
Theoretical Insights Into Multiclass Classification: A High-dimensional Asymptotic View

Christos Thrampoulidis
UC, Santa Barbara
cthrampo@ucsb.edu

Samet Oymak
UC, Riverside
oymak@ece.ucr.edu

Mahdi Soltanolkotabi
University of Southern California
soltanol@usc.edu

Abstract

Contemporary machine learning applications often involve classification tasks with many classes. Despite their extensive use, a precise understanding of the statistical properties and behavior of classification algorithms is still missing, especially in modern regimes where the number of classes is rather large. In this paper, we take a step in this direction by providing the first asymptotically precise analysis of linear multiclass classification. Our theoretical analysis allows us to precisely characterize how the test error varies over different training algorithms, data distributions, problem dimensions as well as number of classes, inter/intra class correlations and class priors. Specifically, our analysis reveals that the classification accuracy is highly distribution-dependent with different algorithms achieving optimal performance for different data distributions and/or training/features sizes. Unlike linear regression/binary classification, the test error in multiclass classification relies on intricate functions of the trained model (e.g., correlation between some of the trained weights) whose asymptotic behavior is difficult to characterize. This challenge is already present in simple classifiers, such as those minimizing a square loss. Our novel theoretical techniques allow us to overcome some of these challenges. The insights gained may pave the way for a precise understanding of other classification algorithms beyond those studied in this paper.

1 Introduction

Multiclass classification is fundamental to a large number of real-world machine learning applications that demand the ability to automatically distinguish between thousands of different classes. Applications include essentially any problem with categorical outputs spanning natural language processing [SVL14], where a seq2seq decoder has to choose the correct word token, reinforcement learning [JGP16, MXSS20], where the agent has to choose the correct action, to recommendation systems, where the model should recommend the correct movie out of many other options. For instance, YouTube’s recommendation system is modeled as an extreme multiclass problem with more than a million classes where each video corresponds to a viable class [CAS16].

The growing list of applications motivate an in-depth exploration of multiclass classification algorithms. Despite their extensive use however, a precise understanding of the statistical properties and behavior of classification algorithms is still missing with many open questions: *What is the total and per class test accuracy? How does this quantity depend on various problem parameters such as data distributions, problem dimensions, etc.? What is the highest test accuracy achievable by any algorithm? What is the best algorithm for each scenario? Which algorithm achieves the highest accuracy on rare or minority classes? How does the answer to the above question change in modern regimes where the number of classes is large?*

Asymptotic analysis in modern high-dimensional regimes where the number of training data and feature sizes grow in tandem with each other provides a promising setting for precisely quantifying the accuracy of classification algorithms as a function of problem variables and resolving the questions above. However, despite the rich literature on precise high-dimensional estimation and more recently

binary classification, multiclass classification is an under-explored venue possibly due to the difficulty of capturing the intricate dependencies between the classes even for relatively simple linear classifiers.

Contributions. We initiate a precise asymptotic study of linear multiclass classification in the modern high-dimensional regime, where the sizes of the training data and of the feature vectors grow large at a proportional rate. A key promise of such a precise analysis is that it allows us to accurately compare between different classification algorithms and data models. Compared to linear regression/binary classification, we identify the following crucial challenge: *the test accuracy in multiclass classification relies on intricate cross-correlations between the trained weights of the classifier*. This has two consequences that drive our analysis. First, in order to obtain sharp asymptotics on the test error of any classifier, it is a prerequisite to precisely quantify the asymptotics of these cross-correlations. Second, the test error does not depend on the correlations in closed-form expressions. Thus, to compare between different classifiers, we need efficient numerical and analytic means to evaluate the test error in terms of the correlation matrices. Interestingly, we show that these challenges are already present in simple classifiers, such as minimizing the square loss, and in stylized distributional settings, such as Gaussian features. Our contributions are as follows:

- We study two different data models: a Gaussian Mixtures Model (GMM) and a Multinomial Logit Model (MLM) with Gaussian features. For each one of them, we provide a precise characterization of total and class-wise test accuracy for three different training algorithms: (i) a least-squares (LS) based classifier, (ii) a weighted least-squares (WLS) based classifier, and (iii) a simple per class averaging (Avg) estimator. For the least-squares based classifiers, we develop a new technique to overcome the technical challenge of characterizing the limiting behavior of the weights' cross-correlations. For the per class averaging classifier, we show that it is Bayes optimal for a GMM with equal priors.
- We discuss efficient means of evaluating the test accuracy as a function of the weights' cross-correlations. This, together with the derived asymptotic formulae for the latter, lead to the first precise high-dimensional characterization of how the total/class-wise accuracy varies for different algorithms, data distributions, problem dimensions as well as number of classes, the inter/intra class correlations and class priors. For special problem geometries, we derive precise conditions on the data distribution and on the relative size of the training set over which each of the two studied algorithms dominates.
- We present and discuss numerical simulations that corroborate our theoretical findings. For instance, with an eye towards making classification algorithms more fair/equitable, we use our precise characterization of the class-wise accuracy to demonstrate how different algorithms behave in the presence of rare/minority classes. We also empirically compare the algorithms studied in this paper to other popular losses such as cross-entropy minimization. This allows us to better understand the performance of various algorithms in modern regimes of large number of classes.

Related Work. There is a classical body of algorithmic work on multiclass classification, e.g., [CS01, LLW04, WW98, BB99, DB94] and several empirical studies of their comparative performance [RK04, Für02, ASS00, PM05]. A more recent extension of this line of work investigates the effect of the loss function in deep neural networks [HYS16, GCOZ17, KS18, BEH20]. Algorithms for extreme multiclass problems with huge number of classes has also been studied in several [CAL13, YHR⁺16, DCO20, RCY⁺19, KMS15] works. On the theory front, numerous works have investigated consistency [Zha04, LLW04, TB07, PSG13, PS16] and finite-sample behavior [KP⁺02, Gue02, ASS00, LLY⁺18, CKMY16, LDBK15, Mau16, LDZK19] of multiclass classification algorithms. Our work differs from this literature in that we are interested in *precise* characterizations of the test accuracy rather than order-wise bounds. Here we focus on linear classifiers, but we consider the modern high-dimensional regime in which both the sample size and the features' dimension are large.

Specifically, our theoretical approach to linear multiclass classification fits in the rapidly growing literature on *sharp* high-dimensional asymptotics of convex optimization-based estimators [Don06, Sto09, OH10, CRPW12, ALMT13, DMM11, BM12, ALMT13, Sto13, OTH13, TOH15, Kar13, EK18, DM16, ORS17, TXH18, TAH18, MM18, WWM19, CM19, HL19, BKRS19, ASH19, JSH20]. Most of this line of work studies linear models and regression problems. More recently there has been a surge of interest in sharp analysis of a variety of methods tailored to binary classification models [TAH15, Hua17, CS18, SC19, MLC19b, MLC19a, KA19, SAH19, TPT20b, DKT19, MRSY19, LS20, MKLZ20, Lol20, TPT20a]. Nevertheless, none of these prior works have yet considered multiclass classification settings. Our paper unveils the salient features of the multiclass setting and shows that corresponding results from the binary setting do not directly apply here. We emphasize that this is the case even for seemingly simple one-vs-all (OVA) classifiers, such as

minimizing the square-loss, that involve training a single binary classifier per class [RK04]. The key technical tool behind our sharp analysis is the convex Gaussian min-max Theorem (CGMT) [TOH15, Sto13]. However, a “naive” application of the CGMT on the original optimization of the classifier does not allow us to compute all the necessary correlations between the classifier’s weights to precisely capture the total/class-wise errors. Instead, our key idea is to formulate an artificial optimization problem, which captures the missing correlations and at the same time conveniently allows us to leverage the CGMT.

Notation. We use $[k]$ to denote $\{1, \dots, k\}$. We use boldface lowercase letters $\mathbf{x}, \mathbf{y}, \boldsymbol{\mu}, \dots$ to denote vectors and boldface uppercase letters $\mathbf{X}, \mathbf{Y}, \mathbf{M}, \dots$ for matrices. We write \mathbf{e}_ℓ for the ℓ -th standard basis vector in \mathbb{R}^k . We also write $\mathbf{I}_k, \mathbf{0}_{k \times k}$ and $\mathbf{1}_k$ for the $k \times k$ identity and all-zeros matrices and the $k \times 1$ all-ones vectors. For a vector $\mathbf{c} \in \mathbb{R}^k$ we write $\arg \max \mathbf{c}$ to denote the index of its largest entry, i.e., $\arg \max \mathbf{c} = \arg \max_{i \in [k]} c_i$. The superscript \dagger denotes pseudoinverse. We use $Q(x)$ for the tail of a standard Gaussian (Q-function). Finally, we reserve variables $G_0, G_1, \dots, G_k \stackrel{iid}{\sim} \mathcal{N}(0, 1)$ to denote i.i.d. standard Gaussians.

2 Problem formulation

We focus on multiclass classification problems with k classes. Specifically, we assume the training data consists of n feature/label pairs $\{(\mathbf{x}_i, Y_i)\}_{i=1}^n$ with $\mathbf{x}_i \in \mathbb{R}^d$ representing the features and $Y_i \in \{1, 2, \dots, k\}$ the associated labels representing one of k classes. It will be convenient to also model the labels as one-hot encoded vectors $\mathbf{y}_i \in \mathbb{R}^k$ representing one of k classes with one-hot encoding, i.e., $\mathbf{y}_i = \mathbf{e}_{Y_i}$. Therefore, when convenient we shall use $\{(\mathbf{x}_i, \mathbf{y}_i)\}_{i=1}^n$ to represent the training data. Throughout, we shall use $\mathbf{X} = [\mathbf{x}_1 \ \mathbf{x}_2 \ \dots \ \mathbf{x}_n] \in \mathbb{R}^{d \times n}$, and $\mathbf{Y} = [\mathbf{y}_1 \ \mathbf{y}_2 \ \dots \ \mathbf{y}_n] \in \mathbb{R}^{k \times n}$, to denote the matrix of features and their labels aggregated into a matrix, respectively. We shall also use $\mathbf{Y}_\ell \in \mathbb{R}^n$ to denote the ℓ -th row of \mathbf{Y} . In our analysis we focus on training linear classifiers. Specifically, we use $\mathbf{W} = [\mathbf{w}_1 \ \mathbf{w}_2 \ \dots \ \mathbf{w}_k]^T \in \mathbb{R}^{k \times d}$ and $\mathbf{b} \in \mathbb{R}^k$ to denote the weights and biases of this linear model, respectively. The overall input-output relationship of the classifier in this case is a function that maps an input vector $\mathbf{x} \in \mathbb{R}^d$ into an output of size k via $\mathbf{x} \mapsto \mathbf{W}\mathbf{x} + \mathbf{b} \in \mathbb{R}^k$, where a training algorithm is used to train the corresponding weights $\mathbf{W} \in \mathbb{R}^{k \times d}$ and biases $\mathbf{b} \in \mathbb{R}^k$. Next we detail the data models and training algorithms that are formally studied in this paper. We end this section by discussing how the test error can be calculated for the different data models.

2.1 Data Models

In our theoretical analysis we assume the training data $\{(\mathbf{x}_i, Y_i)\}_{i=1}^n$ (alternatively $\{(\mathbf{x}_i, \mathbf{y}_i)\}_{i=1}^n$) are generated i.i.d. according to $(\mathbf{x}, Y) / (\mathbf{x}, \mathbf{y})$. We consider two models for the distribution of (\mathbf{x}, \mathbf{y}) which we detail next. In both models we shall use mean/regressor vectors $\{\boldsymbol{\mu}_\ell\}_{\ell=1}^k \in \mathbb{R}^d$ and aggregate them into columns of a matrix of the form $\mathbf{M} := [\boldsymbol{\mu}_1 \ \boldsymbol{\mu}_2 \ \dots \ \boldsymbol{\mu}_k] \in \mathbb{R}^{d \times k}$. In the first model, these vectors represent the mean of the features conditioned on the class, i.e., $\boldsymbol{\mu}_\ell = \mathbb{E}[\mathbf{x} | Y = \ell]$, whereas in the second model these vectors can be viewed as regressor coefficients. We shall refer to $\{\boldsymbol{\mu}_\ell\}_{\ell=1}^k / \mathbf{M}$ as “mean” vectors/matrix in both models. We denote the Grammian matrix of means as $\boldsymbol{\Sigma}_{\boldsymbol{\mu}, \boldsymbol{\mu}} = \mathbf{M}^T \mathbf{M}$. Furthermore, we shall use $\mu_\ell := \|\boldsymbol{\mu}_\ell\|_{\ell_2}$ to denote the norm of the mean vector $\boldsymbol{\mu}_\ell$.

Gaussian Mixture Model (GMM). In this model each example (\mathbf{x}, Y) belongs to class $\ell \in [k]$ with probability π_ℓ , i.e., $\mathbb{P}\{Y = \ell\} = \pi_\ell$. We let $\boldsymbol{\pi} = [\pi_1 \ \pi_2 \ \dots \ \pi_k]^T \in \mathbb{R}^k$ denote the vector of priors which of course obeys $\boldsymbol{\pi} \geq \mathbf{0}$ and $\mathbf{1}^T \boldsymbol{\pi} = 1$. Also, we model the class conditional density of an example in class ℓ with an isotropic Gaussian centered at a mean vector $\boldsymbol{\mu}_\ell$. In particular, we say that a data point (\mathbf{x}, Y) (or its one-hot encoded representation (\mathbf{x}, \mathbf{y})) follows the GMM model when

$$\mathbb{P}\{Y = \ell\} = \pi_\ell \quad \text{and} \quad \mathbf{x} = \boldsymbol{\mu}_Y + \mathbf{z}, \quad \mathbf{z} \sim \mathcal{N}(\mathbf{0}, \sigma^2 \mathbf{I}_d). \quad (2.1)$$

We note that for a training set summarized by the feature and label matrices \mathbf{X} and \mathbf{Y} with columns generated i.i.d. according to the above distribution we have: $\mathbf{X} = \mathbf{M}\mathbf{Y} + \mathbf{Z}$ where $\mathbf{Z} \in \mathbb{R}^{d \times n}$ is a Gaussian noise matrix with i.i.d. $\mathcal{N}(0, \sigma^2)$ entries.

Multinomial Logit Model (MLM). In this model we assume that feature vectors \mathbf{x} are distributed i.i.d. $\mathcal{N}(\mathbf{0}, \mathbf{I}_d)$ and that the conditional density of the class labels is given by the soft-max function. Concretely, we say that a data point (\mathbf{x}, Y) (or its one-hot encoded representation (\mathbf{x}, \mathbf{y})) follows the multinomial logit model when

$$\mathbf{x} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_d) \quad \text{and} \quad \mathbb{P}\{Y = \ell | \mathbf{x}\} = e^{\langle \boldsymbol{\mu}_\ell, \mathbf{x} \rangle} / \sum_{j \in [k]} e^{\langle \boldsymbol{\mu}_j, \mathbf{x} \rangle}. \quad (2.2)$$

2.2 Classification algorithms

As mentioned earlier, in this paper we focus on training linear classifiers of the form $\mathbf{x} \mapsto \mathbf{W}\mathbf{x} + \mathbf{b}$ with $\mathbf{W} \in \mathbb{R}^{k \times d}$ denoting the weights and $\mathbf{b} \in \mathbb{R}^k$ the offset values.

Least-squares (LS). In this approach we train a linear classifier $\mathbf{x} \mapsto \mathbf{W}\mathbf{x} + \mathbf{b}$ via a least-squares fit to the training data: $(\widehat{\mathbf{W}}, \widehat{\mathbf{b}}) := \arg \min_{\mathbf{W}, \mathbf{b}} \frac{1}{2n} \sum_{i=1}^n \|\mathbf{W}\mathbf{x}_i + \mathbf{b} - \mathbf{y}_i\|_{\ell_2}^2 = \frac{1}{2n} \|\mathbf{W}\mathbf{X} + \mathbf{b}\mathbf{1}_n^T - \mathbf{Y}\|_F^2$.

Class averaging (Avg). This approach uses the following weight and offset values $\widehat{\mathbf{W}} := \frac{1}{n} \mathbf{Y}\mathbf{X}^T$ and $\widehat{\mathbf{b}} := \frac{1}{n} \mathbf{Y}\mathbf{1}$. Let n_ℓ be the number of training data from class ℓ then, equivalently, $\widehat{\mathbf{w}}_\ell = \frac{n_\ell}{n} \left(\frac{1}{n_\ell} \sum_{i: Y_i=\ell} \mathbf{x}_i \right)$ and $\widehat{b}_\ell = \frac{n_\ell}{n}$. Therefore, this classifier picks weights according to the empirical mean of features of each class multiplied by the relative frequency of that class and the offset value as the fraction of data points from that class. We note that this algorithm has the same classification performance as the outcome of the ridge-regularized least-squares with infinite regularization.

Weighted Least-squares (WLS). This is a variation of the Least-squares approach where we fit a weighted least squares loss of the form $(\widehat{\mathbf{W}}, \widehat{\mathbf{b}}) := \arg \min_{\mathbf{W}, \mathbf{b}} \frac{1}{2n} \|(\mathbf{W}\mathbf{X} + \mathbf{b}\mathbf{1}_n^T - \mathbf{Y}) \mathbf{D}\|_F^2$. Here, $\mathbf{D} \in \mathbb{R}^{n \times n}$ is a diagonal matrix with the i th diagonal entry equal to $D_{ii} = \omega_\ell$ when the i -th data point is from class ℓ (i.e. $Y_i = \ell$) and $\omega_\ell \geq 0$, $\ell \in [k]$ denote the weights. Aggregating the weights into a vector of the form $\boldsymbol{\omega} = [\omega_1 \ \omega_2 \ \dots \ \omega_k]^T \in \mathbb{R}^k$ we can rewrite \mathbf{D} in the form $\mathbf{D} = \text{diag}(\mathbf{Y}^T \boldsymbol{\omega})$. In this approach the loss associated to data points to class ℓ is weighted by a factor ω_ℓ^2 . For instance, if the class priors are known, a natural choice might be $\omega_\ell = 1/\sqrt{\pi_\ell}$. Such a weighted approach allows the classification algorithm to focus on rare/minority classes which are not well represented in the training data.

Cross-entropy (CE). In this approach the best weight/offset values are determined by fitting a cross entropy loss $(\widehat{\mathbf{W}}, \widehat{\mathbf{b}}) := \arg \min_{\mathbf{W}, \mathbf{b}} \frac{1}{n} \sum_{i=1}^n \log \left(\frac{\sum_{\ell=1}^k e^{\langle \widehat{\mathbf{w}}_\ell, \mathbf{x}_i \rangle + b_\ell}}{e^{\langle \widehat{\mathbf{w}}_{Y_i}, \mathbf{x}_i \rangle + b_{Y_i}}} \right)$. Theoretical analysis for CE is substantially more involved and we defer it to future work. Nevertheless, we compare with this classifier in our numerical simulations.

2.3 Class-wise and total test classification error

Let $\widehat{\mathbf{W}}, \widehat{\mathbf{b}}$ denote the parameters of a trained classifier. Now consider a fresh data sample (\mathbf{x}, Y) generated according to the same distribution as the training data. Once, we have learned the parameters $\widehat{\mathbf{W}}, \widehat{\mathbf{b}}$ of the classifier, the class \widehat{Y} predicted by the classifier is made by a winner takes it all strategy, as follows, $\widehat{Y} = \arg \max_{j \in [k]} \langle \widehat{\mathbf{w}}_j, \mathbf{x} \rangle + \widehat{b}_j$. Therefore, the classification error condition on the true label being c , which we shall refer to as the *class-wise test error*, is equal to

$$\mathbb{P}_{e|c} := \mathbb{P} \{ \widehat{Y} \neq Y | Y = c \} = \mathbb{P} \{ \langle \widehat{\mathbf{w}}_c, \mathbf{x} \rangle + \widehat{b}_c \leq \max_{j \neq c} \langle \widehat{\mathbf{w}}_j, \mathbf{x} \rangle + \widehat{b}_j \}. \quad (2.3)$$

Correspondingly, the *total classification error* is given by

$$\mathbb{P}_e := \mathbb{P} \{ \widehat{Y} \neq Y \} = \mathbb{P} \left\{ \arg \max_{j \in [k]} \{ \langle \widehat{\mathbf{w}}_j, \mathbf{x} \rangle + \widehat{b}_j \} \neq Y \right\} = \mathbb{P} \{ \langle \widehat{\mathbf{w}}_Y, \mathbf{x} \rangle + \widehat{b}_Y \leq \max_{j \neq Y} \langle \widehat{\mathbf{w}}_j, \mathbf{x} \rangle + \widehat{b}_j \}. \quad (2.4)$$

For both the GMM and MLM, the classification error depends on the vector of intercepts $\widehat{\mathbf{b}} \in \mathbb{R}^k$ and the following key “correlation” matrices: $\Sigma_{\mathbf{w}, \mathbf{w}} := \widehat{\mathbf{W}}\widehat{\mathbf{W}}^T$ and $\Sigma_{\mathbf{w}, \boldsymbol{\mu}} := \widehat{\mathbf{W}}\mathbf{M}$.

GMM. In model (2.1), the test error probability is explicitly given by

$$\mathbb{P}_e = \mathbb{P} \left\{ \arg \max \left(\sigma \mathbf{g} + \widehat{\mathbf{b}} + \Sigma_{\mathbf{w}, \boldsymbol{\mu}} \mathbf{e}_Y \right) \neq Y \right\}, \quad \text{where } \mathbf{g} \sim \mathcal{N}(\mathbf{0}, \Sigma_{\mathbf{w}, \mathbf{w}}), \quad (2.5)$$

and Y is independent of \mathbf{g} with probability mass function $\mathbb{P}\{Y = \ell\} = \pi_\ell$, $\ell \in [k]$.

MLM. In model (2.2), the test error probability is explicitly given by

$$\mathbb{P}_e = \mathbb{P} \left\{ \arg \max \left(\mathbf{g} + \widehat{\mathbf{b}} \right) \neq Y(\mathbf{h}) \right\}, \quad \text{where } \begin{bmatrix} \mathbf{g} \\ \mathbf{h} \end{bmatrix} \sim \mathcal{N} \left(\mathbf{0}, \begin{bmatrix} \Sigma_{\mathbf{w}, \mathbf{w}} & \Sigma_{\mathbf{w}, \boldsymbol{\mu}} \\ \Sigma_{\mathbf{w}, \boldsymbol{\mu}}^T & \Sigma_{\boldsymbol{\mu}, \boldsymbol{\mu}} \end{bmatrix} \right), \quad (2.6)$$

and $\mathbb{P}\{Y(\mathbf{h}) = \ell\} = e^{\mathbf{h}_\ell} / \sum_{j \in [k]} e^{\mathbf{h}_j}$, $\ell \in [k]$.

Calculating the class-wise/total misclassification errors. The identities (2.5) and (2.6) (see Section D.1 for a proof) as well as similar ones for the class-wise test error demonstrate that the total/class-wise errors only depend on the correlation matrices $\Sigma_{w,w}$ and $\Sigma_{w,\mu}$, the offset values $\widehat{\mathbf{b}}$ and the the class conditional means. For instance, as we show in the supplementary for GMM the class-wise errors are given by

$$\mathbb{P}_{e|c} = 1 - \mathbb{P}\{\mathbf{S}_c^{1/2} \mathbf{z} \geq \mathbf{t}_c\}, \quad (2.7)$$

where \mathbf{z} is a Gaussian random vector distributed as $\mathcal{N}(\mathbf{0}, \sigma^2 \mathbf{I}_{k-1})$, $\mathbf{S}_c \in \mathbb{R}^{(k-1) \times (k-1)}$ is a symmetric matrix such that its i,j element is given by $[\mathbf{S}_c]_{ij} := \langle \widehat{\mathbf{w}}_c - \widehat{\mathbf{w}}_j, \widehat{\mathbf{w}}_c - \widehat{\mathbf{w}}_i \rangle$ and $\mathbf{t}_c \in \mathbb{R}^{k-1}$ a vector with entries $[\mathbf{t}_c]_i := \langle \widehat{\mathbf{w}}_i - \widehat{\mathbf{w}}_c, \mu_c \rangle + \langle \widehat{\mathbf{b}}_i - \widehat{\mathbf{b}}_c \rangle$. Similarly, based on (2.7) the total classification error in GMM is equal to $\mathbb{P}_e = \sum_{\ell=1}^k \pi_\ell \mathbb{P}_{e|c} = 1 - \sum_{\ell=1}^k \pi_\ell \mathbb{P}\{\mathbf{S}_c^{1/2} \mathbf{z} \geq \mathbf{t}_c\}$. As also detailed in the supplementary, the class-wise/total test errors for MLM similarly depends on quantities of the form $\mathbb{P}\{\mathbf{A}\mathbf{z} \geq \mathbf{t}\}$ with \mathbf{z} a standard Gaussian random vector, \mathbf{A} and \mathbf{t} depending only on correlation matrices, conditional means and classifier offset-values; see Section D.3. There are a variety of algorithmic approaches to calculate $\mathbb{P}\{\mathbf{A}\mathbf{z} \geq \mathbf{t}\}$ once \mathbf{A} and \mathbf{t} are known based on Monte Carlo methods. Analytic bounds on this quantity have also been studied in the literature, e.g., [HH03, SL80]; see more details in Section D.

2.4 High-dimensional regime

This paper derives sharp asymptotic formulae for the class-wise and total classification error of averaging and (weighted) LS algorithms for GMM and MLM. We defer all our proofs to the appendix. All our results hold in the following high-dimensional regime with finite k .

Assumption 1 *We focus on a double asymptotic regime where $n, d \rightarrow \infty$ at a fixed ratio $\gamma = d/n > 0$.*

For the (weighted) least-squares classifier, we focus here in the overdetermined regime $\gamma < 1$. However, our approach is also directly applicable to regularized (or min-norm) LS/WLS in the overparameterized regime $\gamma > 1$.

For a sequence of random variables $\mathcal{X}_{n,d}$ that converges in probability to some constant c in the limit above, we simply write $\mathcal{X}_{n,d} \xrightarrow{P} c$. For a random vector/matrix $\mathbf{v}_{n,d}/\mathbf{V}_{n,d}$ and a deterministic vector/matrix \mathbf{c}/\mathbf{C} , the expressions $\mathbf{v}_{n,d} \xrightarrow{P} \mathbf{c}$ and $\mathbf{V}_{n,d} \xrightarrow{P} \mathbf{C}$ are to be understood entry-wise.

3 Results for Gaussian Mixture Model

In this section we discuss the asymptotics of the intercepts/correlation matrices for the averaging and the LS classifiers for the GMM. The derived formulas can be directly plugged in (2.5) and (2.7) to obtain asymptotics for the total and class-wise test error, respectively. We end this section by also characterizing the Bayes optimal estimator in this model when priors are balanced $\pi_\ell = 1/k, \ell \in [k]$. Additional results on the performance of Weighted LS are deferred to the appendix.

3.1 Class averaging classifier

Proposition 3.1 *Consider data generated according to GMM in an asymptotic regime with any $\gamma > 0$. For the averaging estimator discussed in Section 2.2, the following high-dimensional limits hold*

$$\widehat{\mathbf{b}} \xrightarrow{P} \boldsymbol{\pi}, \quad \Sigma_{w\mu} \xrightarrow{P} \text{diag}(\boldsymbol{\pi}) \cdot \Sigma_{\mu,\mu}, \quad (3.1a)$$

$$\Sigma_{w,w} \xrightarrow{P} \gamma \sigma^2 \cdot \text{diag}(\boldsymbol{\pi}) + \text{diag}(\boldsymbol{\pi}) \cdot \Sigma_{\mu,\mu} \cdot \text{diag}(\boldsymbol{\pi}). \quad (3.1b)$$

The above result allows us to precisely characterize the behavior of the averaging estimator in the high-dimensional regime. Let us consider a few special cases.

Two classes. Consider the special case with two classes with class priors $\pi_1 = 1 - \pi_2 =: \pi$. In this case we can compute the class-wise misclassification probabilities $\mathbb{P}_{e|1}$ and $\mathbb{P}_{e|2}$ explicitly. Specifically using (3.1), we have $\mathbf{S}_1 = \|\pi \boldsymbol{\mu}_1 - (1 - \pi) \boldsymbol{\mu}_2\|_{\ell_2}^2 + \gamma \sigma^2$ and $\mathbf{t}_1 = (1 - 2\pi) + (1 - \pi) \langle \boldsymbol{\mu}_1, \boldsymbol{\mu}_2 \rangle - \pi \|\boldsymbol{\mu}_1\|_{\ell_2}^2$.

Substituting the latter two in (2.7) we arrive at $\mathbb{P}_{e|1} \xrightarrow{P} Q\left(\frac{\pi\|\boldsymbol{\mu}_1\|_{\ell_2}^2 - (1-\pi)(\boldsymbol{\mu}_1, \boldsymbol{\mu}_2) + 2\pi - 1}{\sqrt{\pi\mu_1 - (1-\pi)\mu_2\|_{\ell_2}^2 + \gamma\sigma^2}}\right)$. In the case of equal priors $\pi = \pi_1 = \pi_2 = 1/2$, antipodal and equal energy of the means, i.e., $\boldsymbol{\mu}_1 = -\boldsymbol{\mu}_2$ and $\mu := \|\boldsymbol{\mu}_1\|_{\ell_2} = \|\boldsymbol{\mu}_2\|_{\ell_2}$, we can use the above to conclude that $\mathbb{P}_{e|1} = \mathbb{P}_{e|2} = \frac{1}{2} \mathbb{P}_e = \frac{1}{2} Q\left(\sqrt{\frac{\mu^2}{\mu^2 + \gamma\sigma^2}}\right)$. This formula recovers the result of [MKLZ20] for this special case. Also, as mentioned in [MKLZ20], the formula matches the Bayes optimal error computed in [LM19] for Gaussian mean vectors. This shows that the class averaging method is Bayes optimal in this very simple setting. In Section 3.3, we generalize this result to multiple classes: we show that the average estimator is (asymptotically) Bayes optimal for balanced classes and equal-energy Gaussian means for any $k \geq 2$.

Orthogonal means, equal priors and equal energy. Next we focus on a special case with orthogonal means $\langle \boldsymbol{\mu}_i, \boldsymbol{\mu}_j \rangle = 0$, $i \neq j \in [k]$ of equal energy $\mu^2 := \|\boldsymbol{\mu}_i\|_{\ell_2}^2$ and of equal priors $\pi_i = \pi = 1/k$ for $i \in [k]$. In this case, the class-wise miss-classification error converges to $\mathbb{P}_{e|c} \xrightarrow{P} 1 - \mathbb{P}\{S_c^{1/2} \mathbf{z} > \mathbf{t}\}$, where $\mathbf{S}_c = \pi(\pi\mu^2 + \gamma\sigma^2)(\mathbf{I}_{k-1} + \mathbf{1}_{k-1}\mathbf{1}_{k-1}^T)$ and $\mathbf{t} = -\pi\mu^2\mathbf{1}_{k-1}$. Defining $u_{\text{Avg}} := \frac{\mu^2}{\sigma} \sqrt{\frac{1}{\mu^2 + k\gamma\sigma^2}}$, after some algebraic manipulations the total classification error of the averaging estimator in this case is given by $\mathbb{P}_{e|c} = \mathbb{P}_{e,\text{Avg}} \xrightarrow{P} \mathbb{P}\{G_0 + \max_{j \in [k-1]} G_j \geq u_{\text{Avg}}\}$, where $G_0, \dots, G_{k-1} \stackrel{iid}{\sim} \mathcal{N}(0, 1)$.

3.2 Least-squares classifier

This section focuses on characterizing the intercepts and correlation matrices for the least-squares classifier. To present our results, we assume that the Grammian matrix has eigenvalue decomposition

$$\boldsymbol{\Sigma}_{\boldsymbol{\mu}, \boldsymbol{\mu}} = \mathbf{M}^T \mathbf{M} = \mathbf{V} \boldsymbol{\Sigma}^2 \mathbf{V}^T, \quad \boldsymbol{\Sigma} > \mathbf{0}_{r \times r}, \quad \mathbf{V} \in \mathbb{R}^{k \times r}, \quad r \leq k. \quad (3.2)$$

with $\boldsymbol{\Sigma}$ a diagonal positive-definite matrix and \mathbf{V} an orthonormal matrix obeying $\mathbf{V}^T \mathbf{V} = \mathbf{I}_r$.

Theorem 3.2 *Consider data generated according to GMM in an asymptotic regime with $\gamma < 1$. In addition to (3.2), define the following two positive (semi)-definite matrices: $\mathbf{P} := \text{diag}(\boldsymbol{\pi}) - \boldsymbol{\pi}\boldsymbol{\pi}^T \geq \mathbf{0}_{k \times k}$ and $\boldsymbol{\Delta} := \sigma^2 \mathbf{I}_r + \boldsymbol{\Sigma} \mathbf{V}^T \mathbf{P} \mathbf{V} \boldsymbol{\Sigma} > \mathbf{0}_{r \times r}$. Then, for the least-squares linear classifier $(\widehat{\mathbf{W}}, \widehat{\mathbf{b}})$ the following limits are true asymptotically*

$$\widehat{\mathbf{b}} \xrightarrow{P} \boldsymbol{\pi} - \mathbf{P} \mathbf{V} \boldsymbol{\Sigma} \boldsymbol{\Delta}^{-1} \boldsymbol{\Sigma} \mathbf{V}^T \boldsymbol{\pi}, \quad \boldsymbol{\Sigma}_{\mathbf{w}, \boldsymbol{\mu}} \xrightarrow{P} \mathbf{P} \mathbf{V} \boldsymbol{\Sigma} \boldsymbol{\Delta}^{-1} \boldsymbol{\Sigma} \mathbf{V}^T, \quad (3.3a)$$

$$\boldsymbol{\Sigma}_{\mathbf{w}, \mathbf{w}} \xrightarrow{P} \frac{\gamma}{(1-\gamma)\sigma^2} \mathbf{P} + \mathbf{P} \mathbf{V} \boldsymbol{\Sigma} \boldsymbol{\Delta}^{-1} \left(\boldsymbol{\Delta}^{-1} - \frac{\gamma}{(1-\gamma)\sigma^2} \mathbf{I}_r \right) \boldsymbol{\Sigma} \mathbf{V}^T \mathbf{P}. \quad (3.3b)$$

The above result allows us to precisely characterize the behavior of the least-squares classifier in the high-dimensional regime. In Section H.2, we specialize (3.3) to the case of orthogonal means. Compared to the weight vectors $\widehat{\mathbf{w}}_i, i \in [k]$ of the class averaging classifier that are also (asymptotically) orthogonal when means are orthogonal, this is *not* the case for LS. We show next that these spurious correlations only hurt the classification error when classes are balanced.

Proposition 3.3 *Consider the case of orthogonal, equal energy-means $\boldsymbol{\Sigma}_{\boldsymbol{\mu}, \boldsymbol{\mu}} = \mu \mathbf{I}_k$, balanced priors $\pi_i = 1/k$, $i \in [k]$ and $\gamma < 1$. Setting $u_{\text{LS}} := \frac{\mu^2}{\sigma} \sqrt{\frac{1-\gamma}{\mu^2 + k\gamma\sigma^2}}$, it holds that $\mathbb{P}_{e,\text{LS}} \xrightarrow{P} \mathbb{P}\{G_0 + \max_{j \in [k-1]} G_j \geq u_{\text{LS}}\}$. Specifically, since $u_{\text{LS}} = u_{\text{Avg}} \sqrt{1-\gamma} < u_{\text{Avg}}$, the averaging estimator strictly outperforms LS for all $0 < \gamma < 1$ and $k \geq 2$ in this setting.*

3.3 Bayes estimator for the balanced Gaussian Mixture Model

To check how far the above algorithms are from the lowest misclassification error achievable by any algorithm in this section, we consider a Bayesian setting with Gaussian mean vectors and we derive the Bayes-optimal risk for the case of equal priors. Recall that the Bayes estimator $\hat{Y} = \arg \max_{\ell \in [k]} \mathbb{P}\{Y = \ell \mid \mathbf{X}, \mathbf{Y}, \mathbf{x}\}$ minimizes the risk $\mathbb{P}_e = \mathbb{P}\{\hat{Y} \neq Y\} = \mathbb{E}_{\mathbf{X}, \mathbf{Y}, \mathbf{x}, Y} [\mathbb{1}[\hat{Y} \neq Y]]$.

Proposition 3.4 *Consider $\boldsymbol{\mu}_i \stackrel{iid}{\sim} \mathcal{N}(\mathbf{0}, \frac{\mu^2}{d} \mathbf{I}_d)$ and $\pi_i = 1/k$ for all $i \in [k]$. Set $u_{\text{Bayes}} := \frac{\mu^2}{\sigma} \frac{1}{\sqrt{\mu^2 + k\gamma\sigma^2}}$. Then, the Bayes risk converges to $\mathbb{P}\{G_0 + \max_{\ell \in [k-1]} G_\ell \geq u_{\text{Bayes}}\}$.*

Under Gaussian prior, the means are asymptotically orthogonal and equal-energy. As shown earlier, in this setting, $\mathbb{P}_{e,\text{Avg}} \xrightarrow{P} \mathbb{P}\{G_0 + \max_{\ell \in [k-1]} G_\ell \geq u_{\text{Avg}}\}$. But, $u_{\text{Avg}} = u_{\text{Bayes}}$. Thus, the averaging method is (asymptotically) Bayes optimal for equal-norm, orthogonal means and balanced classes. An analogous result was' derived in [LM19, MKLZ20], but only for binary classification.

4 Results for Multinomial Logit Model

In this section we discuss the asymptotics of the intercepts/correlation matrices for MLM. We present results for arbitrary mean-vectors as well as special cases where the means are mutually orthogonal. Recall the eigenvalue decomposition of the Grammian $\Sigma_{\mu,\mu} = \mathbf{V}\Sigma^2\mathbf{V}^T$ in (3.2). In order to state our results, it is convenient to introduce the following probability vectors in \mathbb{R}^k and \mathbb{R}^{k^2} :

$$\boldsymbol{\pi} := \mathbb{E}\left[\frac{e^{\mathbf{V}\Sigma\mathbf{g}}}{\mathbf{1}_k^T e^{\mathbf{V}\Sigma\mathbf{g}}}\right] \in \mathbb{R}^k \quad \text{and} \quad \boldsymbol{\Pi} := \mathbb{E}\left[\frac{(e^{\mathbf{V}\Sigma\mathbf{g}})(e^{\mathbf{V}\Sigma\mathbf{g}})^T}{(\mathbf{1}_k^T e^{\mathbf{V}\Sigma\mathbf{g}})^2}\right] \in \mathbb{R}^{k \times k}, \quad \text{where } \mathbf{g} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_r). \quad (4.1)$$

Note that $\boldsymbol{\pi}$ and $\boldsymbol{\Pi}$ are the first and second moments of the soft-max mapping of $\mathbf{V}\Sigma\mathbf{g} \sim \mathcal{N}(\mathbf{0}, \Sigma_{\mu,\mu})$. In fact, for the MLM in (2.2) it holds that $\mathbb{P}\{Y = \ell\} = \mathbb{E}[\mathbb{P}\{Y = \ell | \mathbf{x}\}] = \mathbb{E}\left[\frac{e^{\mathbf{e}_\ell^T \mathbf{V}\Sigma\mathbf{g}}}{\mathbf{1}_k^T e^{\mathbf{V}\Sigma\mathbf{g}}}\right] = \boldsymbol{\pi}_\ell$, $\ell \in [k]$ since $\mathbf{M}^T \mathbf{x}$ is distributed as $\mathbf{V}\Sigma\mathbf{g}$. Thus, $\boldsymbol{\pi}$ is the vector of class priors (which explains the slight abuse of notation here in relation to our notation for the class priors of the GMM).

4.1 Class averaging classifier

Proposition 4.1 *Consider data generated according to MLM in an asymptotic regime with any $\gamma > 0$. For the averaging classifier, the following high-dimensional limits hold*

$$\hat{\mathbf{b}} \xrightarrow{P} \boldsymbol{\pi}, \quad \Sigma_{w,\mu} \xrightarrow{P} (\text{diag}(\boldsymbol{\pi}) - \boldsymbol{\Pi}) \cdot \Sigma_{\mu,\mu}, \quad (4.2a)$$

$$\Sigma_{w,w} \xrightarrow{P} \gamma \cdot \text{diag}(\boldsymbol{\pi}) + (\text{diag}(\boldsymbol{\pi}) - \boldsymbol{\Pi}) \cdot \Sigma_{\mu,\mu} (\text{diag}(\boldsymbol{\pi}) - \boldsymbol{\Pi}). \quad (4.2b)$$

Using Gaussian decomposition in (2.6) and checking from (4.2) that $\Sigma_{w,w} - \Sigma_{w,\mu} \Sigma_{\mu,\mu}^\dagger \Sigma_{w,\mu}^T \xrightarrow{P} \gamma \cdot \text{diag}(\boldsymbol{\pi})$ the test error obtains the following explicit form:

$$\mathbb{P}_{e,\text{Avg}} \xrightarrow{P} \mathbb{P}\left\{\arg \max \left\{\sqrt{\gamma} \cdot \text{diag}(\sqrt{\boldsymbol{\pi}}) \cdot \tilde{\mathbf{g}} + (\text{diag}(\boldsymbol{\pi}) - \boldsymbol{\Pi}) \cdot \mathbf{V}\Sigma \cdot \mathbf{g} + \boldsymbol{\pi}\right\} \neq Y(\mathbf{g})\right\}, \quad (4.3)$$

where $\tilde{\mathbf{g}} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_k)$, $\mathbf{g} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_r)$ and $\mathbb{P}\{Y(\mathbf{g}) = c\} = e^{\mathbf{e}_c^T \mathbf{V}\Sigma\mathbf{g}} / \sum_{j \in [k]} e^{\mathbf{e}_j^T \mathbf{V}\Sigma\mathbf{g}}$, $c \in [k]$.

4.2 Least-squares classifier

This section focuses on characterizing the intercepts and correlation matrices for the least-squares classifier. We also use the result to characterize conditions under which LS outperforms averaging.

Theorem 4.2 *Consider data generated according to MLM in an asymptotic regime with $0 < \gamma < 1$. Recall the notation in (4.1). For the LS classifier, the following high-dimensional limits hold.*

$$\hat{\mathbf{b}} \xrightarrow{P} \boldsymbol{\pi}, \quad \Sigma_{w,\mu} \xrightarrow{P} (\text{diag}(\boldsymbol{\pi}) - \boldsymbol{\Pi}) \cdot \Sigma_{\mu,\mu}, \quad (4.4a)$$

$$\Sigma_{w,w} \xrightarrow{P} \frac{\gamma}{1-\gamma} \cdot (\text{diag}(\boldsymbol{\pi}) - \boldsymbol{\pi}\boldsymbol{\pi}^T) + \frac{1-2\gamma}{1-\gamma} \cdot (\text{diag}(\boldsymbol{\pi}) - \boldsymbol{\Pi}) \cdot \Sigma_{\mu,\mu} \cdot (\text{diag}(\boldsymbol{\pi}) - \boldsymbol{\Pi}). \quad (4.4b)$$

It is interesting to observe that (4.4a) is identical to (4.2a). However, the cross-correlations in $\Sigma_{w,w}$ differ. We prove below that this leads to an improved performance of the LS classifier for large sample sizes. First, Theorem 4.2 can be used to check that

$$\Sigma_{w,w} - \Sigma_{w,\mu} \Sigma_{\mu,\mu}^\dagger \Sigma_{w,\mu}^T \xrightarrow{P} \frac{\gamma}{1-\gamma} \left(\text{diag}(\boldsymbol{\pi}) - \boldsymbol{\pi}\boldsymbol{\pi}^T - (\text{diag}(\boldsymbol{\pi}) - \boldsymbol{\Pi}) \Sigma_{\mu,\mu} (\text{diag}(\boldsymbol{\pi}) - \boldsymbol{\Pi}) \right).$$

Thus, the only change in the test-error formula compared to (4.3) is the term $\gamma \cdot \text{diag}(\boldsymbol{\pi})$ substituted by the matrix above.

Proposition 4.3 Assume orthogonal, equal-energy means $\Sigma_{\mu,\mu} = \mu^2 \mathbf{I}_k$, $k \geq 2$. Let $\gamma_* = \frac{\mu^2 k}{(k-1)^2} \left(1 - k \mathbb{E} \left[\frac{e^{2\mu G_1}}{(\sum_{\ell \in [k]} e^{\mu G_\ell})^2} \right] \right)^2 \in (0, 1)$. Then, with probability 1 as $n \rightarrow \infty$, $\mathbb{P}_{e,LS} < \mathbb{P}_{e,Avg} \iff \gamma < \gamma_*$.

5 Numerical Results

This section validates our theory via numerical experiments and provides further insights on multiclass classification. See also Section A for more extensive experiments. We study the class-wise/total test misclassification error in both GMM and MLM for different sample sizes, number of classes and class priors. In line with Section 2.2 we consider four algorithms: (i) Averaging (Avg), (ii) LS, (iii) Weighted LS (WLS) with the i th class weighted by $\omega_\ell^2 = 1/\pi_\ell$, (iv) Cross-Entropy (CE).

Figures 1 and 2 focus on GMM with $k = 9$ classes, $d = 300$ and $\|\mu_i\|_{\ell_2}^2 = 15$. To model different class prior probabilities, we use the distribution $\pi_1 = \pi_2 = \pi_3 = 0.5, \pi_4 = 0.5, \pi_5 = 0.5, \pi_6 = 0.25, \pi_7 = 0.25, \pi_8 = 0.25, \pi_9 = 1/21$. We consider three scenarios: (a) orthogonal means, equal prior ($\pi_i = 1/9$); (b) orthogonal means, different prior; (c) correlated means with pairwise correlation coefficient equal to 0.5 (i.e., $\langle \mu_i, \mu_j \rangle / (\|\mu_i\|_{\ell_2} \|\mu_j\|_{\ell_2}) = 0.5$ for $i \neq j$) and different priors as discussed above. Figure 1 shows the test miss-classification errors as a function of $\gamma := d/n$. In all scenarios our theoretical predictions are a near perfect match to the empirical performance. In scenario (a), class-wise averaging achieves the lowest error as predicted by Proposition 3.4. However, in scenario (b) where the means have different norms the averaging method has higher misclassification error compared with CE, LS and WLS for large sample sizes (small γ). We note that both LS and WLS achieve lower errors compared with CE as the sample size grows. Scenario (c) is similar to (b). However, due to class correlations, the errors are uniformly higher. Figure 2 shows the corresponding class-wise miss-classification errors for the smallest γ in Figure 1 ($\gamma = 0.117$). In scenario (a), errors are equal which is expected given the equal class priors. In scenarios (b) and (c) however, due to different priors, large classes 7,8,9 achieve best accuracy. The performance difference is most visible for the averaging approach. LS mitigates this issue to some extent, while WLS creates the flattest class-wise errors suggesting that it can reduce the miss-classification error on small/minority classes.

Figure 3 focuses on orthogonal classes with varying number of classes k where $\|\mu_i\|_{\ell_2}^2 = 15$ and $d \in \{50, 100, 200\}$ with $kd/n = k\gamma$ fixed at $k\gamma = 20/11$. It plots the ratio of the empirical error probability and our theoretical prediction as k grows until $k = d$. Two observations are worth mentioning here. (1) The accuracy of our predictions noticeably improves as the problem dimension d, n grow as expected given the asymptotic nature of our analysis. Interestingly, the convergence appears to be noticeably faster (as a function of d) for the LS rather than the Averaging classifier. (2) Our theoretical results formally require that k is fixed while d (and n) grow large. Yet, the presented experimental results suggest that they might also hold for large k under the shown scaling. This is a fascinating research question that we believe is worth investigating further.

Figure 4 provides experiments on MLM with $k = 9$ orthogonal classes. Unlike GMM, CE achieves the best performance in MLM. In Figure 4 (a), classes have same norms $\|\mu_i\|_{\ell_2} = 10$, while in Figure 4 (b) we have quadrupled the norms of classes 7,8,9 and doubled the norms of classes 4,5,6. This disparity between the norms seems to help improve the CE accuracy, but hurt LS/averaging accuracy for small γ . Finally, Figure 4 (c) shows the class-wise probability of error associated with (b) for $\gamma = 0.117$ and demonstrates that LS outperforms averaging.

6 Future Directions

This work aims at initiating a precise asymptotic study of multiclass classifiers that provides a promising setting for resolving a rich set of open questions regarding the (comparative) performance of classification algorithms as a function of the involved problem variables. As mentioned, even understanding the statistical performance of one-vs-all multiclass classifiers does not follow directly from the existing literature on binary classifiers. Extending the results of this paper to the one-vs-all logistic and SVM classifiers would allow for a principled comparison among these different choices. A possibly more challenging, albeit mathematically intriguing and practically relevant task, is characterizing the asymptotics of more complicated (non-separable) losses, such as the cross-entropy loss. For this, even characterizing the asymptotic behavior of the correlations $\Sigma_{w,\mu}$

requires new ideas. The previously mentioned study of “extreme multiclass classification” in which the number of classes k is very large is another fascinating direction.

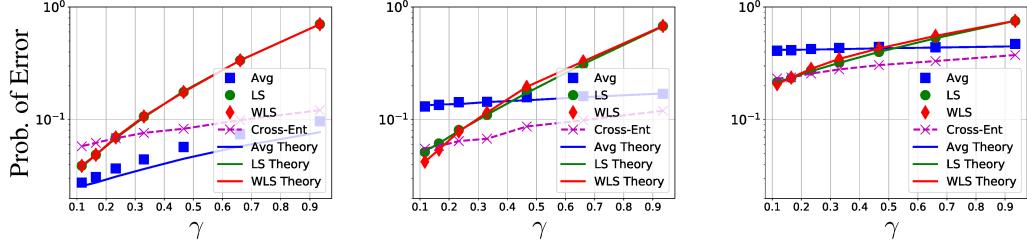


Figure 1: GMM with $k = 9, d = 300$. (a) orthogonal, equal prior, (b) orthogonal, different prior, (c) correlated, different prior.

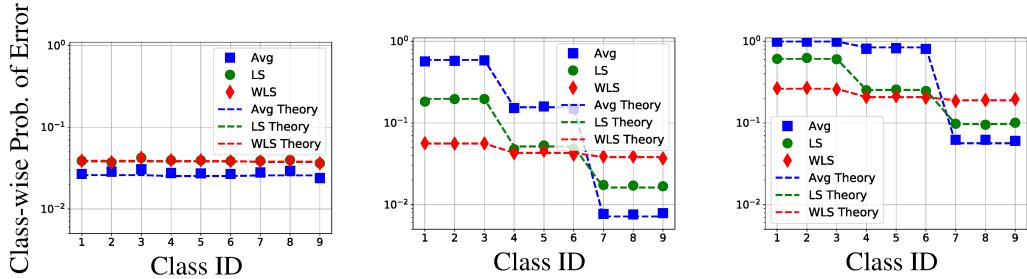


Figure 2: Class-wise probability of errors corresponding to Figure 1 with $\gamma = 0.117$.

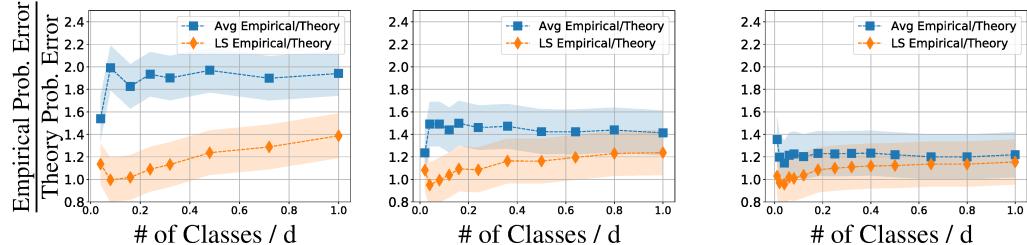


Figure 3: GMM, K/d is varied from 0 to 1 while keeping $K\gamma$ constant for (a) $d = 50$, (b) $d = 100$, (c) $d = 200$.

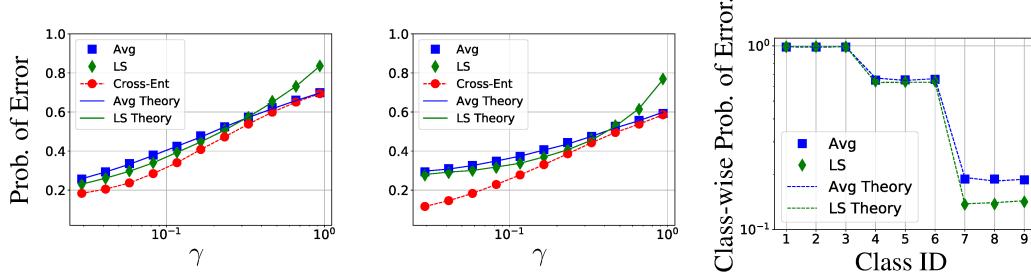


Figure 4: MLM with orthogonal means for (a) equal norms and (b) different norms. (c) Class-wise probability errors for (b).

Acknowledgments

C. Thrampoulidis was partially supported by the NSF under Grant Numbers CCF-2009030 and HDR-1934641. S. Oymak is partially supported by the NSF award CNS-1932254. M. Soltanolkotabi is supported by the Packard Fellowship in Science and Engineering, a Sloan Research Fellowship in Mathematics, an NSF-CAREER under award #1846369, the Air Force Office of Scientific Research Young Investigator Program (AFOSR-YIP) under award #FA9550-18-1-0078, DARPA Learning with Less Labels (LwLL) and FastNICS programs, and NSF-CIF awards #1813877 and #2008443.

Broader Impact

In this paper we develop a precise and asymptotically exact understanding of the statistical behavior of a variety of classification algorithms. In particular we precisely characterize how the total and class-wise accuracy varies under different training algorithms, data distributions, problem dimensions, inter/intra class correlations and class priors. Despite being theoretical/foundational in nature it has potential for broader practical impact. In particular, our precise characterization of class-wise accuracy allows us to understand how different training algorithms impact accuracy of machine learning algorithms on rare/minority classes. Such a precise understanding may help guide the development of more fair/equitable algorithms. On the flip side, such insights may potentially also be used nefariously enabling the marginalization of rare/minority classes by developing algorithms that reduce their class-wise accuracy.

References

- [AG82] Per Kragh Andersen and Richard D Gill. Cox’s regression model for counting processes: a large sample study. *The annals of statistics*, pages 1100–1120, 1982.
- [ALMT13] Dennis Amelunxen, Martin Lotz, Michael B McCoy, and Joel A Tropp. Living on the edge: A geometric theory of phase transitions in convex optimization. *arXiv preprint arXiv:1303.6672*, 2013.
- [ASH19] Ehsan Abbasi, Fariborz Salehi, and Babak Hassibi. Universality in learning from linear measurements. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d’Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems 32*, pages 12372–12382. Curran Associates, Inc., 2019.
- [ASS00] Erin L Allwein, Robert E Schapire, and Yoram Singer. Reducing multiclass to binary: A unifying approach for margin classifiers. *Journal of machine learning research*, 1(Dec):113–141, 2000.
- [BB99] Erin J Bredensteiner and Kristin P Bennett. Multicategory classification by support vector machines. In *Computational Optimization*, pages 53–79. Springer, 1999.
- [BEH20] Anna Sergeevna Bosman, Andries Engelbrecht, and Mardé Helbig. Visualising basins of attraction for the cross-entropy and the squared error neural network loss functions. *Neurocomputing*, 2020.
- [BKRS19] Zhiqi Bu, Jason Klusowski, Cynthia Rush, and Weijie Su. Algorithmic analysis and statistical estimation of slope via approximate message passing. In *Advances in Neural Information Processing Systems*, pages 9361–9371, 2019.
- [BM12] Mohsen Bayati and Andrea Montanari. The lasso risk for gaussian matrices. *Information Theory, IEEE Transactions on*, 58(4):1997–2017, 2012.
- [CAL13] Anna Choromanska, Alekh Agarwal, and John Langford. Extreme multi class classification. In *NIPS Workshop: eXtreme Classification*, submitted, 2013.
- [CAS16] Paul Covington, Jay Adams, and Emre Sargin. Deep neural networks for youtube recommendations. In *Proceedings of the 10th ACM conference on recommender systems*, pages 191–198, 2016.

[CKMY16] Corinna Cortes, Vitaly Kuznetsov, Mehryar Mohri, and Scott Yang. Structured prediction theory based on factor graph complexity. In *Advances in Neural Information Processing Systems*, pages 2514–2522, 2016.

[CM19] Michael Celentano and Andrea Montanari. Fundamental barriers to high-dimensional regression with convex penalties. *arXiv preprint arXiv:1903.10603*, 2019.

[CRPW12] Venkat Chandrasekaran, Benjamin Recht, Pablo A Parrilo, and Alan S Willsky. The convex geometry of linear inverse problems. *Foundations of Computational Mathematics*, 12(6):805–849, 2012.

[CS01] Koby Crammer and Yoram Singer. On the algorithmic implementation of multiclass kernel-based vector machines. *Journal of machine learning research*, 2(Dec):265–292, 2001.

[CS18] Emmanuel J Candès and Pragya Sur. The phase transition for the existence of the maximum likelihood estimate in high-dimensional logistic regression. *arXiv preprint arXiv:1804.09753*, 2018.

[DB94] Thomas G Dietterich and Ghulum Bakiri. Solving multiclass learning problems via error-correcting output codes. *Journal of artificial intelligence research*, 2:263–286, 1994.

[DCO20] Ahmet Demirkaya, Jiasi Chen, and Samet Oymak. Exploring the role of loss functions in multiclass classification. In *2020 54th Annual Conference on Information Sciences and Systems (CISS)*, pages 1–5. IEEE, 2020.

[DKT19] Zeyu Deng, Abla Kammoun, and Christos Thrampoulidis. A model of double descent for high-dimensional binary linear classification. *arXiv preprint arXiv:1911.05822*, 2019.

[DM16] David Donoho and Andrea Montanari. High dimensional robust m-estimation: Asymptotic variance via approximate message passing. *Probability Theory and Related Fields*, 166(3-4):935–969, 2016.

[DMM11] David L Donoho, Arian Maleki, and Andrea Montanari. The noise-sensitivity phase transition in compressed sensing. *Information Theory, IEEE Transactions on*, 57(10):6920–6941, 2011.

[Don06] David L Donoho. Compressed sensing. *Information Theory, IEEE Transactions on*, 52(4):1289–1306, 2006.

[EK18] Noureddine El Karoui. On the impact of predictor geometry on the performance on high-dimensional ridge-regularized generalized robust regression estimators. *Probability Theory and Related Fields*, 170(1-2):95–175, 2018.

[FR13] Simon Foucart and Holger Rauhut. An invitation to compressive sensing. In *A mathematical introduction to compressive sensing*, pages 1–39. Springer, 2013.

[Für02] Johannes Fürnkranz. Round robin classification. *Journal of Machine Learning Research*, 2(Mar):721–747, 2002.

[GCOZ17] Krzysztof Gajowniczek, Leszek J Chmielewski, Arkadiusz Orłowski, and Tomasz Ząbkowski. Generalized entropy cost function in neural networks. In *International Conference on Artificial Neural Networks*, pages 128–136. Springer, 2017.

[Gor88] Yehoram Gordon. *On Milman’s inequality and random subspaces which escape through a mesh in \mathbb{R}^n* . Springer, 1988.

[Gue02] Yann Guermeur. Combining discriminant models with new multi-class svms. *Pattern Analysis & Applications*, 5(2):168–179, 2002.

[HH03] Enkelejd Hashorva and Jürg Hüsler. On multivariate gaussian tails. *Annals of the Institute of Statistical Mathematics*, 55(3):507–522, 2003.

- [HL19] Hong Hu and Yue M Lu. Asymptotics and optimal designs of slope for sparse linear regression. *arXiv preprint arXiv:1903.11582*, 2019.
- [Hua17] Hanwen Huang. Asymptotic behavior of support vector machine for spiked population model. *The Journal of Machine Learning Research*, 18(1):1472–1492, 2017.
- [HYS16] Le Hou, Chen-Ping Yu, and Dimitris Samaras. Squared earth mover’s distance-based loss for training deep neural networks. *arXiv preprint arXiv:1611.05916*, 2016.
- [JGP16] Eric Jang, Shixiang Gu, and Ben Poole. Categorical reparameterization with gumbel-softmax. *arXiv preprint arXiv:1611.01144*, 2016.
- [JSH20] Adel Javanmard, Mahdi Soltanolkotabi, and Hamed Hassani. Precise tradeoffs in adversarial training for linear regression. *arXiv preprint arXiv:2002.10477*, 2020.
- [KA19] A. Kammoun and M.-S. Alouini. On the precise error analysis of support vector machines. *Submitted to IEEE Transactions on information theory*, 2019.
- [Kar13] Noureddine El Karoui. Asymptotic behavior of unregularized and ridge-regularized high-dimensional robust regression estimators: rigorous results. *arXiv preprint arXiv:1311.2445*, 2013.
- [KMS15] Vitaly Kuznetsov, Mehryar Mohri, and U Syed. Rademacher complexity margin bounds for learning with a large number of classes. In *ICML Workshop on Extreme Classification: Learning with a Very Large Number of Labels*, 2015.
- [KP⁺02] Vladimir Koltchinskii, Dmitry Panchenko, et al. Empirical margin distributions and bounding the generalization error of combined classifiers. *The Annals of Statistics*, 30(1):1–50, 2002.
- [KS18] Himanshu Kumar and PS Sastry. Robust loss functions for learning multi-class classifiers. In *2018 IEEE International Conference on Systems, Man, and Cybernetics (SMC)*, pages 687–692. IEEE, 2018.
- [LDBK15] Yunwen Lei, Urun Dogan, Alexander Binder, and Marius Kloft. Multi-class svms: From tighter data-dependent generalization bounds to novel algorithms. In *Advances in Neural Information Processing Systems*, pages 2035–2043, 2015.
- [LDZK19] Yunwen Lei, Ürün Dogan, Ding-Xuan Zhou, and Marius Kloft. Data-dependent generalization bounds for multi-class classification. *IEEE Transactions on Information Theory*, 65(5):2995–3021, 2019.
- [LLW04] Yoonkyung Lee, Yi Lin, and Grace Wahba. Multicategory support vector machines: Theory and application to the classification of microarray data and satellite radiance data. *Journal of the American Statistical Association*, 99(465):67–81, 2004.
- [LLY⁺18] Jian Li, Yong Liu, Rong Yin, Hua Zhang, Lizhong Ding, and Weiping Wang. Multi-class learning: from theory to algorithm. In *Advances in Neural Information Processing Systems*, pages 1586–1595, 2018.
- [LM19] Marc Lelarge and Leo Miolane. Asymptotic bayes risk for gaussian mixture in a semi-supervised setting. *arXiv preprint arXiv:1907.03792*, 2019.
- [Lol20] Panagiotis Lolas. Regularization in high-dimensional regression and classification via random matrix theory. *arXiv preprint arXiv:2003.13723*, 2020.
- [LS20] Tengyuan Liang and Pragya Sur. A precise high-dimensional asymptotic theory for boosting and min- l_1 -norm interpolated classifiers. *arXiv preprint arXiv:2002.01586*, 2020.
- [LT91] Michel Ledoux and Michel Talagrand. *Probability in Banach Spaces: isoperimetry and processes*, volume 23. Springer, 1991.
- [Mau16] Andreas Maurer. A vector-contraction inequality for rademacher complexities. In *International Conference on Algorithmic Learning Theory*, pages 3–17. Springer, 2016.

[MKLZ20] Francesca Mignacco, Florent Krzakala, Yue M Lu, and Lenka Zdeborová. The role of regularization in classification of high-dimensional noisy gaussian mixture. *arXiv preprint arXiv:2002.11544*, 2020.

[MLC19a] X. Mai, Z. Liao, and R. Couillet. A large scale analysis of logistic regression: asymptotic performance and new insights. In *ICASSP*, 2019.

[MLC19b] Xiaoyi Mai, Zhenyu Liao, and Romain Couillet. A large scale analysis of logistic regression: Asymptotic performance and new insights. In *ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 3357–3361. IEEE, 2019.

[MM18] Léo Miolane and Andrea Montanari. The distribution of the lasso: Uniform control over sparse balls and adaptive parameter tuning. *arXiv preprint arXiv:1811.01212*, 2018.

[MRSY19] Andrea Montanari, Feng Ruan, Youngtak Sohn, and Jun Yan. The generalization error of max-margin linear classifiers: High-dimensional asymptotics in the overparametrized regime. *arXiv preprint arXiv:1911.01544*, 2019.

[MXSS20] Jincheng Mei, Chenjun Xiao, Csaba Szepesvari, and Dale Schuurmans. On the global convergence rates of softmax policy gradient methods. *arXiv preprint arXiv:2005.06392*, 2020.

[OH10] Samet Oymak and Babak Hassibi. New null space results and recovery thresholds for matrix rank minimization. *arXiv preprint arXiv:1011.6326*, 2010.

[ORS17] Samet Oymak, Benjamin Recht, and Mahdi Soltanolkotabi. Sharp time–data tradeoffs for linear inverse problems. *IEEE Transactions on Information Theory*, 64(6):4129–4158, 2017.

[OTH13] Samet Oymak, Christos Thrampoulidis, and Babak Hassibi. The squared-error of generalized lasso: A precise analysis. *arXiv preprint arXiv:1311.0830*, 2013.

[PM05] Mahesh Pal and PM Mather. Support vector machines for classification in remote sensing. *International journal of remote sensing*, 26(5):1007–1011, 2005.

[PS16] Bernardo Ávila Pires and Csaba Szepesvári. Multiclass classification calibration functions. *arXiv preprint arXiv:1609.06385*, 2016.

[PSG13] Bernardo Avila Pires, Csaba Szepesvari, and Mohammad Ghavamzadeh. Cost-sensitive multiclass classification risk bounds. In *International Conference on Machine Learning*, pages 1391–1399, 2013.

[RCY⁺19] Ankit Singh Rawat, Jiecao Chen, Felix Xinnan X Yu, Ananda Theertha Suresh, and Sanjiv Kumar. Sampled softmax with random fourier features. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d’Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems 32*, pages 13834–13844. Curran Associates, Inc., 2019.

[RK04] Ryan Rifkin and Aldebaro Klautau. In defense of one-vs-all classification. *Journal of machine learning research*, 5(Jan):101–141, 2004.

[SAH19] Fariborz Salehi, Ehsan Abbasi, and Babak Hassibi. The impact of regularization on high-dimensional logistic regression. *arXiv preprint arXiv:1906.03761*, 2019.

[SC19] Pragya Sur and Emmanuel J Candès. A modern maximum-likelihood theory for high-dimensional logistic regression. *Proceedings of the National Academy of Sciences*, 116(29):14516–14525, 2019.

[SL80] YS Sathe and SR Lingras. A note on the inequalities for tail probability of the multivariate normal distribution. *Communications in Statistics-Theory and Methods*, 9(7):711–715, 1980.

[Sto09] Mihailo Stojnic. Various thresholds for ℓ_1 -optimization in compressed sensing. *arXiv preprint arXiv:0907.3666*, 2009.

- [Sto13] Mihailo Stojnic. A framework to characterize performance of lasso algorithms. *arXiv preprint arXiv:1303.7291*, 2013.
- [SVL14] Ilya Sutskever, Oriol Vinyals, and Quoc V Le. Sequence to sequence learning with neural networks. In *Advances in neural information processing systems*, pages 3104–3112, 2014.
- [TAH15] Christos Thrampoulidis, Ehsan Abbasi, and Babak Hassibi. Lasso with non-linear measurements is equivalent to one with linear measurements. In *Advances in Neural Information Processing Systems*, pages 3420–3428, 2015.
- [TAH18] Christos Thrampoulidis, Ehsan Abbasi, and Babak Hassibi. Precise error analysis of regularized m -estimators in high dimensions. *IEEE Transactions on Information Theory*, 64(8):5592–5628, 2018.
- [TB07] Ambuj Tewari and Peter L Bartlett. On the consistency of multiclass classification methods. *Journal of Machine Learning Research*, 8(May):1007–1025, 2007.
- [TOH15] Christos Thrampoulidis, Samet Oymak, and Babak Hassibi. Regularized linear regression: A precise analysis of the estimation error. In *Proceedings of The 28th Conference on Learning Theory*, pages 1683–1709, 2015.
- [TPT20a] Hossein Taheri, Ramtin Pedarsani, and Christos Thrampoulidis. Fundamental limits of ridge-regularized empirical risk minimization in high dimensions. *arXiv preprint arXiv:2006.08917*, 2020.
- [TPT20b] Hossein Taheri, Ramtin Pedarsani, and Christos Thrampoulidis. Sharp asymptotics and optimal performance for inference in binary models. *arXiv preprint arXiv:2002.07284*, 2020.
- [TXH18] Christos Thrampoulidis, Weiyu Xu, and Babak Hassibi. Symbol error rate performance of box-relaxation decoders in massive mimo. *IEEE Transactions on Signal Processing*, 66(13):3377–3392, 2018.
- [WBW⁺16] Alan Wisler, Visar Berisha, Dennis Wei, Karthikeyan Ramamurthy, and Andreas Spanias. Empirically-estimable multi-class classification bounds. In *2016 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 2594–2598. IEEE, 2016.
- [WW98] Jason Weston and Chris Watkins. Multi-class support vector machines. Technical report, Citeseer, 1998.
- [WWM19] Shuaiwen Wang, Haolei Weng, and Arian Maleki. Does slope outperform bridge regression? *arXiv preprint arXiv:1909.09345*, 2019.
- [YHR⁺16] Ian En-Hsu Yen, Xiangru Huang, Pradeep Ravikumar, Kai Zhong, and Inderjit Dhillon. Pd-sparse: A primal and dual sparse approach to extreme multiclass and multilabel classification. In *International Conference on Machine Learning*, pages 3069–3077, 2016.
- [Zha04] Tong Zhang. Statistical analysis of some multi-category large margin classification methods. *Journal of Machine Learning Research*, 5(Oct):1225–1251, 2004.

Contents

A Additional Numerical Results	17
B Additional Results on Weighted Least-squares classifiers	18
B.1 WLS for GMM	18
B.2 WLS for MLM	19
C Preliminaries	19
C.1 Slepian's inequality	19
C.2 Gaussian integration by parts	20
C.3 Block matrix inversion	20
D Calculating and bounding the missclassification error	20
D.1 Proof of (2.5) and (2.6)	20
D.2 Class-wise and total miss-classification error for GMM	21
D.3 Class-wise and total miss-classification error for MLM	21
D.4 Evaluating and bounding tail probabilities of multivariate Gaussians	22
D.4.1 A special case: Rank-one update of Identity	22
D.4.2 Slepian's bound	22
D.4.3 Simple bounds for GMM	23
E The Class-averaging estimator	24
E.1 Proofs for GMM	24
E.1.1 GMM: Proof of Proposition 3.1	24
E.2 Proofs for MLM	24
E.2.1 Proof of Proposition 4.1	24
E.2.2 Orthogonal means	26
F On the Bayes risk of GMM: Proof of Proposition 3.4	26
G Proof outline for least-squares: key ideas and challenges	28
G.1 Background on the CGMT	30
H Least-squares for GMM	31
H.1 Proof of Theorem 3.2	31
H.1.1 Computing $\Sigma_{w,\mu}$	31
H.1.2 Computing $\Sigma_{w,w}$	34
H.2 Orthogonal means	35
I Least-squares for MLM	36
I.1 Proof of Theorem 4.2	36
I.1.1 Computing $\Sigma_{w,\mu}$	36

I.1.2	Computing $\Sigma_{w,w}$	38
I.2	Orthogonal means and equal-energy	38
I.3	Proof of Proposition 4.3	39
J	Weighted LS for GMM (Proof of Theorem B.1)	42
J.1	Computing $\Sigma_{w,\mu}$	42
J.2	Computing $\Sigma_{w,w}$	48
K	Weighted LS for MLM (Proof of Theorem B.2)	52
K.1	Computing $\Sigma_{w,\mu}$	54
K.3	Computing $\Sigma_{w,w}$	57

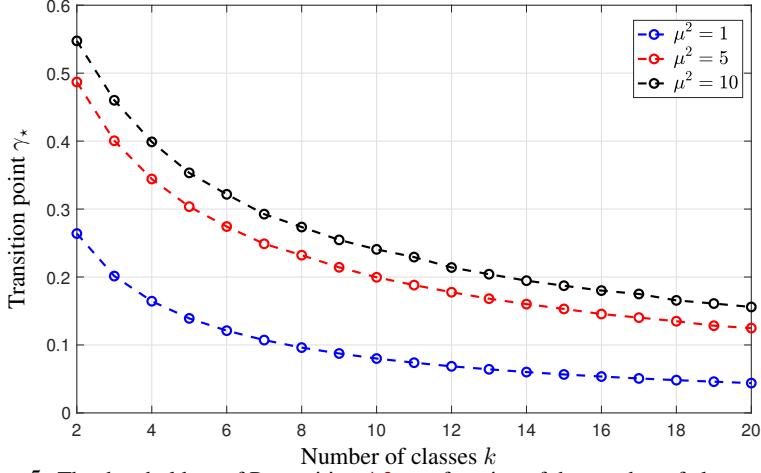


Figure 5: The threshold γ_* of Proposition 4.3 as a function of the number of classes k and the means' energy μ . LS provably outperforms class-averaging for $\gamma < \gamma_*$.

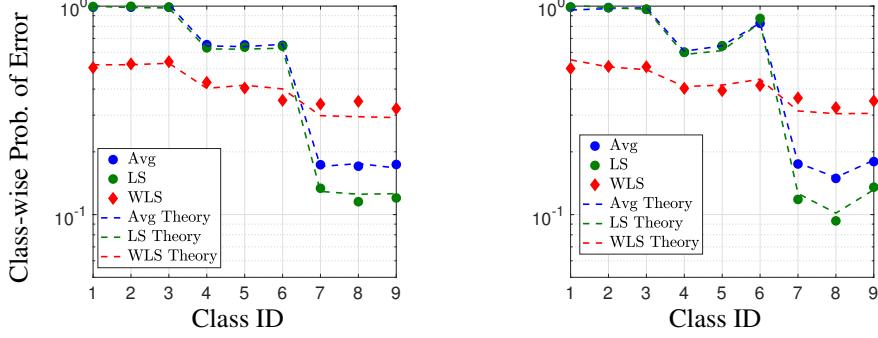


Figure 6: Class-wise probabilities of error for MLM with (a) orthogonal means and (b) correlated means.

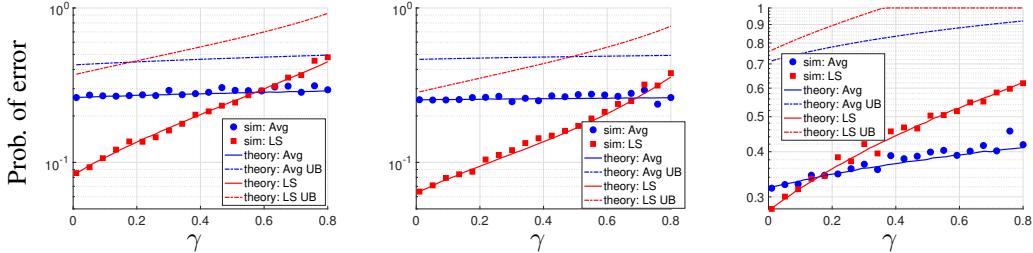


Figure 7: Precise theoretical predictions compared to a theoretical upper bound (UB) obtained by Union bound that does *not* require knowledge of the off-diagonal entries of $\Sigma_{w,w}$. GMM with orthogonal means and (a) equal priors, different strengths; (b) different priors, different strengths; (c) different priors, equal strengths. See text for details.

A Additional Numerical Results

In this section, we provide further numerical experiments.

First, in Figure 5 we investigate the question: *When does least-squares provably outperform averaging?* Our Proposition 4.3 provides a fundamental transition point in sample complexity above which least-squares is provably better than averaging under MLM. In Figure 5, we visualize γ_* as a function

of different number of classes as well as different levels of mean energy. Least-squares outperform averaging in the region below the lines displayed in Figure 5. Our key message is that least-squares work better when the sample complexity is higher and the problem is less noisy. As the number of classes k increase, the problem becomes more difficult/noisy and we require a larger sample complexity to ensure classifier achieves a similar amount of accuracy as small k . Following this intuition, as k increases, γ_* shifts smaller due to larger sample requirement. Similarly energy μ directly controls the noise level of the problem, i.e., larger μ results in a larger signal-to-noise ratio. Thus, as we increase μ , γ_* increases as well because same test accuracy can be achieved with smaller sample size.

Second, Figure 6 provides further experiments on the class-wise probabilities of the MLM model with $k = 9$ classes for $\gamma = 0.1$. Classes 1,2,3 have norms $\|\mu_i\|_{\ell_2} = 15$, while we have quadrupled the norms of classes 7,8,9 and doubled the norms of classes 4,5,6. In scenario (a) the means are orthogonal and in scenario (b) the means are highly correlated. The WLS shown corresponds to the following choice of weights: $\omega_i^2 = 1/\pi_\ell$, $\ell \in [k]$, where π_ℓ is the ℓ th entry of the vector π in (4.1). The theoretical predictions for the class-wise error probabilities are computed using formula (D.7). As was the case for the GMM in Figure 2, we see that WLS creates the flattest class-wise errors.

Finally, in Figure 7 we investigate the following question: *To what extent pairwise class correlations are necessary for performance prediction?* Specifically, we consider a GMM setup with $k = 5$ classes and orthogonal means under three scenarios: (a) $\pi_1 = \pi_2 = \pi_3 = \pi_4 = \pi_5$, $4\|\mu_1\|_{\ell_2} = 4\|\mu_2\|_{\ell_2} = 2\|\mu_3\|_{\ell_2} = 2\|\mu_4\|_{\ell_2} = \|\mu_5\|_{\ell_2} = 4\sqrt{3}$; (b) $4\pi_1 = 4\pi_2 = 2\pi_3 = 2\pi_4 = \pi_5$, $4\|\mu_1\|_{\ell_2} = 4\|\mu_2\|_{\ell_2} = 2\|\mu_3\|_{\ell_2} = 2\|\mu_4\|_{\ell_2} = \|\mu_5\|_{\ell_2} = 4\sqrt{3}$; (c) $4\pi_1 = 4\pi_2 = 2\pi_3 = 2\pi_4 = \pi_5$, $\|\mu_1\|_{\ell_2} = \|\mu_2\|_{\ell_2} = \|\mu_3\|_{\ell_2} = \|\mu_4\|_{\ell_2} = \|\mu_5\|_{\ell_2} = \sqrt{3}$. The solid lines are exact performance predictions based on our theory for averaging and least-squares estimators. The dashed lines are the theoretical upper bounds, which do *not* require the knowledge of cross-correlations between the classes (i.e., off-diagonal entries of $\Sigma_{w,w}$ are unknown). These bounds are calculated by applying a union bound to the class-wise probabilities $\mathbb{P}_{e|c}$, $c \in [k]$ in (D.1) and further appropriately bounding the off-diagonal entries of $\Sigma_{w,w}$ in terms of the self-correlations of the classes, i.e., only the diagonal entries of $\Sigma_{w,w}$. Please see Section D.4.3 for details. Overall, the bounds shown only depend on $\widehat{\mathbf{b}}$, $\Sigma_{w,\mu}$ and $\text{diag}(\Sigma_{w,w})$, which can all be obtained by studying the properties of isolated least-squares on individual classes without understanding their pairwise relations. While this suggests a simpler method to calculate theoretical bounds, there is a visible gap between such upper bounds and exact bounds and this gap is particularly more visible in the third scenario (c), where the bound becomes vacuous for LS. The gap remains visible in scenarios (a) and (b). At least, in these two cases comparing the bounds for averaging and LS to each other reveals the transition in performance gain between the two estimators. However, the cross-point of the curves does not coincide with the true one. This empirical study emphasizes the fact that pairwise correlations are indeed critical for exact asymptotic analysis and naive approaches cannot reproduce, in general, the results of our sharp analysis.

B Additional Results on Weighted Least-squares classifiers

B.1 WLS for GMM

We now focus on characterizing the intercepts/correlation matrices for the WLS classifier.

Theorem B.1 Consider data generated according to GMM and $\gamma < 1$. Consider a weighted LS classifier with weights $\mathbf{D} = \text{diag}(\omega_1, \dots, \omega_k)$ and let η be the unique solution to $\sum_{\ell=1}^k \frac{\pi_\ell \omega_\ell^2}{\omega_\ell^2 + \eta} = \gamma$. Also define $\mathbf{P} := \text{diag}(\widetilde{\pi}) - \widetilde{\pi} \widetilde{\pi}^T \geq \mathbf{0}_{k \times k}$ and $\Delta := \sigma^2 \mathbf{I}_r + \Sigma \mathbf{V}^T \mathbf{P} \Sigma \mathbf{V} > \mathbf{0}_{r \times r}$ with the entries of $\widetilde{\pi}$ given by $\widetilde{\pi}_\ell = \frac{1}{\gamma} \frac{\pi_\ell \omega_\ell^2}{\omega_\ell^2 + \eta}$. Then, for the WLS linear classifier $(\widehat{\mathbf{W}}, \widehat{\mathbf{b}})$ the following asymptotic limits hold

$$\widehat{\mathbf{b}} \xrightarrow{P} \widetilde{\pi} - \mathbf{P} \mathbf{V} \Sigma \Delta^{-1} \Sigma \mathbf{V}^T \widetilde{\pi}, \quad \Sigma_{w,\mu} \xrightarrow{P} \mathbf{P} \mathbf{V} \Sigma \Delta^{-1} \Sigma \mathbf{V}^T, \quad (\text{B.1a})$$

$$\Sigma_{w,w} \xrightarrow{P} \frac{\zeta}{\sigma^2} \mathbf{P} + \mathbf{P} \mathbf{V} \Sigma \Delta^{-1} \left(\Delta^{-1} - \frac{\zeta}{\sigma^2} \mathbf{I}_r \right) \Sigma \mathbf{V}^T \mathbf{P} + \frac{\eta \zeta}{\sigma^2} \mathbf{Q}. \quad (\text{B.1b})$$

Here, $\zeta := \gamma / \left(\eta \sum_{\ell=1}^k \frac{\pi_\ell \omega_\ell^2}{(\omega_\ell^2 + \eta)^2} \right)$ and $\mathbf{Q} \in \mathbb{R}^{k \times k}$ is a known matrix depending on various problem parameters. Its precise value is given in (J.18).

Surprisingly, the effect of the weights is essentially equivalent to adjusting the class priors from π to $\tilde{\pi}$ defined in the theorem (modulo the extra additive term in the cross correlation matrix $\Sigma_{w,w}$). This shows that weighted LS has similar performance to an un-weighted LS applied to a model with different class priors $\tilde{\pi}$. This characterization allows us to precisely understand how different weighting schemes can alter test accuracy for rare/minority classes.

B.2 WLS for MLM

Theorem B.2 predicts the asymptotic performance of *weighted* least-squares for data generated according to MLM.

Theorem B.2 *Consider data generated according to MLM and $\gamma < 1$. Consider a weighted LS classifier with weights $\mathbf{D} = \text{diag}(\omega_1, \dots, \omega_k)$ and let η be the unique solution to $\sum_{\ell=1}^k \frac{\pi_\ell \omega_\ell^2}{\omega_\ell^2 + \eta} = \gamma$. Also define vector $\boldsymbol{\nu} \in \mathbb{R}^k$ with entries given by $\nu_\ell = \frac{1}{\gamma} \frac{\omega_\ell^2}{\omega_\ell^2 + \eta}$ and matrix*

$$\Delta = \mathbb{E}[(\boldsymbol{\nu}^T \mathbf{v}) \mathbf{g} \mathbf{g}^T] - \mathbf{V}^T (\text{diag}(\pi) - \Pi) \boldsymbol{\nu} \boldsymbol{\nu}^T (\text{diag}(\pi) - \Pi) \mathbf{V} \Sigma > \mathbf{0}_{r \times r}, \quad (\text{B.2})$$

where $\mathbf{v} \in \mathbb{R}^k$ is a random vector with entries $V_\ell = e^{\mathbf{e}_\ell^T \mathbf{V} \Sigma \mathbf{g}} / \sum_{\ell' \in [k]} e^{\mathbf{e}_{\ell'}^T \mathbf{V} \Sigma \mathbf{g}}$ for $\mathbf{g} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_r)$. Then, for the WLS linear classifier $(\widehat{\mathbf{W}}, \widehat{\mathbf{b}})$ the following asymptotic limits hold

$$\widehat{\mathbf{b}} \xrightarrow{P} \text{diag}(\boldsymbol{\nu}) \pi - \text{diag}(\boldsymbol{\nu}) (\mathbf{I}_k - \pi \boldsymbol{\nu}^T) (\text{diag}(\pi) - \Pi) \mathbf{V} \Sigma \Delta^{-1} \Sigma \mathbf{V}^T (\text{diag}(\pi) - \Pi) \boldsymbol{\nu}. \quad (\text{B.3a})$$

$$\Sigma_{w,\mu} \xrightarrow{P} \text{diag}(\boldsymbol{\nu}) (\mathbf{I}_k - \pi \boldsymbol{\nu}^T) (\text{diag}(\pi) - \Pi) \mathbf{V} \Sigma \Delta^{-1} \Sigma \mathbf{V}^T. \quad (\text{B.3b})$$

The corresponding formula for the asymptotic limit of the cross-correlation matrix $\Sigma_{w,w}$ is given in (K.38) in Section K.

Of course, the theorem above includes Theorem 4.2 as a special case. Indeed, we show how setting $\omega_\ell = 1$, $\ell \in [k]$ recovers the solution for (un-weighted) LS. First, solving for η simply gives $\eta = \frac{1}{\gamma} - 1$. Thus, $\boldsymbol{\nu} = \mathbf{1}_k$. Also, observe in (4.1) that $(\text{diag}(\pi) - \Pi) \mathbf{1}_k = \mathbf{0}$ and $\mathbf{1}^T \mathbf{v} = 1$. Thus, (B.2) reduces to $\Delta = \mathbb{E}[\mathbf{g} \mathbf{g}^T] = \mathbf{I}_r$. With these, it can be readily checked that (B.3a) and (B.3b) simplify to the expressions in (4.4a).

The term $\mathbb{E}[(\boldsymbol{\nu}^T \mathbf{v}) \mathbf{g} \mathbf{g}^T]$ in (B.2) can be computed using Monte Carlo sampling. It is also possible to slightly simplify the calculations involved using Gaussian integration by parts as shown in Lemma C.3. As mentioned, the formula that predicts $\Sigma_{w,w}$ is given in (K.38). While somewhat more complicated than formulae (B.3), the expression that we provide is also explicit. Numerical simulations shown in Figure 6 in Section A validate the accuracy of the theoretical predictions of the theorem.

C Preliminaries

In this section we gather a few preliminary results that will be used later on in our proofs.

C.1 Slepian's inequality

Lemma C.1 (Slepian's inequality [LT91]) *Let $\mathbf{g} \sim \mathcal{N}(\mathbf{0}, \mathbf{S})$ and $\widetilde{\mathbf{g}} \sim \mathcal{N}(\mathbf{0}, \mathbf{R})$ such that for all $i, j \in [k]$:*

$$S_{ii} = \mathbf{R}_{ii}, \quad \text{and} \quad S_{ij} \geq \mathbf{R}_{ij}.$$

Then, for any $\mathbf{t} \in \mathbb{R}^k$ it holds that

$$\mathbb{P} \left\{ \bigcup_{j \in [k]} \{\mathbf{g}_j \geq \mathbf{t}_j\} \right\} \leq \mathbb{P} \left\{ \bigcup_{j \in [k]} \{\widetilde{\mathbf{g}}_j \geq \mathbf{t}_j\} \right\}.$$

Equivalently, letting $\mathbf{z} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_k)$,

$$1 - \mathbb{P} \{ \mathbf{S}^{1/2} \mathbf{z} \leq \mathbf{t} \} \leq 1 - \mathbb{P} \{ \mathbf{R}^{1/2} \mathbf{z} \leq \mathbf{t} \}.$$

C.2 Gaussian integration by parts

We say that a function $F : \mathbb{R}^m \rightarrow \mathbb{R}$ is of moderate growth if for each $c > 0$,

$$\lim_{\|\mathbf{x}\|_{\ell_2} \rightarrow \infty} F(\mathbf{x}) \exp\left(-c \|\mathbf{x}\|_{\ell_2}^2\right) = 0.$$

The following result is a direct application of Gaussian integration by parts; for instance, see [FR13, Prop. 8.29].

Lemma C.2 (Gaussian integration by parts (GIP)) *Let $\mathbf{g} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_r)$ and function $f : \mathbb{R}^r \rightarrow \mathbb{R}$ such that f and all its first and second order partial derivatives are of moderate growth. Then, the following statements are true:*

- (i) $\mathbb{E}[f(\mathbf{g})\mathbf{g}] = \mathbb{E}[\nabla f(\mathbf{g})]$.
- (ii) $\mathbb{E}[f(\mathbf{g})\mathbf{g}\mathbf{g}^T] = \mathbb{E}[f(\mathbf{g})]\mathbf{I}_r + \mathbb{E}[\nabla^2 f(\mathbf{g})]$.

The following is a corollary of Lemma C.2 applied to the soft-max function.

Lemma C.3 (GIP for the Softmax) *Let $\mathbf{g} \sim \mathcal{N}(\mathbf{0}_r, \mathbf{I}_r)$ and random vector $\mathbf{v} = [V_1, V_2, \dots, V_k]^T$ with entries:*

$$\mathbf{v} = \frac{e^{\mathbf{V}\Sigma\mathbf{g}}}{\mathbf{1}_k^T e^{\mathbf{V}\Sigma\mathbf{g}}}, \quad V_\ell = \frac{e^{\mathbf{e}_\ell^T \mathbf{V}\Sigma\mathbf{g}}}{\sum_{j \in [k]} e^{\mathbf{e}_j^T \mathbf{V}\Sigma\mathbf{g}}}, \quad \ell \in [k]. \quad (\text{C.1})$$

Further recall the notation of $\boldsymbol{\pi}$ and $\boldsymbol{\Pi}$ in (4.1). The following statements are true:

- (i) $\mathbb{E}[\mathbf{v}] = \boldsymbol{\pi}$.
- (ii) For all $i \in [r], \ell \in [k]$,

$$\mathbb{E}[\mathbf{g}_i V_\ell] = (\mathbf{e}_i^T \Sigma \mathbf{V}^T \mathbf{e}_\ell) \boldsymbol{\pi}_\ell - \mathbf{e}_i^T \Sigma \mathbf{V}^T \sum_{j \in [k]} \mathbf{e}_j \boldsymbol{\Pi}_{ij},$$

and in matrix form:

$$\mathbb{E}[\mathbf{v}\mathbf{g}^T] = (\text{diag}(\boldsymbol{\pi}) - \boldsymbol{\Pi}) \mathbf{V}\Sigma.$$

- (iii) For all $\ell \in [k]$ let $s_\ell : \mathbb{R}^k \rightarrow \mathbb{R}$ denote the soft-max function: $s_\ell(\mathbf{x}) = \frac{e^{x_\ell}}{\sum_{i \in [k]} e^{x_i}}$. Then,

$$\mathbb{E}[V_\ell \mathbf{g}\mathbf{g}^T] = \pi_\ell \mathbf{I}_r + \Sigma^T \mathbf{V}^T \mathbb{E}[\nabla^2 s_\ell(\mathbf{V}\Sigma\mathbf{g})] \mathbf{V}\Sigma.$$

C.3 Block matrix inversion

Lemma C.4 (Block matrix inversion) *Let $\mathbf{T} = \begin{bmatrix} \mathbf{A} & \mathbf{b} \\ \mathbf{b}^T & \delta \end{bmatrix}$ be an invertible block matrix. Then*

$$\mathbf{T}^{-1} \begin{bmatrix} \mathbf{f} \\ \epsilon \end{bmatrix} = \begin{bmatrix} \Delta^{-1}(\mathbf{f} - \frac{\epsilon}{\delta} \mathbf{b}) \\ \frac{\epsilon}{\delta} - \frac{1}{\delta} \mathbf{b}^T \Delta^{-1}(\mathbf{f} - \frac{\epsilon}{\delta} \mathbf{b}) \end{bmatrix} \quad (\text{C.2})$$

where $\Delta = \mathbf{A} - \frac{1}{\delta} \mathbf{b}\mathbf{b}^T > \mathbf{0}$ is the Schur complement.

D Calculating and bounding the missclassification error

D.1 Proof of (2.5) and (2.6)

GMM. Starting from (2.4) and using the fact that $\mathbf{x}_i = \boldsymbol{\mu}_Y + \mathbf{z} = \mathbf{M}\mathbf{e}_Y + \mathbf{z}$, $\mathbf{z} \sim \mathcal{N}(\mathbf{0}, \sigma^2 \mathbf{I}_k)$, we have that

$$\mathbb{P}_e = \mathbb{P} \left\{ \arg \max_{j \in [k]} \{ \langle \widehat{\mathbf{w}}_j, \mathbf{M}\mathbf{e}_Y \rangle + \langle \widehat{\mathbf{w}}_j, \mathbf{z} \rangle + \widehat{\mathbf{b}}_j \} \neq Y \right\},$$

or, in matrix-form:

$$\mathbb{P}_e = \mathbb{P} \left\{ \arg \max \{ \widehat{\mathbf{W}} \mathbf{M} \mathbf{e}_Y + \widehat{\mathbf{W}} \mathbf{z} + \widehat{\mathbf{b}} \} \neq Y \right\}.$$

Recall that $\Sigma_{w,\mu} := \widehat{\mathbf{W}} \mathbf{M}$ and note that $\widehat{\mathbf{W}} \mathbf{z}$ is a zero-mean Gaussian vector with covariance matrix $\sigma^2 \widehat{\mathbf{W}} \widehat{\mathbf{W}}^T = \sigma^2 \Sigma_{w,w}$ in order to conclude with the desired formula in (2.5).

MLM. Recall from (2.2) that $\mathbf{x} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_d)$ and Y is distributed such that $\mathbb{P}\{Y = \ell \mid \mathbf{x}\} = e^{\langle \mu_\ell, \mathbf{x} \rangle} / \sum_{j \in [k]} e^{\langle \mu_j, \mathbf{x} \rangle}$. Let $\mathbf{g} = \widehat{\mathbf{W}} \mathbf{x}$ and $\mathbf{h} = \mathbf{M}^T \mathbf{x}$. In this notation, (2.4) becomes

$$\mathbb{P}_e = \mathbb{P} \left\{ \arg \max (\mathbf{g} + \widehat{\mathbf{b}}) \neq Y \right\},$$

with $\mathbb{P}\{Y = \ell \mid \mathbf{x}\} = e^{\mathbf{h}_\ell} / \sum_{j \in [k]} e^{\mathbf{h}_j}$. To complete the proof of (2.6), it is easy to check that $\begin{bmatrix} \mathbf{g} \\ \mathbf{h} \end{bmatrix}$ defined above is jointly Gaussian with zero-mean and covariance matrix $\begin{bmatrix} \Sigma_{w,w} & \Sigma_{w,\mu} \\ \Sigma_{w,\mu}^T & \Sigma_{\mu,\mu} \end{bmatrix}$.

D.2 Class-wise and total miss-classification error for GMM

The class-wise miss-classification error for GMM is given by

$$\mathbb{P}_{e|c} = \mathbb{P} \left(\exists j \neq c : \langle \widehat{\mathbf{w}}_c - \widehat{\mathbf{w}}_j, \mathbf{z} \rangle \leq \langle \widehat{\mathbf{w}}_j - \widehat{\mathbf{w}}_c, \mu_c \rangle + (\widehat{\mathbf{b}}_j - \widehat{\mathbf{b}}_c) \right) \quad (\text{D.1})$$

$$= 1 - \mathbb{P} \left(\forall j \neq c : \langle \widehat{\mathbf{w}}_c - \widehat{\mathbf{w}}_j, \mathbf{z} \rangle \geq \langle \widehat{\mathbf{w}}_j - \widehat{\mathbf{w}}_c, \mu_c \rangle + (\widehat{\mathbf{b}}_j - \widehat{\mathbf{b}}_c) \right), \quad (\text{D.2})$$

where we used that $\mathbf{x} = \mu_c + \mathbf{z}$. Let $\mathbf{S}_c \in \mathbb{R}^{(k-1) \times (k-1)}$ be a symmetric matrix and $\mathbf{t}_c \in \mathbb{R}^{k-1}$ a vector with entries:

$$[\mathbf{t}_c]_j := \langle \widehat{\mathbf{w}}_j - \widehat{\mathbf{w}}_c, \mu_c \rangle + (\widehat{\mathbf{b}}_j - \widehat{\mathbf{b}}_c) \quad j \neq c \in [k] \quad (\text{D.3a})$$

$$[\mathbf{S}_c]_{ij} := \langle \widehat{\mathbf{w}}_c - \widehat{\mathbf{w}}_j, \widehat{\mathbf{w}}_c - \widehat{\mathbf{w}}_i \rangle, \quad i, j \neq c \in [k]. \quad (\text{D.3b})$$

Then, we can rewrite (D.2) as

$$\mathbb{P}_{e|c} := 1 - \mathbb{P} \left\{ \mathbf{S}_c^{1/2} \mathbf{z} \geq \mathbf{t}_c \right\}, \quad (\text{D.4})$$

where the inequality in the rightmost expression applies entry-wise.

Further, by using the law of total probability we have

$$\mathbb{P}_e = \sum_{c=1}^k \pi_c \mathbb{P}_{e|c} = \sum_{c=1}^k \pi_c \left(1 - \mathbb{P} \left\{ \mathbf{S}_c^{1/2} \mathbf{z} \geq \mathbf{t}_c \right\} \right).$$

D.3 Class-wise and total miss-classification error for MLM

In this section, we derive an explicit formula for the class-wise error for MLM. Recall (2.6):

$$\mathbb{P}_e = \mathbb{P} \left\{ \arg \max (\mathbf{g} + \widehat{\mathbf{b}}) \neq Y(\mathbf{h}) \right\}, \quad \text{where } \begin{bmatrix} \mathbf{g} \\ \mathbf{h} \end{bmatrix} \sim \mathcal{N} \left(\mathbf{0}, \begin{bmatrix} \Sigma_{w,w} & \Sigma_{w,\mu} \\ \Sigma_{w,\mu}^T & \Sigma_{\mu,\mu} \end{bmatrix} \right),$$

and $\mathbb{P}\{Y(\mathbf{h}) = \ell\} = e^{\mathbf{h}_\ell} / \sum_{j \in [k]} e^{\mathbf{h}_j}$, $\ell \in [k]$. Using Gaussian decomposition we can write $\mathbf{g} = \widetilde{\mathbf{g}} + \Sigma_{w,\mu} \Sigma_{\mu,\mu}^\dagger \mathbf{h}$ where $\widetilde{\mathbf{g}} \sim \mathcal{N}(\mathbf{0}_k, \Sigma_{w,w} - \Sigma_{w,\mu} \Sigma_{\mu,\mu}^\dagger \Sigma_{w,\mu})$. Using this, we have

$$\begin{aligned} \mathbb{P}_e &= \mathbb{P} \left\{ \arg \max (\widetilde{\mathbf{g}} + \Sigma_{w,\mu} \Sigma_{\mu,\mu}^\dagger \mathbf{h} + \widehat{\mathbf{b}}) \neq Y(\mathbf{h}) \right\} \\ &= \mathbb{P} \left\{ \arg \max (\widetilde{\mathbf{g}} + \Sigma_{w,\mu} \Sigma_{\mu,\mu}^\dagger \mathbf{h} + \widehat{\mathbf{b}}) \neq Y(\mathbf{h}) \right\} \\ &= \sum_{c \in [k]} \mathbb{E}_{\mathbf{h}, \widetilde{\mathbf{g}}} \left[\mathbb{P} \{Y(\mathbf{h}) = c\} \cdot \mathbb{1} \left\{ \arg \max (\widetilde{\mathbf{g}} + \Sigma_{w,\mu} \Sigma_{\mu,\mu}^\dagger \mathbf{h} + \widehat{\mathbf{b}}) \neq c \right\} \right] \\ &= \sum_{c \in [k]} \mathbb{E}_{\mathbf{h}, \widetilde{\mathbf{g}}} \left[\frac{e^{\mathbf{h}_c}}{\sum_{\ell \in [k]} e^{\mathbf{h}_\ell}} \left(1 - \prod_{j \neq c} \mathbb{1} \left\{ \widetilde{\mathbf{g}}_c + [\Sigma_{w,\mu} \Sigma_{\mu,\mu}^\dagger \mathbf{h}]_c + \widehat{\mathbf{b}}_c \geq \widetilde{\mathbf{g}}_j + [\Sigma_{w,\mu} \Sigma_{\mu,\mu}^\dagger \mathbf{h}]_j + \widehat{\mathbf{b}}_j \right\} \right) \right] \\ &= \sum_{c \in [k]} \mathbb{E}_{\mathbf{h}} \left[\frac{e^{\mathbf{h}_c}}{\sum_{\ell \in [k]} e^{\mathbf{h}_\ell}} \left(1 - \mathbb{P}_{\mathbf{z} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_{k-1})} \left\{ \mathbf{S}_c^{1/2} \mathbf{z} \geq \mathbf{t}_c(\mathbf{h}) \right\} \right) \right], \end{aligned} \quad (\text{D.5})$$

where in the last line

$$[\mathbf{t}_c(\mathbf{h})]_j = \widehat{\mathbf{b}}_j - \widehat{\mathbf{b}}_c + [\Sigma_{\mathbf{w}, \mu} \Sigma_{\mu, \mu}^\dagger \mathbf{h}]_j - [\Sigma_{\mathbf{w}, \mu} \Sigma_{\mu, \mu}^\dagger \mathbf{h}]_c, \quad j \neq c \in [k] \quad (\text{D.6a})$$

$$[\mathbf{S}_c]_{i,j} = (\mathbf{e}_c - \mathbf{e}_j)^T (\Sigma_{\mathbf{w}, \mathbf{w}} - \Sigma_{\mathbf{w}, \mu} \Sigma_{\mu, \mu}^\dagger \Sigma_{\mathbf{w}, \mu}) (\mathbf{e}_c - \mathbf{e}_i), \quad i, j \neq c \in [k]. \quad (\text{D.6b})$$

Further recalling the decomposition $\mathbb{P}_e = \sum_c \mathbb{P}\{\widehat{Y} \neq Y | Y = c\} \mathbb{P}\{Y = c\}$ and noting that

$$\mathbb{P}\{Y = c\} = \mathbb{E}_{\mathbf{x}} \left[\mathbb{P}\{Y = c | \mathbf{x}\} \right] = \mathbb{E}_{\mathbf{x}} \left[\frac{1}{(1 + \sum_{j \neq c} e^{(\mu_j - \mu_c)^T \mathbf{x}})} \right] = \pi_c$$

we can see from (D.5) that the class-wise error probabilities can be calculated as follows:

$$\mathbb{P}_{e|c} = \mathbb{P}\{\widehat{Y} \neq Y | Y = c\} = \frac{1}{\pi_c} \mathbb{E}_{\mathbf{h} \sim \mathcal{N}(\mathbf{0}_k, \Sigma_{\mu, \mu})} \left[\frac{e^{\mathbf{h}_c}}{\sum_{\ell \in [k]} e^{\mathbf{h}_\ell}} (1 - \mathbb{P}_{\mathbf{z} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_{k-1})} \{S_c^{1/2} \mathbf{z} \geq \mathbf{t}_c(\mathbf{h})\}) \right], \quad (\text{D.7})$$

where π_c is the c^{th} entry of the vector π in (4.1) and $\mathbf{S}_c, \mathbf{t}_c(\mathbf{h})$ are defined in (D.6).

D.4 Evaluating and bounding tail probabilities of multivariate Gaussians

In Sections D.3 and D.2, we expressed the class-wise probability of missclassification error for both GMM and MLM in the following convenient form for $\mathbf{z} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_{k-1})$,

$$1 - \mathbb{P}\{\mathbf{A}^{1/2} \mathbf{z} \leq \mathbf{t}\} = \mathbb{P}\left\{\bigcup_{i \in [k-1]} \{\tilde{\mathbf{a}}_i^T \mathbf{z} \geq \mathbf{t}_i\}\right\}. \quad (\text{D.8})$$

Here, $\mathbf{A} \geq \mathbf{0} \in \mathbb{R}^{(k-1) \times (k-1)}$, $\mathbf{t} \in \mathbb{R}^{k-1}$ are appropriate coefficient matrices (see (D.4) and (D.5)) and $\tilde{\mathbf{a}}_i$ denotes the i^{th} row of the matrix $\mathbf{A}^{1/2}$. For example, (D.8) maps to (D.4) for $\mathbf{A} \leftarrow \mathbf{S}_c$ and $\mathbf{t} \leftarrow (-\mathbf{t}_c)$.

The formulation above is convenient both in our theoretical analysis, as well as, in simulations. In the rest of this section, we briefly discuss some relevant tools that allow to further simplify or bound expressions in the form of (D.8).

D.4.1 A special case: Rank-one update of Identity

First, we discuss the case where the coefficient matrix \mathbf{A} and vector \mathbf{t} in (D.8) take the special form $\mathbf{A} \propto \mathbf{I} + \mathbf{1}\mathbf{1}^T$ and $\mathbf{t} \propto \mathbf{1}$. This special case appears in some of the stylized symmetric problem settings studied in this paper, such as classification problems with orthogonal and equally-balanced means.

Lemma D.1 *Let $\mathbf{A} = \mathbf{I}_k + \mathbf{1}_k \mathbf{1}_k^T$ and $\mathbf{g} \sim \mathcal{N}(\mathbf{0}, \mathbf{A})$. Then, for any $t \in \mathbb{R}$,*

$$1 - \mathbb{P}\{\mathbf{g} \leq t \mathbf{1}_k\} = \mathbb{P}\{G_0 + \max_{i \in [k]} G_i \geq t\}, \quad G_0, G_1, \dots, G_k \stackrel{\text{iid}}{\sim} \mathcal{N}(0, 1). \quad (\text{D.9})$$

Proof For each $i \in [n]$, we can decompose $\mathbf{g}_i = G_0 + G_i$, where G_0, G_1, \dots, G_k are iid standard normals. Indeed, it can be readily checked from this that $\mathbb{E}[\mathbf{g}_i^2] = 2$ and $\mathbb{E}[\mathbf{g}_i \mathbf{g}_j] = \mathbb{E}[G_0^2] = 1$, $i \neq j$, which is consistent with $\mathbf{g} \sim \mathcal{N}(\mathbf{0}, \mathbf{A})$. Thus, we can write

$$1 - \mathbb{P}\{\mathbf{g} \leq t \mathbf{1}_k\} = \mathbb{P}\left\{\max_{i \in [k]} \mathbf{g}_i \geq t\right\} = \mathbb{P}\{G_0 + \max_{i \in [k]} G_i \geq t\},$$

which completes the proof. ■

D.4.2 Slepian's bound

When the matrix \mathbf{A} does not have the special structure assumed by Lemma D.1, it is not possible in general to provide simple expressions as the one in (D.9). Yet, it might be possible to obtain

upper bounds of the same simple form. Such simple bounds can be useful for theoretical interpretations of otherwise complicated formulae, or can provide efficient means for quick (but, non-tight) implementations.

In this section, we discuss Slepian's inequality (see C.1) as a useful tool in this direction. Assume that $a = \min_{i,j \in [k]} A_{ij} \geq 0$. To begin, note that $\mathbf{A} \geq (\text{diag}(\mathbf{A}) - a\mathbf{I}) + a\mathbf{1}\mathbf{1}^T$, where the inequality holds element-wise and equality is true for the diagonal elements. Then, one can apply Slepian's Lemma C.1 to upper bound the conditional probability of error in (D.8) with the following simple bound:

$$\begin{aligned} 1 - \mathbb{P}\{\mathbf{A}^{1/2}\mathbf{g} \leq \mathbf{t}\} &\leq 1 - \mathbb{P}\left\{((\text{diag}(\mathbf{A}) - a\mathbf{I}) + a\mathbf{1}\mathbf{1}^T)^{1/2}\mathbf{g} \leq \mathbf{t}\right\} \\ &\leq \mathbb{P}\left\{\bigcup_{j \in [k]}\left\{G_0 + G_j\sqrt{[\mathbf{A}]_{jj}/a - 1} \geq [\mathbf{t}]_j/a\right\}\right\}, \quad G_0, G_1, \dots, G_k \stackrel{iid}{\sim} \mathcal{N}(0, 1). \end{aligned}$$

In the second line above, we used the Gaussian decomposition of Lemma D.1.

D.4.3 Simple bounds for GMM

Union bound. Of course, it is also possible to apply (a simpler) union bound to upper bound the tail probability in (D.8). Here, we show explicitly the result of applying union bound to the class-wise error probabilities of the GMM. Specifically, consider (D.1). An application of the union bound leads to the following:

$$\mathbb{P}_{e|c} = \mathbb{P}\left(\exists j \neq c : \langle \widehat{\mathbf{w}}_c - \widehat{\mathbf{w}}_j, \mathbf{z} \rangle \leq \langle \widehat{\mathbf{w}}_j - \widehat{\mathbf{w}}_c, \boldsymbol{\mu}_c \rangle + (\widehat{\mathbf{b}}_j - \widehat{\mathbf{b}}_c)\right) \quad (\text{D.10})$$

$$\begin{aligned} &\leq \sum_{j \neq c} \mathbb{P}_{G \sim \mathcal{N}(0, 1)}\left\{\|\widehat{\mathbf{w}}_j - \widehat{\mathbf{w}}_c\|_{\ell_2} G \leq \langle \widehat{\mathbf{w}}_j - \widehat{\mathbf{w}}_c, \boldsymbol{\mu}_c \rangle + (\widehat{\mathbf{b}}_j - \widehat{\mathbf{b}}_c)\right\} \\ &= \sum_{j \neq c} Q\left(\langle \frac{\widehat{\mathbf{w}}_c - \widehat{\mathbf{w}}_j}{\|\widehat{\mathbf{w}}_c - \widehat{\mathbf{w}}_j\|_{\ell_2}}, \boldsymbol{\mu}_c \rangle + \frac{\widehat{\mathbf{b}}_c - \widehat{\mathbf{b}}_j}{\|\widehat{\mathbf{w}}_c - \widehat{\mathbf{w}}_j\|_{\ell_2}}\right) \\ &= \sum_{j \neq c} Q\left(\frac{-[\mathbf{t}_c]_j}{\sqrt{[\mathbf{S}_c]_{jj}}}\right) \\ &\leq (k-1) \cdot Q(d_{\min}), \end{aligned} \quad (\text{D.11})$$

where in (D.11) $\mathbf{S}_c, \mathbf{t}_c$ are defined in (D.3) and in the last line we denote $d_{\min} := \min_{j \neq c} \{-[\mathbf{t}_c]_j / \sqrt{[\mathbf{S}_c]_{jj}}\}$.

Union bound without knowledge of cross-correlations $\langle \mathbf{w}_i, \mathbf{w}_j \rangle, i \neq j$. It is worth noting that the upper bound in (D.11) requires knowledge of the cross-correlations $\langle \widehat{\mathbf{w}}_j, \widehat{\mathbf{w}}_c \rangle, j \neq c$, i.e., of the off-diagonal entries of $\boldsymbol{\Sigma}_{\mathbf{w}, \mathbf{w}}$. Thankfully, our analysis allows predicting these values. For comparison, we ask whether it is possible to further upper bound the class-wise error probability if only the diagonal entries of $\boldsymbol{\Sigma}_{\mathbf{w}, \mathbf{w}}$ (i.e., the norms $\|\widehat{\mathbf{w}}_j\|_{\ell_2}, j \in [k]$) were known. A simple answer to this question is as follows. Observe that $\sqrt{[\mathbf{S}_c]_{jj}} = \|\widehat{\mathbf{w}}_c - \widehat{\mathbf{w}}_j\|_{\ell_2} \leq \|\widehat{\mathbf{w}}_c\|_{\ell_2} + \|\widehat{\mathbf{w}}_j\|_{\ell_2}$. Thus, if $[\mathbf{t}_c]_j < 0$ then the j th term in (D.11) is further upper bounded by $Q\left(\frac{-[\mathbf{t}_c]_j}{\|\widehat{\mathbf{w}}_c\|_{\ell_2} + \|\widehat{\mathbf{w}}_j\|_{\ell_2}}\right)$:

$$\mathbb{P}_{e|c} \leq \begin{cases} \sum_{j \neq c} Q\left(\frac{-[\mathbf{t}_c]_j}{\|\widehat{\mathbf{w}}_c\|_{\ell_2} + \|\widehat{\mathbf{w}}_j\|_{\ell_2}}\right) & \text{if } \mathbf{t}_c \leq \mathbf{0}, \\ 1 & \text{otherwise.} \end{cases}$$

Unfortunately, this bound becomes non-trivial for the class-wise probability of error only if

$$[\mathbf{t}_c]_j < 0 \Leftrightarrow \langle \widehat{\mathbf{w}}_c, \boldsymbol{\mu}_c \rangle + \widehat{\mathbf{b}}_c \geq \langle \widehat{\mathbf{w}}_j, \boldsymbol{\mu}_c \rangle + \widehat{\mathbf{b}}_j \text{ for all } j \neq c.$$

Intuitively, this assumes a regime where the weight vector $\widehat{\mathbf{w}}_c$ corresponding to class c aligns better with the corresponding mean vector $\boldsymbol{\mu}_c$ than the rest of the weight vectors $\widehat{\mathbf{w}}_j, j \neq c$. This emphasizes the important role of the cross-correlation matrix $\boldsymbol{\Sigma}_{\mathbf{w}, \mathbf{w}}$ (including the off-diagonals) for accurate performance prediction. For an illustration, we have implemented this bound and have compared it to our sharp predictions in Figure 7.

Oracle lower bound. For completeness, we briefly discuss an oracle lower bound for the class-wise probability of error in GMM. Specifically, assume that the means $\boldsymbol{\mu}_i, i \in [n]$ are known. Then

the problem of classifying a new sample \mathbf{x} is a k -ary hypothesis testing problem with Gaussian conditionals. Denote $\mathbb{P}_{\text{genie, Bayes}}$ the Bayes error of this hypothesis testing problem. Clearly $\mathbb{P}_{\text{genie, Bayes}}$ is a lower bound on the error of any classifier that is trained on data. The paper [WBW⁺16] further lower bounds $\mathbb{P}_{\text{genie, Bayes}}$ in terms of the Bayesian probability of errors between every two classes as follows:

$$\begin{aligned} \mathbb{P}_e \geq \mathbb{P}_{\text{genie, Bayes}} &\geq \frac{2}{k} \sum_{i \neq j} \pi_i \mathbb{P}_{\text{genie, ij}} \\ &= \frac{2}{k} \sum_{i \neq j} \frac{\pi_i}{\pi_i + \pi_j} \left\{ \pi_i \cdot Q\left(\frac{\|\boldsymbol{\mu}_i - \boldsymbol{\mu}_j\|_{\ell_2}}{2} + \frac{\log(\pi_i/\pi_j)}{\|\boldsymbol{\mu}_i - \boldsymbol{\mu}_j\|_{\ell_2}} \right) + \pi_j \cdot Q\left(\frac{\|\boldsymbol{\mu}_i - \boldsymbol{\mu}_j\|_{\ell_2}}{2} - \frac{\log(\pi_i/\pi_j)}{\|\boldsymbol{\mu}_i - \boldsymbol{\mu}_j\|_{\ell_2}} \right) \right\}. \end{aligned} \quad (\text{D.12})$$

where $\mathbb{P}_{\text{genie, ij}}$ is the Bayesian error between classes i and j with priors $\frac{\pi_i}{\pi_i + \pi_j}$ and $\frac{\pi_j}{\pi_i + \pi_j}$. For the last equality we have used the well-known formula for the Bayesian probability of binary Gaussian hypothesis testing. In the case of equal-priors the genie lower bound above simplifies to

$$\mathbb{P}_e \geq \frac{2}{k} \sum_{i \neq j} \pi_i \cdot Q\left(\frac{\|\boldsymbol{\mu}_i - \boldsymbol{\mu}_j\|_{\ell_2}}{2} \right). \quad (\text{D.13})$$

Note that, in contrast to (D.13), our analysis allows for precise evaluations of the missclassification error \mathbb{P}_e .

E The Class-averaging estimator

E.1 Proofs for GMM

E.1.1 GMM: Proof of Proposition 3.1

The first statement (3.1a) follows directly from the fact that $\frac{1}{n} \mathbf{1}^T \mathbf{Y}_i = \frac{n_i}{n} \xrightarrow{P} \pi_i$. For the next two statements note that

$$\widehat{\mathbf{w}}_i = \frac{1}{n} \mathbf{M} \mathbf{Y} \mathbf{Y}_i + \frac{1}{n} \mathbf{Z} \mathbf{Y}_i = \frac{1}{n} \sum_{j=1}^k \boldsymbol{\mu}_j (\mathbf{Y}_j^T \mathbf{Y}_i) + \frac{1}{n} \mathbf{Z} \mathbf{Y}_i = \frac{\|\mathbf{Y}_i\|_{\ell_2}^2}{n} \boldsymbol{\mu}_i + \frac{1}{n} \mathbf{Z} \mathbf{Y}_i, \quad (\text{E.1})$$

where in the last line we used orthogonality of the rows \mathbf{Y}_j of the matrix \mathbf{Y} :

$$\langle \mathbf{Y}_i, \mathbf{Y}_j \rangle = 0, \forall i \neq j \in [k]. \quad (\text{E.2})$$

To conclude simply use the facts that for all $i \in [k]$:

- (i) $\frac{\|\mathbf{Y}_i\|_{\ell_2}^2}{n} = \frac{n_i}{n} \xrightarrow{P} \pi_i$.
- (ii) $\mathbf{Z} \mathbf{Y}_i \sim \sigma \|\mathbf{Y}_i\|_{\ell_2} \mathbf{g}_i$ with $\mathbf{g}_i \stackrel{iid}{\sim} \mathcal{N}(\mathbf{0}, \mathbf{I}_d)$ because of (E.2).
- (iii) $\frac{\|\mathbf{g}_i\|_{\ell_2}}{\sqrt{n}} \xrightarrow{P} \sqrt{\gamma}$ and $\frac{1}{\sqrt{n}} \langle \mathbf{g}_i, \boldsymbol{\mu}_j \rangle \xrightarrow{P} 0$.

E.2 Proofs for MLM

E.2.1 Proof of Proposition 4.1

Let us define $\mathbf{g} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_r)$ and random vector $\mathbf{v} = [V_1, V_2, \dots, V_k]^T$ with entries:

$$\mathbf{v} = \frac{e^{\mathbf{V} \boldsymbol{\Sigma} \mathbf{g}}}{\mathbf{1}_k^T e^{\mathbf{V} \boldsymbol{\Sigma} \mathbf{g}}}, \quad V_i = \frac{e^{\mathbf{e}_i^T \mathbf{V} \boldsymbol{\Sigma} \mathbf{g}}}{\sum_{j \in [k]} e^{\mathbf{e}_j^T \mathbf{V} \boldsymbol{\Sigma} \mathbf{g}}}, \quad i \in [k]. \quad (\text{E.3})$$

We will prove the following three statements:

$$\widehat{\mathbf{b}} \xrightarrow{P} \mathbb{E}[\mathbf{v}] \quad (\text{E.4a})$$

$$\boldsymbol{\Sigma}_{\mathbf{w}, \boldsymbol{\mu}} \xrightarrow{P} \mathbb{E}[\mathbf{v} \mathbf{g}^T] \boldsymbol{\Sigma} \mathbf{V}^T \quad (\text{E.4b})$$

$$\boldsymbol{\Sigma}_{\mathbf{w}, \mathbf{w}} \xrightarrow{P} \gamma \cdot \text{diag}(\mathbb{E}[\mathbf{v}]) + \mathbb{E}[\mathbf{v} \mathbf{g}^T] \cdot \mathbb{E}[\mathbf{g} \mathbf{v}^T]. \quad (\text{E.4c})$$

These lead to (4.2) using Lemma C.3. Therefore, in what follows, we prove (E.4)

For the intercepts $\widehat{\mathbf{b}}_\ell$, $\ell \in [k]$ it holds that

$$\widehat{\mathbf{b}}_\ell = \frac{1}{n} \mathbf{1}^T \mathbf{Y}_\ell \xrightarrow{P} \mathbb{P}\{Y = \ell\} = \mathbb{E} \left[\frac{e^{\mathbf{h}_\ell}}{\sum_{j \in [k]} e^{\mathbf{h}_j}} \right],$$

where $\mathbf{h} \sim \mathcal{N}(\mathbf{0}_k, \Sigma_{\mu\mu})$. To deduce the first statement in (E.4a), note that $\mathbf{h} \stackrel{(D)}{=} \mathbf{V}\Sigma\mathbf{g}$.

Continuing with the vectors $\widehat{\mathbf{w}}_\ell$, $\ell \in [k]$, recall that $\mathbf{w}_\ell = \frac{1}{n} \mathbf{X} \mathbf{Y}_\ell = \frac{1}{n} \sum_{i \in [n]} \mathbf{x}_{i\ell} [\mathbf{Y}_\ell]_{i\ell}$. Consider the singular decomposition

$$\mathbf{M} = \mathbf{U}\Sigma\mathbf{V}^T = [\mathbf{u}_1 \quad \mathbf{u}_2 \quad \dots \quad \mathbf{u}_r] \text{diag}(\sigma_1, \sigma_2, \dots, \sigma_r) \begin{bmatrix} \mathbf{v}_1^T \\ \mathbf{v}_2^T \\ \vdots \\ \mathbf{v}_r^T \end{bmatrix},$$

with $\mathbf{U} \in \mathbb{R}^{d \times r}$, $\Sigma \in \mathbb{R}^{r \times r}$, and $\mathbf{V} \in \mathbb{R}^{k \times r}$ where $r = \text{rank}(\mathbf{M}) \leq k$. Decompose $\mathbf{X} \in \mathbb{R}^{d \times n}$ as $\mathbf{X} = \mathbf{U}\mathbf{U}^T \mathbf{X} + \mathbf{P}^\perp \mathbf{X}$ with $\mathbf{P}^\perp = \mathbf{I}_d - \mathbf{U}\mathbf{U}^T$. With this notation we compute

$$\begin{aligned} \langle \mathbf{w}_\ell, \mu_c \rangle &= \frac{1}{n} \sum_{i \in [n]} \mathbf{x}_i^T \mu_c [\mathbf{Y}_\ell]_i = \frac{1}{n} \sum_{j=1}^r \sum_{i \in [n]} (\mathbf{x}_i^T \mathbf{u}_j) \cdot (\mu_c^T \mathbf{u}_j) [\mathbf{Y}_\ell]_i + \frac{1}{n} \sum_{i \in [n]} \cdot (\mu_c^T \mathbf{P}^\perp \mathbf{x}_i) [\mathbf{Y}_\ell]_i \\ &\xrightarrow{P} \sum_{j=1}^r (e_c^T \mathbf{V} \Sigma \mathbf{e}_j) \mathbb{E} \left[\mathbf{g}_j \frac{(e^{\mathbf{e}_\ell^T \mathbf{V} \Sigma \mathbf{g}})}{\sum_{\ell' \in [k]} e^{\mathbf{e}_{\ell'}^T \mathbf{V} \Sigma \mathbf{g}}} \right] \\ &= \sum_{j=1}^r \mathbb{E}[V_\ell \mathbf{g}_j] (e_j^T \Sigma \mathbf{V}^T \mathbf{e}_c). \end{aligned} \tag{E.5}$$

Here, we have recognized that for every $i \in [n]$: $\mathbf{U}^T \mathbf{x}_i \sim \mathbf{g}$, and also, conditioned on \mathbf{x}_i : $[\mathbf{Y}_\ell]_i \sim \text{Bern}\left(e^{\mu_\ell^T \mathbf{x}_i} / \sum_{\ell' \in [k]} e^{\mu_{\ell'}^T \mathbf{x}_i}\right)$ and $\mu_\ell^T \mathbf{x}_i = e_\ell^T \mathbf{V} \Sigma \mathbf{U}^T \mathbf{x}_i \sim e_\ell^T \mathbf{V} \Sigma \mathbf{g}$, $\ell \in [k]$. This shows the second statement in (E.4b) when expressed in matrix form.

We proceed similarly with the proof of the last statement in (E.4c) as follows:

$$\begin{aligned} \langle \mathbf{w}_\ell, \mathbf{w}_c \rangle &= \frac{1}{n^2} \sum_{i_\ell \in [n], i_c \in [n]} \mathbf{x}_{i_\ell}^T \mathbf{x}_{i_c} [\mathbf{Y}_\ell]_{i_\ell} [\mathbf{Y}_c]_{i_c} \\ &= \frac{1}{n^2} \sum_{j=1}^r \sum_{i_\ell \in [n], i_c \in [n]} (\mathbf{x}_{i_\ell}^T \mathbf{u}_j) (\mathbf{x}_{i_c}^T \mathbf{u}_j) [\mathbf{Y}_\ell]_{i_\ell} [\mathbf{Y}_c]_{i_c} + \frac{1}{n^2} \sum_{i_\ell \in [n], i_c \in [n]} (\mathbf{P}^\perp \mathbf{x}_{i_\ell})^T (\mathbf{P}^\perp \mathbf{x}_{i_c}) [\mathbf{Y}_\ell]_{i_\ell} [\mathbf{Y}_c]_{i_c} \end{aligned}$$

For $i_\ell = i_c = i \in [n]$ note that

$$\frac{1}{n^2} \sum_{i \in [n]} \sum_{j=1}^r (\mathbf{x}_i^T \mathbf{u}_j) (\mathbf{x}_i^T \mathbf{u}_j) [\mathbf{Y}_\ell]_i [\mathbf{Y}_c]_i \xrightarrow{P} 0,$$

while, for $i_\ell \neq i_c$,

$$\begin{aligned} \frac{1}{n^2} \sum_{i_\ell \neq i_c \in [n]} \sum_{j=1}^r (\mathbf{x}_{i_\ell}^T \mathbf{u}_j) [\mathbf{Y}_\ell]_{i_\ell} (\mathbf{x}_{i_c}^T \mathbf{u}_j) [\mathbf{Y}_c]_{i_c} &\xrightarrow{P} \sum_{j=1}^r \mathbb{E} \left[\mathbf{g}_j \cdot \frac{(e^{\mathbf{e}_\ell^T \mathbf{V} \Sigma \mathbf{g}})}{\sum_{\ell' \in [k]} e^{\mathbf{e}_{\ell'}^T \mathbf{V} \Sigma \mathbf{g}}} \right] \mathbb{E} \left[\mathbf{g}_j \cdot \frac{(e^{\mathbf{e}_c^T \mathbf{V} \Sigma \mathbf{g}})}{\sum_{\ell' \in [k]} e^{\mathbf{e}_{\ell'}^T \mathbf{V} \Sigma \mathbf{g}}} \right] \\ &= \sum_{j=1}^r \mathbb{E}[V_\ell \mathbf{g}_j] \mathbb{E}[\mathbf{g}_j V_c] = e_\ell^T \mathbf{E}[\mathbf{v} \mathbf{g}^T] \cdot \mathbf{E}[\mathbf{g} \mathbf{v}^T] e_c^T \end{aligned}$$

Furthermore,

$$\frac{1}{n} \sum_{i \in [n]} \|\mathbf{P}^\perp \mathbf{x}_i\|_{\ell_2}^2 [\mathbf{Y}_\ell]_i [\mathbf{Y}_c]_i^2 \xrightarrow{P} \gamma \cdot \mathbf{1}_{\ell, c} \cdot \mathbb{E} \left[\left(\frac{e^{\mathbf{e}_\ell^T \mathbf{V} \Sigma \mathbf{g}}}{\sum_{\ell' \in [k]} e^{\mathbf{e}_{\ell'}^T \mathbf{V} \Sigma \mathbf{g}}} \right) \right] = \gamma \cdot e_\ell^T \text{diag}(\mathbb{E}[\mathbf{v}]) e_c.$$

Combining the last two displays results in (E.4b), as desired.

E.2.2 Orthogonal means

Here, we specialize the general result of Proposition 4.1 to the special case of orthogonal means: $\langle \boldsymbol{\mu}_i, \boldsymbol{\mu}_j \rangle = 0, \forall i \neq j$. Recall the notation $\mu_i = \|\boldsymbol{\mu}_i\|_{\ell_2}, i \in [k]$. Then, in this case the parameters in (4.1) are simply given by the following

$$\pi_i := \mathbb{E} \left[\frac{e^{\mu_i G_i}}{\sum_{\ell \in [k]} e^{\mu_\ell G_\ell}} \right], i \in [k] \quad \text{and} \quad \boldsymbol{\Pi}_{ij} := \mathbb{E} \left[\frac{e^{\mu_i G_i} e^{\mu_j G_j}}{\left(\sum_{\ell \in [k]} e^{\mu_\ell G_\ell} \right)^2} \right], i, j \in [k]. \quad (\text{E.6})$$

Specifically, (4.3) can be equivalently expressed as

$$\mathbb{P}_{e, \text{Avg}} \xrightarrow{P} \mathbb{P} \left(\arg \max_{\ell \in [k]} \{ \gamma \cdot \text{diag}(\boldsymbol{\pi}) \cdot \tilde{\mathbf{g}} + (\text{diag}(\boldsymbol{\pi}) - \boldsymbol{\Pi}) \cdot \boldsymbol{\Sigma} \mathbf{g} \} \neq Y(\mathbf{g}) \right) \quad (\text{E.7})$$

$$= \mathbb{P} \left(\bigcup_{j \neq Y} \{ \gamma \cdot \boldsymbol{\pi}_\ell \cdot \tilde{\mathbf{g}}_\ell \geq \gamma \cdot \boldsymbol{\pi}_Y \cdot \tilde{\mathbf{g}}_Y + (\mathbf{e}_Y - \mathbf{e}_\ell)^T (\text{diag}(\boldsymbol{\pi}) - \boldsymbol{\Pi}) \boldsymbol{\Sigma} \mathbf{g} + (\boldsymbol{\pi}_Y - \boldsymbol{\pi}_\ell) \} \right), \quad (\text{E.8})$$

where $\mathbf{g}, \tilde{\mathbf{g}} \stackrel{iid}{\sim} \mathcal{N}(\mathbf{0}, \mathbf{I}_k)$, $\mathbb{P}(Y(\mathbf{g}) = c) = \frac{e^{\mu_c \mathbf{g}_c}}{\sum_{\ell \in [k]} e^{\mu_\ell \mathbf{g}_\ell}}$ and $\boldsymbol{\Sigma} = \text{diag}(\mu_1, \dots, \mu_k)$.

F On the Bayes risk of GMM: Proof of Proposition 3.4

Without loss of generality in this proof we assume $\sigma = 1$. The general result follows by simply replacing (μ, σ) with $(\frac{\mu}{\sigma}, 1)$ and using the proof for $\sigma = 1$. Recall that the feature vectors $\mathbf{x}_1, \dots, \mathbf{x}_n$ of the training data set are given by:

$$\mathbf{x}_i = \mathbf{M} \mathbf{y}_i + \mathbf{z}_i, \quad i \in [n],$$

where the matrix of means $\mathbf{M} \in \mathbb{R}^{d \times k}$ has iid Gaussian entries with variance μ^2/d , $\mathbf{z}_i \stackrel{iid}{\sim} \mathcal{N}(\mathbf{0}, \mathbf{I}_d)$ and $\mathbf{y}_i \stackrel{iid}{\sim} \text{Unif}(\mathbf{e}_1, \dots, \mathbf{e}_k)$ with \mathbf{e}_j denoting the j^{th} canonical vector in \mathbb{R}^k . By definition here the Bayes estimator is the maximum-likelihood (ML) estimator. By applying the law of total probability and by successive application of the Bayes rule we have the following chain of reformulations of the ML:

$$\begin{aligned} \hat{y}_{n+1} &= \arg \max_{\mathbf{e}_j, j \in [k]} P \{ \mathbf{y} = \mathbf{e}_j \mid \mathbf{X}, \mathbf{Y}, \mathbf{x}_{n+1} \} \\ &= \arg \max_{\mathbf{e}_j, j \in [k]} \int P \{ \mathbf{y} = \mathbf{e}_j \mid \mathbf{M}, \mathbf{X}, \mathbf{Y}, \mathbf{x}_{n+1} \} P \{ \mathbf{M} \mid \mathbf{X}, \mathbf{Y}, \mathbf{x}_{n+1} \} d\mathbf{M} \\ &= \arg \max_{\mathbf{e}_j, j \in [k]} \int \frac{P \{ \mathbf{x}_{n+1} \mid \mathbf{y} = \mathbf{e}_j, \mathbf{M}, \mathbf{X}, \mathbf{Y} \} \cdot P \{ \mathbf{y} = \mathbf{e}_j \mid \mathbf{M}, \mathbf{X}, \mathbf{Y} \}}{P \{ \mathbf{x}_{n+1} \mid \mathbf{M}, \mathbf{X}, \mathbf{Y} \}} P \{ \mathbf{M} \mid \mathbf{X}, \mathbf{Y}, \mathbf{x}_{n+1} \} d\mathbf{M} \\ &= \arg \max_{\mathbf{e}_j, j \in [k]} \int P \{ \mathbf{x}_{n+1} \mid \mathbf{y} = \mathbf{e}_j, \mathbf{M} \} \frac{P \{ \mathbf{M} \mid \mathbf{X}, \mathbf{Y}, \mathbf{x}_{n+1} \}}{P \{ \mathbf{x}_{n+1} \mid \mathbf{M}, \mathbf{X}, \mathbf{Y} \}} d\mathbf{M} \end{aligned} \quad (\text{F.1})$$

$$= \arg \max_{\mathbf{e}_j, j \in [k]} \int P \{ \mathbf{x}_{n+1} \mid \mathbf{y} = \mathbf{e}_j, \mathbf{M} \} \frac{P \{ \mathbf{M} \mid \mathbf{X}, \mathbf{Y} \}}{P \{ \mathbf{x}_{n+1} \mid \mathbf{X}, \mathbf{Y} \}} d\mathbf{M} \quad (\text{F.2})$$

$$= \arg \max_{\mathbf{e}_j, j \in [k]} \int P \{ \mathbf{x}_{n+1} \mid \mathbf{y} = \mathbf{e}_j, \mathbf{M} \} P \{ \mathbf{M} \mid \mathbf{X}, \mathbf{Y} \} d\mathbf{M} \quad (\text{F.3})$$

To arrive in (F.1) we used that $P(\mathbf{y} = \mathbf{e}_j \mid \mathbf{M}, \mathbf{X}, \mathbf{Y}) = \pi_j, \forall j \in [k]$ and $P(\mathbf{x}_{n+1} \mid \mathbf{y} = \mathbf{e}_j, \mathbf{M}, \mathbf{X}, \mathbf{Y}) = P(\mathbf{x}_{n+1} \mid \mathbf{y} = \mathbf{e}_j, \mathbf{M})$. Also, (F.2) follows by recognizing that $P(\mathbf{x}_{n+1} \mid \mathbf{X}, \mathbf{Y}) > 0$ is independent of the variable of integration \mathbf{M} and of the optimization variable j . For the same reasons, in (F.3) we have ignored the normalizing term $P(\mathbf{X} \mid \mathbf{Y})$.

Recalling that $\mathbf{z}_{n+1} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_d)$, we have that $P(\mathbf{x}_{n+1} \mid \mathbf{y} = \mathbf{e}_j, \mathbf{M}) \propto \exp \left(-\|\mathbf{x}_{n+1} - \mathbf{e}_j\|_{\ell_2}^2 / 2 \right)$ where \propto hides constant positive terms. Moreover, the posterior probability of the mean matrix given

the training data is given by

$$\begin{aligned} P(\mathbf{X} | \mathbf{M}, \mathbf{Y}) \cdot P(\mathbf{M}) &\propto \exp\left(-\frac{\|\mathbf{X} - \mathbf{M}\mathbf{Y}\|_{\ell_2}^2}{2}\right) \cdot \exp\left(-\frac{\|\mathbf{M}\|_{\ell_2}^2}{2(\mu^2/d)}\right) \\ &\propto \prod_{c=1}^k \left\{ \exp\left(-\frac{\|\boldsymbol{\mu}_c\|_{\ell_2}^2}{2(\mu^2/d)}\right) \cdot \prod_{i \in \mathcal{C}_c} \exp\left(-\frac{\|\mathbf{x}_i - \boldsymbol{\mu}_c\|_{\ell_2}^2}{2}\right) \right\}, \end{aligned} \quad (\text{F.4})$$

where we denote by \mathcal{C}_c the collection of training samples that belong to class $c \in [k]$, i.e. $\mathcal{C}_c = \{i \in [n] \mid \mathbf{y}_i = \mathbf{e}_c\}$.

With these the objective function of the ML rule in (F.3) becomes:

$$\hat{\mathbf{y}}_{n+1} = \arg \max_{j \in [k]} \mathcal{I}(j, \mathcal{C}_j \cap \{n+1\}) \cdot \prod_{\substack{c=1 \\ c \neq j}}^k \mathcal{I}(c, \mathcal{C}_c), \quad (\text{F.5})$$

where for $\ell \in [k]$ and a subset $\mathcal{A} \subset [n+1]$ we denote

$$\mathcal{I}(\ell, \mathcal{A}) := \int d\boldsymbol{\mu}_\ell \exp\left(-\frac{\|\boldsymbol{\mu}_\ell\|_{\ell_2}^2}{2(\mu^2/d)}\right) \cdot \exp\left(-\sum_{i \in \mathcal{A}} \frac{\|\mathbf{x}_i - \boldsymbol{\mu}_\ell\|_{\ell_2}^2}{2}\right).$$

By completing the squares and invoking a gaussian integral it can be shown that

$$\begin{aligned} \mathcal{I}(\ell, \mathcal{A}) &:= \sqrt{\frac{(d/\mu^2 + |\mathcal{A}|)}{(2\varpi)^d}} \exp\left(-\frac{\left(1 - \frac{1}{d/\mu^2 + |\mathcal{A}|}\right)}{2} \sum_{i \in \mathcal{A}} \|\mathbf{x}_i\|_{\ell_2}^2 + \frac{1}{2(d/\mu^2 + |\mathcal{A}|)} \sum_{i \in \mathcal{A}} \langle \mathbf{x}_i, \sum_{\substack{j \in \mathcal{A} \\ j \neq i}} \mathbf{x}_j \rangle\right) \\ &:= \sqrt{\frac{(d/\mu^2 + |\mathcal{A}|)}{(2\varpi)^d}} \exp\left(-\frac{1}{2\left(\frac{d}{\mu^2} + |\mathcal{A}|\right)} \left(\left(\frac{d}{\mu^2} + |\mathcal{A}| - 1\right) \sum_{i \in \mathcal{A}} \|\mathbf{x}_i\|_{\ell_2}^2 - \sum_{i \in \mathcal{A}} \langle \mathbf{x}_i, \sum_{\substack{j \in \mathcal{A} \\ j \neq i}} \mathbf{x}_j \rangle\right)\right) \\ &:= \sqrt{\frac{(d/\mu^2 + |\mathcal{A}|)}{(2\varpi)^d}} \exp\left(-\frac{1}{2\left(\frac{d/n}{\mu^2} + |\mathcal{A}|/n\right)} \left(\left(\frac{d/n}{\mu^2} + \frac{|\mathcal{A}|}{n} - \frac{1}{n}\right) \sum_{i \in \mathcal{A}} \|\mathbf{x}_i\|_{\ell_2}^2 - \frac{1}{n} \sum_{i \in \mathcal{A}} \langle \mathbf{x}_i, \sum_{\substack{j \in \mathcal{A} \\ j \neq i}} \mathbf{x}_j \rangle\right)\right). \end{aligned}$$

Using this in (F.5) we have that

$$\hat{\mathbf{y}}_{n+1} = \arg \max_{j \in [k]} \mathcal{I}(j) \cdot \exp\left(-\frac{1}{2\left(\frac{d/n}{\mu^2} + \frac{n_j+1}{n}\right)} \left(\left(\frac{d/n}{\mu^2} + \frac{n_j}{n}\right) \|\mathbf{x}_{n+1}\|_{\ell_2}^2 - \frac{2}{n} \langle \mathbf{x}_{n+1}, \sum_{\ell \in \mathcal{C}_j} \mathbf{x}_\ell \rangle\right)\right), \quad (\text{F.6})$$

where $\xi(n_c) := \frac{d/n}{\mu^2} + \frac{n_c}{n}$, $c \in [k]$ and

$$\mathcal{I}(j) := \left\{ \prod_{\substack{c=1 \\ c \neq j}}^k e^{-\frac{1}{2\xi(n_c)} \left((\xi(n_c) - \frac{1}{n}) \sum_{i \in \mathcal{C}_c} \|\mathbf{x}_i\|_{\ell_2}^2 - \frac{1}{n} \sum_{i \in \mathcal{C}_c} \langle \mathbf{x}_i, \sum_{\ell \in \mathcal{C}_c \setminus \{j\}} \mathbf{x}_\ell \rangle \right)} \right\} \cdot e^{-\frac{1}{2(\xi(n_j) + \frac{1}{n})} \left(\xi(n_j) \sum_{i \in \mathcal{C}_j} \|\mathbf{x}_i\|_{\ell_2}^2 - \frac{1}{n} \sum_{i \in \mathcal{C}_j} \langle \mathbf{x}_i, \sum_{\ell \in \mathcal{C}_j \setminus \{j\}} \mathbf{x}_\ell \rangle \right)}.$$

We conclude that

$$\begin{aligned} \hat{\mathbf{y}}_{n+1} &= \arg \max_{j \in [k]} \log(\mathcal{I}(j)) - \frac{1}{2\left(\frac{d/n}{\mu^2} + \frac{n_j+1}{n}\right)} \left\{ \left(\frac{d/n}{\mu^2} + \frac{n_j}{n}\right) \|\mathbf{x}_{n+1}\|_{\ell_2}^2 - \frac{2}{n} \langle \mathbf{x}_{n+1}, \sum_{\ell \in \mathcal{C}_j} \mathbf{x}_\ell \rangle \right\} \\ &= \arg \max_{j \in [k]} \log(\mathcal{I}(j)) + \frac{1}{2\left(\frac{d/n}{\mu^2} + \frac{n_j+1}{n}\right)} \left\{ \frac{2}{n} \langle \mathbf{x}_{n+1}, \sum_{\ell \in \mathcal{C}_j} \mathbf{x}_\ell \rangle \right\}. \end{aligned} \quad (\text{F.7})$$

Next, we evaluate the objective in (F.7) in the asymptotic limit $n, d \rightarrow \infty, n/d = \gamma$. First, since $n_c/n \xrightarrow{P} \pi$, note that $\mathcal{I}(j) - \mathcal{I}(\ell) \xrightarrow{P} 0$ for all $\ell, j \in [k]$. Moreover, note that

$$\begin{aligned} \frac{1}{n} \langle \mathbf{x}_{n+1}, \sum_{\ell \in \mathcal{C}_j} \mathbf{x}_\ell \rangle &= \frac{1}{n} \langle M\mathbf{y}_{n+1} + \mathbf{z}_{n+1}, n_j \boldsymbol{\mu}_j + \sum_{\ell \in \mathcal{C}_j} \mathbf{z}_\ell \rangle \\ &= \frac{n_j}{n} \langle M\mathbf{y}_{n+1}, \boldsymbol{\mu}_j \rangle + \frac{n_j}{n} \langle \mathbf{z}_{n+1}, \boldsymbol{\mu}_j \rangle + \frac{1}{n} \sum_{\ell \in \mathcal{C}_j} \langle M\mathbf{y}_{n+1}, \mathbf{z}_\ell \rangle + \frac{1}{n} \sum_{\ell \in \mathcal{C}_j} \langle \mathbf{z}_{n+1}, \mathbf{z}_\ell \rangle \end{aligned} \quad (\text{F.8})$$

For each one of the four terms in (F.8), we have the following by the CLT:

$$\begin{aligned} \frac{n_j}{n} \langle M\mathbf{y}_{n+1}, \boldsymbol{\mu}_j \rangle &\xrightarrow{P} \pi\mu^2 \langle \mathbf{y}_{n+1}, \mathbf{e}_j \rangle \\ \frac{n_j}{n} \langle \mathbf{z}_{n+1}, \boldsymbol{\mu}_j \rangle &\xrightarrow{(D)} \mathcal{N}(0, \pi^2 r^2) \\ \frac{1}{n} \sum_{\ell \in \mathcal{C}_j} \langle M\mathbf{y}_{n+1}, \mathbf{z}_\ell \rangle &\xrightarrow{P} 0 \\ \frac{1}{n} \sum_{\ell \in \mathcal{C}_j} \langle \mathbf{z}_{n+1}, \mathbf{z}_\ell \rangle &\xrightarrow{(D)} \mathcal{N}(0, \pi\gamma), \end{aligned}$$

where in the last line we used the fact that $\frac{1}{\sqrt{n_j}} \sum_{\ell \in \mathcal{C}_j} \frac{\langle \mathbf{z}_{n+1}, \mathbf{z}_\ell \rangle}{\sqrt{n}} \xrightarrow{(D)} \mathcal{N}(0, \gamma)$.

Therefore, in the asymptotic limit, the Bayes estimator is the solution to:

$$\hat{\mathbf{y}}_{n+1} = \arg \max_{\mathbf{e}_j, j \in [k]} \pi\mu^2 \langle \mathbf{y}_{n+1}, \mathbf{e}_j \rangle + \sqrt{\pi(\pi\mu^2 + \gamma)} G_j, \quad G_1, \dots, G_k \stackrel{iid}{\sim} \mathcal{N}(0, 1). \quad (\text{F.9})$$

As such, the probability of error is

$$\begin{aligned} \mathbb{P}_e &= \mathbb{P} \{ \hat{\mathbf{y}}_{n+1} \neq \mathbf{y}_{n+1} \} = \mathbb{P} \left\{ \pi\mu^2 + \sqrt{\pi(\pi\mu^2 + \gamma)} G_0 \leq \max_{\ell \in [k-1]} \sqrt{\pi(\pi\mu^2 + \gamma)} G_\ell \right\} \\ &= \mathbb{P} \left\{ G_0 + \max_{\ell \in [k-1]} G_\ell \geq \mu^2 \sqrt{\frac{\pi}{\pi\mu^2 + \gamma}} \right\}. \end{aligned} \quad (\text{F.10})$$

G Proof outline for least-squares: key ideas and challenges

In this section, we provide a proof sketch for the analysis of the multiclass least-squares (LS) classifier. Specifically, we discuss our approach towards specifying the high-dimensional limits of the key quantities needed to evaluate the classification error: \mathbf{b} , $\boldsymbol{\Sigma}_{\mathbf{w}, \boldsymbol{\mu}}$, and, $\boldsymbol{\Sigma}_{\mathbf{w}, \mathbf{w}}$. For simplicity, we focus here on the performance of the LS classifier GMM. We note that our proofs for the MLM and the Weighted Least-Squares (WLS) classifiers follow the same general strategy, but in some parts require more involved and intricate analysis and derivations. Our proof follows the following general steps; see the appendix for complete details and derivations.

Step I: Decomposing the loss across classes. Recall from Section 2.2 that the multiclass LS classifier produces a linear classifier $\mathbf{x} \mapsto \mathbf{W}\mathbf{x} + \mathbf{b}$ via a least-squares fit to the training data:

$$(\widehat{\mathbf{W}}, \widehat{\mathbf{b}}) := \frac{1}{2n} \|\mathbf{W}\mathbf{X} + \mathbf{b}\mathbf{1}_n^T - \mathbf{Y}\|_F^2. \quad (\text{G.1})$$

Notice that the objective function above is separable. That is,

$$\frac{1}{2n} \|\mathbf{W}\mathbf{X} + \mathbf{b}\mathbf{1}_n^T - \mathbf{Y}\|_F^2 = \frac{1}{2n} \sum_{\ell=1}^k \|\mathbf{X}^T \mathbf{w}_\ell + b_\ell \mathbf{1}_n - \mathbf{Y}_\ell\|_{\ell_2}^2.$$

Hence, for each $\ell \in [k]$,

$$(\widehat{\mathbf{w}}_\ell, \widehat{b}_\ell) = \arg \min_{\mathbf{w}_\ell, b_\ell} \frac{1}{2n} \|\mathbf{X}^T \mathbf{w}_\ell + b_\ell \mathbf{1}_n - \mathbf{Y}_\ell\|_{\ell_2}^2. \quad (\text{G.2})$$

This decomposition is convenient for analysis as it is easier to compute the statistical properties of the simple single-output LS in (G.2) compared to the multi-output objective in (G.1). Indeed, as we show, in Step III, this simplification will eventually allow us to compute the high-dimensional behavior of the following key quantities for all $\ell \in [k]$: (i) the intercept \widehat{b}_ℓ , (ii) the mean-correlations $\langle \widehat{w}_\ell, \mu_c \rangle$, $c \in [k]$, (iii) the norm $\|\widehat{w}_\ell\|_{\ell_2}$.

Step II: Reduction to an Auxiliary Optimization (AO) problem via CGMT. To calculate the high-dimensional statistical behavior of (G.2) we use the Convex Gaussian min-max Theorem (CGMT) [Sto13, TOH15] framework. We provide a brief introduction of the CGMT machinery in Section G.1. Roughly stated, this framework allows us to replace a *Primary Optimization* (PO) problem of the form (G.2) with an *Auxiliary Optimization* (AO) problem that is simpler to analyze, but is predictive of the behavior of the latter. For instance, for the PO in (G.2) in the GMM, after some algebraic manipulations, the AO problem takes the form

$$\frac{1}{2} \left(\min_{\mathbf{w}_\ell, b_\ell} \frac{1}{\sqrt{n}} \|\sigma \|\mathbf{w}_\ell\|_{\ell_2} \mathbf{g} + \mathbf{Y}^T \mathbf{M}^T \mathbf{w}_\ell + b_\ell \mathbf{1}_n - \mathbf{Y}_\ell \|_{\ell_2} + \frac{1}{\sqrt{n}} \sigma \mathbf{h}^T \mathbf{w}_\ell \right)_+^2, \quad (\text{G.3})$$

where $(x)_+ := \max(0, x)$ and $\mathbf{g} \in \mathbb{R}^n$ and $\mathbf{h} \in \mathbb{R}^d$ are two independent Gaussian random vectors distributed as $\mathcal{N}(\mathbf{0}, \mathbf{I}_n)$ and $\mathcal{N}(\mathbf{0}, \mathbf{I}_d)$.

Step III: Simplification of the AO and computing $\Sigma_{w, \mu}$ and \mathbf{b} . In this step we carry out a series of intricate calculations to further simplify (G.3) and characterize its various asymptotic properties. At a high-level, we follow the principled machinery introduced in [TOH15, TAH18], organizing our analysis in three intermediate steps: (a) Scalarization; (b) Convergence analysis; and (c) Deterministic analysis. We note that each one of these intermediate steps for the multiclass setting is more involved than in previously considered regression and binary classification settings. The detailed derivations are deferred to the Appendix H.1. At the end of this analysis step, we have computed the high-dimensional behavior of the intercepts \widehat{b}_ℓ , $\ell \in [k]$, the mean-correlations $\langle \widehat{w}_\ell, \mu_c \rangle$, $\ell, c \in [k]$, the norms $\|\widehat{w}_\ell\|_{\ell_2}$, $\ell \in [k]$ and the LS training loss $\|\mathbf{X}^T \widehat{\mathbf{w}}_\ell + \widehat{b}_\ell \mathbf{1}_n - \mathbf{Y}_\ell\|_{\ell_2}$. In particular, for GMM these calculations allow us to conclude the following limits for all $\ell \in [k]$:

$$\widehat{b}_\ell \xrightarrow{P} \pi_\ell (1 - (\mathbf{e}_\ell - \boldsymbol{\pi})^T \mathbf{V} \boldsymbol{\Sigma} \boldsymbol{\Delta}^{-1} \boldsymbol{\Sigma} \mathbf{V}^T \boldsymbol{\pi}), \quad \mathbf{M}^T \widehat{\mathbf{w}}_\ell \xrightarrow{P} \pi_\ell \mathbf{V} \boldsymbol{\Sigma} \boldsymbol{\Delta}^{-1} \boldsymbol{\Sigma} \mathbf{V}^T (\mathbf{e}_\ell - \boldsymbol{\pi}), \quad (\text{G.4})$$

and

$$\|\mathbf{w}_\ell\|_{\ell_2}^2 \xrightarrow{P} \frac{\gamma}{(1-\gamma)\sigma^2} \pi_\ell (1 - \pi_\ell) + \pi_\ell^2 (\mathbf{e}_\ell - \boldsymbol{\pi})^T \mathbf{V} \boldsymbol{\Sigma} \boldsymbol{\Delta}^{-1} \left(\boldsymbol{\Delta}^{-1} - \frac{\gamma}{(1-\gamma)\sigma^2} \mathbf{I}_r \right) \boldsymbol{\Sigma} \mathbf{V}^T (\mathbf{e}_\ell - \boldsymbol{\pi}), \quad (\text{G.5})$$

where $\boldsymbol{\Delta} := \sigma^2 \mathbf{I}_r + \boldsymbol{\Sigma} \mathbf{V}^T \mathbf{P} \boldsymbol{\Sigma} \mathbf{V} > \mathbf{0}_{r \times r}$ and $\mathbf{P} := \text{diag}(\boldsymbol{\pi}) - \boldsymbol{\pi} \boldsymbol{\pi}^T$.

Expressing (G.4) in matrix form leads to (3.3a) in Theorem 3.2. Thus, it remains to prove (3.3b), i.e., to determine the high-dimensional limit of $\Sigma_{w, w}$. Note that (G.5) already determines the diagonal entries of $\Sigma_{w, w}$. However, thus far, our analysis treats the optimization of each classifier $\widehat{\mathbf{w}}_\ell$, $\ell \in [k]$ independently and provides no information for the cross-correlation $\langle \widehat{\mathbf{w}}_\ell, \widehat{\mathbf{w}}_c \rangle$, $\ell \neq c \in [k]$.

Step IV: Computing $\Sigma_{w, w}$ and capturing cross-correlations. The final and most involved part of our analysis is characterizing the asymptotic behavior of $\Sigma_{w, w}$. To see why this is particularly challenging note that the reduction from (G.1) to (G.2) “breaks” the dependence of all $\widehat{\mathbf{w}}_1, \widehat{\mathbf{w}}_2, \dots, \widehat{\mathbf{w}}_k$ on the *same* feature matrix \mathbf{X} . Capturing this dependence is crucial in determining the “cross-correlations” $\langle \widehat{\mathbf{w}}_\ell, \widehat{\mathbf{w}}_c \rangle$, $\ell \neq c$. As noted in Section 2.3 the matrix $\Sigma_{w, w}$ is needed to calculate the class-wise and total miss-classification errors. Unfortunately, the CGMT is *not* directly applicable to the multi-output LS optimization in (G.1). Our idea to circumvent this challenge builds on the following simple observation: the vector $\widehat{\mathbf{w}}_{\ell, c} = \widehat{\mathbf{w}}_\ell + \widehat{\mathbf{w}}_c$ is itself the solution to another simple single-output LS problem.

Lemma G.1 *For $\ell \neq c \in [k]$, let $\widehat{\mathbf{w}}_\ell$, $\widehat{\mathbf{w}}_c$ be the ℓ and c -th row of $\widehat{\mathbf{W}}$ which is the solution to the multi-output least-squares minimization (G.1). Denote $\widehat{\mathbf{w}}_{\ell, c} := \widehat{\mathbf{w}}_\ell + \widehat{\mathbf{w}}_c$. Then, $\widehat{\mathbf{w}}_{\ell, c}$ is a minimizer in*

the following single-output least-squares problem:

$$\widehat{\mathbf{w}}_{\ell,c} = \arg \min_{\mathbf{w}, b} \frac{1}{2n} \left\| \mathbf{Y}_\ell + \mathbf{Y}_c - \mathbf{X}^T \mathbf{w} - b \mathbf{1}_n \right\|_{\ell_2}^2.$$

Thanks to Lemma G.1, we can use the CGMT to characterize the limiting behavior of $\|\widehat{\mathbf{w}}_\ell + \widehat{\mathbf{w}}_c\|_{\ell_2}$. These calculations are similar to (but, in certain cases, such as for weighted least-squares, more involved than) those in Steps II and III above. Now note that an asymptotic characterization of $\|\widehat{\mathbf{w}}_\ell + \widehat{\mathbf{w}}_c\|_{\ell_2}$ immediately yields the asymptotic characterization of $\langle \widehat{\mathbf{w}}_\ell, \widehat{\mathbf{w}}_c \rangle$ as

$$\langle \widehat{\mathbf{w}}_\ell, \widehat{\mathbf{w}}_c \rangle = \frac{\|\widehat{\mathbf{w}}_\ell + \widehat{\mathbf{w}}_c\|_{\ell_2}^2 - \|\widehat{\mathbf{w}}_\ell\|_{\ell_2}^2 - \|\widehat{\mathbf{w}}_c\|_{\ell_2}^2}{2}, \quad (\text{G.6})$$

and $\|\widehat{\mathbf{w}}_\ell\|_{\ell_2}, \|\widehat{\mathbf{w}}_c\|_{\ell_2}$ are already computed in Step IV (cf. (G.5)). For the GMM, the analysis in this step allow us to calculate the asymptotic behavior of $\Sigma_{w,w}$ as promised in (3.3b) in Theorem 3.2:

$$\Sigma_{w,w} \xrightarrow{P} \frac{\gamma}{(1-\gamma)\sigma^2} \mathbf{P} + \mathbf{P} \mathbf{V} \Sigma \Delta^{-1} \left(\Delta^{-1} - \frac{\gamma}{(1-\gamma)\sigma^2} \mathbf{I}_r \right) \Sigma \mathbf{V}^T \mathbf{P}.$$

G.1 Background on the CGMT

The CGMT is an extension of Gordon's Gaussian min-max inequality (GMT) [Gor88]. In the context of high-dimensional inference problems, Gordon's inequality was first successfully used in the study of sharp phase-transitions in noiseless Compressed Sensing [Sto09, CRPW12, ALMT13, Sto09]. More recently, [Sto13] (see also [ALMT13, Sec. 10.3]) discovered that Gordon's inequality is essentially tight for certain convex problems. A concrete and general formulation of this idea was given by [TOH15] and was called the CGMT.

In order to summarize the essential ideas, consider the following two Gaussian processes:

$$X_{\mathbf{w},\mathbf{u}} := \mathbf{u}^T \mathbf{G} \mathbf{w} + \psi(\mathbf{w}, \mathbf{u}), \quad (\text{G.7a})$$

$$Y_{\mathbf{w},\mathbf{u}} := \|\mathbf{w}\|_{\ell_2} \mathbf{g}^T \mathbf{u} + \|\mathbf{u}\|_{\ell_2} \mathbf{h}^T \mathbf{w} + \psi(\mathbf{w}, \mathbf{u}), \quad (\text{G.7b})$$

where: $\mathbf{G} \in \mathbb{R}^{n \times d}$, $\mathbf{g} \in \mathbb{R}^n$, $\mathbf{h} \in \mathbb{R}^d$, they all have entries iid Gaussian; the sets $\mathcal{S}_w \subset \mathbb{R}^d$ and $\mathcal{S}_u \subset \mathbb{R}^n$ are compact; and, $\psi : \mathbb{R}^d \times \mathbb{R}^n \rightarrow \mathbb{R}$. For these two processes, define the following (random) min-max optimization programs, which are referred to as the *primary optimization* (PO) problem and the *auxiliary optimization* AO:

$$\Phi(\mathbf{G}) = \min_{\mathbf{w} \in \mathcal{S}_w} \max_{\mathbf{u} \in \mathcal{S}_u} X_{\mathbf{w},\mathbf{u}}, \quad (\text{G.8a})$$

$$\phi(\mathbf{g}, \mathbf{h}) = \min_{\mathbf{w} \in \mathcal{S}_w} \max_{\mathbf{u} \in \mathcal{S}_u} Y_{\mathbf{w},\mathbf{u}}. \quad (\text{G.8b})$$

If the sets \mathcal{S}_w and \mathcal{S}_u are convex and *bounded*, and ψ is continuous *convex-concave* on $\mathcal{S}_w \times \mathcal{S}_u$, then, for any $\nu \in \mathbb{R}$ and $t > 0$, it holds [TOH15, Thm. 3]:

$$\mathbb{P}(|\Phi(\mathbf{G}) - \nu| > t) \leq 2 \mathbb{P}(|\phi(\mathbf{g}, \mathbf{h}) - \nu| > t). \quad (\text{G.9})$$

In words, concentration of the optimal cost of the AO problem around q^* implies concentration of the optimal cost of the corresponding PO problem around the same value q^* . Asymptotically, if we can show that $\phi(\mathbf{g}, \mathbf{h}) \xrightarrow{P} q^*$, then we can conclude that $\Phi(\mathbf{G}) \xrightarrow{P} q^*$. Moreover, starting from (G.9) and under appropriate strict convexity conditions, the CGMT shows that concentration of the optimal solution of the AO problem implies concentration of the optimal solution of the PO around the same value. For example, if minimizers of (G.8b) satisfy $\|\mathbf{w}_\phi(\mathbf{g}, \mathbf{h})\|_{\ell_2} \xrightarrow{P} \alpha^* > 0$, then, the same holds true for the minimizers of (G.8a): $\|\mathbf{w}_\Phi(\mathbf{G})\|_{\ell_2} \xrightarrow{P} \alpha^*$. Thus, one can analyze the AO to infer corresponding properties of the PO, the premise being of course that the former is simpler to handle than the latter.

In [TAH18], the authors introduce a principled machinery that allows to (a) express a quite general family of convex inference optimization problems in the form of the PO and (b) properly analyze the corresponding AO. In particular, the analysis of the AO is performed in three intermediate steps. First, the (random) optimization over vector variables is simplified to an easier optimization over only few scalar variables, termed the "scalarized AO". After the scalarization step, it is possible to establish (uniform) convergence of the scalarized AO to a deterministic min-max optimization problem over only a few scalar variables. The convergence step is followed by the analysis of the latter deterministic problem, which leads to the desired asymptotic characterizations.

H Least-squares for GMM

H.1 Proof of Theorem 3.2

H.1.1 Computing $\Sigma_{w,\mu}$

The LS classifier solves:

$$\begin{aligned} \min_{\mathbf{W} \in \mathbb{R}^{k \times d}, \mathbf{b} \in \mathbb{R}^k} \frac{1}{2n} \|\mathbf{W}\mathbf{X} + \mathbf{b}\mathbf{1}_n^T - \mathbf{Y}\|_F^2 &= \sum_{\ell=1}^k \min_{\mathbf{w}_\ell, b_\ell} \frac{1}{2n} \|\mathbf{X}^T \mathbf{w}_\ell + b_\ell \mathbf{1}_n - \mathbf{Y}_\ell\|_{\ell_2}^2 \\ &= \sum_{\ell=1}^k \min_{\mathbf{w}_\ell, b_\ell} \frac{1}{2n} \|\mathbf{Y}^T \mathbf{M}^T \mathbf{w}_\ell + \mathbf{Z}^T \mathbf{w}_\ell + b_\ell \mathbf{1}_n - \mathbf{Y}_\ell\|_{\ell_2}^2. \end{aligned}$$

Define

$$\mathcal{L}_{PO}(\mathbf{w}_\ell, b_\ell) := \frac{1}{2n} \|\mathbf{Y}^T \mathbf{M}^T \mathbf{w}_\ell + \mathbf{Z}^T \mathbf{w}_\ell + b_\ell \mathbf{1}_n - \mathbf{Y}_\ell\|_{\ell_2}^2. \quad (\text{H.1})$$

Identifying the AO. To continue further note that by duality we have

$$\min_{\mathbf{w}_\ell, b_\ell} \mathcal{L}_{PO}(\mathbf{w}_\ell, b_\ell) = \min_{\mathbf{w}_\ell, b_\ell} \max_{\mathbf{s}} \frac{1}{n} \left(\mathbf{s}^T \mathbf{Y}^T \mathbf{M}^T \mathbf{w}_\ell + \mathbf{s}^T \mathbf{Z}^T \mathbf{w}_\ell + b_\ell \mathbf{s}^T \mathbf{1}_n - \mathbf{s}^T \mathbf{Y}_\ell - \frac{\|\mathbf{s}\|_{\ell_2}^2}{2} \right).$$

Note that the above is jointly convex in $(\mathbf{w}_\ell, b_\ell)$ and concave in \mathbf{s} and the Gaussian matrix \mathbf{Z} is independent of everything else. Thus, the objective is in the form of (G.7a) and so we consider the corresponding Auxiliary Optimization (AO) problem:

$$\min_{\mathbf{w}_\ell, b_\ell} \max_{\mathbf{s}} \frac{1}{n} \left(\mathbf{s}^T \mathbf{Y}^T \mathbf{M}^T \mathbf{w}_\ell + \sigma \|\mathbf{w}_\ell\|_{\ell_2} \mathbf{g}^T \mathbf{s} + \sigma \|\mathbf{s}\|_{\ell_2} \mathbf{h}^T \mathbf{w}_\ell + b_\ell \mathbf{s}^T \mathbf{1}_n - \mathbf{s}^T \mathbf{Y}_\ell - \frac{\|\mathbf{s}\|_{\ell_2}^2}{2} \right),$$

where $\mathbf{g} \in \mathbb{R}^n$ and $\mathbf{h} \in \mathbb{R}^d$ are independent Gaussian random vectors with i.i.d. $\mathcal{N}(0, 1)$ entries. Maximizing over the direction of \mathbf{s} and setting its norm $\beta = \|\mathbf{s}\|_{\ell_2}$ we arrive at

$$\begin{aligned} \min_{\mathbf{w}_\ell, b_\ell} \max_{\beta \geq 0} \frac{1}{n} \left(\beta \left(\sigma \|\mathbf{w}_\ell\|_{\ell_2} \mathbf{g} + \mathbf{Y}^T \mathbf{M}^T \mathbf{w}_\ell + b_\ell \mathbf{1}_n - \mathbf{Y}_\ell \right)_{\ell_2} + \beta \sigma \mathbf{h}^T \mathbf{w}_\ell - \frac{\beta^2}{2} \right) \\ = \min_{\mathbf{w}_\ell, b_\ell} \frac{1}{2n} \left(\left(\sigma \|\mathbf{w}_\ell\|_{\ell_2} \mathbf{g} + \mathbf{Y}^T \mathbf{M}^T \mathbf{w}_\ell + b_\ell \mathbf{1}_n - \mathbf{Y}_\ell \right)_{\ell_2} + \sigma \mathbf{h}^T \mathbf{w}_\ell \right)_+^2 \\ = \frac{1}{2} \left(\min_{\mathbf{w}_\ell, b_\ell} \frac{1}{\sqrt{n}} \left(\sigma \|\mathbf{w}_\ell\|_{\ell_2} \mathbf{g} + \mathbf{Y}^T \mathbf{M}^T \mathbf{w}_\ell + b_\ell \mathbf{1}_n - \mathbf{Y}_\ell \right)_{\ell_2} + \frac{1}{\sqrt{n}} \sigma \mathbf{h}^T \mathbf{w}_\ell \right)_+^2 \end{aligned}$$

Scalarization of the AO. For convenience, define

$$\bar{\phi}_{AO,\ell} := \min_{\mathbf{w}_\ell, b_\ell} \frac{1}{\sqrt{n}} \left(\sigma \|\mathbf{w}_\ell\|_{\ell_2} \mathbf{g} + \mathbf{Y}^T \mathbf{M}^T \mathbf{w}_\ell + b_\ell \mathbf{1}_n - \mathbf{Y}_\ell \right)_{\ell_2} + \frac{\sigma}{\sqrt{n}} \mathbf{h}^T \mathbf{w}_\ell. \quad (\text{H.2})$$

To continue, consider the singular value decomposition

$$\mathbf{M} = \mathbf{U} \boldsymbol{\Sigma} \mathbf{V}^T = [\mathbf{u}_1 \quad \mathbf{u}_2 \quad \dots \quad \mathbf{u}_r] \text{diag}(\sigma_1, \sigma_2, \dots, \sigma_r) \begin{bmatrix} \mathbf{v}_1^T \\ \mathbf{v}_2^T \\ \vdots \\ \mathbf{v}_r^T \end{bmatrix}, \quad (\text{H.3})$$

with $\mathbf{U} \in \mathbb{R}^{d \times r}$, $\boldsymbol{\Sigma} \in \mathbb{R}^{r \times r}$, and $\mathbf{V} \in \mathbb{R}^{k \times r}$ where $r = \text{rank}(\mathbf{M}) \leq k$. We further decompose \mathbf{w}_ℓ in its projections on the orthogonal columns $\mathbf{u}_1, \dots, \mathbf{u}_r$ of \mathbf{U} :

$$\mathbf{w}_\ell = \sum_{i=1}^r \alpha_i \mathbf{u}_i + \alpha_0 \mathbf{w}_\ell^\perp,$$

where $\|\mathbf{w}_\ell^\perp\|_{\ell_2} = 1$ and $\mathbf{U}^T \mathbf{w}_\ell^\perp = \mathbf{0}$, $\alpha_0 \geq 0$ and we denote

$$\alpha_i := \mathbf{u}_i^T \mathbf{w}_\ell, i \in [r]. \quad (\text{H.4})$$

We also define $\boldsymbol{\alpha} = [\alpha_1 \ \alpha_2 \ \dots \ \alpha_k]^T$. In this notation, we have

$$\begin{aligned} \bar{\phi}_{AO,\ell}(\mathbf{g}, \mathbf{h}) &:= \min_{\alpha_0 \geq 0, \boldsymbol{\alpha} \in \mathbb{R}^r, b_\ell} \frac{1}{\sqrt{n}} \left\| \sigma \sqrt{\alpha_0^2 + \|\boldsymbol{\alpha}\|_{\ell_2}^2} \mathbf{g} + \mathbf{Y}^T \mathbf{V} \boldsymbol{\Sigma} \boldsymbol{\alpha} + b_\ell \mathbf{1}_n - \mathbf{Y}_\ell \right\|_{\ell_2} \\ &\quad + \sum_{i=1}^r \alpha_i \sigma \frac{\mathbf{h}^T \mathbf{u}_i}{\sqrt{n}} + \frac{\alpha_0 \sigma}{\sqrt{n}} \min_{\mathbf{w}_\ell^\perp} (\mathbf{h}^T \mathbf{w}_\ell^\perp) \\ &= \min_{\alpha_0 \geq 0, \boldsymbol{\alpha} \in \mathbb{R}^r, b_\ell} \frac{1}{\sqrt{n}} \left\| \sigma \sqrt{\alpha_0^2 + \|\boldsymbol{\alpha}\|_{\ell_2}^2} \mathbf{g} + \mathbf{Y}^T \mathbf{V} \boldsymbol{\Sigma} \boldsymbol{\alpha} + b_\ell \mathbf{1}_n - \mathbf{Y}_\ell \right\|_{\ell_2} \\ &\quad + \sum_{i=1}^r \alpha_i \sigma \frac{\mathbf{h}^T \mathbf{u}_i}{\sqrt{n}} - \alpha_0 \sigma \frac{\|\mathbf{h}^\perp\|_{\ell_2}}{\sqrt{n}}, \end{aligned} \quad (\text{H.5})$$

where in the second line we denote \mathbf{h}^\perp the projection of \mathbf{h} onto the complement subspace of the span of $\mathbf{u}_1, \dots, \mathbf{u}_r$ and we recalled that $\|\mathbf{w}_\ell^\perp\|_{\ell_2} = 1$ and $\langle \mathbf{w}_\ell^\perp, \mathbf{u}_i \rangle = 0, i \in [r]$.

Convergence of the AO. First, note that

$$\begin{aligned} \frac{1}{n} \left\| \mathbf{Y}^T \mathbf{V} \boldsymbol{\Sigma} \boldsymbol{\alpha} + b_\ell \mathbf{1}_n - \mathbf{Y}_\ell \right\|_{\ell_2}^2 &= \frac{1}{n} \left\| \mathbf{Y}^T (\mathbf{V} \boldsymbol{\Sigma} \boldsymbol{\alpha} - \mathbf{e}_\ell) + b_\ell \mathbf{1}_n \right\|_{\ell_2}^2 \\ &= \frac{1}{n} \left\| \mathbf{Y}^T (\mathbf{V} \boldsymbol{\Sigma} \boldsymbol{\alpha} - \mathbf{e}_\ell) \right\|_{\ell_2}^2 + b_\ell^2 + \frac{2}{n} b_\ell \mathbf{1}_n^T \mathbf{Y}^T (\mathbf{V} \boldsymbol{\Sigma} \boldsymbol{\alpha} - \mathbf{e}_\ell) \\ &= \text{trace} \left((\mathbf{V} \boldsymbol{\Sigma} \boldsymbol{\alpha} - \mathbf{e}_\ell)^T \text{diag} \left(\frac{n_1}{n}, \frac{n_2}{n}, \dots, \frac{n_k}{n} \right) (\mathbf{V} \boldsymbol{\Sigma} \boldsymbol{\alpha} - \mathbf{e}_\ell) \right) \\ &\quad + b_\ell^2 + 2b_\ell \left[\frac{n_1}{n} \quad \frac{n_2}{n} \quad \dots \quad \frac{n_k}{n} \right] (\mathbf{V} \boldsymbol{\Sigma} \boldsymbol{\alpha} - \mathbf{e}_\ell) \end{aligned}$$

Thus

$$\begin{aligned} \frac{1}{n} \left\| \mathbf{Y}^T \mathbf{V} \boldsymbol{\Sigma} \boldsymbol{\alpha} + b_\ell \mathbf{1}_n - \mathbf{Y}_\ell \right\|_{\ell_2}^2 &\xrightarrow{P} \text{trace} \left((\mathbf{V} \boldsymbol{\Sigma} \boldsymbol{\alpha} - \mathbf{e}_\ell)^T \text{diag}(\boldsymbol{\pi}) (\mathbf{V} \boldsymbol{\Sigma} \boldsymbol{\alpha} - \mathbf{e}_\ell) \right) \\ &\quad + b_\ell^2 + 2b_\ell \boldsymbol{\pi}^T (\mathbf{V} \boldsymbol{\Sigma} \boldsymbol{\alpha} - \mathbf{e}_\ell) \\ &= \boldsymbol{\alpha}^T (\boldsymbol{\Sigma} \mathbf{V}^T \text{diag}(\boldsymbol{\pi}) \boldsymbol{\Sigma}) \boldsymbol{\alpha} - 2\pi_\ell \boldsymbol{\alpha}^T \boldsymbol{\Sigma} \mathbf{V}^T \mathbf{e}_\ell + 2b_\ell \boldsymbol{\alpha}^T \boldsymbol{\Sigma} \mathbf{V}^T \boldsymbol{\pi} \\ &\quad + b_\ell^2 - 2b_\ell \pi_\ell + \pi_\ell \end{aligned}$$

At this point, observe that we have reduced the AO to an optimization problem over only $r + 2$ scalar variables. Using the law of large numbers, the fact that $\|\mathbf{h}^\perp\|_{\ell_2}$ concentrates around $\sqrt{d-r}$ and $(d-r)/n \xrightarrow{P} \gamma$, as well as the limit calculation above, it is not hard to see that for fixed α_0, b_ℓ and $\boldsymbol{\alpha} = [\alpha_1, \dots, \alpha_r]^T \in \mathbb{R}^r$, the objective function in (H.5) converges to the following:

$$\begin{aligned} \mathcal{D}_\ell(\alpha_0, \boldsymbol{\alpha}, b_\ell) &:= \sqrt{\alpha_0^2 \sigma^2 + \boldsymbol{\alpha}^T (\sigma^2 \mathbf{I}_r + \boldsymbol{\Sigma} \mathbf{V}^T \text{diag}(\boldsymbol{\pi}) \boldsymbol{\Sigma}) \boldsymbol{\alpha} - 2\boldsymbol{\alpha}^T (\pi_\ell \boldsymbol{\Sigma} \mathbf{V}^T \mathbf{e}_\ell - b_\ell \boldsymbol{\Sigma} \mathbf{V}^T \boldsymbol{\pi}) + b_\ell^2 - 2b_\ell \pi_\ell + \pi_\ell} \\ &\quad - \alpha_0 \sigma \sqrt{\gamma}, \end{aligned} \quad (\text{H.6})$$

We will show in the next paragraph that the argument inside the square-root in (H.6) is a convex quadratic over $(\alpha_0, \boldsymbol{\alpha}, b_\ell)$ (see (H.19)). Thus, the function $\mathcal{D}_\ell(\alpha_0, \boldsymbol{\alpha}, b_\ell)$ is jointly convex. Using uniform convergence of convex functions over compact sets [AG82, Cor.. II.1], we arrive at

$$\bar{\phi}_{AO,\ell}(\mathbf{g}, \mathbf{h}) \xrightarrow{P} \min_{\alpha_0 \geq 0, \boldsymbol{\alpha}, b_\ell} \mathcal{D}_\ell(\alpha_0, \boldsymbol{\alpha}, b_\ell). \quad (\text{H.7})$$

Deterministic Analysis. Here, we analyze the deterministic scalar minimization on the RHS of (H.7). Define

$$\mathbf{A} := \begin{bmatrix} \sigma^2 \mathbf{I}_r + \boldsymbol{\Sigma} \mathbf{V}^T \text{diag}(\boldsymbol{\pi}) \boldsymbol{\Sigma} & \boldsymbol{\Sigma} \mathbf{V}^T \boldsymbol{\pi} \\ \boldsymbol{\pi}^T \boldsymbol{\Sigma} \boldsymbol{\Sigma} & 1 \end{bmatrix} \quad \text{and} \quad \mathbf{c}_\ell = \begin{bmatrix} \boldsymbol{\Sigma} \mathbf{V}^T \mathbf{e}_\ell \\ 1 \end{bmatrix}, \quad (\text{H.8})$$

and observe that we can write

$$\mathcal{D}_\ell(\alpha_0, \boldsymbol{\alpha}, b_\ell) = \sqrt{\alpha_0^2 \sigma^2 + \pi_\ell + [\boldsymbol{\alpha}^T \ b_\ell] \mathbf{A} \begin{bmatrix} \boldsymbol{\alpha} \\ b_\ell \end{bmatrix} - 2\pi_\ell \mathbf{c}_\ell^T \begin{bmatrix} \boldsymbol{\alpha} \\ b_\ell \end{bmatrix} - \alpha_0 \sigma \sqrt{\gamma}}. \quad (\text{H.9})$$

First, note that the matrix \mathbf{A} is positive definite. This can be checked by computing the Schur complement of \mathbf{A} :

$$\Delta := \sigma^2 \mathbf{I}_r + \Sigma \mathbf{V}^T \mathbf{P} \mathbf{V} \Sigma := \sigma^2 \mathbf{I}_r + \Sigma \mathbf{V}^T (\text{diag}(\boldsymbol{\pi}) - \boldsymbol{\pi} \boldsymbol{\pi}^T) \mathbf{V} \Sigma > \mathbf{0}_{r \times r}. \quad (\text{H.10})$$

Positive definiteness above holds because $\mathbf{P} := (\text{diag}(\boldsymbol{\pi}) - \boldsymbol{\pi} \boldsymbol{\pi}^T) \succeq \mathbf{0}_{k \times k}$. Thus the term under the square-root in (H.9) is a strictly convex quadratic. Thus, \mathcal{D}_ℓ is jointly convex in its arguments.

To simplify the RHS of (H.7) we proceed by minimizing $\mathcal{D}_\ell(\alpha_0, \boldsymbol{\alpha}, b_\ell)$ over $(\boldsymbol{\alpha}, b_\ell)$ which from Lemma (C.3) is equal to

$$\begin{aligned} \begin{bmatrix} \widehat{\boldsymbol{\alpha}} \\ \widehat{b}_\ell \end{bmatrix} &= \pi_\ell \mathbf{A}^{-1} \mathbf{c}_\ell = \pi_\ell \begin{bmatrix} \mathbf{I} & \mathbf{0} \\ -\boldsymbol{\pi}^T \mathbf{V} \Sigma & 1 \end{bmatrix} \begin{bmatrix} \Delta^{-1} & \mathbf{0} \\ \mathbf{0}^T & 1 \end{bmatrix} \begin{bmatrix} \mathbf{I} & -\Sigma \mathbf{V}^T \boldsymbol{\pi} \\ \mathbf{0}^T & 1 \end{bmatrix} \begin{bmatrix} \Sigma \mathbf{V}^T \mathbf{e}_\ell \\ 1 \end{bmatrix} \\ &= \pi_\ell \begin{bmatrix} \mathbf{I} & \mathbf{0} \\ -\boldsymbol{\pi}^T \mathbf{V} \Sigma & 1 \end{bmatrix} \begin{bmatrix} \Delta^{-1} & \mathbf{0} \\ \mathbf{0}^T & 1 \end{bmatrix} \begin{bmatrix} -\Sigma \mathbf{V}^T (\boldsymbol{\pi} - \mathbf{e}_\ell) \\ 1 \end{bmatrix} \\ &= \pi_\ell \begin{bmatrix} -\Delta^{-1} \Sigma \mathbf{V}^T (\boldsymbol{\pi} - \mathbf{e}_\ell) \\ 1 + \pi^T \mathbf{V} \Sigma \Delta^{-1} \Sigma \mathbf{V}^T (\boldsymbol{\pi} - \mathbf{e}_\ell) \end{bmatrix}. \end{aligned} \quad (\text{H.11})$$

Thus, the minimum value attained is

$$-\pi_\ell^2 \begin{bmatrix} -(\boldsymbol{\pi} - \mathbf{e}_\ell)^T \mathbf{V} \Sigma & 1 \end{bmatrix} \begin{bmatrix} \Delta^{-1} & \mathbf{0} \\ \mathbf{0}^T & 1 \end{bmatrix} \begin{bmatrix} -\Sigma \mathbf{V}^T (\boldsymbol{\pi} - \mathbf{e}_\ell) \\ 1 \end{bmatrix} = -\pi_\ell^2 \left(1 + (\boldsymbol{\pi} - \mathbf{e}_\ell)^T \mathbf{V} \Sigma \Delta^{-1} \Sigma \mathbf{V}^T (\boldsymbol{\pi} - \mathbf{e}_\ell) \right).$$

Using the above, (H.7) reduces to

$$\bar{\phi}_{AO,\ell}(\mathbf{g}, \mathbf{h}) \xrightarrow{P} \min_{\alpha_0 \geq 0} \sqrt{\alpha_0^2 \sigma^2 + \pi_\ell - \pi_\ell^2 \left(1 + (\boldsymbol{\pi} - \mathbf{e}_\ell)^T \mathbf{V} \Sigma \Delta^{-1} \Sigma \mathbf{V}^T (\boldsymbol{\pi} - \mathbf{e}_\ell) \right)} - \alpha_0 \sigma \sqrt{\gamma}. \quad (\text{H.12})$$

Setting the derivative with respect to α_0 to zero we arrive at

$$\frac{\alpha_0 \sigma^2}{\sqrt{\alpha_0^2 \sigma^2 + \pi_\ell - \pi_\ell^2 \left(1 + (\boldsymbol{\pi} - \mathbf{e}_\ell)^T \mathbf{V} \Sigma \Delta^{-1} \Sigma \mathbf{V}^T (\boldsymbol{\pi} - \mathbf{e}_\ell) \right)}} = \sigma \sqrt{\gamma}.$$

Thus,

$$\widehat{\alpha}_0 = \frac{1}{\sigma} \sqrt{\frac{\gamma}{1 - \gamma}} \sqrt{\pi_\ell (1 - \pi_\ell) - \pi_\ell^2 (\boldsymbol{\pi} - \mathbf{e}_\ell)^T \mathbf{V} \Sigma \Delta^{-1} \Sigma \mathbf{V}^T (\boldsymbol{\pi} - \mathbf{e}_\ell)}. \quad (\text{H.13})$$

Plugging the latter into (H.12) we arrive at

$$\bar{\phi}_{AO,\ell}(\mathbf{g}, \mathbf{h}) \xrightarrow{P} \sqrt{1 - \gamma} \sqrt{\pi_\ell (1 - \pi_\ell) - \pi_\ell^2 (\boldsymbol{\pi} - \mathbf{e}_\ell)^T \mathbf{V} \Sigma \Delta^{-1} \Sigma \mathbf{V}^T (\boldsymbol{\pi} - \mathbf{e}_\ell)}.$$

Asymptotic predictions. First, from (H.11) the bias term converges as follows:

$$\widehat{b}_\ell \xrightarrow{P} \pi_\ell \left(1 + \boldsymbol{\pi}^T \mathbf{V} \Sigma \Delta^{-1} \Sigma \mathbf{V}^T (\boldsymbol{\pi} - \mathbf{e}_\ell) \right).$$

Thus,

$$\widehat{\mathbf{b}} \xrightarrow{P} \text{diag}(\boldsymbol{\pi}) \left(\mathbf{1}_k + (\boldsymbol{\pi} \mathbf{1}_k^T - \mathbf{I}_k) \mathbf{V} \Sigma \Delta^{-1} \Sigma \mathbf{V}^T \boldsymbol{\pi} \right).$$

Recall from (H.4) that $\boldsymbol{\alpha} = \mathbf{U}^T \mathbf{w}_\ell$. Thus, the correlations $\langle \boldsymbol{\mu}_i, \mathbf{w}_\ell \rangle$, $i \in [k]$ converge as follows:

$$\mathbf{M}^T \mathbf{w}_\ell = \mathbf{V} \Sigma \mathbf{U}^T \mathbf{w}_\ell \xrightarrow{P} \mathbf{V} \Sigma \widehat{\boldsymbol{\alpha}} = -\pi_\ell \mathbf{V} \Sigma \Delta^{-1} \Sigma \mathbf{V}^T (\boldsymbol{\pi} - \mathbf{e}_\ell). \quad (\text{H.14})$$

Here, convergence applies element-wise to the entries of the involved random vectors. Moreover, from the analysis above we can predict the limit of the norm $\|\mathbf{w}_\ell\|_{\ell_2}$. For this, note that $\|\mathbf{w}_\ell\|_{\ell_2}^2 = \widehat{\alpha}_0^2 + \widehat{\boldsymbol{\alpha}}^T \widehat{\boldsymbol{\alpha}}$. Thus,

$$\|\mathbf{w}_\ell\|_{\ell_2}^2 \xrightarrow{P} \frac{\gamma}{(1 - \gamma) \sigma^2} \pi_\ell (1 - \pi_\ell) + \pi_\ell^2 (\boldsymbol{\pi} - \mathbf{e}_\ell)^T \mathbf{V} \Sigma \Delta^{-1} \left(\Delta^{-1} - \frac{\gamma}{(1 - \gamma) \sigma^2} \mathbf{I}_r \right) \Sigma \mathbf{V}^T (\boldsymbol{\pi} - \mathbf{e}_\ell). \quad (\text{H.15})$$

H.1.2 Computing $\Sigma_{w,w}$

In the previous section we used the CGMT to predict the bias \widehat{b}_ℓ , the correlations $\langle \mu_i, \widehat{w}_\ell \rangle$, $i[k]$ and the norm $\|\widehat{w}_\ell\|_{\ell_2}$ for all $\ell \in [k]$ members of the multi-output classifier. Here, we show how to compute the limits of the cross-correlations $\langle \widehat{w}_\ell, \widehat{w}_j \rangle, \ell \neq j \in [k]$.

Lemma H.1 *For $\ell \neq j \in [k]$, let \widehat{w}_ℓ \widehat{w}_j be solutions to the least-squares minimization*

$$(\widehat{w}_\ell, \widehat{w}_j, \widehat{b}_\ell, \widehat{b}_j) = \arg \min_{\mathbf{w}_\ell, \mathbf{w}_j, b_\ell, b_j} \left\{ \frac{1}{2n} \|\mathbf{Y}_\ell - \mathbf{X}^T \mathbf{w}_\ell - b_\ell \mathbf{1}_n\|_{\ell_2}^2 + \frac{1}{2n} \|\mathbf{Y}_j - \mathbf{X}^T \mathbf{w}_j - b_j \mathbf{1}_n\|_{\ell_2}^2 \right\}.$$

Denote $\widehat{w}_{\ell,j} := \widehat{w}_\ell + \widehat{w}_j$ and $\widehat{b}_{\ell,j} := \widehat{b}_\ell + \widehat{b}_j$. Then, $(\widehat{w}_{\ell,j}, \widehat{b}_{\ell,j})$ is a minimizer in the following least-squares problem:

$$(\widehat{w}_{\ell,j}, \widehat{b}_{\ell,j}) = \arg \min_{\mathbf{w}, b} \frac{1}{2n} \|\mathbf{Y}_\ell + \mathbf{Y}_j - \mathbf{X}^T \mathbf{w} - b \mathbf{1}_n\|_{\ell_2}^2. \quad (\text{H.16})$$

Proof Clearly the minimization in (H.16) is convex. Thus, it suffices to prove that $\widehat{w}_\ell + \widehat{w}_j$ satisfies the KKT conditions. First, by optimality of \widehat{w}_ℓ , we have that

$$\mathbf{X} (\mathbf{Y}_\ell - \mathbf{X}^T \widehat{w}_\ell - \widehat{b}_\ell \mathbf{1}_n) = 0$$

Similarly, for \widehat{w}_j :

$$\mathbf{X} (\mathbf{Y}_j - \mathbf{X}^T \widehat{w}_j - \widehat{b}_j \mathbf{1}_n) = 0.$$

Adding the equations on the above displays we find that

$$\mathbf{X} (\mathbf{Y}_\ell + \mathbf{Y}_j - \mathbf{X}^T (\widehat{w}_j + \widehat{w}_\ell) - (\widehat{b}_j + \widehat{b}_\ell) \mathbf{1}_n) = 0.$$

Recognize that this coincides with the optimality condition for (H.16). Thus, the proof is complete. ■

Thanks to Lemma H.1, we can use the CGMT to characterize the limiting behavior of $\|\widehat{w}_\ell + \widehat{w}_j\|_{\ell_2}$. Observe that this immediately gives the limit of $\langle \widehat{w}_\ell, \widehat{w}_j \rangle$ since

$$\langle \widehat{w}_\ell, \widehat{w}_j \rangle = \frac{\|\widehat{w}_\ell + \widehat{w}_j\|_{\ell_2}^2 - \|\widehat{w}_\ell\|_{\ell_2}^2 - \|\widehat{w}_j\|_{\ell_2}^2}{2}. \quad (\text{H.17})$$

The analysis of (H.16) is very similar to that of (H.1); thus, most details are omitted. Similar to (H.5) we can relate (H.16) with the following AO problem:

$$\begin{aligned} \bar{\phi}_{AO, \ell, j}(\mathbf{g}, \mathbf{h}) := & \min_{\beta_0 \geq 0, \beta \in \mathbb{R}^r, b_{\ell,j}} \frac{1}{\sqrt{n}} \left\| \sigma \sqrt{\beta_0^2 + \|\beta\|_{\ell_2}^2} \mathbf{g} + \mathbf{Y}^T \mathbf{V} \Sigma \beta + b_{\ell,j} \mathbf{1}_n - \mathbf{Y}_\ell - \mathbf{Y}_j \right\|_{\ell_2} \\ & + \sigma \sum_{i=1}^r \beta_i \frac{\mathbf{h}^T \mathbf{u}_i}{\sqrt{n}} - \sigma \beta_0 \frac{\|\mathbf{h}^\perp\|_{\ell_2}}{\sqrt{n}}, \end{aligned} \quad (\text{H.18})$$

where we have decomposed

$$\mathbf{w}_{\ell,j} = \sum_{i=1}^r \beta_i \mathbf{u}_i + \beta_0 \mathbf{w}_{\ell,j}^\perp,$$

with $\|\mathbf{w}_{\ell,j}^\perp\|_{\ell_2} = 1$ and $\mathbf{U}^T \mathbf{w}_{\ell,j}^\perp = \mathbf{0}_r$.

Using a calculation similar to the one leading to (H.9) we can show that (H.18) converges point-wise in $\beta_0, \beta = [\beta_1, \dots, \beta_r], b_{\ell,j}$ to the following:

$$\mathcal{D}_\ell(\beta_0, \beta, b_{\ell,j}) = \sqrt{\beta_0^2 \sigma^2 + \pi_\ell + \pi_j + [\beta^T \ b_{\ell,j}] \mathbf{A} \begin{bmatrix} \beta \\ b_{\ell,j} \end{bmatrix} - 2 \mathbf{d}_{\ell,j}^T \begin{bmatrix} \beta \\ b_{\ell,j} \end{bmatrix} - \beta_0 \sigma \sqrt{\gamma}}, \quad (\text{H.19})$$

where \mathbf{A} is as in (H.8) and we have further defined

$$\mathbf{d}_{\ell,j} := \begin{bmatrix} \pi_\ell \Sigma \mathbf{V}^T \mathbf{e}_\ell + \pi_j \Sigma \mathbf{V}^T \mathbf{e}_j \\ \pi_\ell + \pi_j \end{bmatrix}.$$

Thus, similar to (H.11) we can compute the minimizer of the deterministic objective in (H.19):

$$\begin{bmatrix} \widehat{\boldsymbol{\beta}} \\ \widehat{b}_{\ell,j} \end{bmatrix} = \begin{bmatrix} -\boldsymbol{\Delta}^{-1} \boldsymbol{\Sigma} \boldsymbol{V}^T (\pi_\ell (\boldsymbol{\pi} - \boldsymbol{e}_\ell) + \pi_j (\boldsymbol{\pi} - \boldsymbol{e}_j)) \\ \pi_\ell + \pi_j + \boldsymbol{\pi}^T \boldsymbol{V} \boldsymbol{\Sigma} \boldsymbol{\Delta}^{-1} \boldsymbol{\Sigma} \boldsymbol{V}^T (\pi_\ell (\boldsymbol{\pi} - \boldsymbol{e}_\ell) + \pi_j (\boldsymbol{\pi} - \boldsymbol{e}_j)) \end{bmatrix}, \quad (\text{H.20})$$

and

$$\widehat{\beta}_0 = \frac{1}{\sigma} \sqrt{\frac{\gamma}{1-\gamma}} \sqrt{\pi_\ell + \pi_j - (\pi_\ell + \pi_j)^2 - (\pi_\ell (\boldsymbol{\pi} - \boldsymbol{e}_\ell) + \pi_j (\boldsymbol{\pi} - \boldsymbol{e}_j))^T \boldsymbol{V} \boldsymbol{\Sigma} \boldsymbol{\Delta}^{-1} \boldsymbol{\Sigma} \boldsymbol{V}^T (\pi_\ell (\boldsymbol{\pi} - \boldsymbol{e}_\ell) + \pi_j (\boldsymbol{\pi} - \boldsymbol{e}_j))}, \quad (\text{H.21})$$

where recall that $\boldsymbol{\Delta}$ is as in (H.10).

From the CGMT, we have that $\|\widehat{\boldsymbol{w}}_\ell + \widehat{\boldsymbol{w}}_j\|_{\ell_2}^2 \xrightarrow{P} \widehat{\beta}_0^2 + \|\boldsymbol{\beta}\|_{\ell_2}^2$. Combining this with the calculations above, we conclude that

$$\begin{aligned} \|\widehat{\boldsymbol{w}}_\ell + \widehat{\boldsymbol{w}}_j\|_{\ell_2}^2 &\xrightarrow{P} \frac{\gamma}{(1-\gamma)\sigma^2} (\pi_\ell + \pi_j) (1 - \pi_\ell - \pi_j) \\ &+ (\pi_\ell (\boldsymbol{\pi} - \boldsymbol{e}_\ell) + \pi_j (\boldsymbol{\pi} - \boldsymbol{e}_j))^T \boldsymbol{V} \boldsymbol{\Sigma} \boldsymbol{\Delta}^{-1} \left(\boldsymbol{\Delta}^{-1} - \frac{\gamma}{(1-\gamma)\sigma^2} \boldsymbol{I}_r \right) \boldsymbol{\Sigma} \boldsymbol{V}^T (\pi_\ell (\boldsymbol{\pi} - \boldsymbol{e}_\ell) + \pi_j (\boldsymbol{\pi} - \boldsymbol{e}_j)) \end{aligned} \quad (\text{H.22})$$

Finally, using (H.22) and (H.15) in (H.17) it follows that

$$\langle \boldsymbol{w}_\ell, \boldsymbol{w}_j \rangle \xrightarrow{P} \pi_\ell \pi_j \left(-\frac{\gamma}{(1-\gamma)\sigma^2} + (\boldsymbol{\pi} - \boldsymbol{e}_\ell)^T \boldsymbol{V} \boldsymbol{\Sigma} \boldsymbol{\Delta}^{-1} \left(\boldsymbol{\Delta}^{-1} - \frac{\gamma}{(1-\gamma)\sigma^2} \boldsymbol{I}_r \right) \boldsymbol{\Sigma} \boldsymbol{V}^T (\boldsymbol{\pi} - \boldsymbol{e}_j) \right). \quad (\text{H.23})$$

H.2 Orthogonal means

Here, we specialize the asymptotic predictions of Theorem 3.2 to the case of orthogonal means $\langle \boldsymbol{\mu}_i, \boldsymbol{\mu}_j \rangle = 0$, $i \neq j$.

Corollary H.2 (Orthogonal means) *Consider the case of orthogonal means, i.e. $\langle \boldsymbol{\mu}_i, \boldsymbol{\mu}_j \rangle = 0$, $\forall i \neq j$ and $\gamma < 1$ with Euclidean norms given by $\mu_i = \|\boldsymbol{\mu}_i\|_{\ell_2}$. Define the following parameters for $i \in [k]$:*

$$\rho_i := \pi_i \sigma^2 / (\sigma^2 + \pi_i \mu_i^2) \quad \text{and} \quad \beta_i = \rho_i \sigma^2 / (\sigma^2 - \sum_{i=1}^k \pi_i \rho_i \mu_i^2).$$

Then, the following asymptotic limits hold for the least-squares classifier, for all $i, j \in [k]$:

$$\widehat{b}_i \xrightarrow{P} \beta_i, \quad \langle \widehat{\boldsymbol{w}}_i, \boldsymbol{\mu}_j \rangle \xrightarrow{P} \frac{1}{\sigma} (\mathbb{1}_{ij} - \beta_i) \rho_j \mu_j, \quad (\text{H.24a})$$

$$\langle \widehat{\boldsymbol{w}}_i, \widehat{\boldsymbol{w}}_j \rangle \xrightarrow{P} \frac{1}{\sigma^4} \beta_i \beta_j \sum_{\ell=1}^k \rho_\ell^2 \mu_\ell^2 - \frac{1}{\sigma^4} \beta_i \rho_j^2 \mu_j^2 - \frac{1}{\sigma^4} \beta_j \rho_i^2 \mu_i^2 - \frac{\gamma \beta_i \rho_j}{(1-\gamma)\sigma^2} + \frac{\mathbb{1}_{ij}}{\sigma^2} \left(\frac{\gamma}{(1-\gamma)} \rho_i + \frac{1}{\sigma^2} \rho_i^2 \mu_i^2 \right) \quad (\text{H.24b})$$

Furthermore, if the means have equal norms $\mu := \mu_i$ and the classes are balanced: $\pi_i = 1/k$, $i \in [k]$, then, setting $u_{\text{LS}} := \frac{\mu^2}{\sigma} \sqrt{\frac{1-\gamma}{\mu^2 + k\gamma\sigma^2}}$, it holds that

$$\mathbb{P}_e = \mathbb{P} \left\{ G_0 + \max_{j \in [k-1]} G_j \geq u_{\text{LS}} \right\}, \quad G_0, G_1, \dots, G_{k-1} \stackrel{iid}{\sim} \mathcal{N}(0, 1). \quad (\text{H.25})$$

Proof This is a direct corollary of Theorem 3.2. Indeed, (H.24) can be derived from (3.3) after substituting $\boldsymbol{V} = \boldsymbol{I}_k$, $\boldsymbol{\Sigma} = \text{diag}(\mu_1, \mu_2, \dots, \mu_k)$ and some algebra steps that we omit for brevity.

Instead, we outline below how to conclude (H.25) from (H.24). Assume that $\mu_i = \mu$, $\forall i \in [k]$ and $\pi_i = \pi = 1/k$, $\forall i \in [k]$. Recall from (2.7) that $\mathbb{P}(\text{error} | \boldsymbol{y} = \boldsymbol{e}_c) = 1 - \mathbb{P}(\boldsymbol{S}_c^{1/2} \boldsymbol{z} > \boldsymbol{t})$, and using (H.24) it can be checked that

$$\boldsymbol{S}_c = \frac{\pi}{1 + \pi \mu^2} \left(\frac{\pi \mu^2}{1 + \pi \mu^2} + \frac{\gamma}{1 - \gamma} \right) (\boldsymbol{I}_k + \mathbf{1}_k \mathbf{1}_k^T) \quad \text{and} \quad \boldsymbol{t} = -\frac{\pi \mu^2}{1 + \pi \mu^2} \mathbf{1}.$$

Thus, setting

$$u_{\text{LS}} := \mu^2 \sqrt{\frac{\pi}{\pi\mu^2 + \left(\frac{\gamma}{1-\gamma}\right)(1+\pi\mu^2)}} = \mu^2 \sqrt{\frac{1-\gamma}{\mu^2 + \gamma/\pi}}, \quad (\text{H.26})$$

and applying Lemma (D.1), the probability of error is given by the advertised expression. \blacksquare

I Least-squares for MLM

I.1 Proof of Theorem 4.2

I.1.1 Computing $\Sigma_{w,\mu}$

Assume that \mathbf{X}, \mathbf{Y} are generated from the MLM.

Fix any $\ell \in [k]$. The classifier parameters $\widehat{\mathbf{w}}_\ell, \widehat{b}_\ell$ minimize the following objective function $\mathcal{L}_{PO}(\mathbf{w}_\ell, b_\ell) := \frac{1}{2n} \|\mathbf{X}^T \mathbf{w}_\ell + b_\ell \mathbf{1}_n - \mathbf{Y}_\ell\|_{\ell_2}^2$.

Identifying the AO. To continue further note that by duality we have

$$\min_{\mathbf{w}_\ell, b_\ell} \mathcal{L}_{PO}(\mathbf{w}_\ell, b_\ell) = \min_{\mathbf{w}_\ell, b_\ell} \max_{\mathbf{s}} \frac{1}{n} \left(\mathbf{s}^T \mathbf{X}^T \mathbf{w}_\ell + b_\ell \mathbf{s}^T \mathbf{1}_n - \mathbf{s}^T \mathbf{Y}_\ell - \frac{\|\mathbf{s}\|_{\ell_2}^2}{2} \right), \quad (\text{I.1})$$

and the optimization is jointly convex in $(\mathbf{w}_\ell, b_\ell)$ and concave in \mathbf{s} . Here, note that \mathbf{Y}_ℓ depends on the Gaussian matrix \mathbf{X} . Thus, before applying the CGMT, we need to break this dependence as follows. Consider the singular value decomposition

$$\mathbf{M} = \mathbf{U} \mathbf{\Sigma} \mathbf{V}^T = [\mathbf{u}_1 \quad \mathbf{u}_2 \quad \dots \quad \mathbf{u}_r] \text{diag}(\sigma_1, \sigma_2, \dots, \sigma_r) \begin{bmatrix} \mathbf{v}_1^T \\ \mathbf{v}_2^T \\ \vdots \\ \mathbf{v}_r^T \end{bmatrix}, \quad (\text{I.2})$$

with $\mathbf{U} \in \mathbb{R}^{d \times r}$, $\mathbf{\Sigma} \in \mathbb{R}^{r \times r}$, and $\mathbf{V} \in \mathbb{R}^{k \times r}$ where $r = \text{rank}(\mathbf{M}) \leq k$. For every $i \in [n]$, we decompose \mathbf{x}_i in its projection on the subspace spanned orthogonal columns $\mathbf{u}_1, \dots, \mathbf{u}_r$ as follows:

$$\mathbf{x}_i = \mathbf{U} \mathbf{U}^T \mathbf{X}_i + \mathbf{P}^\perp \mathbf{X}_i = \mathbf{U} \tilde{\mathbf{g}}_i + \mathbf{P}^\perp \mathbf{x}_i,$$

where $\mathbf{P}^\perp = \mathbf{I}_r - \mathbf{U} \mathbf{U}^T$, and we denote

$$\tilde{\mathbf{G}} := [\tilde{\mathbf{g}}_1 \quad \tilde{\mathbf{g}}_2 \quad \dots \quad \tilde{\mathbf{g}}_n], \quad \tilde{\mathbf{g}}_i := \mathbf{U}^T \mathbf{x}_i \in \mathbb{R}^r, i \in [n]. \quad (\text{I.3})$$

Recalling that $\mathbf{x}_i \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_d)$ note that

$$\tilde{\mathbf{g}}_i \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_r) \quad \text{and} \quad \tilde{\mathbf{g}}_i \perp \mathbf{P}^\perp \mathbf{x}_i. \quad (\text{I.4})$$

Further recall that for all $i \in [n]$, conditioned on \mathbf{x}_i

$$[\mathbf{Y}_\ell]_i \sim \text{Bern} \left(\frac{e^{\mu_\ell^T \mathbf{x}_i}}{\sum_{\ell' \in [k]} e^{\mu_{\ell'}^T \mathbf{x}_i}} \right) \sim \text{Bern} \left(\frac{e^{\mathbf{e}_\ell^T \mathbf{V} \Sigma \tilde{\mathbf{g}}_i}}{\sum_{\ell' \in [k]} e^{\mathbf{e}_{\ell'}^T \mathbf{V} \Sigma \tilde{\mathbf{g}}_i}} \right), \quad (\text{I.5})$$

where we used (I.3) and the SVD decomposition of \mathbf{M} . In this notation, we can rewrite the PO as follows:

$$\min_{\mathbf{w}_\ell, b_\ell} \max_{\mathbf{s}} \frac{1}{n} \left(\mathbf{s}^T \mathbf{X}^T \mathbf{P}^\perp \mathbf{w}_\ell + \mathbf{s}^T \tilde{\mathbf{G}}^T \mathbf{U}^T \mathbf{w}_\ell + b_\ell \mathbf{s}^T \mathbf{1}_n - \mathbf{s}^T \mathbf{Y}_\ell - \frac{\|\mathbf{s}\|_{\ell_2}^2}{2} \right)$$

From (I.4) and (I.5) notice that \mathbf{Y}_ℓ depends only on $\tilde{\mathbf{G}}$ and $\tilde{\mathbf{G}}$ is independent of $\mathbf{X}^T \mathbf{P}^\perp$. Therefore, the corresponding Auxiliary Optimization (AO) problem becomes

$$\min_{\mathbf{w}_\ell, b_\ell} \max_{\mathbf{s}} \frac{1}{n} \left(\|\mathbf{P}^\perp \mathbf{w}_\ell\|_{\ell_2} \mathbf{g}^T \mathbf{s} + \|\mathbf{s}\|_{\ell_2} \mathbf{h}^T \mathbf{P}^\perp \mathbf{w}_\ell + \mathbf{s}^T \tilde{\mathbf{G}}^T \mathbf{U}^T \mathbf{w}_\ell + b_\ell \mathbf{s}^T \mathbf{1}_n - \mathbf{s}^T \mathbf{Y}_\ell - \frac{\|\mathbf{s}\|_{\ell_2}^2}{2} \right), \quad (\text{I.6})$$

where $\mathbf{g} \in \mathbb{R}^n$ and $\mathbf{h} \in \mathbb{R}^d$ are iid Gaussian vectors independent of everything else.

Scalarization of the AO. Maximizing over the direction of \mathbf{s} and denoting its norm $\beta = \|\mathbf{s}\|_{\ell_2} \geq 0$ we arrive at

$$\begin{aligned} \min_{\mathbf{w}_\ell, b_\ell} \max_{\beta \geq 0} \frac{1}{n} & \left(\beta \|\mathbf{P}^\perp \mathbf{w}_\ell\|_{\ell_2} \mathbf{g} + \mathbf{G}^T \mathbf{U}^T \mathbf{w}_\ell + b_\ell \mathbf{1}_n - \mathbf{Y}_\ell \|_{\ell_2} + \beta \mathbf{h}^T \mathbf{P}^\perp \mathbf{w}_\ell - \frac{\beta^2}{2} \right) \\ &= \min_{\mathbf{w}_\ell, b_\ell} \frac{1}{2n} \left(\|\mathbf{P}^\perp \mathbf{w}_\ell\|_{\ell_2} \mathbf{g} + \tilde{\mathbf{G}}^T \mathbf{U}^T \mathbf{w}_\ell + b_\ell \mathbf{1}_n - \mathbf{Y}_\ell \|_{\ell_2} + \mathbf{h}^T \mathbf{P}^\perp \mathbf{w}_\ell \right)_+^2 \\ &= \frac{1}{2} \left(\min_{\mathbf{w}_\ell, b_\ell} \frac{1}{\sqrt{n}} \|\mathbf{P}^\perp \mathbf{w}_\ell\|_{\ell_2} \mathbf{g} + \tilde{\mathbf{G}}^T \mathbf{U}^T \mathbf{w}_\ell + b_\ell \mathbf{1}_n - \mathbf{Y}_\ell \|_{\ell_2} + \frac{1}{\sqrt{n}} \mathbf{h}^T \mathbf{P}^\perp \mathbf{w}_\ell \right)_+^2 \end{aligned} \quad (\text{I.7})$$

In the remaining, we focus in the inner minimization above. Let us denote

$$\mathbf{a} := \mathbf{U}^T \mathbf{w}_\ell \quad \text{and} \quad \alpha_0 = \|\mathbf{P}^\perp \mathbf{w}_\ell\|_{\ell_2}.$$

Notice that $\mathbf{a} \perp \mathbf{P}^\perp \mathbf{w}_\ell$ and thus the orthogonal decomposition $\mathbf{w}_\ell = \mathbf{U} \mathbf{a} + \mathbf{P}^\perp \mathbf{w}_\ell$. With this observation, we can optimize over the direction of $\mathbf{P}^T \mathbf{w}_\ell$ in (I.7) by aligning it with $-\mathbf{P}^T \mathbf{h}$. With this, the minimization in (I.7) reduces to the following

$$\min_{\mathbf{a}, \alpha_0 \geq 0, b_\ell} \frac{1}{\sqrt{n}} \|\alpha_0 \mathbf{g} + \tilde{\mathbf{G}}^T \mathbf{a} + b_\ell \mathbf{1}_n - \mathbf{Y}_\ell\|_{\ell_2} - \alpha_0 \frac{1}{\sqrt{n}} \|\mathbf{P}^\perp \mathbf{h}\|_{\ell_2}. \quad (\text{I.8})$$

Convergence of the AO. First, we argue on point-wise convergence of the objective function in (I.8). Fix \mathbf{a}, α_0 and \mathbf{b}_ℓ . From the WLLN, $\frac{1}{\sqrt{n}} \|\mathbf{P}^\perp \mathbf{h}\|_{\ell_2} \xrightarrow{P} \sqrt{\gamma}$ and

$$\frac{1}{n} \|\alpha_0 \mathbf{g} + \tilde{\mathbf{G}}^T \mathbf{a} + b_\ell \mathbf{1}_n - \mathbf{Y}_\ell\|_{\ell_2}^2 = \frac{1}{n} \sum_{i=1}^n (\alpha_0 \mathbf{g}_i + \mathbf{a}^T \tilde{\mathbf{g}}_i + b_\ell - [\mathbf{Y}_\ell]_i)^2 \xrightarrow{P} \mathbb{E}[(\alpha_0 G_0 + \mathbf{a}^T \mathbf{g} + \mathbf{b}_\ell - Y_\ell)^2], \quad (\text{I.9})$$

where the expectation is over $\mathbf{g} \sim \mathcal{N}(\mathbf{0}_r, \mathbf{I}_r)$ (with some abuse of notation) and

$$Y_\ell \sim \text{Bern}(V_\ell) \quad \text{and} \quad V_\ell = \frac{e^{\mathbf{e}_\ell^T \mathbf{V} \Sigma \mathbf{g}}}{\sum_{\ell'=1}^k e^{\mathbf{e}_{\ell'}^T \mathbf{V} \Sigma \mathbf{g}}}. \quad (\text{I.10})$$

Therefore, point-wise on \mathbf{a}, α_0 and \mathbf{b}_ℓ , the objective of the AO converges to

$$\mathcal{D}_\ell(\alpha_0, \mathbf{a}, b_\ell) := \sqrt{\mathbb{E}[(\alpha_0 G_0 + \mathbf{a}^T \mathbf{g} + \mathbf{b}_\ell - Y_\ell)^2]} - \alpha_0 \sqrt{\gamma}. \quad (\text{I.11})$$

Next, with an argument based on convexity and compactness similar to that in ‘‘Convergence analysis of the AO’’ in Section H it can be argued that the convergence above is uniform. Thus,

$$(\text{I.8}) \xrightarrow{P} \min_{\alpha_0 \geq 0, \mathbf{a}, b_\ell} \mathcal{D}_\ell(\alpha_0, \mathbf{a}, b_\ell). \quad (\text{I.12})$$

Deterministic analysis of the AO. Here, we solve the deterministic minimization problem in (I.12). Optimization over b_ℓ is straightforward. By setting

$$b_\ell = \mathbb{E}[Y_\ell] = \mathbb{E}[V_\ell],$$

we now have to optimize

$$\min_{\alpha_0 \geq 0, \mathbf{a}} \sqrt{\alpha_0^2 + \mathbb{E}[(\mathbf{a}^T \mathbf{g} - Y_\ell)^2] - (\mathbb{E}[V_\ell])^2} - \alpha_0 \sqrt{\gamma}. \quad (\text{I.13})$$

By direct differentiation and first-order optimality, we compute the optimal values as follows:

$$\tilde{a}_j = \mathbb{E}[\mathbf{g}_j Y_\ell] = \mathbb{E}[\mathbf{g}_j V_\ell], \quad j \in [r], \quad (\text{I.14})$$

$$\tilde{\alpha}_0^2 = \frac{\gamma}{1-\gamma} \left(\text{Var}[Y_\ell] - \sum_{j=1}^r (\mathbb{E}[\mathbf{g}_j V_\ell])^2 \right) = \frac{\gamma}{1-\gamma} \left(\mathbb{E}[V_\ell] - (\mathbb{E}[V_\ell])^2 - \sum_{j=1}^r (\mathbb{E}[\mathbf{g}_j V_\ell])^2 \right). \quad (\text{I.15})$$

Asymptotic Predictions. From the analysis above, we conclude with the following limits about the solution $\widehat{\mathbf{b}}_\ell, \widehat{\mathbf{w}}_\ell$ of the PO:

$$\widehat{\mathbf{b}}_\ell \xrightarrow{P} \mathbb{E}[V_\ell] \quad (\text{I.16a})$$

$$\langle \mu_c, \widehat{\mathbf{w}}_\ell \rangle \xrightarrow{P} \mathbf{e}_c^T \mathbf{V} \Sigma \mathbb{E}[\mathbf{g} V_\ell], \quad c \in [k] \quad (\text{I.16b})$$

$$\|\widehat{\mathbf{w}}_\ell\|_{\ell_2}^2 \xrightarrow{P} \sum_{j=1}^r (\mathbb{E}[\mathbf{g}_j V_\ell])^2 + \frac{\gamma}{1-\gamma} \left(\mathbb{E}[V_\ell] - (\mathbb{E}[V_\ell])^2 - \sum_{j=1}^r (\mathbb{E}[\mathbf{g}_j V_\ell])^2 \right) \quad (\text{I.16c})$$

$$= \frac{\gamma}{1-\gamma} \left(\mathbb{E}[V_\ell] - (\mathbb{E}[V_\ell])^2 \right) + \frac{1-2\gamma}{1-\gamma} \sum_{j=1}^r (\mathbb{E}[\mathbf{g}_j V_\ell])^2. \quad (\text{I.16d})$$

Recall the notation in (4.1). Note that $\mathbb{E}[V_\ell] = \pi_\ell$. Moreover, using Gaussian integration by parts Lemma C.3, it can be shown that $\mathbb{E}[V_\ell \mathbf{g}] = \Sigma \mathbf{V}^T (\text{diag}(\pi) - \Pi) \mathbf{e}_\ell$. Using these and writing in matrix form, we arrive at (4.4a).

I.1.2 Computing $\Sigma_{w,w}$

Here, we prove (4.4b). Specifically, we compute the correlations $\langle \widehat{\mathbf{w}}_\ell, \widehat{\mathbf{w}}_c \rangle$, $\ell \neq c \in [k]$ by following the strategy of Section H.1.2. Specifically, in view of Lemma H.1 we need to study the following PO:

$$\min_{\mathbf{w}, \mathbf{b}} \max_{\mathbf{s}} \frac{1}{n} \left(\mathbf{s}^T \mathbf{X}^T \mathbf{w} + \mathbf{b} \mathbf{s}^T \mathbf{1}_n - \mathbf{s}^T (\mathbf{Y}_\ell + \mathbf{Y}_c) - \frac{\|\mathbf{s}\|_{\ell_2}^2}{2} \right) \quad (\text{I.17})$$

which is minimized by $\widehat{\mathbf{w}}_\ell + \widehat{\mathbf{w}}_c$. Thus the analysis will lead us to an asymptotic formula for $\|\widehat{\mathbf{w}}_\ell + \widehat{\mathbf{w}}_c\|_{\ell_2}$. This when combined with the formulae for $\|\widehat{\mathbf{w}}_\ell\|_{\ell_2}$ and $\|\widehat{\mathbf{w}}_c\|_{\ell_2}$ in (I.16d) will give the desired.

The analysis of (I.17) is almost identical to the analysis of (I.1) in the previous section. Specifically, without repeating all the details for brevity, it can be shown that the AO of (I.17) converges to the following (cf. (I.11)):

$$\mathcal{D}_\ell(\alpha_0, \boldsymbol{\alpha}, b_\ell) := \sqrt{\mathbb{E}[(\alpha_0 G_0 + \mathbf{a}^T \mathbf{g} + \mathbf{b}_\ell - Y_{\ell,c})^2]} - \alpha_0 \sqrt{\gamma}, \quad (\text{I.18})$$

where as before $G_0 \sim \mathcal{N}(0, 1)$, $\mathbf{g} \sim \mathcal{N}(\mathbf{0}_r, \mathbf{I}_r)$, only now (I.10) is modified to:

$$Y_{\ell,c} \sim \text{Bern}(V_c + V_\ell) \quad \text{and as before:} \quad V_\ell = \frac{e^{\mathbf{e}_\ell^T \mathbf{V} \Sigma \mathbf{g}}}{\sum_{\ell'=1}^k e^{\mathbf{e}_{\ell'}^T \mathbf{V} \Sigma \mathbf{g}}}. \quad (\text{I.19})$$

With these, it can be shown that

$$\|\widehat{\mathbf{w}}_\ell + \widehat{\mathbf{w}}_c\|_{\ell_2}^2 \xrightarrow{P} \sum_{j=1}^r (\mathbb{E}[\mathbf{g}_j (V_c + V_\ell)])^2 + \frac{\gamma}{1-\gamma} \left(\mathbb{E}[V_c + V_\ell] - (\mathbb{E}[V_c + V_\ell])^2 - \sum_{j=1}^r (\mathbb{E}[\mathbf{g}_j (V_c + V_\ell)])^2 \right)$$

Combining this with (I.16d), we conclude that for $\ell \neq c \in [k]$:

$$\langle \widehat{\mathbf{w}}_\ell, \widehat{\mathbf{w}}_c \rangle \xrightarrow{P} \frac{1-2\gamma}{1-\gamma} \sum_{j=1}^r \mathbb{E}[\mathbf{g}_j V_c] \mathbb{E}[\mathbf{g}_j V_\ell] - \frac{\gamma}{1-\gamma} \mathbb{E}[V_c] \mathbb{E}[V_\ell]. \quad (\text{I.20})$$

This shows (4.4b) after applying Gaussian integration by parts and expressing it in matrix form; see Lemma C.3.

I.2 Orthogonal means and equal-energy

Here, we use Theorem 4.2 to prove that, in contrast to the GMM, in the MLM under orthogonal and equal-energy means: LS outperforms the averaging classifier for large enough sample sizes. Assuming orthogonal means of equal energy μ :

$$\boldsymbol{\pi} = \boldsymbol{\pi}_1 \mathbf{1}_k = (1/k) \mathbf{1}_k, \quad (\text{I.21})$$

$$\Pi = (\Pi_{11} - \Pi_{12}) \mathbf{I}_k + \Pi_{12} \mathbf{1}_k \mathbf{1}_k^T \quad \text{with} \quad \Pi_{12} = \frac{1-k^2 \Pi_{11}^2}{k(k-1)} \quad \text{and} \quad \Pi_{11} = \mathbb{E} \left[\frac{e^{2\mu G_1}}{\left(\sum_{\ell \in [k]} e^{\mu G_\ell} \right)^2} \right].$$

Then,

$$\Sigma_{\mathbf{w}, \mathbf{w}} - \Sigma_{\mathbf{w}, \mu} \Sigma_{\mu, \mu}^{-1} \Sigma_{\mathbf{w}, \mu}^T \xrightarrow{P} \frac{\gamma}{1 - \gamma} \cdot (p \mathbf{I}_k - q \mathbf{1}_k \mathbf{1}_k^T), \quad (\text{I.22})$$

where we defined

$$p := \boldsymbol{\pi}_1 - \mu^2 (\boldsymbol{\pi}_1 - \boldsymbol{\Pi}_{11} + \boldsymbol{\Pi}_{12})^2 \text{ and } q := (\boldsymbol{\pi}_1^2 + \boldsymbol{\Pi}_{12}^2 \mu^2 k - 2\mu^2 \boldsymbol{\Pi}_{12} (\boldsymbol{\pi}_1 - \boldsymbol{\Pi}_{11} + \boldsymbol{\Pi}_{12})). \quad (\text{I.23})$$

Thus, similar to (4.3) and with the same notation,

$$\mathbb{P}_{e, \text{LS}} \xrightarrow{P} \mathbb{P} \left\{ \arg \max_{\ell \in [k]} \left\{ \sqrt{\frac{\gamma}{1 - \gamma}} \cdot (p \mathbf{I}_k - q \mathbf{1}_k \mathbf{1}_k^T)^{1/2} \cdot \tilde{\mathbf{g}} + \mu \cdot (\text{diag}(\boldsymbol{\pi}) - \boldsymbol{\Pi}) \mathbf{g} \right\} \neq Y(\mathbf{g}) \right\}. \quad (\text{I.24})$$

In (I.24) (as well as in (4.3)), note that the matrices multiplying $\tilde{\mathbf{g}}$ and \mathbf{g} have all the form of a rank one update of a (scaled) identity matrix. It turns out that we can exploit this structure to simplify the formulae for the test error even further. Importantly, this lets us directly compare $\mathbb{P}_{e, \text{LS}}$ and $P_{e, \text{Avg}}$ of the two classifiers. These are detailed in Section I.3.

I.3 Proof of Proposition 4.3

In (E.7) and (I.24), we showed the following limits for orthogonal means of equal-energy $\mu > 0$:

$$\begin{aligned} P_{e, \text{Avg}} &\xrightarrow{P} \mathbb{P} \left(\arg \max_{\ell \in [k]} \left\{ \gamma \cdot \boldsymbol{\pi}_1 \cdot \mathbf{I}_k \cdot \tilde{\mathbf{g}} + \mu \left((\boldsymbol{\pi}_1 - \boldsymbol{\Pi}_{11}) \cdot \mathbf{I}_k + \boldsymbol{\Pi}_{12} \mathbf{1}_k \mathbf{1}_k^T \right) \cdot \mathbf{g} \right\} \neq Y(\mathbf{g}) \right) \\ \mathbb{P}_{e, \text{LS}} &\xrightarrow{P} \mathbb{P} \left(\arg \max_{\ell \in [k]} \left\{ \sqrt{\frac{\gamma}{1 - \gamma}} \cdot (p \mathbf{I}_k - q \mathbf{1}_k \mathbf{1}_k^T)^{1/2} \cdot \tilde{\mathbf{g}} + \mu \left((\boldsymbol{\pi}_1 - \boldsymbol{\Pi}_{11}) \cdot \mathbf{I}_k + \boldsymbol{\Pi}_{12} \mathbf{1}_k \mathbf{1}_k^T \right) \cdot \mathbf{g} \right\} \neq Y(\mathbf{g}) \right), \end{aligned}$$

where $\mathbb{P}(Y(\mathbf{g}) = \ell) = \frac{e^{\mu \mathbf{g}_\ell}}{\sum_{j \in [k]} e^{\mu \mathbf{g}_j}}$ and we have further used (I.21) and the notation in (I.23).

We compare the expression on the RHS in the above display by applying Lemma I.1 below with the following substitutions

$$\begin{aligned} \mathbf{g} &\leftarrow \tilde{\mathbf{g}}, \quad \mathbf{h} \leftarrow \mathbf{g}, \quad c(\mathbf{h}) \leftarrow Y(\mathbf{g}) \\ p_2 &\leftarrow \frac{\gamma}{1 - \gamma} (\boldsymbol{\pi}_1 - \mu^2 (\boldsymbol{\pi}_1 - \boldsymbol{\Pi}_{11} + \boldsymbol{\Pi}_{12})^2), \quad q_2 \leftarrow \frac{\gamma}{1 - \gamma} (\boldsymbol{\pi}_1^2 + \boldsymbol{\Pi}_{12}^2 \mu^2 k - 2\mu^2 \boldsymbol{\Pi}_{12} (\boldsymbol{\pi}_1 - \boldsymbol{\Pi}_{11} + \boldsymbol{\Pi}_{12})), \\ p_1 &\leftarrow \gamma \boldsymbol{\pi}_1, \quad q_1 \leftarrow 0. \end{aligned}$$

This shows that with probability 1, $\mathbb{P}_{e, \text{LS}} < \mathbb{P}_{e, \text{Avg}}$ if and only if $p_2 < p_1 \Leftrightarrow \gamma < \gamma_* = \mu^2 (\boldsymbol{\pi}_1 - \boldsymbol{\Pi}_{11} + \boldsymbol{\Pi}_{12})^2 / \boldsymbol{\pi}_1$. To retrieve (??), recall that $\boldsymbol{\pi}_1 = 1/k$ and $k \boldsymbol{\Pi}_{11} + (k^2 - k) \boldsymbol{\Pi}_{12} = 1$. The only thing left to prove is that $\gamma_* < 1$. To see this note that $p_2 > 0$ from positive semi-definiteness of the Schur matrix in (I.23). It takes simple algebra to conclude that $p_2 > 0 \implies \gamma_* < 1$.

Lemma I.1 *Let $k \geq 2$, $\mathbf{g} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_k)$, $\mathbf{h} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_k)$, and discrete random variable $c(\mathbf{h})$ such that $\mathbb{P}(c(\mathbf{h}) = \ell) = e^{\mathbf{h}_\ell} / \sum_{j \in [k]} e^{\mathbf{h}_j}$. Consider the function $F : \mathbb{R}_{>0} \times \mathbb{R} \rightarrow [0, 1]$ defined as follows*

$$F(p, q) = \mathbb{P} \left(\arg \max \left\{ (p \mathbf{I}_k - q \mathbf{1}_k \mathbf{1}_k^T)^{1/2} \mathbf{g} + (\alpha \mathbf{I}_k - \beta \mathbf{1}_k \mathbf{1}_k^T)^{1/2} \mathbf{h} \right\} \neq c(\mathbf{h}) \right),$$

such that $p \mathbf{I}_k - q \mathbf{1}_k \mathbf{1}_k^T > 0$ and fixed $\alpha \mathbf{I}_k - \beta \mathbf{1}_k \mathbf{1}_k^T > 0$. Then, the following statements are true.

1. $F(p, q) = \mathbb{P} \left(\arg \max \left\{ \sqrt{p} \cdot \mathbf{g} + \sqrt{\alpha} \mathbf{h} \right\} \neq c(\mathbf{h}) \right).$
2. For $0 < p_2 < p_1$ and any $q_1 < \frac{p_1}{k}$, $q_2 < \frac{p_2}{k}$, it holds that $F(p_2, q_2) < F(p_1, q_1)$.

Proof Fix any $p > 0$, $q \leq \frac{p}{k}$. Denote $\mathbf{T} := (p \mathbf{I}_k - q \mathbf{1}_k \mathbf{1}_k^T)^{1/2}$ and $\mathbf{S} := (\alpha \mathbf{I}_k - \beta \mathbf{1}_k \mathbf{1}_k^T)^{1/2}$ for convenience. It can be checked that $\mathbf{T} := \left(\sqrt{p} \mathbf{I}_k + \frac{\sqrt{p} - qk - \sqrt{p}}{k} \mathbf{1}_k \mathbf{1}_k^T \right)$ and $\mathbf{S} := \left(\sqrt{\alpha} \mathbf{I}_k + \frac{\sqrt{\alpha} - \beta k - \sqrt{\alpha}}{k} \mathbf{1}_k \mathbf{1}_k^T \right)$. From these, it follows directly that

$$F(p, q) = \mathbb{P} \left(\arg \max \left\{ \sqrt{p} \cdot \mathbf{g} + \sqrt{\alpha} \mathbf{h} \right\} \neq c(\mathbf{h}) \right).$$

This shows the first statement.

Next, we show the second statement. Using the distribution of $c(\mathbf{h})$ and symmetry we have the following chain of equalities:

$$\begin{aligned}
1 - F(p, q) &= \mathbb{P} \left\{ \arg \max_{j \in [k]} \left\{ \sqrt{p} \cdot \mathbf{g}_j + \sqrt{\alpha} \mathbf{h}_j \right\} = c(\mathbf{h}) \right\} \\
&= k \cdot \mathbb{E} \left[\frac{e^{\mathbf{h}_k}}{\sum_{j \in [k]} e^{\mathbf{h}_j}} \cdot \mathbb{1} \left\{ \arg \max \left\{ \sqrt{p} \cdot \mathbf{g}_j + \sqrt{\alpha} \mathbf{h}_j \right\} = k \right\} \right] \\
&= k \cdot \mathbb{E} \left[\frac{e^{\mathbf{h}_k}}{\sum_{j \in [k]} e^{\mathbf{h}_j}} \cdot \prod_{j \in [k-1]} \mathbb{1} \left\{ \sqrt{p} \cdot \mathbf{g}_j + \sqrt{\alpha} \mathbf{h}_j < \sqrt{p} \cdot \mathbf{g}_k + \sqrt{\alpha} \mathbf{h}_k \right\} \right] \\
&= k \cdot \mathbb{E} \left[\frac{e^{\mathbf{h}_k}}{\sum_{j \in [k]} e^{\mathbf{h}_j}} \cdot \prod_{j \in [k-1]} \mathbb{1} \left\{ \mathbf{g}_j < \mathbf{g}_k + \frac{\sqrt{\alpha} \mathbf{h}_k - \sqrt{\alpha} \mathbf{h}_j}{\sqrt{p}} \right\} \right] \\
&= k \cdot \mathbb{E} \left[\frac{e^{\mathbf{h}_k}}{\sum_{j \in [k]} e^{\mathbf{h}_j}} \cdot \prod_{j \in [k-1]} Q \left(\mathbf{g}_k + \frac{\sqrt{\alpha} \mathbf{h}_j - \sqrt{\alpha} \mathbf{h}_k}{\sqrt{p}} \right) \right] =: k \cdot G(\sqrt{p}), \quad (\text{I.25})
\end{aligned}$$

where in the last line we used the rotational symmetry of the Gaussian distribution:

$$\mathbb{P} \left\{ \mathbf{g}_j < \mathbf{g}_k + \frac{\sqrt{\alpha} \mathbf{h}_k - \sqrt{\alpha} \mathbf{h}_j}{\sqrt{p}} \mid \mathbf{h}_1, \dots, \mathbf{h}_k \right\} = \mathbb{P} \left\{ \mathbf{g}_j > \mathbf{g}_k + \frac{\sqrt{\alpha} \mathbf{h}_j - \sqrt{\alpha} \mathbf{h}_k}{\sqrt{p}} \mid \mathbf{h}_1, \dots, \mathbf{h}_k \right\},$$

and the fact that $\mathbf{g}_1, \dots, \mathbf{g}_{k-1}$ are independent.

Next, we will show that the function $\mathcal{G}(\cdot)$ defined above is strictly decreasing in $(0, \infty)$. Towards this goal, using $Q'(x) = -\frac{1}{\sqrt{2\pi}} e^{-x^2/2} = -\phi(x)$ and using the shorthand

$$H_{kj} = \mathbf{h}_j - \mathbf{h}_k, \quad j \in [k],$$

we may compute the derivative of \mathcal{G} at any $s > 0$ as follows:

$$\begin{aligned}
\frac{d\mathcal{G}(s)}{ds} &= \sum_{i \in [k-1]} \mathbb{E} \left[\frac{e^{\mathbf{h}_k}}{\sum_{j \in [k]} e^{\mathbf{h}_j}} \cdot \phi \left(\mathbf{g}_k + \frac{\sqrt{\alpha} H_{ki}}{s} \right) \cdot \frac{\sqrt{\alpha} H_{ki}}{s^2} \cdot \prod_{j \neq i \in [k-1]} Q \left(\mathbf{g}_k + \frac{\sqrt{\alpha} H_{kj}}{s} \right) \right] \\
&= \sum_{i \in [k-1]} \mathbb{E} \left[\frac{1}{\sum_{j \in [k]} e^{H_{kj}}} \cdot \phi \left(\mathbf{g}_k + \frac{\sqrt{\alpha} H_{ki}}{s} \right) \cdot \frac{\sqrt{\alpha} H_{ki}}{s^2} \cdot \prod_{j \neq i \in [k-1]} Q \left(\mathbf{g}_k + \frac{\sqrt{\alpha} H_{kj}}{s} \right) \right] \\
&= \sum_{i \in [k-1]} \frac{\sqrt{\alpha}}{s^2} \mathbb{E} \left[H_{ki} \cdot \mathcal{A}_i \left(\mathbf{g}_k, \{H_{kj}\}_{j \in [k-1]} \right) \right], \quad (\text{I.26})
\end{aligned}$$

where in the last line we have defined

$$\mathcal{A}_i \left(\mathbf{g}_k, \{H_{kj}\}_{j \in [k-1]} \right) := \frac{1}{1 + \sum_{j \in [k-1]} e^{H_{kj}}} \cdot \phi \left(\mathbf{g}_k + \frac{\sqrt{\alpha} H_{ki}}{s} \right) \cdot \prod_{j \neq i \in [k-1]} Q \left(\mathbf{g}_k + \frac{\sqrt{\alpha} H_{kj}}{s} \right),$$

Next, we use Gaussian integration by parts (GIBP) to further simplify the expression in (I.26). Fix any $i \in [k-1]$. Then, by (GIBP):

$$A_i := \mathbb{E} \left[H_{ki} \cdot \mathcal{A}_i \left(\mathbf{g}_k, \{H_{kj}\}_{j \in [k-1]} \right) \right] \quad (\text{I.27})$$

$$\begin{aligned}
&= \mathbb{E} [H_{ki}^2] \mathbb{E} \left[\frac{d}{dH_{ki}} \mathcal{A}_i \left(\mathbf{g}_k, \{H_{kj}\}_{j \in [k-1]} \right) \right] + \sum_{\substack{\ell \in [k-1] \\ \ell \neq i}} \mathbb{E} [H_{ki} \cdot H_{k\ell}] \mathbb{E} \left[\frac{d}{dH_{k\ell}} \mathcal{A}_i \left(\mathbf{g}_k, \{H_{kj}\}_{j \in [k-1]} \right) \right] \\
&= 2 \underbrace{\mathbb{E} \left[\frac{d}{dH_{ki}} \mathcal{A}_i \left(\mathbf{g}_k, \{H_{kj}\}_{j \in [k-1]} \right) \right]}_{\text{TermI}} + \underbrace{\sum_{\substack{\ell \in [k-1] \\ \ell \neq i}} \mathbb{E} \left[\frac{d}{dH_{k\ell}} \mathcal{A}_i \left(\mathbf{g}_k, \{H_{kj}\}_{j \in [k-1]} \right) \right]}_{\text{TermII}}, \quad (\text{I.28})
\end{aligned}$$

where in the second line, we used the fact that $\mathbf{h} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_k)$ to compute

$$\mathbb{E} [H_{ki}^2] = 2, \quad \text{and} \quad \mathbb{E} [H_{ki} H_{k\ell}] = 1, \quad \ell \neq i, \quad \ell \in [k-1].$$

We now compute the derivatives in (I.28). First, for any $\ell \in [k-1]$, $\ell \neq i$,

$$\begin{aligned} \frac{d\mathcal{A}_i(\mathbf{g}_k, \{H_{kj}\}_{j \in [k-1]})}{dH_{k\ell}} &= -\frac{e^{H_{k\ell}}}{(1 + \sum_{j \in [k-1]} e^{H_{kj}})^2} \cdot \phi\left(\mathbf{g}_k + \frac{\sqrt{\alpha} H_{ki}}{s}\right) \cdot \prod_{j \neq i \in [k-1]} Q\left(\mathbf{g}_k + \frac{\sqrt{\alpha} H_{kj}}{s}\right) =: \text{TermII(a)}_\ell \\ &- \frac{\sqrt{\alpha}}{s} \cdot \frac{1}{1 + \sum_{j \in [k-1]} e^{H_{kj}}} \cdot \phi\left(\mathbf{g}_k + \frac{\sqrt{\alpha} H_{ki}}{s}\right) \cdot \phi\left(\mathbf{g}_k + \frac{\sqrt{\alpha} H_{k\ell}}{s}\right) \cdot \prod_{j \neq (i, \ell) \in [k-1]} Q\left(\mathbf{g}_k + \frac{\sqrt{\alpha} H_{kj}}{s}\right) =: \text{TermII(b)}_\ell \\ &= \text{TermII(a)}_\ell + \text{TermII(b)}_\ell \end{aligned} \quad (\text{I.29})$$

Thus,

$$\text{TermII} = \sum_{\ell \neq i \in [k-1]} \mathbb{E}[\text{TermII(a)}_\ell] + \mathbb{E}[\text{TermII(b)}_\ell] =: \mathbb{E}[\text{TermII(a)}] + N_i, \quad (\text{I.30})$$

where we defined

$$N_i = -\frac{\sqrt{\alpha}}{s} \cdot \mathbb{E} \left[\frac{\phi\left(\mathbf{g}_k + \frac{\sqrt{\alpha} H_{ki}}{s}\right)}{1 + \sum_{j \in [k-1]} e^{H_{kj}}} \cdot \sum_{\ell \neq i \in [k-1]} \left\{ \phi\left(\mathbf{g}_k + \frac{\sqrt{\alpha} H_{k\ell}}{s}\right) \cdot \prod_{j \neq (i, \ell) \in [k-1]} Q\left(\mathbf{g}_k + \frac{\sqrt{\alpha} H_{kj}}{s}\right) \right\} \right] < 0. \quad (\text{I.31})$$

and we remark for later use that

$$\mathbb{E}[\text{TermII(a)}] = \sum_{\ell \neq i \in [k-1]} \mathbb{E}[\text{TermII(a)}_\ell] < 0. \quad (\text{I.32})$$

Second, it holds that

$$\begin{aligned} \frac{d\mathcal{A}_i(\mathbf{g}_k, \{H_{kj}\}_{j \in [k-1]})}{dH_{ki}} &= -\frac{e^{H_{ki}}}{(1 + \sum_{j \in [k-1]} e^{H_{kj}})^2} \cdot \phi\left(\mathbf{g}_k + \frac{\sqrt{\alpha} H_{ki}}{s}\right) \cdot \prod_{j \neq i \in [k-1]} Q\left(\mathbf{g}_k + \frac{\sqrt{\alpha} H_{kj}}{s}\right) \\ &+ \frac{d\phi\left(\mathbf{g}_k + \frac{\sqrt{\alpha} H_{ki}}{s}\right)}{dH_{ki}} \cdot \frac{1}{1 + \sum_{j \in [k-1]} e^{H_{kj}}} \cdot \prod_{j \neq i \in [k-1]} Q\left(\mathbf{g}_k + \frac{\sqrt{\alpha} H_{kj}}{s}\right) \\ &= -\frac{e^{H_{ki}}}{(1 + \sum_{j \in [k-1]} e^{H_{kj}})^2} \cdot \phi\left(\mathbf{g}_k + \frac{\sqrt{\alpha} H_{ki}}{s}\right) \cdot \prod_{j \neq i \in [k-1]} Q\left(\mathbf{g}_k + \frac{\sqrt{\alpha} H_{kj}}{s}\right) =: \text{TermI(a)} \\ &- \frac{\sqrt{\alpha}}{s} \left(\mathbf{g}_k + \frac{\sqrt{\alpha} H_{ki}}{s} \right) \phi\left(\mathbf{g}_k + \frac{\sqrt{\alpha} H_{ki}}{s}\right) \cdot \frac{1}{1 + \sum_{j \in [k-1]} e^{H_{kj}}} \cdot \prod_{j \neq i \in [k-1]} Q\left(\mathbf{g}_k + \frac{\sqrt{\alpha} H_{kj}}{s}\right) =: \text{TermI(b)}, \\ &= \text{TermI(a)} + \text{TermI(b)}, \end{aligned} \quad (\text{I.33})$$

where in the penultimate line we used the fact that $\phi'(x) = -x\phi(x)$. Consider the two terms in (I.33). Clearly,

$$\mathbb{E}[\text{TermI(a)}] < 0. \quad (\text{I.34})$$

For the second term we observe that:

$$\mathbb{E}[\text{TermI(b)}] = -\frac{\sqrt{\alpha}}{s} \mathbb{E} \left[\left(\mathbf{g}_k + \frac{\sqrt{\alpha} H_{ki}}{s} \right) \cdot \mathcal{A}_i(\mathbf{g}_k, \{H_{kj}\}_{j \in [k-1]}) \right] \quad (\text{I.35})$$

$$\begin{aligned} &= -\frac{\alpha}{s^2} \cdot \mathbb{E}[H_{ki} \cdot \mathcal{A}_i(\mathbf{g}_k, \{H_{kj}\}_{j \in [k-1]})] - \frac{\sqrt{\alpha}}{s} \mathbb{E}[\mathbf{g}_k \cdot \mathcal{A}_i(\mathbf{g}_k, \{H_{kj}\}_{j \in [k-1]})] \\ &= -\frac{\alpha}{s^2} \cdot A_i - \frac{\sqrt{\alpha}}{s} \mathbb{E}[\mathbf{g}_k \cdot \mathcal{A}_i(\mathbf{g}_k, \{H_{kj}\}_{j \in [k-1]})]. \end{aligned} \quad (\text{I.36})$$

Moreover, using again GIBP, $\mathbb{E}[\mathbf{g}_k^2] = 1$, $\mathbb{E}[\mathbf{g}_k H_{kj}] = 0$, $j \in [k]$ and the fact that $\phi'(x) = -x\phi(x)$,

$$\begin{aligned}
\mathbb{E}[\mathbf{g}_k \cdot \mathcal{A}_i(\mathbf{g}_k, \{H_{kj}\}_{j \in [k-1]})] &= \mathbb{E}\left[\frac{d\mathcal{A}_i(\mathbf{g}_k, \{H_{kj}\}_{j \in [k]})}{d\mathbf{g}_k}\right] \\
&= -\mathbb{E}\left[\left(\mathbf{g}_k + \frac{\sqrt{\alpha} H_{ki}}{s}\right) \cdot \mathcal{A}_i(\mathbf{g}_k, \{H_{kj}\}_{j \in [k-1]})\right] \\
&\quad - \sum_{\ell \neq i \in [k-1]} \mathbb{E}\left[\phi\left(\mathbf{g}_k + \frac{\sqrt{\alpha} H_{k\ell}}{s}\right) \cdot \frac{1}{1 + \sum_{j \in [k-1]} e^{H_{k\ell}}} \cdot \phi\left(\mathbf{g}_k + \frac{\sqrt{\alpha} H_{ki}}{s}\right) \cdot \prod_{j \neq (i, \ell) \in [k-1]} Q\left(\mathbf{g}_k + \frac{\sqrt{\alpha} H_{kj}}{s}\right)\right] \\
&= \frac{s}{\sqrt{\alpha}} \mathbb{E}[\text{TermI(b)}] + \frac{s}{\sqrt{\alpha}} N_i,
\end{aligned} \tag{I.37}$$

where, we have recalled (I.35) and (I.31). Using (I.37) in (I.36), we find that

$$\mathbb{E}[\text{TermI(b)}] = -\frac{\alpha}{s^2} \cdot A_i - \mathbb{E}[\text{TermI(b)}] - N_i \implies \mathbb{E}[\text{TermI(b)}] = -\frac{\alpha}{2s^2} \cdot A_i - \frac{N_i}{2}. \tag{I.38}$$

We are now ready to put things together:

$$\begin{aligned}
A_i &= 2 \cdot \text{TermI} + \text{TermII} \quad \text{by (I.28)} \\
&= 2 \mathbb{E}[\text{TermI(a)}] + 2 \mathbb{E}[\text{TermI(b)}] + \mathbb{E}[\text{TermII(a)}] + N_i \quad \text{by (I.30)} \\
&= 2 \mathbb{E}[\text{TermI(a)}] - \frac{\alpha}{s^2} \cdot A_i - N_i + \mathbb{E}[\text{TermII(a)}] + N_i \quad \text{by (I.38)} \\
\implies A_i &= \frac{s^2}{s^2 + \alpha} (2 \mathbb{E}[\text{TermI(a)}] + \mathbb{E}[\text{TermII(a)}]) \\
&< 0 \quad \text{by (I.34) and (I.32).}
\end{aligned}$$

From this, (I.26) and (I.27), we have shown that \mathcal{G} is strictly decreasing in $(0, \infty)$. Recalling the definition of \mathcal{G} in (I.25), this implies that $F(p, q)$ is strictly increasing in $p > 0$, as desired to complete the proof. \blacksquare

J Weighted LS for GMM (Proof of Theorem B.1)

J.1 Computing $\Sigma_{w, \mu}$

The WLS estimator solves:

$$\begin{aligned}
\min_{\mathbf{W} \in \mathbb{R}^{k \times d}, \mathbf{b} \in \mathbb{R}^k} \quad & \frac{1}{2n} \|(\mathbf{W}\mathbf{X} + \mathbf{b}\mathbf{1}_n^T - \mathbf{Y})\mathbf{D}\|_F^2 = \sum_{\ell=1}^k \min_{\mathbf{w}_\ell, b_\ell} \frac{1}{2n} \|\mathbf{D}(\mathbf{X}^T \mathbf{w}_\ell + b_\ell \mathbf{1}_n - \mathbf{Y}_\ell)\|_{\ell_2}^2 \\
&= \sum_{\ell=1}^k \min_{\mathbf{w}_\ell, b_\ell} \frac{1}{2n} \|\mathbf{D}(\mathbf{Y}^T \mathbf{M}^T \mathbf{w}_\ell + \mathbf{Z}^T \mathbf{w}_\ell + b_\ell \mathbf{1}_n - \mathbf{Y}_\ell)\|_{\ell_2}^2.
\end{aligned}$$

Define

$$\mathcal{L}_{PO}(\mathbf{w}_\ell, b_\ell) := \frac{1}{2n} \|\mathbf{D}(\mathbf{Y}^T \mathbf{M}^T \mathbf{w}_\ell + \mathbf{Z}^T \mathbf{w}_\ell + b_\ell \mathbf{1}_n - \mathbf{Y}_\ell)\|_{\ell_2}^2. \tag{J.1}$$

Identifying the AO. By duality we have

$$\min_{\mathbf{w}_\ell, b_\ell} \mathcal{L}_{PO}(\mathbf{w}_\ell, b_\ell) = \min_{\mathbf{u}, \mathbf{w}_\ell, b_\ell} \max_s \frac{1}{n} \left(s^T \mathbf{D} \mathbf{Y}^T \mathbf{M}^T \mathbf{w}_\ell + s^T \mathbf{D} \mathbf{Z}^T \mathbf{w}_\ell + b_\ell s^T \mathbf{D} \mathbf{1}_n - s^T \mathbf{D} \mathbf{Y}_\ell - s^T \mathbf{u} + \frac{\|\mathbf{u}\|_{\ell_2}^2}{2} \right)$$

Note that the above is jointly convex in $(\mathbf{u}, \mathbf{w}_\ell, b_\ell)$ and concave in s . Thus, we consider the Auxiliary Optimization (AO) problem

$$\min_{\mathbf{u}, \mathbf{w}_\ell, b_\ell} \max_s \frac{1}{n} \left(s^T \mathbf{D} \mathbf{Y}^T \mathbf{M}^T \mathbf{w}_\ell + \sigma \|\mathbf{w}_\ell\|_{\ell_2} \mathbf{g}^T \mathbf{D} \mathbf{s} + \sigma \|\mathbf{D} \mathbf{s}\|_{\ell_2} \mathbf{h}^T \mathbf{w}_\ell + b_\ell s^T \mathbf{D} \mathbf{1}_n - s^T \mathbf{D} \mathbf{Y}_\ell - s^T \mathbf{u} + \frac{\|\mathbf{u}\|_{\ell_2}^2}{2} \right),$$

where $\mathbf{g} \in \mathbb{R}^n$ and $\mathbf{h} \in \mathbb{R}^d$ are independent Gaussian random vectors with i.i.d. $\mathcal{N}(0, 1)$ entries. Moreover, we carry out a change of variable $\mathbf{s} \rightarrow \mathbf{D}\mathbf{s}$ to arrive at

$$\min_{\mathbf{u}, \mathbf{w}_\ell, b_\ell} \max_{\mathbf{s}} \frac{1}{n} \left(\mathbf{s}^T \mathbf{Y}^T \mathbf{M}^T \mathbf{w}_\ell + \sigma \|\mathbf{w}_\ell\|_{\ell_2} \mathbf{g}^T \mathbf{s} + \sigma \|\mathbf{s}\|_{\ell_2} \mathbf{h}^T \mathbf{w}_\ell + b_\ell \mathbf{s}^T \mathbf{1}_n - \mathbf{s}^T \mathbf{Y}_\ell - \mathbf{s}^T \mathbf{D}^{-1} \mathbf{u} + \frac{\|\mathbf{u}\|_{\ell_2}^2}{2} \right).$$

Simplification of the AO. Maximizing over the direction of \mathbf{s} and setting its norm $\beta = \|\mathbf{s}\|_{\ell_2}$ above we arrive at

$$\begin{aligned} & \min_{\mathbf{u}, \mathbf{w}_\ell, b_\ell} \max_{\beta \geq 0} \max_{\mathbf{s}: \|\mathbf{s}\|_{\ell_2}=1} \frac{1}{n} \left(\beta \mathbf{s}^T \mathbf{Y}^T \mathbf{M}^T \mathbf{w}_\ell + \sigma \beta \|\mathbf{w}_\ell\|_{\ell_2} \mathbf{g}^T \mathbf{s} + \sigma \beta \mathbf{h}^T \mathbf{w}_\ell + b_\ell \beta \mathbf{s}^T \mathbf{1}_n - \beta \mathbf{s}^T \mathbf{Y}_\ell - \beta \mathbf{s}^T \mathbf{D}^{-1} \mathbf{u} + \frac{\|\mathbf{u}\|_{\ell_2}^2}{2} \right) \\ &= \min_{\mathbf{u}, \mathbf{w}_\ell, b_\ell} \max_{\beta \geq 0} \frac{1}{n} \left(\beta \|\mathbf{Y}^T \mathbf{M}^T \mathbf{w}_\ell + \sigma \|\mathbf{w}_\ell\|_{\ell_2} \mathbf{g} + b_\ell \mathbf{1}_n - \mathbf{Y}_\ell - \mathbf{D}^{-1} \mathbf{u} \right)_{\ell_2} + \sigma \beta \mathbf{h}^T \mathbf{w}_\ell + \frac{\|\mathbf{u}\|_{\ell_2}^2}{2} \\ &= \min_{\mathbf{u}, b_\ell} \max_{\beta \geq 0} \min_{\mathbf{w}_\ell} \frac{1}{n} \left(\beta \|\mathbf{Y}^T \mathbf{M}^T \mathbf{w}_\ell + \sigma \|\mathbf{w}_\ell\|_{\ell_2} \mathbf{g} + b_\ell \mathbf{1}_n - \mathbf{Y}_\ell - \mathbf{D}^{-1} \mathbf{u} \right)_{\ell_2} + \sigma \beta \mathbf{h}^T \mathbf{w}_\ell + \frac{\|\mathbf{u}\|_{\ell_2}^2}{2}. \end{aligned}$$

To continue, consider the singular value decomposition

$$\mathbf{M} = \mathbf{U} \Sigma \mathbf{V}^T = [\mathbf{u}_1 \quad \mathbf{u}_2 \quad \dots \quad \mathbf{u}_r] \text{diag}(\sigma_1, \sigma_2, \dots, \sigma_r) \begin{bmatrix} \mathbf{v}_1^T \\ \mathbf{v}_2^T \\ \vdots \\ \mathbf{v}_r^T \end{bmatrix}, \quad (\text{J.2})$$

with $r := \text{rank}(\mathbf{M}) \leq k$ and define the variable $\boldsymbol{\alpha} = \mathbf{U}^T \mathbf{w}_\ell$ and $\boldsymbol{\alpha}_\perp = \mathbf{U}_\perp^T \mathbf{w}_\ell$ where \mathbf{U}_\perp is the orthogonal complement of the columns of \mathbf{U} . With these definitions the above optimization problem reduces to

$$\min_{\mathbf{u}, b_\ell} \max_{\beta \geq 0} \min_{\boldsymbol{\alpha}} \min_{\boldsymbol{\alpha}_\perp} \frac{1}{n} \left(\beta \left\| \mathbf{Y}^T \mathbf{V} \Sigma \boldsymbol{\alpha} + \sigma \sqrt{\|\boldsymbol{\alpha}\|_{\ell_2}^2 + \|\boldsymbol{\alpha}_\perp\|_{\ell_2}^2} \mathbf{g} + b_\ell \mathbf{1}_n - \mathbf{Y}_\ell - \mathbf{D}^{-1} \mathbf{u} \right\|_{\ell_2} + \sigma \beta \mathbf{h}^T \mathbf{U} \boldsymbol{\alpha} + \sigma \beta \mathbf{h}^T \mathbf{U}_\perp \boldsymbol{\alpha}_\perp + \frac{\|\mathbf{u}\|_{\ell_2}^2}{2} \right)$$

Decomposing the optimization over $\boldsymbol{\alpha}_\perp$ in terms of its direction and norm $\alpha_0 = \|\boldsymbol{\alpha}_\perp\|_{\ell_2}$ we arrive at

$$\min_{\mathbf{u}, b_\ell} \max_{\beta \geq 0} \min_{\boldsymbol{\alpha}} \min_{\alpha_0 \geq 0} \frac{1}{n} \left(\beta \left\| \mathbf{Y}^T \mathbf{V} \Sigma \boldsymbol{\alpha} + \sigma \sqrt{\|\boldsymbol{\alpha}\|_{\ell_2}^2 + \alpha_0^2} \mathbf{g} + b_\ell \mathbf{1}_n - \mathbf{Y}_\ell - \mathbf{D}^{-1} \mathbf{u} \right\|_{\ell_2} + \sigma \beta \mathbf{h}^T \mathbf{U} \boldsymbol{\alpha} - \sigma \alpha_0 \beta \|\mathbf{U}_\perp^T \mathbf{h}\|_{\ell_2} + \frac{\|\mathbf{u}\|_{\ell_2}^2}{2} \right)$$

Since $\mathbf{U}^T \mathbf{h}$ is $r \leq k$ dimensional in our asymptotic regime the term $\frac{\mathbf{h}^T \mathbf{U} \boldsymbol{\alpha}}{n}$ can be ignored. Also replacing β with β/\sqrt{n} we thus arrive at

$$\begin{aligned} & \min_{\mathbf{u}, b_\ell} \max_{\beta \geq 0} \min_{\boldsymbol{\alpha}} \min_{\alpha_0 \geq 0} \frac{\beta}{\sqrt{n}} \left\| \mathbf{Y}^T \mathbf{V} \Sigma \boldsymbol{\alpha} + \sigma \sqrt{\|\boldsymbol{\alpha}\|_{\ell_2}^2 + \alpha_0^2} \mathbf{g} + b_\ell \mathbf{1}_n - \mathbf{Y}_\ell - \mathbf{D}^{-1} \mathbf{u} \right\|_{\ell_2} - \frac{1}{\sqrt{n}} \sigma \alpha_0 \beta \|\mathbf{U}_\perp^T \mathbf{h}\|_{\ell_2} + \frac{\|\mathbf{u}\|_{\ell_2}^2}{2n} \\ &= \min_{\mathbf{u}, b_\ell} \max_{\beta \geq 0} \min_{\boldsymbol{\alpha}} \min_{\alpha_0 \geq 0} \min_{\tau \geq 0} \frac{\beta}{2n\tau} \left\| \mathbf{Y}^T \mathbf{V} \Sigma \boldsymbol{\alpha} + \sigma \sqrt{\|\boldsymbol{\alpha}\|_{\ell_2}^2 + \alpha_0^2} \mathbf{g} + b_\ell \mathbf{1}_n - \mathbf{Y}_\ell - \mathbf{D}^{-1} \mathbf{u} \right\|_{\ell_2}^2 + \frac{\beta\tau}{2} \\ & \quad - \frac{1}{\sqrt{n}} \sigma \alpha_0 \beta \|\mathbf{U}_\perp^T \mathbf{h}\|_{\ell_2} + \frac{\|\mathbf{u}\|_{\ell_2}^2}{2n} \\ &= \min_{b_\ell} \max_{\beta \geq 0} \min_{\boldsymbol{\alpha}} \min_{\alpha_0 \geq 0} \min_{\tau \geq 0} \min_{\mathbf{u}} \frac{\beta}{2n\tau} \left\| \mathbf{Y}^T \mathbf{V} \Sigma \boldsymbol{\alpha} + \sigma \sqrt{\|\boldsymbol{\alpha}\|_{\ell_2}^2 + \alpha_0^2} \mathbf{g} + b_\ell \mathbf{1}_n - \mathbf{Y}_\ell - \mathbf{D}^{-1} \mathbf{u} \right\|_{\ell_2}^2 + \frac{\beta\tau}{2} \\ & \quad - \frac{1}{\sqrt{n}} \sigma \alpha_0 \beta \|\mathbf{U}_\perp^T \mathbf{h}\|_{\ell_2} + \frac{\|\mathbf{u}\|_{\ell_2}^2}{2n}. \end{aligned}$$

Setting the derivative with respect to \mathbf{u} to zero we arrive at

$$\mathbf{u} = \frac{\beta}{\tau} \mathbf{D}^{-1} \left(\mathbf{I} + \frac{\beta}{\tau} \mathbf{D}^{-2} \right)^{-1} \left(\mathbf{Y}^T \mathbf{V} \Sigma \boldsymbol{\alpha} + \sigma \sqrt{\|\boldsymbol{\alpha}\|_{\ell_2}^2 + \alpha_0^2} \mathbf{g} + b_\ell \mathbf{1}_n - \mathbf{Y}_\ell \right).$$

Plugging the latter into the above the AO simplifies to

$$\min_{b_\ell} \max_{\beta \geq 0} \min_{\alpha} \min_{\alpha_0 \geq 0} \min_{\tau \geq 0} \frac{\beta}{2\tau n} \text{trace} \left(\mathbf{t}^T \left(\mathbf{I} + \frac{\beta}{\tau} \mathbf{D}^{-2} \right)^{-1} \mathbf{t} \right) - \frac{1}{\sqrt{n}} \sigma \alpha_0 \beta \| \mathbf{U}_\perp^T \mathbf{h} \|_{\ell_2} + \frac{\beta \tau}{2}$$

where

$$\mathbf{t} := \mathbf{Y}^T \mathbf{V} \Sigma \alpha + \sigma \sqrt{\| \alpha \|_{\ell_2}^2 + \alpha_0^2} \mathbf{g} + b_\ell \mathbf{1}_n - \mathbf{Y}_\ell = \mathbf{Y}^T (\mathbf{V} \Sigma \alpha - \mathbf{e}_\ell) + \sigma \sqrt{\| \alpha \|_{\ell_2}^2 + \alpha_0^2} \mathbf{g} + b_\ell \mathbf{1}_n$$

To continue note that in our asymptotic regime we have

$$\frac{1}{\sqrt{n}} \| \mathbf{U}_\perp^T \mathbf{h} \|_{\ell_2} \xrightarrow{P} \sqrt{\gamma}$$

and the cross terms can be ignored so that in an asymptotic sense

$$\begin{aligned} & \frac{1}{n} \text{trace} \left(\mathbf{t}^T \left(\mathbf{I} + \frac{\beta}{\tau} \mathbf{D}^{-2} \right)^{-1} \mathbf{t} \right) \\ &= \frac{1}{n} \sigma^2 \left(\| \alpha \|_{\ell_2}^2 + \alpha_0^2 \right) \text{trace} \left(\left(\mathbf{I} + \frac{\beta}{\tau} \mathbf{D}^{-2} \right)^{-1} \right) \\ &+ \frac{1}{n} \left(\mathbf{Y}^T (\mathbf{V} \Sigma \alpha - \mathbf{e}_\ell) + b_\ell \mathbf{1}_n \right)^T \left(\mathbf{I} + \frac{\beta}{\tau} \mathbf{D}^{-2} \right)^{-1} \left(\mathbf{Y}^T (\mathbf{V} \Sigma \alpha - \mathbf{e}_\ell) + b_\ell \mathbf{1}_n \right) \end{aligned}$$

Therefore we arrive at

$$\begin{aligned} & \min_{b_\ell} \max_{\beta \geq 0} \min_{\alpha} \min_{\tau \geq 0} \min_{\alpha_0 \geq 0} \frac{\beta}{2\tau n} \sigma^2 \left(\| \alpha \|_{\ell_2}^2 + \alpha_0^2 \right) \text{trace} \left(\left(\mathbf{I} + \frac{\beta}{\tau} \mathbf{D}^{-2} \right)^{-1} \right) - \sigma \alpha_0 \beta \sqrt{\gamma} + \frac{\beta \tau}{2} \\ &+ \frac{\beta}{2\tau n} \left(\mathbf{Y}^T (\mathbf{V} \Sigma \alpha - \mathbf{e}_\ell) + b_\ell \mathbf{1}_n \right)^T \left(\mathbf{I} + \frac{\beta}{\tau} \mathbf{D}^{-2} \right)^{-1} \left(\mathbf{Y}^T (\mathbf{V} \Sigma \alpha - \mathbf{e}_\ell) + b_\ell \mathbf{1}_n \right), \end{aligned}$$

which can be rewritten in the form

$$\begin{aligned} & \min_{b_\ell} \max_{\beta \geq 0} \min_{\alpha} \min_{\tau \geq 0} \frac{\beta}{2\tau n} \sigma^2 \| \alpha \|_{\ell_2}^2 \text{trace} \left(\left(\mathbf{I} + \frac{\beta}{\tau} \mathbf{D}^{-2} \right)^{-1} \right) + \frac{\beta \tau}{2} \\ &+ \frac{\beta}{2\tau n} \left(\mathbf{Y}^T (\mathbf{V} \Sigma \alpha - \mathbf{e}_\ell) + b_\ell \mathbf{1}_n \right)^T \left(\mathbf{I} + \frac{\beta}{\tau} \mathbf{D}^{-2} \right)^{-1} \left(\mathbf{Y}^T (\mathbf{V} \Sigma \alpha - \mathbf{e}_\ell) + b_\ell \mathbf{1}_n \right) \\ &+ \frac{\beta}{2\tau n} \sigma^2 \alpha_0^2 \text{trace} \left(\left(\mathbf{I} + \frac{\beta}{\tau} \mathbf{D}^{-2} \right)^{-1} \right) - \alpha_0 \sigma \beta \sqrt{\gamma}. \end{aligned}$$

To continue further we shall assume $\mathbf{D} = \text{diag}(\mathbf{Y}^T \boldsymbol{\omega})$. Note that in this case

$$\frac{1}{n} \text{trace} \left(\left(\mathbf{I} + \frac{\beta}{\tau} \mathbf{D}^{-2} \right)^{-1} \right) = \frac{1}{n} \sum_{i=1}^n \frac{(\mathbf{y}_i^T \boldsymbol{\omega})^2}{(\mathbf{y}_i^T \boldsymbol{\omega})^2 + \frac{\beta}{\tau}} = \sum_{\ell=1}^k \frac{n_\ell}{n} \frac{\omega_\ell^2}{\omega_\ell^2 + \frac{\beta}{\tau}} \xrightarrow{P} \sum_{\ell=1}^k \frac{\pi_\ell \omega_\ell^2}{\omega_\ell^2 + \frac{\beta}{\tau}}.$$

Also,

$$\begin{aligned} & \frac{1}{n} \left(\mathbf{Y}^T (\mathbf{V} \Sigma \alpha - \mathbf{e}_\ell) + b_\ell \mathbf{1}_n \right)^T \left(\mathbf{I} + \frac{\beta}{\tau} \mathbf{D}^{-2} \right)^{-1} \left(\mathbf{Y}^T (\mathbf{V} \Sigma \alpha - \mathbf{e}_\ell) + b_\ell \mathbf{1}_n \right) \\ &= \frac{1}{n} (\mathbf{V} \Sigma \alpha - \mathbf{e}_\ell)^T \mathbf{Y} \left(\mathbf{I} + \frac{\beta}{\tau} \mathbf{D}^{-2} \right)^{-1} \mathbf{Y}^T (\mathbf{V} \Sigma \alpha - \mathbf{e}_\ell) \\ &+ \frac{2}{n} b_\ell \mathbf{1}_n^T \left(\mathbf{I} + \frac{\beta}{\tau} \mathbf{D}^{-2} \right)^{-1} \mathbf{Y}^T (\mathbf{V} \Sigma \alpha - \mathbf{e}_\ell) + \frac{b_\ell^2}{n} \text{trace} \left(\left(\mathbf{I} + \frac{\beta}{\tau} \mathbf{D}^{-2} \right)^{-1} \right) \\ &\xrightarrow{P} (\mathbf{V} \Sigma \alpha - \mathbf{e}_\ell)^T \text{diag} \left(\frac{\pi_1 \omega_1^2}{\omega_1^2 + \frac{\beta}{\tau}}, \frac{\pi_2 \omega_2^2}{\omega_2^2 + \frac{\beta}{\tau}}, \dots, \frac{\pi_k \omega_k^2}{\omega_k^2 + \frac{\beta}{\tau}} \right) (\mathbf{V} \Sigma \alpha - \mathbf{e}_\ell) \\ &+ 2b_\ell \begin{bmatrix} \frac{\pi_1 \omega_1^2}{\omega_1^2 + \frac{\beta}{\tau}} & \frac{\pi_2 \omega_2^2}{\omega_2^2 + \frac{\beta}{\tau}} & \dots & \frac{\pi_k \omega_k^2}{\omega_k^2 + \frac{\beta}{\tau}} \end{bmatrix} (\mathbf{V} \Sigma \alpha - \mathbf{e}_\ell) + b_\ell^2 \left(\sum_{\ell=1}^k \frac{\pi_\ell \omega_\ell^2}{\omega_\ell^2 + \frac{\beta}{\tau}} \right). \end{aligned}$$

Next define

$$\begin{aligned} \mathbf{A}(\eta) &:= \begin{bmatrix} \sigma^2(\boldsymbol{\pi}^T \boldsymbol{\nu}(\eta)) \mathbf{I} + \boldsymbol{\Sigma} \mathbf{V}^T \text{diag}(\boldsymbol{\pi}) \text{diag}(\boldsymbol{\nu}(\eta)) \mathbf{V} \boldsymbol{\Sigma} & \boldsymbol{\Sigma} \mathbf{V}^T \text{diag}(\boldsymbol{\nu}(\eta)) \boldsymbol{\pi} \\ \boldsymbol{\pi}^T \text{diag}(\boldsymbol{\nu}(\eta)) \mathbf{V} \boldsymbol{\Sigma} & \boldsymbol{\pi}^T \boldsymbol{\nu}(\eta) \end{bmatrix} \\ \mathbf{c}_\ell &:= \begin{bmatrix} \boldsymbol{\Sigma} \mathbf{V}^T \mathbf{e}_\ell \\ 1 \end{bmatrix}, \end{aligned} \quad (\text{J.3})$$

where

$$\boldsymbol{\nu}(\eta) = \frac{1}{\gamma} \begin{bmatrix} \frac{\omega_1^2}{\omega_1^2 + \eta} \\ \frac{\omega_2^2}{\omega_2^2 + \eta} \\ \vdots \\ \frac{\omega_k^2}{\omega_k^2 + \eta} \end{bmatrix}. \quad (\text{J.4})$$

We thus arrive at

$$\begin{aligned} \min_{\boldsymbol{\alpha}} \min_{b_\ell} \max_{\alpha_0 \geq 0} \max_{\beta \geq 0} \min_{\tau \geq 0} & \frac{\gamma\beta}{2\tau} \left(\pi_\ell \boldsymbol{\nu}_\ell \left(\frac{\beta}{\tau} \right) + [\boldsymbol{\alpha}^T \ b_\ell] \mathbf{A} \left(\frac{\beta}{\tau} \right) [\boldsymbol{\alpha}]_{b_\ell} - 2\pi_\ell \boldsymbol{\nu}_\ell \left(\frac{\beta}{\tau} \right) \mathbf{c}_\ell^T [\boldsymbol{\alpha}]_{b_\ell} \right) \\ & + \frac{\gamma\beta}{2\tau} \sigma^2 \left(\boldsymbol{\pi}^T \boldsymbol{\nu} \left(\frac{\beta}{\tau} \right) \right) \alpha_0^2 - \alpha_0 \sigma \beta \sqrt{\gamma} + \frac{\beta\tau}{2}. \end{aligned}$$

Deterministic Analysis of the AO. Setting the derivative of the above with respect to α_0 to zero we arrive at

$$\frac{\gamma\beta}{\tau} \sigma^2 \left(\boldsymbol{\pi}^T \boldsymbol{\nu} \left(\frac{\beta}{\tau} \right) \right) \alpha_0 - \sigma \beta \sqrt{\gamma} = 0 \quad \Rightarrow \quad \alpha_0 = \frac{\tau}{\sigma \sqrt{\gamma} \left(\boldsymbol{\pi}^T \boldsymbol{\nu} \left(\frac{\beta}{\tau} \right) \right)}.$$

Note that the above objective has the form

$$f \left(\frac{\beta}{\tau} \right) - \alpha_0 \sigma \beta \sqrt{\gamma} + \frac{\beta\tau}{2}$$

with

$$f(\eta) := \frac{\eta\gamma}{2} \left(\pi_\ell \boldsymbol{\nu}_\ell(\eta) + [\boldsymbol{\alpha}^T \ b_\ell] \mathbf{A}(\eta) [\boldsymbol{\alpha}]_{b_\ell} - 2\pi_\ell \boldsymbol{\nu}_\ell(\eta) \mathbf{c}_\ell^T [\boldsymbol{\alpha}]_{b_\ell} \right) + \frac{\gamma\eta}{2} \sigma^2 (\boldsymbol{\pi}^T \boldsymbol{\nu}(\eta)) \alpha_0^2.$$

Thus setting the derivatives with respect to β and τ to zero, we have

$$\frac{1}{\tau} f' \left(\frac{\beta}{\tau} \right) - \alpha_0 \sigma \sqrt{\gamma} + \frac{\tau}{2} = 0 \quad \Rightarrow \quad f' \left(\frac{\beta}{\tau} \right) - \alpha_0 \sigma \sqrt{\gamma} \tau + \frac{\tau^2}{2} = 0 \quad \Rightarrow \quad f' \left(\frac{\beta}{\tau} \right) = \tau^2 \left(\frac{1}{\boldsymbol{\pi}^T \boldsymbol{\nu} \left(\frac{\beta}{\tau} \right)} - \frac{1}{2} \right)$$

and

$$-\frac{\beta}{\tau^2} f' \left(\frac{\beta}{\tau} \right) + \frac{\beta}{2} = 0 \quad \Rightarrow \quad \tau^2 = 2f' \left(\frac{\beta}{\tau} \right).$$

Combining the latter two we conclude that $\boldsymbol{\pi}^T \boldsymbol{\nu} \left(\frac{\beta}{\tau} \right) = 1$. Thus, $\eta = \frac{\beta}{\tau}$ is the solution to $\boldsymbol{\pi}^T \boldsymbol{\nu}(\eta) = 1$. To calculate τ and hence α_0 we calculate f' which is equal to

$$\begin{aligned} f'(\eta) &= \frac{\gamma}{2} \left(\pi_\ell \boldsymbol{\nu}_\ell(\eta) + [\boldsymbol{\alpha}^T \ b_\ell] \mathbf{A}(\eta) [\boldsymbol{\alpha}]_{b_\ell} - 2\pi_\ell \boldsymbol{\nu}_\ell(\eta) \mathbf{c}_\ell^T [\boldsymbol{\alpha}]_{b_\ell} \right) + \frac{\gamma}{2} \sigma^2 (\boldsymbol{\pi}^T \boldsymbol{\nu}(\eta)) \alpha_0^2 + \frac{\gamma\eta}{2} \sigma^2 \alpha_0^2 (\boldsymbol{\pi}^T \boldsymbol{\nu}'(\eta)) \\ &+ \frac{\gamma\eta}{2} \left(\pi_\ell \boldsymbol{\nu}'_\ell(\eta) + [\boldsymbol{\alpha}^T \ b_\ell] \mathbf{A}'(\eta) [\boldsymbol{\alpha}]_{b_\ell} - 2\pi_\ell \boldsymbol{\nu}'_\ell(\eta) \mathbf{c}_\ell^T [\boldsymbol{\alpha}]_{b_\ell} \right), \end{aligned}$$

where

$$\boldsymbol{\nu}'(\eta) = -\frac{1}{\gamma} \begin{bmatrix} \frac{\omega_1^2}{(\omega_1^2 + \eta)^2} \\ \frac{\omega_2^2}{(\omega_2^2 + \eta)^2} \\ \vdots \\ \frac{\omega_k^2}{(\omega_k^2 + \eta)^2} \end{bmatrix}$$

$$\mathbf{A}'(\eta) := \begin{bmatrix} \sigma^2 (\boldsymbol{\pi}^T \boldsymbol{\nu}'(\eta)) \mathbf{I} + \boldsymbol{\Sigma} \mathbf{V}^T \text{diag}(\boldsymbol{\pi}) \text{diag}(\boldsymbol{\nu}'(\eta)) \mathbf{V} \boldsymbol{\Sigma} & \boldsymbol{\Sigma} \mathbf{V}^T \text{diag}(\boldsymbol{\nu}'(\eta)) \boldsymbol{\pi} \\ \boldsymbol{\pi}^T \text{diag}(\boldsymbol{\nu}'(\eta)) \mathbf{V} \boldsymbol{\Sigma} & \boldsymbol{\pi}^T \boldsymbol{\nu}'(\eta) \end{bmatrix}$$

$$\mathbf{c}_\ell := \begin{bmatrix} \boldsymbol{\Sigma} \mathbf{V}^T \mathbf{e}_\ell \\ 1 \end{bmatrix}.$$

Now note that at the optimal point we have

$$f'(\eta) = \frac{\tau^2}{2} = \frac{\gamma}{2} \sigma^2 (\boldsymbol{\pi}^T \boldsymbol{\nu}(\eta))^2 \alpha_0^2.$$

Thus from the above we can conclude that

$$\begin{aligned} \alpha_0^2 &= -\frac{1}{\eta \sigma^2 (\boldsymbol{\pi}^T \boldsymbol{\nu}'(\eta))} \left(\pi_\ell \boldsymbol{\nu}_\ell(\eta) + [\boldsymbol{\alpha}^T \ b_\ell] \mathbf{A}(\eta) \begin{bmatrix} \boldsymbol{\alpha} \\ b_\ell \end{bmatrix} - 2\pi_\ell \boldsymbol{\nu}_\ell(\eta) \mathbf{c}_\ell^T \begin{bmatrix} \boldsymbol{\alpha} \\ b_\ell \end{bmatrix} \right) \\ &\quad - \frac{1}{\sigma^2 (\boldsymbol{\pi}^T \boldsymbol{\nu}'(\eta))} \left(\pi_\ell \boldsymbol{\nu}'_\ell(\eta) + [\boldsymbol{\alpha}^T \ b_\ell] \mathbf{A}'(\eta) \begin{bmatrix} \boldsymbol{\alpha} \\ b_\ell \end{bmatrix} - 2\pi_\ell \boldsymbol{\nu}'_\ell(\eta) \mathbf{c}_\ell^T \begin{bmatrix} \boldsymbol{\alpha} \\ b_\ell \end{bmatrix} \right). \end{aligned}$$

Thus the AO optimization problem reduces to

$$\min_{b_\ell} \min_{\boldsymbol{\alpha}} \frac{\eta \gamma}{2} \left(\pi_\ell \boldsymbol{\nu}_\ell + [\boldsymbol{\alpha}^T \ b_\ell] \mathbf{A} \begin{bmatrix} \boldsymbol{\alpha} \\ b_\ell \end{bmatrix} - 2\pi_\ell \boldsymbol{\nu}_\ell \mathbf{c}_\ell^T \begin{bmatrix} \boldsymbol{\alpha} \\ b_\ell \end{bmatrix} \right),$$

where η is the solution to

$$\sum_{\ell=1}^k \frac{\pi_\ell \omega_\ell^2}{\omega_\ell^2 + \eta} = \gamma,$$

and

$$\begin{aligned} \mathbf{A}(\eta) &:= \begin{bmatrix} \sigma^2 \mathbf{I}_r + \boldsymbol{\Sigma} \mathbf{V}^T \text{diag}(\boldsymbol{\pi} \odot \boldsymbol{\nu}) \mathbf{V} \boldsymbol{\Sigma} & \boldsymbol{\Sigma} \mathbf{V}^T (\boldsymbol{\pi} \odot \boldsymbol{\nu}) \\ (\boldsymbol{\pi} \odot \boldsymbol{\nu})^T \mathbf{V} \boldsymbol{\Sigma} & 1 \end{bmatrix} \\ \mathbf{c}_\ell &:= \begin{bmatrix} \boldsymbol{\Sigma} \mathbf{V}^T \mathbf{e}_\ell \\ 1 \end{bmatrix}, \end{aligned} \tag{J.5}$$

with

$$\boldsymbol{\nu} := \frac{1}{\gamma} \begin{bmatrix} \frac{\omega_1^2}{\omega_1^2 + \eta} \\ \frac{\omega_2^2}{\omega_2^2 + \eta} \\ \vdots \\ \frac{\omega_k^2}{\omega_k^2 + \eta} \end{bmatrix}.$$

First, note that the matrix \mathbf{A} is positive definite. This can be checked by computing the Schur complement of \mathbf{A} :

$$\boldsymbol{\Delta} := \sigma^2 \mathbf{I}_r + \boldsymbol{\Sigma} \mathbf{V}^T \mathbf{P} \mathbf{V} \boldsymbol{\Sigma} := \sigma^2 \mathbf{I}_r + \boldsymbol{\Sigma} \mathbf{V}^T \left(\text{diag}(\boldsymbol{\pi} \odot \boldsymbol{\nu}) - (\boldsymbol{\pi} \odot \boldsymbol{\nu}) (\boldsymbol{\pi} \odot \boldsymbol{\nu})^T \right) \mathbf{V} \boldsymbol{\Sigma} > \mathbf{0}_{r \times r}. \tag{J.6}$$

Positive definiteness above holds because $\mathbf{P} := \left(\text{diag}(\boldsymbol{\pi} \odot \boldsymbol{\nu}) - (\boldsymbol{\pi} \odot \boldsymbol{\nu}) (\boldsymbol{\pi} \odot \boldsymbol{\nu})^T \right) \succeq \mathbf{0}_{k \times k}$. Thus the objective is a strictly convex quadratic and is jointly convex in its arguments. We proceed by minimizing the objective over $(\boldsymbol{\alpha}, b_\ell)$ which is equal to

$$\begin{bmatrix} \widehat{\boldsymbol{\alpha}} \\ \widehat{b}_\ell \end{bmatrix} = \pi_\ell \boldsymbol{\nu}_\ell \mathbf{A}^{-1} \mathbf{c}_\ell = \pi_\ell \boldsymbol{\nu}_\ell \begin{bmatrix} -\boldsymbol{\Delta}^{-1} \boldsymbol{\Sigma} \mathbf{V}^T (\boldsymbol{\pi} \odot \boldsymbol{\nu} - \mathbf{e}_\ell) \\ 1 + (\boldsymbol{\pi} \odot \boldsymbol{\nu})^T \mathbf{V} \boldsymbol{\Sigma} \boldsymbol{\Delta}^{-1} \boldsymbol{\Sigma} \mathbf{V}^T (\boldsymbol{\pi} \odot \boldsymbol{\nu} - \mathbf{e}_\ell) \end{bmatrix}. \tag{J.7}$$

Thus, the minimum value attained is

$$\begin{aligned} &-\pi_\ell^2 \nu_\ell^2 \left[-(\boldsymbol{\pi} \odot \boldsymbol{\nu} - \mathbf{e}_\ell)^T \mathbf{V} \boldsymbol{\Sigma} \quad 1 \right] \begin{bmatrix} \boldsymbol{\Delta}^{-1} & \mathbf{0} \\ \mathbf{0}^T & 1 \end{bmatrix} \begin{bmatrix} -\boldsymbol{\Sigma} \mathbf{V}^T (\boldsymbol{\pi} \odot \boldsymbol{\nu} - \mathbf{e}_\ell) \\ 1 \end{bmatrix} \\ &= -\pi_\ell^2 \nu_\ell^2 \left(1 + (\boldsymbol{\pi} \odot \boldsymbol{\nu} - \mathbf{e}_\ell)^T \mathbf{V} \boldsymbol{\Sigma} \boldsymbol{\Delta}^{-1} \boldsymbol{\Sigma} \mathbf{V}^T (\boldsymbol{\pi} \odot \boldsymbol{\nu} - \mathbf{e}_\ell) \right). \end{aligned} \tag{J.8}$$

Thus the objective reduces to

$$\frac{\eta \gamma}{2} \left(\pi_\ell \nu_\ell (1 - \pi_\ell \nu_\ell) - \pi_\ell^2 \nu_\ell^2 (\boldsymbol{\pi} \odot \boldsymbol{\nu} - \mathbf{e}_\ell)^T \mathbf{V} \boldsymbol{\Sigma} \boldsymbol{\Delta}^{-1} \boldsymbol{\Sigma} \mathbf{V}^T (\boldsymbol{\pi} \odot \boldsymbol{\nu} - \mathbf{e}_\ell) \right). \tag{J.9}$$

Therefore,

$$\begin{aligned} \alpha_0^2 &= -\frac{1}{\eta \sigma^2 (\boldsymbol{\pi}^T \boldsymbol{\nu}'(\eta))} \left(\pi_\ell \nu_\ell (1 - \pi_\ell \nu_\ell) - \pi_\ell^2 \nu_\ell^2 (\boldsymbol{\pi} \odot \boldsymbol{\nu} - \mathbf{e}_\ell)^T \mathbf{V} \boldsymbol{\Sigma} \boldsymbol{\Delta}^{-1} \boldsymbol{\Sigma} \mathbf{V}^T (\boldsymbol{\pi} \odot \boldsymbol{\nu} - \mathbf{e}_\ell) \right) \\ &\quad - \frac{1}{\sigma^2 (\boldsymbol{\pi}^T \boldsymbol{\nu}'(\eta))} \left(\pi_\ell \boldsymbol{\nu}'_\ell(\eta) + [\boldsymbol{\alpha}^T \ b_\ell] \mathbf{A}'(\eta) \begin{bmatrix} \boldsymbol{\alpha} \\ b_\ell \end{bmatrix} - 2\pi_\ell \boldsymbol{\nu}'_\ell(\eta) \mathbf{c}_\ell^T \begin{bmatrix} \boldsymbol{\alpha} \\ b_\ell \end{bmatrix} \right), \end{aligned}$$

where

$$\begin{bmatrix} \widehat{\boldsymbol{\alpha}} \\ \widehat{b}_\ell \end{bmatrix} = \pi_\ell \nu_\ell \begin{bmatrix} -\boldsymbol{\Delta}^{-1} \boldsymbol{\Sigma} \boldsymbol{V}^T (\boldsymbol{\pi} \odot \boldsymbol{\nu} - \mathbf{e}_\ell) \\ 1 + (\boldsymbol{\pi} \odot \boldsymbol{\nu})^T \boldsymbol{V} \boldsymbol{\Sigma} \boldsymbol{\Delta}^{-1} \boldsymbol{\Sigma} \boldsymbol{V}^T (\boldsymbol{\pi} \odot \boldsymbol{\nu} - \mathbf{e}_\ell) \end{bmatrix} \quad (\text{J.10})$$

and

$$\begin{aligned} \boldsymbol{\nu}'(\eta) &= -\frac{1}{\gamma} \begin{bmatrix} \frac{\omega_1^2}{(\omega_1^2 + \eta)^2} \\ \frac{\omega_2^2}{(\omega_2^2 + \eta)^2} \\ \vdots \\ \frac{\omega_k^2}{(\omega_k^2 + \eta)^2} \end{bmatrix} \\ \mathbf{A}'(\eta) &:= \begin{bmatrix} \sigma^2 (\boldsymbol{\pi}^T \boldsymbol{\nu}'(\eta)) \mathbf{I} + \boldsymbol{\Sigma} \boldsymbol{V}^T \text{diag}(\boldsymbol{\pi}) \text{diag}(\boldsymbol{\nu}'(\eta)) \boldsymbol{V} \boldsymbol{\Sigma} & \boldsymbol{\Sigma} \boldsymbol{V}^T \text{diag}(\boldsymbol{\nu}'(\eta)) \boldsymbol{\pi} \\ \boldsymbol{\pi}^T \text{diag}(\boldsymbol{\nu}'(\eta)) \boldsymbol{V} \boldsymbol{\Sigma} & \boldsymbol{\pi}^T \boldsymbol{\nu}'(\eta) \end{bmatrix}. \end{aligned}$$

To continue note that

$$\begin{aligned} \pi_\ell \boldsymbol{\nu}'_\ell(\eta) + [\boldsymbol{\alpha}^T \quad b_\ell] \mathbf{A}'(\eta) \begin{bmatrix} \boldsymbol{\alpha} \\ b_\ell \end{bmatrix} &- 2\pi_\ell \boldsymbol{\nu}'_\ell(\eta) \mathbf{c}_\ell^T \begin{bmatrix} \boldsymbol{\alpha} \\ b_\ell \end{bmatrix} \\ &= \pi_\ell \boldsymbol{\nu}'_\ell(\eta) + \pi_\ell^2 \nu_\ell^2 \mathbf{c}_\ell^T \mathbf{A}^{-1} \mathbf{A}'(\eta) \mathbf{A}^{-1} \mathbf{c}_\ell - 2\pi_\ell^2 \nu_\ell' \nu_\ell \mathbf{c}_\ell^T \mathbf{A}^{-1} \mathbf{c}_\ell. \end{aligned}$$

Thus,

$$\begin{aligned} \alpha_0^2 &= -\frac{1}{\eta \sigma^2 (\boldsymbol{\pi}^T \boldsymbol{\nu}'(\eta))} \left(\pi_\ell \nu_\ell (1 - \pi_\ell \nu_\ell) - \pi_\ell^2 \nu_\ell^2 (\boldsymbol{\pi} \odot \boldsymbol{\nu} - \mathbf{e}_\ell)^T \boldsymbol{V} \boldsymbol{\Sigma} \boldsymbol{\Delta}^{-1} \boldsymbol{\Sigma} \boldsymbol{V}^T (\boldsymbol{\pi} \odot \boldsymbol{\nu} - \mathbf{e}_\ell) \right) \\ &\quad - \frac{1}{\sigma^2 (\boldsymbol{\pi}^T \boldsymbol{\nu}'(\eta))} \left(\pi_\ell \boldsymbol{\nu}'_\ell(\eta) + \pi_\ell^2 \nu_\ell^2 \mathbf{c}_\ell^T \mathbf{A}^{-1} \mathbf{A}'(\eta) \mathbf{A}^{-1} \mathbf{c}_\ell - 2\pi_\ell^2 \nu_\ell' \nu_\ell \mathbf{c}_\ell^T \mathbf{A}^{-1} \mathbf{c}_\ell \right) \\ &= \frac{\gamma}{\eta \sigma^2 \left(\sum_{\ell=1}^k \frac{\pi_\ell \omega_\ell^2}{(\omega_\ell^2 + \eta)^2} \right)} \left(\pi_\ell \nu_\ell (1 - \pi_\ell \nu_\ell) - \pi_\ell^2 \nu_\ell^2 (\boldsymbol{\pi} \odot \boldsymbol{\nu} - \mathbf{e}_\ell)^T \boldsymbol{V} \boldsymbol{\Sigma} \boldsymbol{\Delta}^{-1} \boldsymbol{\Sigma} \boldsymbol{V}^T (\boldsymbol{\pi} \odot \boldsymbol{\nu} - \mathbf{e}_\ell) \right) \\ &\quad + \frac{\gamma}{\sigma^2 \left(\sum_{\ell=1}^k \frac{\pi_\ell \omega_\ell^2}{(\omega_\ell^2 + \eta)^2} \right)} \left(\pi_\ell \boldsymbol{\nu}'_\ell(\eta) + \pi_\ell^2 \nu_\ell^2 \mathbf{c}_\ell^T \mathbf{A}^{-1} \mathbf{A}'(\eta) \mathbf{A}^{-1} \mathbf{c}_\ell - 2\pi_\ell^2 \nu_\ell' \nu_\ell \mathbf{c}_\ell^T \mathbf{A}^{-1} \mathbf{c}_\ell \right) \\ &:= \frac{\zeta}{\sigma^2} \left(\pi_\ell \nu_\ell (1 - \pi_\ell \nu_\ell) - \pi_\ell^2 \nu_\ell^2 (\boldsymbol{\pi} \odot \boldsymbol{\nu} - \mathbf{e}_\ell)^T \boldsymbol{V} \boldsymbol{\Sigma} \boldsymbol{\Delta}^{-1} \boldsymbol{\Sigma} \boldsymbol{V}^T (\boldsymbol{\pi} \odot \boldsymbol{\nu} - \mathbf{e}_\ell) \right) \\ &\quad + \frac{\zeta \eta}{\sigma^2} \left(\pi_\ell \boldsymbol{\nu}'_\ell(\eta) + \pi_\ell^2 \nu_\ell^2 \mathbf{c}_\ell^T \mathbf{A}^{-1} \mathbf{A}'(\eta) \mathbf{A}^{-1} \mathbf{c}_\ell - 2\pi_\ell^2 \nu_\ell' \nu_\ell \mathbf{c}_\ell^T \mathbf{A}^{-1} \mathbf{c}_\ell \right) \end{aligned}$$

where $\zeta := \frac{\gamma}{\eta \left(\sum_{\ell=1}^k \frac{\pi_\ell \omega_\ell^2}{(\omega_\ell^2 + \eta)^2} \right)}$.

Asymptotic predictions. First, from (J.7) the bias term converges as follows:

$$\widehat{b}_\ell \xrightarrow{P} \pi_\ell \nu_\ell \left(1 + (\boldsymbol{\pi} \odot \boldsymbol{\nu})^T \boldsymbol{V} \boldsymbol{\Sigma} \boldsymbol{\Delta}^{-1} \boldsymbol{\Sigma} \boldsymbol{V}^T (\boldsymbol{\pi} \odot \boldsymbol{\nu} - \mathbf{e}_\ell) \right).$$

Thus,

$$\widehat{\mathbf{b}} \xrightarrow{P} (\mathbf{I}_k - \mathbf{P} \boldsymbol{V} \boldsymbol{\Sigma} \boldsymbol{\Delta}^{-1} \boldsymbol{\Sigma} \boldsymbol{V}^T) (\boldsymbol{\pi} \odot \boldsymbol{\nu}).$$

Recall that $\boldsymbol{\alpha} = \mathbf{U}^T \mathbf{w}_\ell$. Thus, the correlations $\langle \boldsymbol{\mu}_i, \mathbf{w}_\ell \rangle$, $i \in [k]$ converge as follows:

$$\mathbf{M}^T \mathbf{w}_\ell = \boldsymbol{V} \boldsymbol{\Sigma} \mathbf{U}^T \mathbf{w}_\ell \xrightarrow{P} \boldsymbol{V} \boldsymbol{\Sigma} \widehat{\boldsymbol{\alpha}} = -\pi_\ell \nu_\ell \boldsymbol{V} \boldsymbol{\Sigma} \boldsymbol{\Delta}^{-1} \boldsymbol{\Sigma} \boldsymbol{V}^T (\boldsymbol{\pi} \odot \boldsymbol{\nu} - \mathbf{e}_\ell). \quad (\text{J.11})$$

Here, convergence applies element-wise to the entries of the involved random vectors. Moreover, from the analysis above we can predict the limit of the norm $\|\mathbf{w}_\ell\|_{\ell_2}$. For this, note that $\|\mathbf{w}_\ell\|_{\ell_2}^2 = \widehat{\alpha}_0^2 + \widehat{\boldsymbol{\alpha}}^T \widehat{\boldsymbol{\alpha}}$. Thus,

$$\begin{aligned} \|\mathbf{w}_\ell\|_{\ell_2}^2 &\xrightarrow{P} \frac{\zeta}{\sigma^2} \pi_\ell \nu_\ell (1 - \pi_\ell \nu_\ell) + \pi_\ell^2 \nu_\ell^2 (\boldsymbol{\pi} \odot \boldsymbol{\nu} - \mathbf{e}_\ell)^T \boldsymbol{V} \boldsymbol{\Sigma} \boldsymbol{\Delta}^{-1} \left(\boldsymbol{\Delta}^{-1} - \frac{\zeta}{\sigma^2} \mathbf{I}_r \right) \boldsymbol{\Sigma} \boldsymbol{V}^T (\boldsymbol{\pi} \odot \boldsymbol{\nu} - \mathbf{e}_\ell) \\ &\quad + \frac{\eta \zeta}{\sigma^2} \left(\pi_\ell \boldsymbol{\nu}'_\ell(\eta) + \pi_\ell^2 \nu_\ell^2 \mathbf{c}_\ell^T \mathbf{A}^{-1} \mathbf{A}' \mathbf{A}^{-1} \mathbf{c}_\ell - 2\pi_\ell^2 \nu_\ell' \nu_\ell \mathbf{c}_\ell^T \mathbf{A}^{-1} \mathbf{c}_\ell \right). \end{aligned} \quad (\text{J.12})$$

J.2 Computing $\Sigma_{w,w}$

In the previous section we used the CGMT to predict the bias \widehat{b}_ℓ , the correlations $\langle \mu_i, \widehat{w}_\ell \rangle$, $i[k]$ and the norm $\|\widehat{w}_\ell\|_{\ell_2}$ for all $\ell \in [k]$ members of the multi-output classifier. Here, we show how to compute the limits of the cross-correlations $\langle \widehat{w}_\ell, \widehat{w}_j \rangle$, $\ell \neq j \in [k]$.

Lemma J.1 For $\ell \neq j \in [k]$ let \widehat{w}_ℓ \widehat{w}_j be solutions to the least-squares minimization (??), i.e.,

$$\begin{aligned} & (\widehat{w}_\ell, \widehat{w}_j, \widehat{b}_\ell, \widehat{b}_j) \\ &= \arg \min_{\mathbf{w}_\ell, \mathbf{w}_j, b_\ell, b_j} \left\{ \frac{1}{2n} \|\mathbf{D}(\mathbf{Y}_\ell - \mathbf{X}^T \mathbf{w}_\ell - b_\ell \mathbf{1}_n)\|_{\ell_2}^2 + \frac{1}{2n} \|\mathbf{D}(\mathbf{Y}_j - \mathbf{X}^T \mathbf{w}_j - b_j \mathbf{1}_n)\|_{\ell_2}^2 \right\}. \end{aligned}$$

Denote $\widehat{w}_{\ell,j} := \widehat{w}_\ell + \widehat{w}_j$ and $\widehat{b}_{\ell,j} := \widehat{b}_\ell + \widehat{b}_j$. Then, $(\widehat{w}_{\ell,j}, \widehat{b}_{\ell,j})$ is a minimizer in the following least-squares problem:

$$(\widehat{w}_{\ell,j}, \widehat{b}_{\ell,j}) = \arg \min_{\mathbf{w}, b} \frac{1}{2n} \|\mathbf{D}(\mathbf{Y}_\ell + \mathbf{Y}_j - \mathbf{X}^T \mathbf{w} - b \mathbf{1}_n)\|_{\ell_2}^2 \quad (\text{J.13})$$

Proof Clearly the minimization in (J.13) is convex. Thus, it suffices to prove that $\widehat{w}_\ell + \widehat{w}_j$ satisfies the KKT conditions. First, by optimality of \widehat{w}_ℓ , we have that

$$\mathbf{X} \mathbf{D}^2 (\mathbf{Y}_\ell - \mathbf{X}^T \widehat{w}_\ell - \widehat{b}_\ell \mathbf{1}_n) = 0$$

Similarly, for \widehat{w}_j :

$$\mathbf{X} \mathbf{D}^2 (\mathbf{Y}_j - \mathbf{X}^T \widehat{w}_j - \widehat{b}_j \mathbf{1}_n) = 0.$$

Adding the equations on the above displays we find that

$$\mathbf{X} \mathbf{D}^2 (\mathbf{Y}_\ell + \mathbf{Y}_j - \mathbf{X}^T (\widehat{w}_j + \widehat{w}_\ell) - (\widehat{b}_j + \widehat{b}_\ell) \mathbf{1}_n) = 0.$$

Recognize that this coincides with the optimality condition for (J.13). Thus, the proof is complete. ■

Thanks to Lemma J.1, we can use the CGMT to characterize the limiting behavior of $\|\widehat{w}_\ell + \widehat{w}_j\|_{\ell_2}$. Observe that this immediately gives the limit of $\langle \widehat{w}_\ell, \widehat{w}_j \rangle$ since

$$\langle \widehat{w}_\ell, \widehat{w}_j \rangle = \frac{\|\widehat{w}_\ell + \widehat{w}_j\|_{\ell_2}^2 - \|\widehat{w}_\ell\|_{\ell_2}^2 - \|\widehat{w}_j\|_{\ell_2}^2}{2}. \quad (\text{J.14})$$

The analysis of (J.13) is very similar to that of (H.1). In particular we use the following decomposition

$$\mathbf{w}_{\ell,j} = \sum_{i=1}^r \beta_i \mathbf{u}_i + \beta_0 \mathbf{w}_{\ell,j}^\perp,$$

with $\|\mathbf{w}_{\ell,j}^\perp\|_{\ell_2} = 1$ and $\mathbf{U}^T \mathbf{w}_{\ell,j}^\perp = \mathbf{0}_r$. This allows us to arrive at

$$\begin{aligned} & \min_{\beta} \min_{\mathbf{b}_{\ell,j}} \max_{\alpha_0 \geq 0} \max_{\beta \geq 0} \min_{\tau \geq 0} \\ & \frac{\gamma \beta}{\tau} \left(\pi_\ell \boldsymbol{\nu}_\ell \left(\frac{\beta}{\tau} \right) + \pi_j \boldsymbol{\nu}_j \left(\frac{\beta}{\tau} \right) + [\boldsymbol{\beta}^T \mathbf{b}_{\ell,j}] \mathbf{A} \left(\frac{\beta}{\tau} \right) \left[\begin{matrix} \boldsymbol{\beta} \\ \mathbf{b}_{\ell,j} \end{matrix} \right] - 2 \left(\pi_\ell \boldsymbol{\nu}_\ell \left(\frac{\beta}{\tau} \right) \mathbf{c}_\ell + \pi_j \boldsymbol{\nu}_j \left(\frac{\beta}{\tau} \right) \mathbf{c}_j \right)^T \left[\begin{matrix} \boldsymbol{\beta} \\ \mathbf{b}_{\ell,j} \end{matrix} \right] \right) \\ & + \frac{\gamma \beta}{2\tau} \sigma^2 \left(\boldsymbol{\pi}^T \boldsymbol{\nu} \left(\frac{\beta}{\tau} \right) \right) \beta_0^2 - \beta_0 \sigma \beta \sqrt{\gamma} + \frac{\beta \tau}{2}, \end{aligned}$$

where $\mathbf{A}(\eta)$ and \mathbf{c}_ℓ are as in (J.3) and $\boldsymbol{\nu}(\eta)$ is as in (J.4).

Setting the derivative of the above with respect to α_0 to zero we arrive at

$$\frac{\gamma \beta}{\tau} \sigma^2 \left(\boldsymbol{\pi}^T \boldsymbol{\nu} \left(\frac{\beta}{\tau} \right) \right) \beta_0 - \sigma \beta \sqrt{\gamma} = 0 \quad \Rightarrow \quad \beta_0 = \frac{\tau}{\sigma \sqrt{\gamma} \left(\boldsymbol{\pi}^T \boldsymbol{\nu} \left(\frac{\beta}{\tau} \right) \right)}.$$

Note that the above objective has the form

$$g \left(\frac{\beta}{\tau} \right) - \beta_0 \sigma \beta \sqrt{\gamma} + \frac{\beta \tau}{2},$$

with

$$g(\eta) := \frac{\eta\gamma}{2} \left(\pi_\ell \boldsymbol{\nu}_\ell(\eta) + \pi_j \boldsymbol{\nu}_j(\eta) + [\boldsymbol{\beta}^T \ b_{\ell,j}] \mathbf{A}(\eta) \begin{bmatrix} \boldsymbol{\beta} \\ b_{\ell,j} \end{bmatrix} - 2(\pi_\ell \boldsymbol{\nu}_\ell(\eta) \mathbf{c}_\ell + \pi_j \boldsymbol{\nu}_j(\eta) \mathbf{c}_j)^T \begin{bmatrix} \boldsymbol{\beta} \\ b_{\ell,j} \end{bmatrix} \right) + \frac{\gamma\eta}{2} \sigma^2 (\boldsymbol{\pi}^T \boldsymbol{\nu}(\eta)) \beta_0^2.$$

Thus, the derivatives with respect to β and τ to zero we have

$$\frac{1}{\tau} g' \left(\frac{\beta}{\tau} \right) - \beta_0 \sigma \sqrt{\gamma} + \frac{\tau}{2} = 0 \quad \Rightarrow \quad g' \left(\frac{\beta}{\tau} \right) - \beta_0 \sigma \sqrt{\gamma} \tau + \frac{\tau^2}{2} = 0 \quad \Rightarrow \quad g' \left(\frac{\beta}{\tau} \right) = \tau^2 \left(\frac{1}{\boldsymbol{\pi}^T \boldsymbol{\nu} \left(\frac{\beta}{\tau} \right)} - \frac{1}{2} \right)$$

and

$$-\frac{\beta}{\tau^2} g' \left(\frac{\beta}{\tau} \right) + \frac{\beta}{2} = 0 \quad \Rightarrow \quad \tau^2 = 2g' \left(\frac{\beta}{\tau} \right).$$

Combining the latter two we conclude that $\boldsymbol{\pi}^T \boldsymbol{\nu} \left(\frac{\beta}{\tau} \right) = 1$. Thus, $\eta = \frac{\beta}{\tau}$ is the solution to $\boldsymbol{\pi}^T \boldsymbol{\nu}(\eta) = 1$. To calculate τ and hence β_0 we calculate g' which is equal to

$$g'(\eta) = \frac{\gamma}{2} \left(\pi_\ell \boldsymbol{\nu}_\ell(\eta) + \pi_j \boldsymbol{\nu}_j(\eta) + [\boldsymbol{\beta}^T \ b_{\ell,j}] \mathbf{A}(\eta) \begin{bmatrix} \boldsymbol{\beta} \\ b_{\ell,j} \end{bmatrix} - 2(\pi_\ell \boldsymbol{\nu}_\ell(\eta) \mathbf{c}_\ell + \pi_j \boldsymbol{\nu}_j(\eta) \mathbf{c}_j)^T \begin{bmatrix} \boldsymbol{\beta} \\ b_{\ell,j} \end{bmatrix} \right) + \frac{\gamma}{2} \sigma^2 (\boldsymbol{\pi}^T \boldsymbol{\nu}(\eta)) \beta_0^2 + \frac{\gamma\eta}{2} \sigma^2 \beta_0^2 (\boldsymbol{\pi}^T \boldsymbol{\nu}'(\eta)) + \frac{\gamma\eta}{2} \left(\pi_\ell \boldsymbol{\nu}'_\ell(\eta) + \pi_j \boldsymbol{\nu}'_j(\eta) + [\boldsymbol{\beta}^T \ b_{\ell,j}] \mathbf{A}'(\eta) \begin{bmatrix} \boldsymbol{\beta} \\ b_{\ell,j} \end{bmatrix} - 2(\pi_\ell \boldsymbol{\nu}'_\ell(\eta) \mathbf{c}_\ell + \pi_j \boldsymbol{\nu}'_j(\eta) \mathbf{c}_j)^T \begin{bmatrix} \boldsymbol{\beta} \\ b_{\ell,j} \end{bmatrix} \right).$$

Here, we have

$$\boldsymbol{\nu}'(\eta) = -\frac{1}{\gamma} \begin{bmatrix} \frac{\omega_1^2}{(\omega_1^2 + \eta)^2} \\ \frac{\omega_2^2}{(\omega_2^2 + \eta)^2} \\ \dots \\ \frac{\omega_k^2}{(\omega_k^2 + \eta)^2} \end{bmatrix}$$

$$\mathbf{A}'(\eta) := \begin{bmatrix} \sigma^2 (\boldsymbol{\pi}^T \boldsymbol{\nu}'(\eta)) \mathbf{I} + \boldsymbol{\Sigma} \mathbf{V}^T \text{diag}(\boldsymbol{\pi}) \text{diag}(\boldsymbol{\nu}'(\eta)) \mathbf{V} \boldsymbol{\Sigma} & \boldsymbol{\Sigma} \mathbf{V}^T \text{diag}(\boldsymbol{\nu}'(\eta)) \boldsymbol{\pi} \\ \boldsymbol{\pi}^T \text{diag}(\boldsymbol{\nu}'(\eta)) \mathbf{V} \boldsymbol{\Sigma} & \boldsymbol{\pi}^T \boldsymbol{\nu}'(\eta) \end{bmatrix}$$

$$\mathbf{c}_\ell := \begin{bmatrix} \boldsymbol{\Sigma} \mathbf{V}^T \mathbf{e}_\ell \\ 1 \end{bmatrix}.$$

Now note that at the optimal point we have

$$g'(\eta) = \frac{\tau^2}{2} = \frac{\gamma}{2} \sigma^2 (\boldsymbol{\pi}^T \boldsymbol{\nu}(\eta)) \beta_0^2.$$

Thus from the above we can conclude that

$$\beta_0^2 = -\frac{1}{\eta \sigma^2 (\boldsymbol{\pi}^T \boldsymbol{\nu}'(\eta))} \left(\pi_\ell \boldsymbol{\nu}_\ell(\eta) + \pi_j \boldsymbol{\nu}_j(\eta) + [\boldsymbol{\beta}^T \ b_{\ell,j}] \mathbf{A}(\eta) \begin{bmatrix} \boldsymbol{\beta} \\ b_{\ell,j} \end{bmatrix} - 2(\pi_\ell \boldsymbol{\nu}_\ell(\eta) \mathbf{c}_\ell + \pi_j \boldsymbol{\nu}_j(\eta) \mathbf{c}_j)^T \begin{bmatrix} \boldsymbol{\beta} \\ b_{\ell,j} \end{bmatrix} \right) - \frac{1}{\sigma^2 (\boldsymbol{\pi}^T \boldsymbol{\nu}'(\eta))} \left(\pi_\ell \boldsymbol{\nu}'_\ell(\eta) + \pi_j \boldsymbol{\nu}'_j(\eta) + [\boldsymbol{\beta}^T \ b_{\ell,j}] \mathbf{A}'(\eta) \begin{bmatrix} \boldsymbol{\beta} \\ b_{\ell,j} \end{bmatrix} - 2(\pi_\ell \boldsymbol{\nu}'_\ell(\eta) \mathbf{c}_\ell + \pi_j \boldsymbol{\nu}'_j(\eta) \mathbf{c}_j)^T \begin{bmatrix} \boldsymbol{\beta} \\ b_{\ell,j} \end{bmatrix} \right).$$

Thus, the AO problem reduces to

$$\min_{\mathbf{b}_{\ell,j}} \min_{\boldsymbol{\beta}} \frac{\eta\gamma}{2} \left(\pi_\ell \boldsymbol{\nu}_\ell + \pi_j \boldsymbol{\nu}_j + [\boldsymbol{\beta}^T \ b_{\ell,j}] \mathbf{A} \begin{bmatrix} \boldsymbol{\beta} \\ b_{\ell,j} \end{bmatrix} - 2(\pi_\ell \boldsymbol{\nu}_\ell \mathbf{c}_\ell + \pi_j \boldsymbol{\nu}_j \mathbf{c}_j)^T \begin{bmatrix} \boldsymbol{\beta} \\ b_{\ell,j} \end{bmatrix} \right),$$

where η is the solution to

$$\sum_{\ell=1}^k \frac{\pi_\ell \omega_\ell^2}{\omega_\ell^2 + \eta} = \gamma.$$

Thus, similar to (J.10) we can compute the minimizer of the deterministic

$$\begin{aligned} \begin{bmatrix} \widehat{\beta} \\ b_{\ell,j} \end{bmatrix} &= \begin{bmatrix} -\Delta^{-1} \Sigma V^T (\pi_\ell \nu_\ell (\boldsymbol{\pi} \odot \boldsymbol{\nu} - \mathbf{e}_\ell) + \pi_j \nu_j (\boldsymbol{\pi} \odot \boldsymbol{\nu} - \mathbf{e}_j)) \\ \pi_\ell \nu_\ell + \pi_j \nu_j + (\boldsymbol{\pi} \odot \boldsymbol{\nu})^T V \Sigma \Delta^{-1} \Sigma V^T (\pi_\ell \nu_\ell (\boldsymbol{\pi} \odot \boldsymbol{\nu} - \mathbf{e}_\ell) + \pi_j \nu_j (\boldsymbol{\pi} \odot \boldsymbol{\nu} - \mathbf{e}_j)) \end{bmatrix} \\ &= \mathbf{A}^{-1} (\pi_\ell \boldsymbol{\nu}_\ell \mathbf{c}_\ell + \pi_j \boldsymbol{\nu}_j \mathbf{c}_j) \end{aligned} \quad (\text{J.15})$$

and

$$\begin{aligned} \boldsymbol{\nu}'(\eta) &= -\frac{1}{\gamma} \begin{bmatrix} \frac{\omega_1^2}{(\omega_1^2 + \eta)^2} \\ \frac{\omega_2^2}{(\omega_2^2 + \eta)^2} \\ \dots \\ \frac{\omega_k^2}{(\omega_k^2 + \eta)^2} \end{bmatrix} \\ \mathbf{A}'(\eta) &:= \begin{bmatrix} \sigma^2 (\boldsymbol{\pi}^T \boldsymbol{\nu}'(\eta)) \mathbf{I} + \Sigma V^T \text{diag}(\boldsymbol{\pi}) \text{diag}(\boldsymbol{\nu}'(\eta)) V \Sigma & \Sigma V^T \text{diag}(\boldsymbol{\nu}'(\eta)) \boldsymbol{\pi} \\ \boldsymbol{\pi}^T \text{diag}(\boldsymbol{\nu}'(\eta)) V \Sigma & \boldsymbol{\pi}^T \boldsymbol{\nu}'(\eta) \end{bmatrix}. \end{aligned}$$

To continue note that

$$\begin{aligned} \pi_\ell \boldsymbol{\nu}'_\ell(\eta) + \pi_j \boldsymbol{\nu}'_j(\eta) + [\beta^T \ b_{\ell,j}] \mathbf{A}'(\eta) \begin{bmatrix} \beta \\ b_{\ell,j} \end{bmatrix} &- 2 (\pi_\ell \boldsymbol{\nu}'_\ell(\eta) \mathbf{c}_\ell + \pi_j \boldsymbol{\nu}'_j(\eta) \mathbf{c}_j)^T \begin{bmatrix} \beta \\ b_{\ell,j} \end{bmatrix} \\ &= \pi_\ell \boldsymbol{\nu}'_\ell(\eta) + \pi_j \boldsymbol{\nu}'_j(\eta) + (\pi_\ell \boldsymbol{\nu}'_\ell(\eta) \mathbf{c}_\ell + \pi_j \boldsymbol{\nu}'_j(\eta) \mathbf{c}_j)^T \mathbf{A}^{-1} \mathbf{A}'(\eta) \mathbf{A}^{-1} (\pi_\ell \boldsymbol{\nu}'_\ell(\eta) \mathbf{c}_\ell + \pi_j \boldsymbol{\nu}'_j(\eta) \mathbf{c}_j) \\ &\quad - 2 (\pi_\ell \boldsymbol{\nu}'_\ell(\eta) \mathbf{c}_\ell + \pi_j \boldsymbol{\nu}'_j(\eta) \mathbf{c}_j)^T \mathbf{A}^{-1} (\pi_\ell \boldsymbol{\nu}'_\ell(\eta) \mathbf{c}_\ell + \pi_j \boldsymbol{\nu}'_j(\eta) \mathbf{c}_j). \end{aligned}$$

Thus,

$$\begin{aligned} \beta_0^2 &= -\frac{1}{\eta \sigma^2 (\boldsymbol{\pi}^T \boldsymbol{\nu}'(\eta))} \left(\pi_\ell \boldsymbol{\nu}_\ell(\eta) + \pi_j \boldsymbol{\nu}_j(\eta) + [\beta^T \ b_{\ell,j}] \mathbf{A}(\eta) \begin{bmatrix} \beta \\ b_{\ell,j} \end{bmatrix} - 2 (\pi_\ell \boldsymbol{\nu}_\ell(\eta) \mathbf{c}_\ell + \pi_j \boldsymbol{\nu}_j(\eta) \mathbf{c}_j)^T \begin{bmatrix} \beta \\ b_{\ell,j} \end{bmatrix} \right) \\ &\quad - \frac{1}{\sigma^2 (\boldsymbol{\pi}^T \boldsymbol{\nu}'(\eta))} \left(\pi_\ell \boldsymbol{\nu}'_\ell(\eta) + \pi_j \boldsymbol{\nu}'_j(\eta) + [\beta^T \ b_{\ell,j}] \mathbf{A}'(\eta) \begin{bmatrix} \beta \\ b_{\ell,j} \end{bmatrix} - 2 (\pi_\ell \boldsymbol{\nu}'_\ell(\eta) \mathbf{c}_\ell + \pi_j \boldsymbol{\nu}'_j(\eta) \mathbf{c}_j)^T \begin{bmatrix} \beta \\ b_{\ell,j} \end{bmatrix} \right) \\ &= \frac{\zeta}{\sigma^2} \left(\pi_\ell \nu_\ell + \pi_j \nu_j - (\pi_\ell \nu_\ell + \pi_j \nu_j)^2 \right. \\ &\quad \left. - (\pi_\ell \nu_\ell (\boldsymbol{\pi} \odot \boldsymbol{\nu} - \mathbf{e}_\ell) + \pi_j \nu_j (\boldsymbol{\pi} \odot \boldsymbol{\nu} - \mathbf{e}_j))^T V \Sigma \Delta^{-1} \Sigma V^T (\pi_\ell \nu_\ell (\boldsymbol{\pi} \odot \boldsymbol{\nu} - \mathbf{e}_\ell) + \pi_j \nu_j (\boldsymbol{\pi} \odot \boldsymbol{\nu} - \mathbf{e}_j)) \right) \\ &\quad + \frac{\zeta \eta}{\sigma^2} \left(\pi_\ell \boldsymbol{\nu}'_\ell(\eta) + \pi_j \boldsymbol{\nu}'_j(\eta) + (\pi_\ell \boldsymbol{\nu}_\ell(\eta) \mathbf{c}_\ell + \pi_j \boldsymbol{\nu}_j(\eta) \mathbf{c}_j)^T \mathbf{A}^{-1} \mathbf{A}'(\eta) \mathbf{A}^{-1} (\pi_\ell \boldsymbol{\nu}_\ell(\eta) \mathbf{c}_\ell + \pi_j \boldsymbol{\nu}_j(\eta) \mathbf{c}_j) \right. \\ &\quad \left. - 2 (\pi_\ell \boldsymbol{\nu}'_\ell(\eta) \mathbf{c}_\ell + \pi_j \boldsymbol{\nu}'_j(\eta) \mathbf{c}_j)^T \mathbf{A}^{-1} (\pi_\ell \boldsymbol{\nu}_\ell(\eta) \mathbf{c}_\ell + \pi_j \boldsymbol{\nu}_j(\eta) \mathbf{c}_j) \right), \end{aligned}$$

where $\zeta := \frac{\gamma}{\eta \left(\sum_{\ell=1}^k \frac{\pi_\ell \omega_\ell^2}{(\omega_\ell^2 + \eta)^2} \right)}$. From the CGMT, we have that $\|\widehat{\mathbf{w}}_\ell + \widehat{\mathbf{w}}_j\|_{\ell_2}^2 \xrightarrow{P} \widehat{\beta}_0^2 + \|\beta\|_{\ell_2}^2$. Combining this with the calculations above, we conclude that

$$\begin{aligned} \|\widehat{\mathbf{w}}_\ell + \widehat{\mathbf{w}}_j\|_{\ell_2}^2 &\xrightarrow{P} \frac{\zeta}{\sigma^2} (\pi_\ell \nu_\ell + \pi_j \nu_j) (1 - \pi_\ell \nu_\ell - \pi_j \nu_j) \\ &\quad + (\pi_\ell \nu_\ell (\boldsymbol{\pi} \odot \boldsymbol{\nu} - \mathbf{e}_\ell) + \pi_j \nu_j (\boldsymbol{\pi} \odot \boldsymbol{\nu} - \mathbf{e}_j))^T V \Sigma \Delta^{-1} \left(\Delta^{-1} - \frac{\eta}{\sigma^2} I_r \right) \Sigma V^T (\pi_\ell \nu_\ell (\boldsymbol{\pi} \odot \boldsymbol{\nu} - \mathbf{e}_\ell) + \pi_j \nu_j (\boldsymbol{\pi} \odot \boldsymbol{\nu} - \mathbf{e}_j)) \\ &\quad + \frac{\zeta \eta}{\sigma^2} \left(\pi_\ell \boldsymbol{\nu}'_\ell(\eta) + \pi_j \boldsymbol{\nu}'_j(\eta) + (\pi_\ell \boldsymbol{\nu}_\ell(\eta) \mathbf{c}_\ell + \pi_j \boldsymbol{\nu}_j(\eta) \mathbf{c}_j)^T \mathbf{A}^{-1} \mathbf{A}'(\eta) \mathbf{A}^{-1} (\pi_\ell \boldsymbol{\nu}_\ell(\eta) \mathbf{c}_\ell + \pi_j \boldsymbol{\nu}_j(\eta) \mathbf{c}_j) \right. \\ &\quad \left. - 2 (\pi_\ell \boldsymbol{\nu}'_\ell(\eta) \mathbf{c}_\ell + \pi_j \boldsymbol{\nu}'_j(\eta) \mathbf{c}_j)^T \mathbf{A}^{-1} (\pi_\ell \boldsymbol{\nu}_\ell(\eta) \mathbf{c}_\ell + \pi_j \boldsymbol{\nu}_j(\eta) \mathbf{c}_j) \right). \end{aligned} \quad (\text{J.16})$$

Finally, using (J.16) and (J.12) in (J.14) it follows that

$$\begin{aligned}
\langle \mathbf{w}_\ell, \mathbf{w}_j \rangle &\xrightarrow{P} \\
&\pi_\ell \nu_\ell \pi_j \nu_j \left(-\frac{\zeta}{\sigma^2} + (\boldsymbol{\pi} \odot \boldsymbol{\nu} - \mathbf{e}_\ell)^T \mathbf{V} \boldsymbol{\Sigma} \boldsymbol{\Delta}^{-1} \left(\boldsymbol{\Delta}^{-1} - \frac{\zeta}{\sigma^2} \mathbf{I}_r \right) \boldsymbol{\Sigma} \mathbf{V}^T (\boldsymbol{\pi} \odot \boldsymbol{\nu} - \mathbf{e}_j) \right) \\
&+ \frac{\zeta \eta}{\sigma^2} \pi_\ell \nu_\ell \pi_j \nu_j \mathbf{c}_j^T \mathbf{A}^{-1} \mathbf{A}' \mathbf{A}^{-1} \mathbf{c}_\ell - \frac{\zeta \eta}{\sigma^2} \pi_\ell \pi_j (\nu_\ell \nu'_j + \nu'_\ell \nu_j) \mathbf{c}_j^T \mathbf{A}^{-1} \mathbf{c}_\ell \\
&= \pi_\ell \nu_\ell \pi_j \nu_j \left(-\frac{\zeta}{\sigma^2} + (\boldsymbol{\pi} \odot \boldsymbol{\nu} - \mathbf{e}_\ell)^T \mathbf{V} \boldsymbol{\Sigma} \boldsymbol{\Delta}^{-1} \left(\boldsymbol{\Delta}^{-1} - \frac{\zeta}{\sigma^2} \mathbf{I}_r \right) \boldsymbol{\Sigma} \mathbf{V}^T (\boldsymbol{\pi} \odot \boldsymbol{\nu} - \mathbf{e}_j) \right) \\
&+ \frac{\zeta \eta}{\sigma^2} \pi_\ell \nu_\ell \pi_j \nu_j \left(\mathbf{e}_j^T \begin{bmatrix} \boldsymbol{\Sigma} \mathbf{V}^T \\ \mathbf{1}^T \end{bmatrix}^T \mathbf{A}^{-1} \mathbf{A}' \mathbf{A}^{-1} \begin{bmatrix} \boldsymbol{\Sigma} \mathbf{V}^T \\ \mathbf{1}^T \end{bmatrix} \mathbf{e}_\ell \right) \\
&- \frac{\eta \zeta}{\sigma^2} \pi_\ell \pi_j (\nu_\ell \nu'_j + \nu'_\ell \nu_j) \left(\mathbf{e}_j^T \begin{bmatrix} \boldsymbol{\Sigma} \mathbf{V}^T \\ \mathbf{1}^T \end{bmatrix}^T \mathbf{A}^{-1} \begin{bmatrix} \boldsymbol{\Sigma} \mathbf{V}^T \\ \mathbf{1}^T \end{bmatrix} \mathbf{e}_\ell \right) \tag{J.17}
\end{aligned}$$

Putting everything together we arrive at

$$\boldsymbol{\Sigma}_{\mathbf{w}, \mathbf{w}} \xrightarrow{P} \frac{\zeta}{\sigma^2} \mathbf{P} + \mathbf{P} \mathbf{V} \boldsymbol{\Sigma} \boldsymbol{\Delta}^{-1} \left(\boldsymbol{\Delta}^{-1} - \frac{\zeta}{\sigma^2} \mathbf{I}_r \right) \boldsymbol{\Sigma} \mathbf{V}^T \mathbf{P} + \frac{\zeta \eta}{\sigma^2} \mathbf{Q},$$

where

$$\begin{aligned}
\mathbf{Q} := &\text{diag}(\boldsymbol{\pi} \odot \boldsymbol{\nu}') + \text{diag}(\boldsymbol{\pi} \odot \boldsymbol{\nu}) \begin{bmatrix} \boldsymbol{\Sigma} \mathbf{V}^T \\ \mathbf{1}^T \end{bmatrix}^T (\mathbf{A}^{-1} \mathbf{A}' \mathbf{A}^{-1}) \begin{bmatrix} \boldsymbol{\Sigma} \mathbf{V}^T \\ \mathbf{1}^T \end{bmatrix} \text{diag}(\boldsymbol{\pi} \odot \boldsymbol{\nu}) \\
&- \text{diag}(\boldsymbol{\pi} \odot \boldsymbol{\nu}') \begin{bmatrix} \boldsymbol{\Sigma} \mathbf{V}^T \\ \mathbf{1}^T \end{bmatrix}^T \mathbf{A}^{-1} \begin{bmatrix} \boldsymbol{\Sigma} \mathbf{V}^T \\ \mathbf{1}^T \end{bmatrix} \text{diag}(\boldsymbol{\pi} \odot \boldsymbol{\nu}) - \text{diag}(\boldsymbol{\pi} \odot \boldsymbol{\nu}) \begin{bmatrix} \boldsymbol{\Sigma} \mathbf{V}^T \\ \mathbf{1}^T \end{bmatrix}^T \mathbf{A}^{-1} \begin{bmatrix} \boldsymbol{\Sigma} \mathbf{V}^T \\ \mathbf{1}^T \end{bmatrix} \text{diag}(\boldsymbol{\pi} \odot \boldsymbol{\nu}')
\end{aligned}$$

and as mentioned earlier

$$\begin{aligned}
\mathbf{A}' := &\mathbf{A}'(\eta) := \begin{bmatrix} \sigma^2 (\boldsymbol{\pi}^T \boldsymbol{\nu}'(\eta)) \mathbf{I} + \boldsymbol{\Sigma} \mathbf{V}^T \text{diag}(\boldsymbol{\pi}) \text{diag}(\boldsymbol{\nu}'(\eta)) \mathbf{V} \boldsymbol{\Sigma} & \boldsymbol{\Sigma} \mathbf{V}^T \text{diag}(\boldsymbol{\nu}'(\eta)) \boldsymbol{\pi} \\ \boldsymbol{\pi}^T \text{diag}(\boldsymbol{\nu}'(\eta)) \mathbf{V} \boldsymbol{\Sigma} & \boldsymbol{\pi}^T \boldsymbol{\nu}'(\eta) \end{bmatrix} \\
\mathbf{A} := &\mathbf{A}(\eta) := \begin{bmatrix} \sigma^2 (\boldsymbol{\pi}^T \boldsymbol{\nu}(\eta)) \mathbf{I} + \boldsymbol{\Sigma} \mathbf{V}^T \text{diag}(\boldsymbol{\pi}) \text{diag}(\boldsymbol{\nu}(\eta)) \mathbf{V} \boldsymbol{\Sigma} & \boldsymbol{\Sigma} \mathbf{V}^T \text{diag}(\boldsymbol{\nu}(\eta)) \boldsymbol{\pi} \\ \boldsymbol{\pi}^T \text{diag}(\boldsymbol{\nu}(\eta)) \mathbf{V} \boldsymbol{\Sigma} & \boldsymbol{\pi}^T \boldsymbol{\nu}(\eta) \end{bmatrix}.
\end{aligned}$$

Let us end by simplifying \mathbf{Q} to this aim

$$\begin{aligned}
\mathbf{A}^{-1} \begin{bmatrix} \boldsymbol{\Sigma} \mathbf{V}^T \\ \mathbf{1}^T \end{bmatrix} \text{diag}(\boldsymbol{\pi} \odot \boldsymbol{\nu}) &= \begin{bmatrix} \mathbf{I} & 0 \\ -\tilde{\boldsymbol{\pi}}^T \mathbf{V} \boldsymbol{\Sigma} & 1 \end{bmatrix} \begin{bmatrix} \boldsymbol{\Delta}^{-1} & 0 \\ \mathbf{0}^T & 1 \end{bmatrix} \begin{bmatrix} \mathbf{I} & -\boldsymbol{\Sigma} \mathbf{V}^T \tilde{\boldsymbol{\pi}} \\ \mathbf{0}^T & 1 \end{bmatrix} \begin{bmatrix} \boldsymbol{\Sigma} \mathbf{V}^T \\ \mathbf{1}^T \end{bmatrix} \text{diag}(\boldsymbol{\pi} \odot \boldsymbol{\nu}) \\
&= \begin{bmatrix} \mathbf{I} & 0 \\ -\tilde{\boldsymbol{\pi}}^T \mathbf{V} \boldsymbol{\Sigma} & 1 \end{bmatrix} \begin{bmatrix} \boldsymbol{\Delta}^{-1} & 0 \\ \mathbf{0}^T & 1 \end{bmatrix} \begin{bmatrix} \boldsymbol{\Sigma} \mathbf{V}^T (\mathbf{I} - \tilde{\boldsymbol{\pi}} \mathbf{1}^T) \\ \mathbf{1}^T \end{bmatrix} \text{diag}(\boldsymbol{\pi} \odot \boldsymbol{\nu}) \\
&= \begin{bmatrix} \mathbf{I} & 0 \\ -\tilde{\boldsymbol{\pi}}^T \mathbf{V} \boldsymbol{\Sigma} & 1 \end{bmatrix} \begin{bmatrix} \boldsymbol{\Delta}^{-1} \boldsymbol{\Sigma} \mathbf{V}^T (\mathbf{I} - \tilde{\boldsymbol{\pi}} \mathbf{1}^T) \\ \mathbf{1}^T \end{bmatrix} \text{diag}(\boldsymbol{\pi} \odot \boldsymbol{\nu}) \\
&= \begin{bmatrix} \boldsymbol{\Delta}^{-1} \boldsymbol{\Sigma} \mathbf{V}^T (\mathbf{I} - \tilde{\boldsymbol{\pi}} \mathbf{1}^T) \\ -\tilde{\boldsymbol{\pi}}^T \mathbf{V} \boldsymbol{\Sigma} \boldsymbol{\Delta}^{-1} \boldsymbol{\Sigma} \mathbf{V}^T (\mathbf{I} - \tilde{\boldsymbol{\pi}} \mathbf{1}^T) + \mathbf{1}^T \end{bmatrix} \text{diag}(\boldsymbol{\pi} \odot \boldsymbol{\nu}) \\
&= \begin{bmatrix} \boldsymbol{\Delta}^{-1} \boldsymbol{\Sigma} \mathbf{V}^T (\text{diag}(\tilde{\boldsymbol{\pi}}) - \tilde{\boldsymbol{\pi}} \tilde{\boldsymbol{\pi}}^T) \\ -\tilde{\boldsymbol{\pi}}^T \mathbf{V} \boldsymbol{\Sigma} \boldsymbol{\Delta}^{-1} \boldsymbol{\Sigma} \mathbf{V}^T (\text{diag}(\tilde{\boldsymbol{\pi}}) - \tilde{\boldsymbol{\pi}} \tilde{\boldsymbol{\pi}}^T) + \tilde{\boldsymbol{\pi}}^T \end{bmatrix}
\end{aligned}$$

Thus, defining $\tilde{\boldsymbol{\pi}}' = \boldsymbol{\pi} \odot \boldsymbol{\nu}'$ we have

$$\begin{aligned}
\text{diag}(\boldsymbol{\pi} \odot \boldsymbol{\nu}') \begin{bmatrix} \boldsymbol{\Sigma} \mathbf{V}^T \\ \mathbf{1}^T \end{bmatrix}^T \mathbf{A}^{-1} \begin{bmatrix} \boldsymbol{\Sigma} \mathbf{V}^T \\ \mathbf{1}^T \end{bmatrix} \text{diag}(\boldsymbol{\pi} \odot \boldsymbol{\nu}) &= \text{diag}(\boldsymbol{\pi} \odot \boldsymbol{\nu}') \begin{bmatrix} \boldsymbol{\Sigma} \mathbf{V}^T \\ \mathbf{1}^T \end{bmatrix}^T \begin{bmatrix} \boldsymbol{\Delta}^{-1} \boldsymbol{\Sigma} \mathbf{V}^T (\text{diag}(\tilde{\boldsymbol{\pi}}) - \tilde{\boldsymbol{\pi}} \tilde{\boldsymbol{\pi}}^T) \\ -\tilde{\boldsymbol{\pi}}^T \mathbf{V} \boldsymbol{\Sigma} \boldsymbol{\Delta}^{-1} \boldsymbol{\Sigma} \mathbf{V}^T (\text{diag}(\tilde{\boldsymbol{\pi}}) - \tilde{\boldsymbol{\pi}} \tilde{\boldsymbol{\pi}}^T) + \tilde{\boldsymbol{\pi}}^T \end{bmatrix} \\
&= (\text{diag}(\tilde{\boldsymbol{\pi}}') - \tilde{\boldsymbol{\pi}}' \tilde{\boldsymbol{\pi}}^T) \mathbf{V} \boldsymbol{\Sigma} \boldsymbol{\Delta}^{-1} \boldsymbol{\Sigma} \mathbf{V}^T (\text{diag}(\tilde{\boldsymbol{\pi}}) - \tilde{\boldsymbol{\pi}} \tilde{\boldsymbol{\pi}}^T) + \tilde{\boldsymbol{\pi}}' \tilde{\boldsymbol{\pi}}^T
\end{aligned}$$

Using the above and recalling $\tilde{\pi}' = \pi \odot \nu'$ we arrive at

$$\begin{aligned} \mathbf{Q} = & \text{diag}(\tilde{\pi}') \\ & + \left[\begin{array}{c} \Delta^{-1} \Sigma V^T (\text{diag}(\tilde{\pi}) - \tilde{\pi} \tilde{\pi}^T) \\ - \tilde{\pi}^T V \Sigma \Delta^{-1} \Sigma V^T (\text{diag}(\tilde{\pi}) - \tilde{\pi} \tilde{\pi}^T) + \tilde{\pi}^T \end{array} \right]^T \mathbf{A}' \left[\begin{array}{c} \Delta^{-1} \Sigma V^T (\text{diag}(\tilde{\pi}) - \tilde{\pi} \tilde{\pi}^T) \\ - \tilde{\pi}^T V \Sigma \Delta^{-1} \Sigma V^T (\text{diag}(\tilde{\pi}') - \tilde{\pi} \tilde{\pi}^T) + \tilde{\pi}^T \end{array} \right] \\ & - (\text{diag}(\tilde{\pi}') - \tilde{\pi}' \tilde{\pi}^T) V \Sigma \Delta^{-1} \Sigma V^T (\text{diag}(\tilde{\pi}) - \tilde{\pi} \tilde{\pi}^T) - \tilde{\pi}' \tilde{\pi}^T \\ & - (\text{diag}(\tilde{\pi}) - \tilde{\pi} \tilde{\pi}^T) V \Sigma \Delta^{-1} \Sigma V^T (\text{diag}(\tilde{\pi}') - \tilde{\pi} \tilde{\pi}'^T) - \tilde{\pi} \tilde{\pi}'^T \end{aligned} \quad (\text{J.18})$$

where

$$\begin{aligned} \mathbf{A}' &= \begin{bmatrix} \sigma^2 (\tilde{\pi}'^T \mathbf{1}) \mathbf{I} + \Sigma V^T \text{diag}(\tilde{\pi}') V \Sigma & \Sigma V^T \tilde{\pi}' \\ \tilde{\pi}'^T V \Sigma & \tilde{\pi}'^T \mathbf{1} \end{bmatrix} \\ \mathbf{A} &= \begin{bmatrix} \sigma^2 \mathbf{I} + \Sigma V^T \text{diag}(\tilde{\pi}) V \Sigma & \Sigma V^T \tilde{\pi} \\ \tilde{\pi}^T V \Sigma & 1 \end{bmatrix}. \end{aligned}$$

Using the above the cross-correlation matrix $\Sigma_{\mathbf{w}, \mathbf{w}}$ is given by

$$\Sigma_{\mathbf{w}, \mathbf{w}} \xrightarrow{P} \frac{\zeta}{\sigma^2} \mathbf{P} + \mathbf{P} V \Sigma \Delta^{-1} \left(\Delta^{-1} - \frac{\zeta}{\sigma^2} \mathbf{I}_r \right) \Sigma V^T \mathbf{P} + \frac{\zeta \eta}{\sigma^2} \mathbf{Q}.$$

K Weighted LS for MLM (Proof of Theorem B.2)

Let $\mathbf{D} := \mathbf{D}^{(n)} := \text{diag}(D_1, \dots, D_n)$ be a diagonal matrix with non-zero diagonal entries. In particular, assume that the diagonal entries of \mathbf{D} are distributed $D_i \stackrel{iid}{\sim} D$ where the random variable D may depend on the entries of the matrix of response variables \mathbf{Y} . Here, we focus on the following setting:

$$\mathbf{D} = \sum_{j \in [k]} \text{diag}(\omega_j \mathbf{Y}_j), \quad \omega_j \geq 0, \quad j \in [k]. \quad (\text{K.1})$$

Specifically, for (K.1), we have $D_i \stackrel{iid}{\sim} D$ with $D = \omega_\ell Y_\ell + \sum_{i \neq \ell \in [k]} \omega_i Y_i$, where for all $c \in [k]$:

$$\mathbb{P}([Y_1, Y_2, \dots, Y_k]^T = \mathbf{e}_c) = V_c = \frac{e^{\mathbf{e}_c^T V \Sigma g}}{\sum_{\ell'=1}^k e^{\mathbf{e}_{\ell'}^T V \Sigma g}}, \quad (\text{K.2})$$

$\mathbf{M} \mathbf{M}^T = \mathbf{V} \Sigma^2 \mathbf{V}^T$, and $\mathbf{g} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_r)$.

With these, we consider the weighted least-squares (WLS) solution for $\ell \in [k]$:

$$(\hat{\mathbf{w}}_\ell, \hat{b}) = \arg \min_{\mathbf{w}, b} \mathcal{L}_{PO}(\mathbf{w}, b) := \frac{1}{2n} \left\| \mathbf{D} (\mathbf{X}^T \mathbf{w} + b \mathbf{1}_n - \mathbf{Y}_\ell) \right\|_{\ell_2}^2,$$

where \mathbf{D} is as in (K.1). In fact, it is convenient to rewrite the above as follows:

$$(\hat{\mathbf{w}}_\ell, \hat{b}) = \arg \min_{\mathbf{w}, b, \mathbf{u}} \max_{\mathbf{s}} \frac{1}{n} \left(\mathbf{s}^T \mathbf{D} \mathbf{X}^T \mathbf{w} + b \mathbf{s}^T \mathbf{D} \mathbf{1}_n \mathbf{s}^T \mathbf{D} \mathbf{Y}_\ell - \mathbf{s}^T \mathbf{u} + \frac{\|\mathbf{u}\|_{\ell_2}^2}{2} \right). \quad (\text{K.3})$$

Identifying the AO. The PO in (K.3) is very similar to (I.1). In particular, following step by step the same decomposition trick as in Section H.1.1, it can be shown that the AO corresponding to (K.3) becomes (cf. (I.6))

$$\min_{\mathbf{w}_\ell, b_\ell, \mathbf{u}} \max_{\mathbf{s}} \frac{1}{n} \left(\left\| \mathbf{P}^\perp \mathbf{w}_\ell \right\|_{\ell_2} \mathbf{g}^T \mathbf{D} \mathbf{s} + \left\| \mathbf{D} \mathbf{s} \right\|_{\ell_2} \mathbf{h}^T \mathbf{P}^\perp \mathbf{w}_\ell + \mathbf{s}^T \mathbf{D} \tilde{\mathbf{G}}^T \mathbf{U}^T \mathbf{w}_\ell + b_\ell \mathbf{s}^T \mathbf{D} \mathbf{1}_n - \mathbf{s}^T \mathbf{D} \mathbf{Y}_\ell - \mathbf{u}^T \mathbf{s} + \frac{\|\mathbf{u}\|_{\ell_2}^2}{2} \right),$$

where we use the same notation as in Section H.1.1 for $\mathbf{P}^\perp, \mathbf{U}, \tilde{\mathbf{G}}, \mathbf{g}$ and \mathbf{h} . Recall also the relation of \mathbf{Y}_ℓ to $\tilde{\mathbf{G}}$ in (I.5).

Scalarization of the AO. We start the process of simplifying the AO by setting $\beta := \|\mathbf{D}\mathbf{s}\|_{\ell_2}/\sqrt{n}$ and optimizing over the direction of $\mathbf{D}\mathbf{s}$ to equivalently write the AO as

$$\min_{\mathbf{w}_\ell, \mathbf{b}_\ell, \mathbf{u}} \max_{\beta \geq 0} \frac{1}{\sqrt{n}} \left(\beta \|\mathbf{P}^\perp \mathbf{w}_\ell\|_{\ell_2} \mathbf{g} + \tilde{\mathbf{G}}^T \mathbf{U}^T \mathbf{w}_\ell + b_\ell \mathbf{1}_n - \mathbf{Y}_\ell - \mathbf{D}^{-1} \mathbf{u} \right) + \frac{\|\mathbf{u}\|_{\ell_2}^2}{2n}, \quad (\text{K.4})$$

Next, focus on the minimization over \mathbf{w}_ℓ . Let us denote

$$\mathbf{a} := \mathbf{U}^T \mathbf{w}_\ell \quad \text{and} \quad \alpha_0 = \|\mathbf{P}^\perp \mathbf{w}_\ell\|_{\ell_2}.$$

Notice that $\mathbf{a} \perp \mathbf{P}^\perp \mathbf{w}_\ell$ and thus the orthogonal decomposition $\mathbf{w}_\ell = \mathbf{U}\mathbf{a} + \mathbf{P}^\perp \mathbf{w}_\ell$. With this observation, note that the optimal direction of $\mathbf{P}^T \mathbf{w}_\ell$ in (K.4) aligns with $\mathbf{P}^T \mathbf{h}$ for all values of β . Therefore, (K.4) reduces to

$$\min_{\mathbf{a}, \alpha_0 \geq 0, \mathbf{b}_\ell, \mathbf{u}} \max_{\beta \geq 0} \frac{1}{\sqrt{n}} \left(\beta \|\alpha_0 \mathbf{g} + \tilde{\mathbf{G}}^T \mathbf{a} + b_\ell \mathbf{1}_n - \mathbf{Y}_\ell - \mathbf{D}^{-1} \mathbf{u}\|_{\ell_2} - \beta \alpha_0 \|\mathbf{P}^\perp \mathbf{h}\|_{\ell_2} \right) + \frac{\|\mathbf{u}\|_{\ell_2}^2}{2n}, \quad (\text{K.5})$$

Continuing let us denote $\mathbf{t} := \alpha_0 \mathbf{g} + \tilde{\mathbf{G}}^T \mathbf{a} + b_\ell \mathbf{1}_n - \mathbf{Y}_\ell$ for convenience and rewrite $\|\mathbf{t} - \mathbf{D}^{-1} \mathbf{u}\|_{\ell_2}$ as follows

$$\frac{\|\mathbf{t} - \mathbf{D}^{-1} \mathbf{u}\|_{\ell_2}}{\sqrt{n}} = \min_{\tau > 0} \frac{\tau}{2} + \frac{\|\mathbf{t} - \mathbf{D}^{-1} \mathbf{u}\|_{\ell_2}^2}{2\tau n}.$$

Note that the resulting minimization is convex in \mathbf{u} and concave in β . Also, by considering the bounded AO (such that β is bounded; see [DKT19, Sec. A]), we can flip the order of min-max and optimize over \mathbf{u} first. In particular, \mathbf{u} minimizes the following strictly convex quadratic

$$\min_{\mathbf{u}} \left\{ \frac{1}{n} \left(\frac{\beta}{2\tau} \|\mathbf{D}^{-1} \mathbf{u}\|_{\ell_2} + \frac{1}{2} \|\mathbf{u}\|_{\ell_2}^2 - \frac{\beta}{\tau} \mathbf{t}^T \mathbf{D}^{-1} \mathbf{u} \right) \right\} = \frac{1}{2n} \mathbf{u}^T \left(\frac{\beta}{\tau} \mathbf{D}^{-2} + \mathbf{I}_n \right) \mathbf{u} - \frac{\beta}{\tau n} \mathbf{t}^T \mathbf{D}^{-1} \mathbf{u}.$$

In particular,

$$\mathbf{u} = \frac{\beta}{\tau} \left(\frac{\beta}{\tau} \mathbf{D}^{-2} + \mathbf{I} \right)^{-1} \mathbf{D}^{-1} \mathbf{t} = \left(\mathbf{D}^{-1} + \frac{\tau}{\beta} \mathbf{D} \right)^{-1} (\alpha_0 \mathbf{g} + \tilde{\mathbf{G}}^T \mathbf{a} + b_\ell \mathbf{1}_n - \mathbf{Y}_\ell)$$

Putting things together, the new objective function of (K.5) becomes

$$\begin{aligned} \min_{\mathbf{a}, \alpha_0 \geq 0, \mathbf{b}_\ell, \tau > 0} \max_{\beta \geq 0} \mathcal{R}(\mathbf{a}, \alpha_0, b_\ell, \tau, \beta) \\ \text{where } \mathcal{R}(\mathbf{a}, \alpha_0, b_\ell, \tau, \beta) := \frac{\beta \tau}{2n} + \frac{\beta}{2\tau n} \|\mathbf{t}\|_{\ell_2}^2 - \frac{\beta}{2\tau n} \mathbf{t}^T \left(\mathbf{I} + \frac{\tau}{\beta} \mathbf{D}^2 \right)^{-1} \mathbf{t} - \frac{\beta \alpha_0}{\sqrt{n}} \|\mathbf{P}^\perp \mathbf{h}\|_{\ell_2}. \end{aligned} \quad (\text{K.6})$$

Convergence of the AO After having simplified the AO into an optimization problem over $r + 4$ variables, we are ready to study its asymptotic behavior. First, we argue on point-wise convergence of \mathcal{R} in (K.6). Fix $\mathbf{a}, \alpha_0, \mathbf{b}_\ell, \tau$ and β . From the WLLN, $\frac{1}{\sqrt{n}} \|\mathbf{P}^\perp \mathbf{h}\|_{\ell_2} \xrightarrow{P} \sqrt{\gamma}$ and as in (I.9)

$$\frac{1}{n} \|\mathbf{t}\|_{\ell_2}^2 = \frac{1}{n} \sum_{i=1}^n (\alpha_0 \mathbf{g}_i + \mathbf{a}^T \tilde{\mathbf{g}}_i + b_\ell - [\mathbf{Y}_\ell]_i)^2 \xrightarrow{P} \mathbb{E} \left[(\alpha_0 G_0 + \mathbf{a}^T \mathbf{g} + \mathbf{b}_\ell - Y_\ell)^2 \right],$$

where the expectation is over $\mathbf{g} \sim \mathcal{N}(\mathbf{0}_r, \mathbf{I}_r)$ (with some abuse of notation) and

$$Y_\ell \sim \text{Bern}(V_\ell) \quad \text{and} \quad V_\ell = \frac{e^{\mathbf{e}_\ell^T \mathbf{V} \Sigma \mathbf{g}}}{\sum_{\ell'=1}^r e^{\mathbf{e}_{\ell'}^T \mathbf{V} \Sigma \mathbf{g}}}. \quad (\text{K.7})$$

Furthermore,

$$\frac{1}{n} \mathbf{t}^T \left(\mathbf{I} + \frac{\tau}{\beta} \mathbf{D}^2 \right)^{-1} \mathbf{t} = \frac{1}{n} \sum_{i=1}^n \frac{(\alpha_0 \mathbf{g}_i + \mathbf{a}^T \tilde{\mathbf{g}}_i + b_\ell - [\mathbf{Y}_\ell]_i)^2}{1 + \frac{\tau}{\beta} d_i^2} \xrightarrow{P} \mathbb{E} \left[\frac{(\alpha_0 G_0 + \mathbf{a}^T \mathbf{g} + \mathbf{b}_\ell - Y_\ell)^2}{1 + \frac{\tau}{\beta} D^2} \right]$$

Therefore, point-wise on $\mathbf{a}, \alpha_0, \mathbf{b}_\ell, \tau$ and β , the objective \mathcal{R} of the AO converges to

$$\begin{aligned}\mathcal{D}_\ell(\alpha_0, \boldsymbol{\alpha}, \mathbf{b}_\ell, \tau, \beta) &:= \frac{\beta\tau}{2} + \frac{\beta}{2\tau} \mathbb{E} \left[(\alpha_0 G_0 + \mathbf{a}^T \mathbf{g} + \mathbf{b}_\ell - Y_\ell)^2 \right] - \frac{\beta}{2\tau} \mathbb{E} \left[\frac{(\alpha_0 G_0 + \mathbf{a}^T \mathbf{g} + \mathbf{b}_\ell - Y_\ell)^2}{1 + \frac{\tau}{\beta} D^2} \right] - \beta \alpha_0 \sqrt{\gamma} \\ &= \frac{\beta\tau}{2} + \frac{1}{2} \mathbb{E} \left[\frac{D^2 (\alpha_0 G_0 + \mathbf{a}^T \mathbf{g} + \mathbf{b}_\ell - Y_\ell)^2}{1 + \frac{\tau}{\beta} D^2} \right] - \beta \alpha_0 \sqrt{\gamma} \\ &= \frac{\beta\tau}{2} + \frac{1}{2} \mathbb{E} \left[\frac{(\alpha_0 G_0 + \mathbf{a}^T \mathbf{g} + \mathbf{b}_\ell - Y_\ell)^2}{D^{-2} + (\tau/\beta)} \right] - \beta \alpha_0 \sqrt{\gamma}.\end{aligned}\tag{K.8}$$

We note that the function above is jointly convex in $(\alpha_0, \boldsymbol{\alpha}, \mathbf{b}_\ell, \tau)$ and concave in β .

K.1 Computing $\Sigma_{w,\mu}$

It can be checked that the first order optimality conditions of $\mathcal{D}_\ell(\alpha_0, \boldsymbol{\alpha}, \mathbf{b}_\ell, \tau, \beta)$ with respect to β and $\tau > 0$ are given as follows:

$$\beta^2 = \mathbb{E} \left[\frac{(\alpha_0 G_0 + \mathbf{a}^T \mathbf{g} + \mathbf{b}_\ell - Y_\ell)^2}{(D^{-2} + (\tau/\beta))^2} \right] \quad \text{or} \quad \beta = 0,\tag{K.9}$$

$$\alpha_0 \sqrt{\gamma} = \frac{\tau}{2} + \frac{\tau}{2\beta^2} \cdot \mathbb{E} \left[\frac{(\alpha_0 G_0 + \mathbf{a}^T \mathbf{g} + \mathbf{b}_\ell - Y_\ell)^2}{(D^{-2} + (\tau/\beta))^2} \right].\tag{K.10}$$

Thus, at optimality either $\beta = 0$ or $\tau = \alpha_0 \sqrt{\gamma}$. In what follows, consider the solution $\tau = \alpha_0 \sqrt{\gamma}$. We will show that this leads to the true saddle point of \mathcal{D} .

Moreover, by denoting $\eta := \frac{\beta}{\tau}$ and recalling from (K.7) that $Y_\ell = \text{Bern}(V_\ell)$, we can express $\mathcal{D}_\ell(\alpha_0, \boldsymbol{\alpha}, \mathbf{b}_\ell, \tau, \beta)$ as follows

$$\frac{\beta\tau}{2} + \frac{\alpha_0^2}{2} \mathbb{E} \left[\frac{1}{D^{-2} + 1/\eta} \right] - \beta \alpha_0 \sqrt{\gamma} + \frac{1}{2} [\mathbf{a}^T \mathbf{b}_\ell] \cdot \mathbf{A}(\eta) \cdot [\mathbf{a}] - \mathbf{c}_\ell^T(\eta) [\mathbf{b}_\ell] + \frac{1}{2} \mathbb{E} \left[\frac{Y_\ell^2}{D^{-2} + 1/\eta} \right],$$

where

$$\mathbf{A}(\eta) := \begin{bmatrix} \mathbb{E} \left[\frac{\mathbf{g} \mathbf{g}^T}{D^{-2} + 1/\eta} \right] & \mathbb{E} \left[\frac{\mathbf{g}}{D^{-2} + 1/\eta} \right] \\ \mathbb{E} \left[\frac{\mathbf{g}^T}{D^{-2} + 1/\eta} \right] & \mathbb{E} \left[\frac{1}{D^{-2} + 1/\eta} \right] \end{bmatrix}\tag{K.11a}$$

$$\mathbf{c}_\ell(\eta) := \left[\mathbb{E} \left[\frac{\mathbf{g} Y_\ell}{D^{-2} + 1/\eta} \right] \right]\tag{K.11b}$$

we have the following first-order optimality conditions for α_0 , \mathbf{a} and \mathbf{b}_ℓ :

$$\begin{bmatrix} \mathbf{a} \\ \mathbf{b}_\ell \end{bmatrix} = \mathbf{A}^{-1}(\eta) \cdot \mathbf{c}_\ell(\eta)\tag{K.12}$$

$$\alpha_0 = \beta \sqrt{\gamma} / \mathbb{E} \left[\frac{1}{D^{-2} + 1/\eta} \right].\tag{K.13}$$

Rearranging (K.13) and using $\tau = \alpha_0 \sqrt{\gamma}$ gives the following equation for η :

$$\frac{\alpha_0 \sqrt{\gamma}}{\beta} \mathbb{E} \left[\frac{1}{D^{-2} + 1/\eta} \right] = \gamma \stackrel{\tau = \alpha_0 \sqrt{\gamma}}{\implies} \mathbb{E} \left[\frac{1/\eta}{D^{-2} + 1/\eta} \right] = \gamma.\tag{K.14}$$

Thus, the optimal values of \mathbf{a} and \mathbf{b}_ℓ are found by (K.12) for η the positive solution of the equation in (K.14). To solve for α_0 , we combine (K.13) and (K.9) which leads to

$$\alpha_0^2 \left(\gamma \eta^2 - \mathbb{E} \left[\left(\frac{1}{D^{-2} + 1/\eta} \right)^2 \right] \right) = \mathbb{E} \left[\frac{(\mathbf{a}^T \mathbf{g} + \mathbf{b}_\ell - Y_\ell)^2}{(D^{-2} + 1/\eta)^2} \right],\tag{K.15}$$

where we have also used the RHS of (K.14). Next, we specialize these findings to the special structure of the weighting matrix \mathbf{D} in (K.1).

Applying weighting (K.1). Assume (K.1) holds. In this case, Equation (K.14) that determines the value of $\eta > 0$ becomes

$$F(\eta) := \sum_{i \in [k]} \frac{\pi_i \omega_i^2}{\omega_i^2 + \eta} = \gamma, \quad (\text{K.16})$$

where we have recalled the notation in (4.1) $\pi_i := \mathbb{E}[V_i] > 0$, $i \in [k]$. It can be easily checked by direct differentiation that $\eta \mapsto F$ is strictly decreasing in $(0, \infty)$. Also, using $\sum_{i \in [k]} \pi_i = 1$ the range of F in $(0, \infty)$ is $(0, 1)$. Thus, it follows that (K.16) has a unique solution for all $\gamma \in (0, 1)$.

Also, in this case we can write (K.11) in the following more convenient form:

$$\mathbf{A}(\eta) := \sum_{i \in [k]} \underbrace{\left(\frac{\omega_i^2 \eta}{\omega_i^2 + \eta} \right) \mathbb{E} \left[\begin{bmatrix} \mathbf{g} \\ 1 \end{bmatrix} \begin{bmatrix} \mathbf{g}^T & 1 \end{bmatrix} V_i \right]}_{=: \tilde{\mathbf{A}}_i} \quad (\text{K.17})$$

$$\mathbf{c}_\ell(\eta) := \underbrace{\left(\frac{\omega_\ell^2 \eta}{\omega_\ell^2 + \eta} \right) \mathbb{E} \left[\begin{bmatrix} \mathbf{g} \\ 1 \end{bmatrix} V_\ell \right]}_{=: \tilde{\mathbf{c}}_\ell}. \quad (\text{K.18})$$

For convenience let us define vectors $\boldsymbol{\nu} := \boldsymbol{\nu}(\eta)$, $\tilde{\boldsymbol{\pi}} = \tilde{\boldsymbol{\pi}}(\eta) \in \mathbb{R}^k$ with entries:

$$\tilde{\pi}_i := \pi_i \left(\frac{1}{\gamma} \cdot \frac{\omega_i^2}{\omega_i^2 + \eta} \right) =: \pi_i \cdot \nu_i \quad (\text{K.19})$$

Because of (K.16), notice that $\tilde{\boldsymbol{\pi}}$ is a probability vector, i.e.

$$\tilde{\boldsymbol{\pi}}^T \mathbf{1}_k = \boldsymbol{\pi}^T \boldsymbol{\nu} = 1.$$

With the notation above, it holds

$$\begin{aligned} \mathbf{A}(\eta) &= \gamma \cdot \eta \cdot \begin{bmatrix} \sum_{i \in [k]} \nu_i \cdot \mathbb{E}[V_i \mathbf{g} \mathbf{g}^T] & \sum_{i \in [k]} \nu_i \cdot \mathbb{E}[V_i \mathbf{g}] \\ \sum_{i \in [k]} \nu_i \cdot \mathbb{E}[V_i \mathbf{g}^T] & 1 \end{bmatrix} \\ &= \gamma \cdot \eta \cdot \begin{bmatrix} \sum_{i \in [k]} \nu_i \cdot \mathbb{E}[V_i \mathbf{g} \mathbf{g}^T] & \Sigma \mathbf{V}^T (\text{diag}(\boldsymbol{\pi}) - \boldsymbol{\Pi}) \boldsymbol{\nu} \\ \boldsymbol{\nu}^T (\text{diag}(\boldsymbol{\pi}) - \boldsymbol{\Pi}) \mathbf{V} \Sigma & 1 \end{bmatrix} \end{aligned} \quad (\text{K.20})$$

$$= \gamma \cdot \eta \cdot \begin{bmatrix} \mathbb{E}[(\boldsymbol{\nu}^T \mathbf{v}) \mathbf{g} \mathbf{g}^T] & \Sigma \mathbf{V}^T (\text{diag}(\boldsymbol{\pi}) - \boldsymbol{\Pi}) \boldsymbol{\nu} \\ \boldsymbol{\nu}^T (\text{diag}(\boldsymbol{\pi}) - \boldsymbol{\Pi}) \mathbf{V} \Sigma & 1 \end{bmatrix} \quad (\text{K.21})$$

$$\begin{aligned} \mathbf{c}_\ell(\eta) &= \gamma \cdot \eta \cdot \begin{bmatrix} \boldsymbol{\nu}_\ell \mathbb{E}[V_\ell \mathbf{g}] \\ \tilde{\pi}_\ell \end{bmatrix} \\ &= \gamma \cdot \eta \cdot \begin{bmatrix} \Sigma \mathbf{V}^T (\text{diag}(\boldsymbol{\pi}) - \boldsymbol{\Pi}) \boldsymbol{\nu}_\ell \mathbf{e}_\ell \\ \tilde{\pi}_\ell \end{bmatrix} \end{aligned} \quad (\text{K.22})$$

where we have also used the fact that $\mathbb{E}[V_i \mathbf{g}] = \Sigma \mathbf{V}^T (\text{diag}(\boldsymbol{\pi}) - \boldsymbol{\Pi}) \mathbf{e}_i$, $i \in [k]$ and recalled the notation

$$\mathbf{v} = [V_1, \dots, V_k]^T.$$

Using (K.30) and (K.31), we conclude from (K.12) the following expressions for \mathbf{a} and \mathbf{b} :

$$\mathbf{a} = \boldsymbol{\Delta}^{-1} \Sigma \mathbf{V}^T (\text{diag}(\boldsymbol{\pi}) - \boldsymbol{\Pi}) \cdot \boldsymbol{\nu}_\ell \cdot (\mathbf{e}_\ell - \pi_\ell \boldsymbol{\nu}), \quad (\text{K.23})$$

$$\mathbf{b}_\ell = \tilde{\pi}_\ell - \boldsymbol{\nu}^T (\text{diag}(\boldsymbol{\pi}) - \boldsymbol{\Pi}) \mathbf{V} \Sigma \boldsymbol{\Delta}^{-1} \Sigma \mathbf{V}^T (\text{diag}(\boldsymbol{\pi}) - \boldsymbol{\Pi}) \cdot \boldsymbol{\nu}_\ell \cdot (\mathbf{e}_\ell - \pi_\ell \boldsymbol{\nu}), \quad (\text{K.24})$$

where we defined

$$\boldsymbol{\Delta} = \mathbb{E}[(\boldsymbol{\nu}^T \mathbf{v}) \mathbf{g} \mathbf{g}^T] - \Sigma \mathbf{V}^T (\text{diag}(\boldsymbol{\pi}) - \boldsymbol{\Pi}) \boldsymbol{\nu} \boldsymbol{\nu}^T (\text{diag}(\boldsymbol{\pi}) - \boldsymbol{\Pi}) \mathbf{V} \Sigma > \mathbf{0}_{r \times r}. \quad (\text{K.25})$$

Finally, we show how to compute α_0 using (K.15). The RHS in (K.15) can be computed as

$$\begin{aligned} & \sum_{i \neq \ell \in [k]} \frac{[\mathbf{a}^T \quad \mathbf{b}_\ell] \widetilde{\mathbf{A}}_i \begin{bmatrix} \mathbf{a} \\ \mathbf{b}_\ell \end{bmatrix}}{(\omega_i^{-2} + 1/\eta)^2} + \frac{[\mathbf{a}^T \quad \mathbf{b}_\ell] \widetilde{\mathbf{A}}_\ell \begin{bmatrix} \mathbf{a} \\ \mathbf{b}_\ell \end{bmatrix} - 2[\mathbf{a}^T \quad \mathbf{b}_\ell] \widetilde{\mathbf{c}}_\ell + \boldsymbol{\pi}_\ell}{(\omega_\ell^{-2} + 1/\eta)^2} \\ &= \eta^2 \cdot \gamma^2 \cdot \left\{ [\mathbf{a}^T \quad \mathbf{b}_\ell] \left(\sum_{i \in [k]} \boldsymbol{\nu}_i^2 \widetilde{\mathbf{A}}_i \right) \begin{bmatrix} \mathbf{a} \\ \mathbf{b}_\ell \end{bmatrix} - 2[\mathbf{a}^T \quad \mathbf{b}_\ell] \boldsymbol{\nu}_\ell^2 \widetilde{\mathbf{c}}_\ell + \boldsymbol{\pi}_\ell \boldsymbol{\nu}_\ell^2 \right\}, \end{aligned}$$

where $\mathbf{a}, \mathbf{b}_\ell$ are as in (K.23) and (K.24). Also, note that

$$\mathbb{E} \left[\left(\frac{1}{D^{-2} + 1/\eta} \right)^2 \right] = \eta^2 \sum_{i \in [k]} \frac{\boldsymbol{\pi}_i \omega_i^4}{(\omega_i^2 + \eta)^2} = \eta^2 \cdot \gamma^2 \cdot \boldsymbol{\pi}^T \text{diag}(\boldsymbol{\nu}) \boldsymbol{\nu} = \eta^2 \cdot \gamma^2 \cdot \widetilde{\boldsymbol{\pi}}^T \boldsymbol{\nu}.$$

Put together, we have the following expression for α_0 :

$$\begin{aligned} \alpha_0^2 &= \frac{1}{(1/\gamma - \widetilde{\boldsymbol{\pi}}^T \boldsymbol{\nu})} \cdot \left\{ [\mathbf{a}^T \quad \mathbf{b}_\ell] \left(\sum_{i \in [k]} \boldsymbol{\nu}_i^2 \widetilde{\mathbf{A}}_i \right) \begin{bmatrix} \mathbf{a} \\ \mathbf{b}_\ell \end{bmatrix} - 2[\mathbf{a}^T \quad \mathbf{b}_\ell] \boldsymbol{\nu}_\ell^2 \widetilde{\mathbf{c}}_\ell + \boldsymbol{\pi}_\ell \boldsymbol{\nu}_\ell^2 \right\} \\ &= \frac{1}{(1/\gamma - \widetilde{\boldsymbol{\pi}}^T \boldsymbol{\nu})} \cdot \left\{ [\mathbf{a}^T \quad \mathbf{b}_\ell] \mathbf{A}' \begin{bmatrix} \mathbf{a} \\ \mathbf{b}_\ell \end{bmatrix} - 2[\mathbf{a}^T \quad \mathbf{b}_\ell] \begin{bmatrix} \boldsymbol{\Sigma} \mathbf{V}^T (\text{diag}(\boldsymbol{\pi}) - \boldsymbol{\Pi}) \boldsymbol{\nu}_\ell^2 \mathbf{e}_\ell \\ \widetilde{\boldsymbol{\pi}}_\ell \cdot \boldsymbol{\nu}_\ell \end{bmatrix} + \widetilde{\boldsymbol{\pi}}_\ell \cdot \boldsymbol{\nu}_\ell \right\}, \end{aligned} \quad (\text{K.26})$$

where $\mathbf{a}, \mathbf{b}_\ell$ are as in (K.23), (K.24) and we have also defined

$$\mathbf{A}' = \begin{bmatrix} \mathbb{E}[(\boldsymbol{\nu}^T \text{diag}(\boldsymbol{\nu}) \mathbf{v}) \mathbf{g} \mathbf{g}^T] & \boldsymbol{\Sigma} \mathbf{V}^T (\text{diag}(\boldsymbol{\pi}) - \boldsymbol{\Pi}) \text{diag}(\boldsymbol{\nu}) \boldsymbol{\nu} \\ \boldsymbol{\nu}^T \text{diag}(\boldsymbol{\nu}) (\text{diag}(\boldsymbol{\pi}) - \boldsymbol{\Pi}) \mathbf{V} \boldsymbol{\Sigma} & \boldsymbol{\nu}^T \text{diag}(\boldsymbol{\nu}) \boldsymbol{\pi} \end{bmatrix}. \quad (\text{K.27})$$

Asymptotic Predictions. Writing (K.24) in vector form we find that

$$\widehat{\mathbf{b}} \xrightarrow{P} \widetilde{\boldsymbol{\pi}} - \text{diag}(\boldsymbol{\nu}) (\mathbf{I}_k - \boldsymbol{\pi} \boldsymbol{\nu}^T) (\text{diag}(\boldsymbol{\pi}) - \boldsymbol{\Pi}) \mathbf{V} \boldsymbol{\Sigma} \boldsymbol{\Delta}^{-1} \boldsymbol{\Sigma} \mathbf{V}^T (\text{diag}(\boldsymbol{\pi}) - \boldsymbol{\Pi}) \boldsymbol{\nu}. \quad (\text{K.28})$$

Also, recalling that $\mathbf{e}_\ell^T \boldsymbol{\Sigma}_{\mathbf{w}, \boldsymbol{\mu}} = \widehat{\mathbf{w}}_\ell^T \mathbf{U} \boldsymbol{\Sigma} \mathbf{V}^T \xrightarrow{P} \mathbf{a}^T \boldsymbol{\Sigma} \mathbf{V}^T$ and using (K.23):

$$\boldsymbol{\Sigma}_{\mathbf{w}, \boldsymbol{\mu}} \xrightarrow{P} \text{diag}(\boldsymbol{\nu}) (\mathbf{I}_k - \boldsymbol{\pi} \boldsymbol{\nu}^T) (\text{diag}(\boldsymbol{\pi}) - \boldsymbol{\Pi}) \mathbf{V} \boldsymbol{\Sigma} \boldsymbol{\Delta}^{-1} \boldsymbol{\Sigma} \mathbf{V}^T. \quad (\text{K.29})$$

Finally, for the magnitudes of the weight vectors, recall that $\|\widehat{\mathbf{w}}_\ell\|_{\ell_2}^2 \xrightarrow{P} \|\mathbf{a}\|_{\ell_2}^2 + \alpha_0^2$. Thus, to find the limiting values of the norms, we can combine (K.26) and (K.23)-(K.24). For convenience, we summarize the final expression here. Define the following¹

$$\mathbf{A} := \begin{bmatrix} \mathbb{E}[(\boldsymbol{\nu}^T \mathbf{v}) \mathbf{g} \mathbf{g}^T] & \boldsymbol{\Sigma} \mathbf{V}^T (\text{diag}(\boldsymbol{\pi}) - \boldsymbol{\Pi}) \boldsymbol{\nu} \\ \boldsymbol{\nu}^T (\text{diag}(\boldsymbol{\pi}) - \boldsymbol{\Pi}) \mathbf{V} \boldsymbol{\Sigma} & 1 \end{bmatrix} \quad (\text{K.30})$$

$$\mathbf{c}_\ell := \begin{bmatrix} \boldsymbol{\Sigma} \mathbf{V}^T (\text{diag}(\boldsymbol{\pi}) - \boldsymbol{\Pi}) \boldsymbol{\nu}_\ell \mathbf{e}_\ell \\ \widetilde{\boldsymbol{\pi}}_\ell \end{bmatrix}. \quad (\text{K.31})$$

Further recall the matrix \mathbf{A}' in (K.27).

$$\begin{aligned} \|\widehat{\mathbf{w}}_\ell\|_{\ell_2}^2 &\xrightarrow{P} \|\boldsymbol{\Delta}^{-1} \boldsymbol{\Sigma} \mathbf{V}^T (\text{diag}(\boldsymbol{\pi}) - \boldsymbol{\Pi}) \cdot \boldsymbol{\nu}_\ell \cdot (\mathbf{e}_\ell - \boldsymbol{\pi}_\ell \boldsymbol{\nu})\|_{\ell_2}^2 \\ &+ \frac{1}{(1/\gamma - \widetilde{\boldsymbol{\pi}}^T \boldsymbol{\nu})} \cdot \{ \mathbf{c}_\ell^T \mathbf{A}^{-1} \mathbf{A}' \mathbf{A}^{-1} \mathbf{c}_\ell - 2\boldsymbol{\nu}_\ell \mathbf{c}_\ell^T \mathbf{A}^{-1} \mathbf{c}_\ell + \widetilde{\boldsymbol{\pi}}_\ell \cdot \boldsymbol{\nu}_\ell \}. \end{aligned} \quad (\text{K.32})$$

Remark K.2 Consider the special case $\omega_i = 1$, $i \in [k]$. We show how the above recovers the solution for (un-weighted) LS. First, note that in this case (K.16) simply gives $\eta = \frac{1}{\gamma} - 1$. Thus, $\boldsymbol{\nu} = \mathbf{1}_k$ and $\widetilde{\boldsymbol{\pi}} = \boldsymbol{\pi}$. Also, recall that $(\text{diag}(\boldsymbol{\pi}) - \boldsymbol{\Pi}) \mathbf{1}_k = \mathbf{0}$ and $\mathbf{1}^T \mathbf{v} = 1$. Thus, (K.25) simply gives $\boldsymbol{\Delta} = \mathbb{E}[\mathbf{g} \mathbf{g}^T] = \mathbf{I}_r$. With these, it can be readily checked that (K.28) and (K.23) simplify to the expressions in (4.4a). Similarly, $\mathbf{A} = \mathbf{A}' = \mathbf{I}_{r+1}$ and (K.26) reduces in this case to (I.1.1). For general weight coefficients, such simplifications do not seem possible and one needs to compute the matrix $\mathbb{E}[(\boldsymbol{\nu}^T \mathbf{v}) \mathbf{g} \mathbf{g}^T]$ that appears in the definitions of $\boldsymbol{\Delta}$, \mathbf{A} and \mathbf{A}' . We note that this calculation can be somewhat simplified by applying Gaussian integration by parts similar to lemma C.3.

¹Note the slight abuse of notation compared to the definitions in (K.30) and (K.30). This “renaming” should not be confusing as the constant $\gamma \cdot \eta$ (that is different between the two definitions) cancels when computing $\begin{bmatrix} \mathbf{a} \\ \mathbf{b}_\ell \end{bmatrix} = \mathbf{A}^{-1} \mathbf{c}_\ell$ (see (K.12)).

K.3 Computing $\Sigma_{w,w}$

In this section, we use Lemma H.1 to compute the cross-correlations $\langle \widehat{w}_\ell, \widehat{w}_j \rangle$, $j \neq \ell \in [k]$. Specifically, the analysis of (J.13) is almost identical to the analysis of (K.3) in the previous section. Specifically, without repeating all the details for brevity, it can be shown that the AO of (J.13) converges to $\min_{\mathbf{a}, \alpha_0 \geq 0, b_\ell, \tau > 0} \max_{\beta \geq 0} \mathcal{D}(\mathbf{a}, \alpha_0, b_\ell, \tau, \beta)$ where $\mathcal{D}(\mathbf{a}, \alpha_0, b_\ell, \tau, \beta)$ is as in (K.8) only with Y_ℓ substituted by $Y_{\ell,c}$:

$$Y_{\ell,c} \sim \text{Bern}(V_c + V_\ell) \quad \text{and as before: } V_i = \frac{e^{\mathbf{e}_i^T \mathbf{V} \Sigma \mathbf{g}}}{\sum_{\ell'=1}^r e^{\mathbf{e}_{\ell'}^T \mathbf{V} \Sigma \mathbf{g}}}, \quad i = \ell, c. \quad (\text{K.33})$$

Thus, what changes in the calculations above is in (K.18) and (K.35), where we now have instead

$$\mathbf{c}(\eta) := \left(\frac{1}{\frac{1}{\omega_\ell^2} + 1/\eta} \right) \underbrace{\mathbb{E} \left[\begin{bmatrix} \mathbf{g} \\ 1 \end{bmatrix} V_\ell \right]}_{=: \widetilde{\mathbf{c}}_\ell} + \left(\frac{1}{\frac{1}{\omega_c^2} + 1/\eta} \right) \underbrace{\mathbb{E} \left[\begin{bmatrix} \mathbf{g} \\ 1 \end{bmatrix} V_c \right]}_{=: \widetilde{\mathbf{c}}_c} \quad (\text{K.34})$$

and

$$\begin{aligned} \sum_{i \in \{\ell, c\} \in [k]} \frac{[\mathbf{a}^T \quad \mathbf{b}_\ell] \widetilde{\mathbf{A}}_i \begin{bmatrix} \mathbf{a} \\ \mathbf{b}_\ell \end{bmatrix}}{(\omega_i^{-2} + 1/\eta)^2} + \frac{[\mathbf{a}^T \quad \mathbf{b}_\ell] \widetilde{\mathbf{A}}_\ell \begin{bmatrix} \mathbf{a} \\ \mathbf{b}_\ell \end{bmatrix} - 2[\mathbf{a}^T \quad \mathbf{b}_\ell] \widetilde{\mathbf{c}}_\ell + \boldsymbol{\pi}_\ell}{(\omega_\ell^{-2} + \eta)^2} \\ + \frac{[\mathbf{a}^T \quad \mathbf{b}_\ell] \widetilde{\mathbf{A}}_c \begin{bmatrix} \mathbf{a} \\ \mathbf{b}_\ell \end{bmatrix} - 2[\mathbf{a}^T \quad \mathbf{b}_\ell] \widetilde{\mathbf{c}}_c + \boldsymbol{\pi}_c}{(\omega_c^{-2} + 1/\eta)^2}, \end{aligned} \quad (\text{K.35})$$

respectively. With these and following mutatis-mutandis the steps and the notation in the previous section, we find the following asymptotic expression for the magnitude of $\widehat{w}_\ell + \widehat{w}_c$:

$$\begin{aligned} \|\widehat{w}_\ell + \widehat{w}_c\|_{\ell_2}^2 &\xrightarrow{P} \|\Delta^{-1} \Sigma \mathbf{V}^T (\text{diag}(\boldsymbol{\pi}) - \boldsymbol{\Pi}) \cdot (\boldsymbol{\nu}_\ell \cdot (\mathbf{e}_\ell - \boldsymbol{\pi}_\ell \boldsymbol{\nu}) + \boldsymbol{\nu}_c \cdot (\mathbf{e}_c - \boldsymbol{\pi}_c \boldsymbol{\nu}))\|_{\ell_2}^2 \\ &+ \frac{1}{(1/\gamma - \widetilde{\boldsymbol{\pi}}^T \boldsymbol{\nu})} \cdot \left\{ (\mathbf{c}_\ell + \mathbf{c}_c)^T \mathbf{A}^{-1} \mathbf{A}' \mathbf{A}^{-1} (\mathbf{c}_\ell + \mathbf{c}_c) - 2(\boldsymbol{\nu}_\ell \mathbf{c}_\ell + \boldsymbol{\nu}_c \mathbf{c}_c)^T \mathbf{A}^{-1} (\boldsymbol{\nu}_\ell \mathbf{c}_\ell + \boldsymbol{\nu}_c \mathbf{c}_c) + \widetilde{\boldsymbol{\pi}}_\ell \cdot \boldsymbol{\nu}_\ell + \widetilde{\boldsymbol{\pi}}_c \cdot \boldsymbol{\nu}_c \right\}. \end{aligned}$$

We may now combine this with (K.32) to conclude with the following asymptotic limits for the cross-correlations for all $\ell \neq c \in [k]$:

$$\begin{aligned} \langle \widehat{w}_\ell, \widehat{w}_c \rangle &\xrightarrow{P} \boldsymbol{\nu}_c (\mathbf{e}_c - \boldsymbol{\pi}_c \boldsymbol{\nu})^T (\text{diag}(\boldsymbol{\pi}) - \boldsymbol{\Pi}) \mathbf{V} \Sigma \Delta^{-2} \Sigma \mathbf{V}^T (\text{diag}(\boldsymbol{\pi}) - \boldsymbol{\Pi}) \boldsymbol{\nu}_\ell (\mathbf{e}_\ell - \boldsymbol{\pi}_\ell \boldsymbol{\nu}) \\ &+ \frac{1}{(1/\gamma - \widetilde{\boldsymbol{\pi}}^T \boldsymbol{\nu})} \cdot \left\{ \mathbf{c}_c^T \mathbf{A}^{-1} \mathbf{A}' \mathbf{A}^{-1} \mathbf{c}_\ell - 2\boldsymbol{\nu}_c \boldsymbol{\nu}_\ell \mathbf{c}_c^T \mathbf{A}^{-1} \mathbf{c}_\ell \right\}. \end{aligned} \quad (\text{K.36})$$

$$\begin{aligned} &= \boldsymbol{\nu}_c (\mathbf{e}_c - \boldsymbol{\pi}_c \boldsymbol{\nu})^T (\text{diag}(\boldsymbol{\pi}) - \boldsymbol{\Pi}) \mathbf{V} \Sigma \Delta^{-2} \Sigma \mathbf{V}^T (\text{diag}(\boldsymbol{\pi}) - \boldsymbol{\Pi}) \boldsymbol{\nu}_\ell (\mathbf{e}_\ell - \boldsymbol{\pi}_\ell \boldsymbol{\nu}) \\ &+ \frac{1}{(1/\gamma - \widetilde{\boldsymbol{\pi}}^T \boldsymbol{\nu})} \cdot \left\{ \mathbf{c}_c^T (\mathbf{A}^{-1} \mathbf{A}' \mathbf{A}^{-1} - 2\boldsymbol{\nu}_c \boldsymbol{\nu}_\ell \mathbf{A}^{-1}) \mathbf{c}_\ell \right\}. \end{aligned} \quad (\text{K.37})$$

In matrix form, we have

$$\begin{aligned} \Sigma_{w,w} &\xrightarrow{P} \text{diag}(\boldsymbol{\nu}) (\mathbf{I}_k - \boldsymbol{\pi} \boldsymbol{\nu}^T) (\text{diag}(\boldsymbol{\pi}) - \boldsymbol{\Pi}) \mathbf{V} \Sigma \Delta^{-2} \Sigma \mathbf{V}^T (\text{diag}(\boldsymbol{\pi}) - \boldsymbol{\Pi}) (\mathbf{I}_k - \boldsymbol{\nu} \boldsymbol{\pi}^T) \text{diag}(\boldsymbol{\nu}) \\ &+ \frac{1}{(1/\gamma - \widetilde{\boldsymbol{\pi}}^T \boldsymbol{\nu})} \left\{ [\text{diag}(\boldsymbol{\nu}) (\text{diag}(\boldsymbol{\pi}) - \boldsymbol{\Pi}) \mathbf{V} \Sigma \quad \widetilde{\boldsymbol{\pi}}] \mathbf{A}^{-1} \mathbf{A}' \mathbf{A}^{-1} \left[\begin{array}{c} \Sigma \mathbf{V}^T (\text{diag}(\boldsymbol{\pi}) - \boldsymbol{\Pi}) \text{diag}(\boldsymbol{\nu}) \\ \widetilde{\boldsymbol{\pi}}^T \end{array} \right] \right\} \\ &- 2 \frac{1}{(1/\gamma - \widetilde{\boldsymbol{\pi}}^T \boldsymbol{\nu})} \left\{ \text{diag}(\boldsymbol{\nu}) [\text{diag}(\boldsymbol{\nu}) (\text{diag}(\boldsymbol{\pi}) - \boldsymbol{\Pi}) \mathbf{V} \Sigma \quad \widetilde{\boldsymbol{\pi}}] \mathbf{A}^{-1} \left[\begin{array}{c} \Sigma \mathbf{V}^T (\text{diag}(\boldsymbol{\pi}) - \boldsymbol{\Pi}) \text{diag}(\boldsymbol{\nu}) \\ \widetilde{\boldsymbol{\pi}}^T \end{array} \right] \text{diag}(\boldsymbol{\nu}) \right\} \\ &+ \frac{1}{(1/\gamma - \widetilde{\boldsymbol{\pi}}^T \boldsymbol{\nu})} \left\{ \text{diag}(\boldsymbol{\nu}) \text{diag}(\widetilde{\boldsymbol{\pi}}) \right\}. \end{aligned} \quad (\text{K.38})$$