



HAL
open science

Estimation of multiple mean vectors in high dimension

Gilles Blanchard, Jean-Baptiste Fermanian, Hannah Marienwald

► **To cite this version:**

Gilles Blanchard, Jean-Baptiste Fermanian, Hannah Marienwald. Estimation of multiple mean vectors in high dimension. 2024. hal-04515801

HAL Id: hal-04515801

<https://hal.science/hal-04515801v1>

Preprint submitted on 21 Mar 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

ESTIMATION OF MULTIPLE MEAN VECTORS IN HIGH DIMENSION

G. BLANCHARD^{*,†}, J-B. FERMANIAN^{*,‡}, H. MARIENWALD^{*,§}

ABSTRACT. We endeavour to estimate numerous multi-dimensional means of various probability distributions on a common space based on independent samples. Our approach involves forming estimators through convex combinations of empirical means derived from these samples. We introduce two strategies to find appropriate data-dependent convex combination weights: a first one employing a testing procedure to identify neighbouring means with low variance, which results in a closed-form plug-in formula for the weights, and a second one determining weights via minimization of an upper confidence bound on the quadratic risk. Through theoretical analysis, we evaluate the improvement in quadratic risk offered by our methods compared to the empirical means. Our analysis focuses on a dimensional asymptotics perspective, showing that our methods asymptotically approach an oracle (minimax) improvement as the effective dimension of the data increases. We demonstrate the efficacy of our methods in estimating multiple kernel mean embeddings through experiments on both simulated and real-world datasets.

KEYWORDS. aggregation estimator, effective dimension, high dimension, kernel mean embedding, minimax rate, multiple means estimation.

1. INTRODUCTION

We study the problem of jointly estimating multiple vector means μ_1, \dots, μ_B of distinct probability distributions $\mathbb{P}_1, \dots, \mathbb{P}_B$ over \mathbb{R}^d (an extension to Hilbert spaces is also discussed). The estimation of the means is based on a family of independent sample sets, $X_{\bullet}^{(1)}, \dots, X_{\bullet}^{(B)}$, where each $X_{\bullet}^{(k)}$ with $k \in \llbracket B \rrbracket := \{1, \dots, B\}$ comprises of N_k samples drawn i.i.d. from \mathbb{P}_k . Formally, the joint model is

$$\begin{cases} X_{\bullet}^{(k)} := (X_i^{(k)})_{1 \leq i \leq N_k} \stackrel{\text{i.i.d.}}{\sim} \mathbb{P}_k, k \in \llbracket B \rrbracket; \\ (X_{\bullet}^{(1)}, \dots, X_{\bullet}^{(B)}) \text{ independent.} \end{cases} \quad (1)$$

The distributions are assumed to be at least square-integrable. We refer to a set of samples $X_{\bullet}^{(k)}$ as *bag* and to \mathbb{P}_k as *task*, in line with the domain of multi-task learning.

* Authors contributed equally.

† Université Paris Saclay, Institut Mathématique d'Orsay, France. gilles.blanchard@universite-paris-saclay.fr.

‡ Université Paris Saclay, Institut Mathématique d'Orsay, France. jean-baptiste.fermanian@universite-paris-saclay.fr.

§ BIFOLD, Technische Universität Berlin, Germany. hannah.marienwald@campus.tu-berlin.de.

Our aim is to define estimators $\hat{\mu}_k$ and analyse the risk given by the expected squared distance to the true means μ_k .

Evident candidates are empirical means taken separately for each bag, which we call *naive* estimators. The question we want to tackle is whether it is possible to improve over these individual naive estimators by exploiting similarities between tasks. We propose and study particular estimators $\hat{\mu}_k$ formed by a convex combination of naive estimators of “related” tasks. We insist that absolutely no information about the underlying similarity or task structure is assumed to be known *a priori*. Roughly speaking, we measure relatedness between tasks by estimating the distance between their means.

The goal is to analyse the *relative* risk of the proposed estimators, i.e., the ratio of their risk to that of the corresponding naive estimator. The following questions will guide our estimator construction and analysis:

- (a) what would be the ideal “oracle” convex combination estimator, if some additional *a priori* information about task relatedness were known?
- (b) can an empirical estimator approach the oracle relative risk from the data only, in a suitable asymptotical sense?
- (c) is the oracle relative risk minimax optimal in a suitable asymptotical sense?

Because we focus on the relative risk, the usual asymptotics of the sample size going to infinity is not the most relevant one (though we will assume that the sample sizes are “large enough”). Rather, we will focus on *high-dimensional asymptotics* where the dimension grows large. More precisely, we mean a notion of *effective* dimension rather than ambient space dimension: the effective dimension of a task will be defined from spectral quantities related to its covariance matrix, as is common in high-dimensional statistics.

Motivations for this work. The framework under examination is motivated by scenarios involving large volumes of high-dimensional data. These scenarios typically involve the categorization of independent samples into homogeneous units that may exhibit differences but also varying degrees of similarity. Examples include medical or educational records sourced from different institutions, or purchase histories organised by individual clients on an internet platform. This framework also intersects with the concepts of federated and personalised machine learning (McAuley, 2022; Tan et al., 2022). An application of particular interest within this framework is that of kernel mean embeddings of distributions (Muandet, Fukumizu, Sriperumbudur, and Schölkopf, 2017). This involves estimating means of distributions after a formal mapping of the data into a Hilbert space. Notably, in this context, one anticipates that the effective dimensionality of the mapped data will be high.

Relation to previous work. The problem of estimating multiple means has a long and rich history in statistics, starting in particular with the seminal work of Stein on the eponymous paradox and the James-Stein estimator (James and Stein, 1961), continued with the empirical Bayes point of view on the latter (Efron and Morris, 1972), up to modern considerations on the topic (Brown and Greenshtein, 2009; Jiang

and Zhang, 2009). The topic of “multitask learning” also provides a more recent angle on the problem (Duan and Wang, 2023; Feldman et al., 2014). We defer a detailed discussion to Section 7, but stress that most previous works analysed the *compound* (or cumulated) risk over all tasks and its behaviour in the asymptotics $B \rightarrow \infty$, in a one- or fixed-dimensional setting. By contrast, we will be interested in analyzing the individual risk separately for each task, and in “high dimensional” asymptotics.

We start with a description of the considered setting in Section 2. Sections 3 and 4 introduce two approaches to form convex combination estimators of the means, provide bounds on their relative risks, and a comparison of the two. A minimax analysis for suitable distribution classes is conducted in Section 5. Finally, experiments on artificial and true data are presented in Section 6. All proofs are provided in the Supplemental, wherein Supplemental A contains a list of the used notation for the reader’s convenience.

2. SETTING AND NOTATION

2.1. Loss and risk. We consider the squared norm loss and expected risk

$$L_k(\hat{\mu}_k) := \|\hat{\mu}_k - \mu_k\|^2; \quad R_k(\hat{\mu}_k) := \mathbb{E}[L_k(\hat{\mu}_k)]. \quad (2)$$

of an estimator $\hat{\mu}_k$ for μ_k . The empirical mean $\hat{\mu}_k^{\text{NE}} := \frac{1}{N_k} \sum_{i=1}^{N_k} X_k^{(i)}$, called the *naive estimator*, serves as a reference. Due to the unbiasedness of the naive estimator, its variance is equal to its risk. More specifically, let the *naive risk* be denoted by

$$s_k^2 := R_k(\hat{\mu}_k^{\text{NE}}) = \frac{\text{Tr } \Sigma_k}{N_k}, \quad (3)$$

where Σ_k is the covariance of task k . Then any estimator $\hat{\mu}_k$ is analysed in terms of its *relative risk* to the naive — lower is better :

$$\frac{R_k(\hat{\mu}_k)}{s_k^2}. \quad (4)$$

In contrast to the compound decision setting, our goal is to analyse the relative risk for each task separately. For this reason, the focus is on a specific task, say $k = 1$ and $R_1(\hat{\mu}_1)/s_1^2$ without loss of generality. In Section 5.2 the relative risk averaged over tasks $\frac{1}{B} \sum_{k=1}^B R_k(\hat{\mu}_k)/s_k^2$ is considered.

2.2. High-dimensional asymptotics. Observe from (3) that the naive risk s_1^2 decreases at the parametric rate $\mathcal{O}(N_1^{-1})$. We expect the risk of a competing estimator $\hat{\mu}_1$ to follow the same trend. As a consequence, the role of the sample size will cancel out in the relative risk. In order to state meaningful results, it is necessary to obtain sharp estimates of the other factors in the rate.

To this end, we shift the perspective from a standard asymptotic view point, $N_1 \rightarrow \infty$, to high-dimensional asymptotics, emphasizing the behaviour of the risks as the dimensionality grows. There are different possible definitions of *effective dimensionality*

of a distribution, generally linked to the spectral decay of the covariance matrix and ratios of its Schatten norms. The following ones will be relevant to our analysis:

$$d_k^\bullet := \frac{(\text{Tr } \Sigma_k)^2}{\text{Tr } \Sigma_k^2}, \quad d_k^e := \frac{\text{Tr } \Sigma_k}{\|\Sigma_k\|_\infty}. \quad (5)$$

Observe that in the isotropic setting $\Sigma_k \propto I_d$, the effective dimensions d_k^\bullet and d_k^e coincide with the ambient dimension d , as one would expect. In all cases it holds $1 \leq \sqrt{d_k^\bullet} \leq d_k^e \leq d_k^\bullet \leq d$. In random matrix literature, d^e is sometimes called intrinsic dimension (Hsu et al., 2012; Tropp et al., 2015) or effective rank (Koltchinskii and Lounici, 2016), and $(d^e)^2/d^\bullet$ is known as the numerical or stable rank of Σ (Rudelson and Vershynin, 2007; Tropp et al., 2015). Most notably, we uncover a ‘‘blessing of dimensionality’’ phenomenon: in a nutshell, we will show that the relative risks of our estimators asymptotically approach a suitable notion of oracle relative risk as the (effective) dimensionality increases.

2.3. Distributional assumptions. For our theoretical analysis, we consider the following different possible distributional assumptions:

Assumption (GS, Gaussian setting). For all $k \in \llbracket B \rrbracket$, the distribution \mathbb{P}_k is $\mathcal{N}(\mu_k, \Sigma_k)$.

Assumption (BS, Bounded setting). For all $k \in \llbracket B \rrbracket$, \mathbb{P}_k has support in the ball of radius M centred at 0.

The **(BS)** setting is of particular interest for the application to kernel mean embeddings, for which the assumption of a bounded kernel is very common. All results for **(BS)** are presented in \mathbb{R}^d but can be extended to a separable Hilbert space (up to adequate adaptation of notation).

Supplemental C.4.3 covers another distributional assumption: heavy-tailed distributions with finite fourth moment. These results only hold for some of the proposed estimators (the testing approach, introduced in Section 3) and are, thus, not discussed further in the main text of this paper.

2.4. Simplifying settings. At times we will discuss unrealistic but simplifying settings to help with the exposition or to illuminate our theoretical findings.

Setting (ECSS, Equal Covariance and Sample Sizes). For all $k \in \llbracket B \rrbracket$, $\Sigma_k = \Sigma$ and $N_k = N$, which implies that $s_k^2, d_k^\bullet, d_k^e$ do not depend on k .

Setting (KC, Known Covariances). For all $k \in \llbracket B \rrbracket$, Σ_k is known. Consequently, all derived quantities $\text{Tr } \Sigma_k, \text{Tr } \Sigma_k^2, d_k^\bullet, d_k^e, s_k^2$ are also known.

We will first derive the estimators assuming known covariances **(KC)** but later provide estimates for covariance-related quantities if those are unknown. If the covariances and sample sizes are homogeneous **(ECSS)** the risks are more transparent and interpretable which will help to illuminate our theoretical findings. We insist that the final algorithms neither assume **(KC)** nor **(ECSS)**.

2.5. Naive estimator aggregation. As announced earlier, without loss of generality we focus on estimating task $k = 1$. Furthermore, we focus on estimators which can be written as convex combinations (aggregation) of naive estimators. Let \mathcal{S}_B denote the $(B - 1)$ -dimensional simplex, and $\boldsymbol{\omega} = (\omega_1, \dots, \omega_B) \in \mathcal{S}_B$ be a weight vector, then

$$\hat{\boldsymbol{\mu}}_{\boldsymbol{\omega}} := \sum_{k \in \llbracket B \rrbracket} \omega_k \hat{\boldsymbol{\mu}}_k^{\text{NE}} \quad \text{s.t.} \quad \sum_{k \in \llbracket B \rrbracket} \omega_k = 1 \quad \text{and} \quad \forall k \in \llbracket B \rrbracket : \omega_k \geq 0, \quad (6)$$

whose loss and risk will be abbreviated as $L_1(\boldsymbol{\omega})$ and $R_1(\boldsymbol{\omega})$, respectively. While the weight vector $\boldsymbol{\omega}$ may be data-dependent later, for the present considerations we assume that the weights are *deterministic*. In this case, using independence of the naive estimators and the notation $\Delta_k := \boldsymbol{\mu}_k - \boldsymbol{\mu}_1$, we restate the risk $R_1(\boldsymbol{\omega})$ by its bias-variance decomposition for a fixed $\boldsymbol{\omega}$ as

$$R_1(\boldsymbol{\omega}) = \left\| \sum_{k \in \llbracket B \rrbracket} \omega_k (\boldsymbol{\mu}_k - \boldsymbol{\mu}_1) \right\|^2 + \sum_{k \in \llbracket B \rrbracket} \omega_k^2 s_k^2 = \sum_{k, k' \in \llbracket B \rrbracket} \omega_k \omega_{k'} \langle \Delta_k, \Delta_{k'} \rangle + \sum_{k \in \llbracket B \rrbracket} \omega_k^2 s_k^2, \quad (7)$$

where the first term corresponds to the (squared) bias and the second to the variance. Intuitively, we want to give higher weights to tasks that are close (small task bias $\|\Delta_k\|$) and can be accurately estimated (small naive risk s_k^2). At a first glance, we could set as a goal to find suitable weights $\boldsymbol{\omega}$ that minimise (7); this, however, would require full knowledge of the Gram matrix $(\langle \Delta_k, \Delta_{k'} \rangle)_{k, k' \in \llbracket B \rrbracket}$, in addition to the naive risks s_k^2 . Estimation of the full Gram matrix, accurate enough to approach exact minimization of (7), appears unattainable if the number of tasks B is large and the Gram matrix becomes high-dimensional, which is the scenario we are interested in. For this reason, we will consider optimizing the risk given more limited information, which includes a subset of neighbouring tasks close to the target in relative sense but not their exact position. We define the oracle risk as the minimiser of the worst-case risk of (7) as if this partial information was known to the oracle.

We will consider two strategies to approach that oracle programme from data. In Section 3 we aggregate only means close to the target which are identified by a test procedure. Minimization of an upper bound of the risk yields their weights. In Section 4 we minimise directly an upper confidence bound of the aggregate risk (7) but have to take into account that the means that are further away induce a large uncertainty on the bias term. In both cases, we compare the obtained relative risk to that of the oracle. Additionally, we study the minimax risk under the oracle information in Section 5 and whether the proposed estimators match it.

3. A TESTING APPROACH

A low-risk aggregation estimator (6) combines naive estimations that — at best — provide a reduction in variance but add only a small bias, cf. (7). Our first approach explicitly controls the bias. We aim at identifying a subset of *neighbour* tasks whose means are sufficiently close to the target task. We then restrict the support of the

weights to that subset and form a convex combination of neighbouring naive estimations. This approach and its analysis generalise ideas introduced in Marienwald et al. (2021). Let us first introduce some additional notation.

Definition 1 (τ -neighbouring tasks). Recall the notation $\Delta_k = \mu_k - \mu_1$. For a fixed $\tau > 0$, let $V_\tau \subseteq \llbracket B \rrbracket$ denote the set of all τ -neighbouring tasks (of task 1) as:

$$V_\tau := \{k \in \llbracket B \rrbracket : \|\Delta_k\|^2 \leq \tau s_1^2\}. \quad (8)$$

For $\tau = 0$, for the sake of later notational coherence we define $V_0 := \{1\}$ which deviates from (8) as V_0 does not contain any other tasks $k \neq 1$ even if $\Delta_k = 0$.

Note that this notion of τ -neighbourhood is relative to the naive risk of task 1, and that $1 \in V_\tau$ always holds.

Definition 2 (Relative aggregated variance ν). For a subset $U \subseteq \llbracket B \rrbracket$ of tasks, define their relative aggregated variance (to that of task 1) as:

$$\nu(U) := \frac{s^2(U)}{s_1^2}, \text{ with } s^2(U) := \left(\sum_{k \in U} \frac{1}{s_k^2} \right)^{-1}. \quad (9)$$

Observe that $s^2(U)$ is the variance of the optimal convex combination of unbiased, independent estimators that have different variances s_k^2 — a classical problem of statistics. The quantity $\nu(U)$ is, again, relative to the naive risk of task 1.

The quantity τ can be seen as the worst-case relative bias of a convex combination of their naive estimators for the goal of estimating μ_1 , while $\nu(V_\tau)$ is a best-case relative variance (i.e., all the tasks in V_τ would in fact have mean μ_1). We introduce the following auxiliary function, which will capture an optimal trade-off between these two quantities. It provides a common reference value for the relative risks of our estimators and is of fundamental importance for the remainder of this manuscript.

Definition 3. Define the function $\mathcal{B} : \mathbb{R} \times [0, 1] \rightarrow [0, 1]$ as

$$\mathcal{B}(\tau, \nu) := \left(\frac{\tau}{1 + \tau} \right) + \left(\frac{1}{1 + \tau} \right) \left(\frac{\nu}{1 + \tau(1 - \nu)} \right). \quad (10)$$

Observe that $\mathcal{B}(0, \nu) = \nu$, $\mathcal{B}(\tau, 0) = \frac{\tau}{1 + \tau}$, and \mathcal{B} is increasing in both of its variables.

In the next section, we derive a form of optimal or “oracle” weights for combining naive estimators of tasks belonging to any given subset $V \subseteq V_\tau$, and identify \mathcal{B} as a bound on its relative risk. The following sections (3.2 to 3.4) are concerned with approximating the oracle bound by estimating unknown quantities and using a plug-in principle.

3.1. Oracle procedure. For a fixed $\tau > 0$, assume an oracle provides a set of neighbours V with the guarantee that $V \subseteq V_\tau$ holds. We restrict our attention to convex combinations of naive estimators only in set V , i.e., estimators $\hat{\mu}_\omega$ as in (6) with $\omega \in \mathcal{S}_V$,

the set of convex weights of support included in V . Using the Cauchy-Schwartz inequality in (7) (with $\Delta_1 = 0$), for such aggregated estimators we obtain the risk bound

$$R_1(\boldsymbol{\omega}) \leq \tau s_1^2 (1 - \omega_1)^2 + \sum_{k \in V} \omega_k^2 s_k^2, \quad \text{for all } \boldsymbol{\omega} \in \mathcal{S}_V, \quad (11)$$

which can be optimised for $\boldsymbol{\omega}$. A bound on the oracle relative risk is presented next.

Lemma 1. *Let $\tau > 0$ be fixed. For all $V \subseteq V_\tau$, the weights $\boldsymbol{\omega}_V^* \in \mathcal{S}_V$ that minimise (11) yield the bound*

$$\frac{R_1(\boldsymbol{\omega}_V^*)}{s_1^2} \leq \mathcal{B}(\tau, \nu(V)). \quad (12)$$

The oracle weights $\boldsymbol{\omega}_V^*$ are given by:

$$\omega_{V,k}^*(\tau, \mathbf{s}) = (1 - \lambda) \mathbf{1}\{k = 1\} + \lambda \frac{s^2(V)}{s_k^2}, \quad \text{where } \lambda := \frac{1}{1 + \tau(1 - \nu(V))}. \quad (13)$$

It holds $\mathcal{B}(\tau, \nu(V)) \in [\frac{\tau}{1+\tau}, 1]$, i.e., this bound cannot be better than $\frac{\tau}{1+\tau}$. We will call $\frac{\tau}{1+\tau}$ *best potential improvement* (that can be guaranteed by the oracle bound). The bound on the relative risk depends on the relative neighbourhood size τ and the relative aggregated variance $\nu(V)$. Because \mathcal{B} increases in both variables, small τ and $\nu(V)$ are beneficial. This coincides with what we noted from the bias-variance decomposition (7). If τ is fixed, it is of advantage to consider as many τ -neighbours as possible so that $\nu(V)$ decreases, i.e., to take $V = V_\tau$. On the other hand, reducing the neighbourhood size τ reduces the bias but also leads to a smaller set of neighbours, ergo, a larger relative aggregated variance $\nu(V_\tau)$. Thus, there is a trade-off between both quantities. We may aim at a relative risk close to $\min_{\tau > 0} \mathcal{B}(\tau, V_\tau)$ but for the remainder of this section we assume $\tau > 0$ fixed beforehand.

The following observations enable additional insight into the involved quantities:

- (a) $\mathcal{B}(0, \nu(V)) = \nu(V)$, i.e., when $\tau \searrow 0$, which implies that all tasks in V have the same mean, the bound is given by the relative aggregated variance, as should be expected from the remark following Definition 2.
- (b) $\mathcal{B}(\tau, 0) = \frac{\tau}{1+\tau}$, the best potential improvement is reached when $s^2(V) \searrow 0$. This happens if at least one of the τ -neighbouring means is known with perfect precision and it becomes a “reference point”. This scenario is comparable to the classical James-Stein setting, for which the origin is such a reference point and the James-Stein estimate improves most if the target is close to the origin (see Supplemental B for a detailed discussion). However, $s^2(V) \searrow 0$ also happens when τ -neighbours have a non-zero variance, but their number grows large.
- (c) $\mathcal{B}(\tau, \nu(V))$ remains unchanged if we replace a group of neighbours $V \setminus \{1\}$ by a single τ -neighbour with variance $s^2(V \setminus \{1\})$.

3.2. From an oracle to an empirical procedure. In practice, the oracle information about the relative neighbours is unavailable. However, we can hope to approach the oracle setting by estimating the set of τ -neighbours V_τ and their risks s_k^2 . We will

assume that such estimates are independent of the samples used to compute $(\widehat{\mu}_k^{\text{NE}})_{k \in \llbracket B \rrbracket}$. (To this end, one might resort to sample splitting.) The independence assumption of estimates is emphasised by a tilde notation: $(\widetilde{V}, \widetilde{\mathbf{s}}^2)$.

The simplest is to plug in such estimates into the oracle formula (13). The next proposition quantifies how the relative risk of the plug-in procedure can be bounded, provided the estimation error is.

Proposition 1. *Let $\tau > 0$ be fixed. Assume $\widetilde{V} \subseteq \llbracket B \rrbracket$, $\widetilde{\mathbf{s}}^2 = (\widetilde{s}_k^2)_{k \in \llbracket B \rrbracket} \in \mathbb{R}_+^B$ are possibly random but independent of the samples in model (1). Let V^* be some deterministic reference set, such that $1 \in V^*$. Let $(\widetilde{V}, \widetilde{\mathbf{s}}^2)$ be plugged in for (V, \mathbf{s}^2) into (13), giving rise to weight vector $\widetilde{\omega}$. Conditionally to the event*

$$\begin{cases} V^* \subseteq \widetilde{V} \subseteq V_\tau, \\ |\widetilde{s}_k^2 - s_k^2| \leq \eta s_k^2, \text{ for all } k \in \widetilde{V}, \text{ and some } \eta \in [0, 1), \end{cases} \quad (14)$$

it holds

$$\frac{R_1(\widetilde{\omega})}{s_1^2} \leq \left(\frac{1 + \eta}{1 - \eta} \right) \mathcal{B}(\tau, \nu(\widetilde{V})) \leq \left(\frac{1 + \eta}{1 - \eta} \right) \mathcal{B}(\tau, \nu(V^*)). \quad (15)$$

Comparing the oracle relative risk bound (12) with that of the empirical procedure (15), note first the requirement that all estimated neighbours are τ -neighbours ($\widetilde{V} \subseteq V_\tau$); secondly, the oracle risk is deteriorated by two factors: the excess factor $(1 + \eta)/(1 - \eta) \geq 1$ which quantifies what we lose due to estimation of the neighbours' risks; and the replacement of the set of true neighbours by the smaller set V^* , under the requirement that $V^* \subseteq \widetilde{V}$ holds. To summarise, we expect the risk of the empirical procedure to be close to the oracle risk if (1) the relative estimation error η for naive risks is small, and (2) we can guarantee the ‘‘sandwiching’’ property $V^* \subseteq \widetilde{V} \subseteq V_\tau$, with V^* as large as possible; typically we would be satisfied with $V^* = V_{(1-\varepsilon)\tau}$ for a small ε .

The next sections will introduce such estimates and the fulfillment of event (14) under certain conditions, starting with the estimation of neighbour tasks.

3.3. Finding neighbours (known covariances). For now let us assume **(KC)**; we will generalise to unknown covariances in the next section. Accordingly, the naive risks s_k^2 are known, so that $\eta = 0$ in the context of (15), and we focus on the estimation of the set of neighbours. We assume that we are doing so using independent ‘‘tilde’’ data $(\widetilde{X}_\bullet^{(k)})_{k \in \llbracket B \rrbracket}$ which are drawn from (1) but independent of $(X_\bullet^{(k)})_{k \in \llbracket B \rrbracket}$ (e.g., using sample splitting). For clarity $X_\bullet^{(k)}$ and $\widetilde{X}_\bullet^{(k)}$ are assumed to be of the same size N_k . Given the first requirement $\widetilde{V} \subseteq V_\tau$, it is natural to think of \widetilde{V} as the output of a multiple test procedure (for which the null hypothesis for task k is *not* being a τ -neighbour, i.e., $\|\Delta_k\| > \tau s_1^2$).

Our approach is based on recent results for two-sample mean vector testing (Blanchard and Fermanian, 2023). Assume $N_k \geq 2$ for all $k \in \llbracket B \rrbracket$. For $k \in \llbracket B \rrbracket \setminus \{1\}$, we

form an unbiased estimator for $\|\Delta_k\|^2$ based on the U-statistics

$$\tilde{U}_k := \sum_{\ell \in \{1, k\}} \sum_{\substack{i, j=1 \\ i \neq j}}^{N_\ell} \frac{\langle \tilde{X}_i^{(\ell)}, \tilde{X}_j^{(\ell)} \rangle}{N_\ell(N_\ell - 1)} - 2 \sum_{i=1}^{N_1} \sum_{j=1}^{N_k} \frac{\langle \tilde{X}_i^{(1)}, \tilde{X}_j^{(k)} \rangle}{N_1 N_k}. \quad (16)$$

The following proposition is a direct consequence of Blanchard and Fermanian (2023):

Proposition 2. *Assume (GS), (KC) hold and let $\alpha \in (0, 1)$, $\tau > 0$ be fixed. Let $\tilde{T}_k^{(\tau)}$ be given by*

$$\tilde{T}_k^{(\tau)} := \mathbf{1} \left\{ \tilde{U}_k \leq \tau s_1^2 \right\}. \quad (17)$$

Put for $k \in \llbracket B \rrbracket$

$$\tau_{\min}^k := 32 \left(\frac{1}{\sqrt{d_1^\bullet}} + \frac{s_k^2/s_1^2}{\sqrt{d_k^\bullet}} \right) \log(8\alpha^{-1}), \quad (18)$$

then it holds:

$$\text{if } \|\mu_1 - \mu_k\|^2 > \tau_k^+ s_1^2 : \quad \mathbb{P} \left[\tilde{T}_k^{(\tau)} = 1 \right] \leq \alpha; \quad (19)$$

$$\text{if } \|\mu_1 - \mu_k\|^2 \leq \tau_k^- s_1^2 : \quad \mathbb{P} \left[\tilde{T}_k^{(\tau)} = 0 \right] \leq \alpha. \quad (20)$$

where $\tau_k^\pm = \left(\sqrt{\tau} \pm \sqrt{\tau_{\min}^k} \right)_+^2$.

Equations (19)-(20) can be understood as controls of the type I/II error level for the test of $\|\Delta_k\|^2 > \tau_k^+ s_1^2$ versus the alternative $\|\Delta_k\|^2 \leq \tau_k^- s_1^2$. It is possible to make the original null hypothesis $\|\Delta_k\|^2 > \tau s_1^2$ appear through notation translation ($\sqrt{\tau} \leftarrow \sqrt{\tau_k^-}$, $\sqrt{\tau_k^+} \leftarrow \sqrt{\tau}$, if we assume additionally $\tau \geq \tau_{\min}^k$). We prefer to keep the above more symmetric form, also because the rejection set (17) has a simple form, used in practice.

The test is able to identify mean differences very accurately relative to the target threshold τs_1^2 if $\tau \gg \tau_{\min}^k$. Formula (18) highlights the crucial role of the effective dimensionality for this minimal threshold of reliable detection. In the simplified (ECSS) setting, this threshold is simply of order $1/\sqrt{d_1^\bullet}$. This reflects the known phenomenon that testing is more reliable than estimation in high dimensions; distances that can be detected might be of smaller order than the typical estimation error. For fixed τ and increasing dimension, the inconclusive gap between the null and the alternative vanishes with increasing dimension — a desirable property given the sandwiching property that we aim for (see (14)).

In general non-(ECSS) configurations, we still want to keep τ_{\min}^k small of order $1/\sqrt{d_1^\bullet}$. In view of the second term in (18), this suggests to only consider tasks with $s_k^2/\sqrt{d_k^\bullet} \leq \varsigma s_1^2/\sqrt{d_1^\bullet}$ for some constant $\varsigma \geq 1$. To this aim, denote the set of tasks

satisfying this criterion as

$$W_{(\varsigma)} := \left\{ k \in \llbracket B \rrbracket : \frac{s_k^2}{\sqrt{d_k^\bullet}} \leq \varsigma \frac{s_1^2}{\sqrt{d_1^\bullet}} \right\} = \left\{ k \in \llbracket B \rrbracket : \frac{\|\Sigma_k\|_2}{N_k} \leq \varsigma \frac{\|\Sigma_1\|_2}{N_1} \right\}, \quad (21)$$

and correspondingly the set of whittled down neighbours as

$$V_{\tau, \varsigma} := V_\tau \cap W_{(\varsigma)}. \quad (22)$$

Note that since we are under **(KC)**, the set $W_{(\varsigma)}$ is assumed to be fully known for now. (We will consider estimating it in the next section.) Then the following corollary makes the obtained sandwiching property explicit:

Corollary 1. *Let $\varsigma \geq 1$ be fixed. Assume **(GS)** and **(KC)** hold and let $\alpha \in (0, 1)$. Then, defining*

$$\tilde{V}_{\tau, \varsigma} := \left\{ k \in \llbracket B \rrbracket : \tilde{T}_k^{(\tau)} = 1 \right\} \cap W_{(\varsigma)}$$

(where $\tilde{T}_k^{(\tau)}$ is as in (17)), with probability at least $1 - \alpha$ it holds

$$V_{\tau^-, \varsigma} \subseteq \tilde{V}_{\tau, \varsigma} \subseteq V_{\tau^+, \varsigma}, \quad (23)$$

$$\text{where } \tau^\pm := \left(\sqrt{\tau} \pm \sqrt{\varsigma \tau_{\min}^\circ} \right)_+, \quad \tau_{\min}^\circ := 64 \log(8B\alpha^{-1}) / \sqrt{d_1^\bullet}.$$

The sandwiching property (23) provides a direct link to Proposition 1. More specifically, Corollary 1 together with Proposition 1 guarantee with high probability that the bound on the relative risk of the plug-in estimate $\hat{\mu}_{\tilde{\omega}}$ of (13) using the estimated set of neighbours $\tilde{V}_{\tau, \varsigma}$ is bounded by $\mathcal{B}(\tau^+, \nu(V_{\tau^-, \varsigma}))$ (recall $\eta = 0$ for now because of **(KC)**, and $\tilde{V}_0 := \{1\}$). Furthermore, for fixed τ , if $d_1^\bullet / (\log B)^2 \rightarrow \infty$ then τ_{\min}° vanishes and it holds $\tau^- \approx \tau \approx \tau^+$. Under **(ECSS)**, we can simply take $\varsigma = 1$ and have $V_{\tau^-, \varsigma} = V_{\tau^-}$, ensuring a relative risk very close to the oracle $\mathcal{B}(\tau, \nu(V_\tau))$. In a general context, there is an additional trade-off through the choice of the constant ς . In both cases, closeness to the oracle relative risk *improves* with increasing effective dimensionality.

3.4. Unknown covariances. In a realistic setting the covariances are unknown, especially in high dimensions. In this section, we estimate all quantities relevant for the fulfilment of Proposition 1, using the same independent “tilde” data $(\tilde{X}_i^{(k)})_{k \in \llbracket B \rrbracket}$ as in the previous section. For simplicity we assume that the sizes N_k of the “tilde” samples are the same as that of the main sample, as we would get by equal-size splitting. Observe that it is not necessary to estimate the full covariance matrices Σ_k , but only scalar quantities related to their Schatten norms. In particular, in the Gaussian setting we have the following result for the natural unbiased estimators of s_k^2 :

Proposition 3. *Let $\tilde{s}_k^2 := \frac{1}{N_k(N_k-1)} \sum_{i=1}^{N_k} \|\tilde{X}_i^{(k)} - \tilde{\mu}_k^{NE}\|^2$, where $\tilde{\mu}_k^{NE} := N_k^{-1} \sum_{i=1}^{N_k} \tilde{X}_i^{(k)}$, and let $\alpha \in (0, 1)$. Assume **(GS)** holds. Then with probability at least $1 - \alpha$:*

$$\forall k \in \llbracket B \rrbracket : \quad \left| \tilde{s}_k^2 - s_k^2 \right| \leq \left(4\sqrt{2} \frac{\log(2B\alpha^{-1})}{\sqrt{d_k^\bullet N_k}} \right) s_k^2. \quad (24)$$

When $N_k \gtrsim \log^2(2B\alpha^{-1})$ for all k , the estimation of s_k^2 has relative accuracy of order $1/\sqrt{d_k^\bullet}$ with probability $(1 - \alpha)$. This finding can be used for the fulfillment of the second requirement of condition (14). It also allows to preserve the qualitative results of Proposition 2 (up to numerical factors) for test (17) wherein \tilde{s}_1^2 is plugged in for s_1^2 . Finally, we also replace $\|\Sigma_k\|_2$ in the definition (21) of set $W_{(\varsigma)}$ by suitable estimators; Proposition 14 in the Supplemental gives the details. It provides a quantitatively precise version of the sandwiching property analogous to (23) with all unknown quantities are replaced by their proposed estimators.

We combine the obtained results in an illustrative example. It shows a fully empirical algorithm that approximates the (whittled down) oracle $B(\tau, V_{\tau, \varsigma})$ (numerical constants are made explicit for concreteness but not meant to be sharp):

Proposition 4. *Assume (GS) holds. Let $\alpha \in (0, 1/3)$. Consider the following plug-in versions of the quantities appearing in (17), (21):*

$$\widetilde{W}_{(\varsigma)} := \left\{ k \in \llbracket B \rrbracket : \frac{\widetilde{Z}_k^{(2)}}{N_k} \leq \varsigma \frac{\widetilde{Z}_1^{(2)}}{N_1} \right\}, \quad \widetilde{T}_k^{(\tau)} := \mathbf{1}\{\widetilde{U}_k \leq \tau \tilde{s}_k^2\}, \quad (25)$$

where \tilde{s}_k^2 as in Prop. 3, and $\widetilde{Z}_k^{(2)}$ estimates $\|\Sigma_k\|_2$ as defined in (59) in the Supplemental. Define the set of estimated τ -neighbours

$$\widetilde{V}_{\tau, \varsigma} := \left\{ k \in \widetilde{W}_{(\varsigma)} : \widetilde{T}_k^{(\tau)} = 1 \right\}. \quad (26)$$

Assume $N_k \geq a(4 + \log(2B\alpha^{-1}))^4$ for all $k \in \llbracket B \rrbracket$, for a big enough numerical constant a ($a = 4400$ works). For fixed $\tau > 0, \varsigma \geq 1$, consider the weights $\tilde{\omega}^\#$ obtained by the modified plug-in $(\widetilde{V}_{\tau, 3\varsigma}, \tilde{\mathbf{s}}^2)$ for (V, \mathbf{s}^2) in (13), where

$$\tilde{\tau} := \left(1 + \frac{1}{60\sqrt{d_1^\bullet}} \right) \left(\sqrt{\tau} + \sqrt{6\varsigma\tilde{\tau}_{\min}^\circ} \right)^2; \quad \tilde{\tau}_{\min}^\circ := \frac{32(\log(8B\alpha^{-1}))}{\sqrt{d_1^\bullet}}; \quad \sqrt{d_1^\bullet} := \frac{N_1\tilde{s}_1^2}{\widetilde{Z}_1^{(2)}}. \quad (27)$$

Then with probability at least $1 - 3\alpha$ over the draw of the “tilde” sample $(\widetilde{X}_\bullet^{(k)})_{k \in \llbracket B \rrbracket}$, it holds

$$\frac{R_1(\tilde{\omega}^\#)}{s_1^2} \leq \left(1 + \frac{1}{10\sqrt{\min_k d_k^\bullet}} \right) \left(1 + \frac{30\sqrt{\varsigma \log(8B\alpha^{-1})}}{(d_1^\bullet)^{\frac{1}{4}}\sqrt{\tau}} \right)^2 \mathcal{B}(\tau, \nu(V_{\tau, \varsigma})),$$

where the expected risk is with respect to the main sample $(X_\bullet^{(k)})_{k \in \llbracket B \rrbracket}$.

3.5. Discussion. To summarise, for fixed values of $\tau, \varsigma, B, (N_k)_{k \in \llbracket B \rrbracket}$, the bound on the relative risk of $\tilde{\omega}^\#$ becomes arbitrarily close to the oracle bound in the high-dimensional asymptotics $d_1^\bullet \rightarrow \infty$. We stress that this applies for *fixed* sample sizes N_k , provided $N_k \gtrsim \log^4 B$. Consequently, the fully empirical procedure is (with high probability) not worse than the naive estimator up to a risk factor very close to 1 (since the oracle bound \mathcal{B} is always less than 1), and potentially performs much better if there are many

true τ -neighbouring tasks (again, as reflected by the oracle factor). The conclusion still holds true if $\tau, \varsigma, B, (N_k)$ vary with d_1^\bullet ($\tau \rightarrow 0$ and/or $B \rightarrow \infty$ being the most interesting situations) provided $\varsigma \log(B)/\sqrt{d_1^\bullet} = o(\tau)$ holds and as $N_k \gtrsim \log^4 B$ as before.

Beyond the Gaussian setting. The results presented above hold under the Gaussian distributional assumptions (**GS**). However, the required components — specifically, concentration of estimators for distances between two means and for Schatten norms of the covariances — can be extended with appropriate modifications to the bounded (**BS**) and heavy-tailed (**HT**) distributional settings. Detailed results are presented in Supplemental C.4 and show the qualitative robustness of our approach beyond the Gaussian setting.

Beyond the testing approach. The testing approach has two flaws: first, the theoretical necessity to partition the data entails a certain loss of efficiency, such as a reduction by a factor of 1/2 when the data is equally split. This consideration has been disregarded in the preceding discussion, where the oracle risk was restricted to the main sample. Second, the issue of parameter selection of τ and ς persists. As previously elucidated, the oracle relative risk \mathcal{B} exhibits a bias-variance trade-off: the aggregated variance decreases with an increase in the number of τ -neighbours, consequently, with the worst-case relative bias τ . Ideally, parameters should be adaptively chosen to strive for optimal oracle improvement $\min_{\tau \geq 0, \varsigma \geq 1} \mathcal{B}(\tau, \nu(V_{\tau, \varsigma}))$. The next section introduces an alternative approach pursuing this objective. Additionally, Section 5 analyses whether the derived bounds are optimal.

4. A “ Q -AGGREGATION” APPROACH

In this section, we propose an alternative approach for forming the weights of the convex combination estimator (6). The weights are found by direct minimization of an upper confidence bound of the risk $R_1(\boldsymbol{\omega})$, i.e.,

$$\hat{\boldsymbol{\omega}} \in \arg \min_{\boldsymbol{\omega} \in S_B} \left(\hat{L}_1(\boldsymbol{\omega}) + u \hat{Q}_1(\boldsymbol{\omega}) \right), \quad (28)$$

where $\hat{L}_1(\boldsymbol{\omega})$ is an unbiased estimate of the risk. The idea of this scheme bears resemblance to Q -aggregation (Lecué and Rigollet, 2014), because the objective function will be a quadratic function of $\boldsymbol{\omega}$. The objective aims at taking into account all individual distances between the bags, rather than selecting those less than a fixed threshold as in the testing approach. The penalization term $\hat{Q}_1(\boldsymbol{\omega})$ shall be a high probability upper bound on the difference between estimated and true *loss* ($\hat{L}_1(\boldsymbol{\omega}) - L_1(\boldsymbol{\omega})$). Observe that the penalization term also depends on the weight vector, since giving more weight to tasks that are further away from the target (large $\|\Delta_k\|$) will result in a larger variability of the risk estimate $\hat{L}_1(\boldsymbol{\omega})$. The parameter u is a calibration constant. Compared to the testing approach, one advantage is that it is not necessary to choose the parameters τ and ς . Furthermore no sample splitting is needed. On the other hand, the procedure

is more computationally demanding since there is no closed form solution to (28). Instead, a solution $\widehat{\boldsymbol{\omega}}$ can be obtained by exponentiated gradient descent on the simplex (Kivinen and Warmuth, 1997).

We present specific choices for $\widehat{L}_1(\boldsymbol{\omega})$, $\widehat{Q}_1(\boldsymbol{\omega})$ and an analysis of the relative risk of the resulting Q -aggregation estimator for **(GS)** in Section 4.1 and for **(BS)** thereafter. In contrast to Lecué and Rigollet (2014), we focus on the effect of the dimension rather than that of the sample size which provides a novel analysis.

4.1. Gaussian setting. Under assumption **(GS)** we propose to use the following estimates to form the Q -aggregation estimator:

$$\widehat{L}_1(\boldsymbol{\omega}) = \left\| \sum_{k=2}^B \omega_k (\widehat{\boldsymbol{\mu}}_i^{\text{NE}} - \widehat{\boldsymbol{\mu}}_1^{\text{NE}}) \right\|^2 + (2\omega_1 - 1) \widehat{s}_1^2, \quad (29)$$

$$\widehat{s}_1^2 := \frac{1}{N_1(N_1 - 1)} \sum_{i=1}^{N_1} \left\| X_i^{(1)} - \widehat{\boldsymbol{\mu}}_1^{\text{NE}} \right\|^2, \quad (30)$$

$$\widehat{Q}_1(\boldsymbol{\omega}) := \sum_{k=2}^B \omega_k \sqrt{\frac{\widehat{q}_k}{N_1}}, \quad \text{where} \quad \widehat{q}_k := \frac{1}{N_1 - 1} \sum_{i=1}^{N_1} \left\langle \widehat{\boldsymbol{\mu}}_1^{\text{NE}} - \widehat{\boldsymbol{\mu}}_k^{\text{NE}}, X_i^{(1)} - \widehat{\boldsymbol{\mu}}_1^{\text{NE}} \right\rangle^2. \quad (31)$$

It can be checked easily that \widehat{s}_1^2 is an unbiased estimator of the naive risk s_1^2 , and that the estimator $\widehat{L}_1(\boldsymbol{\omega})$ is an unbiased estimate of the *conditional* risk $\mathbb{E}[\widehat{L}_1(\boldsymbol{\omega}) - L_1(\boldsymbol{\omega}) | X_{\bullet}^{(k)}, k \geq 2] = 0$. With these choices we establish the following result for the average risk of the Q -aggregation estimator:

Proposition 5. *Assume **(GS)** holds, and let $u_0 \in \mathbb{R}_+$ be fixed such that $\log(17B) \leq u_0 \leq (N_1 - 1)/2$. With $\widehat{L}_1(\boldsymbol{\omega})$ and $\widehat{Q}_1(\boldsymbol{\omega})$ as defined in (29),(31), let*

$$\widehat{\boldsymbol{\omega}} \in \arg \min_{\boldsymbol{\omega} \in \mathcal{S}_B} \left(\widehat{L}_1(\boldsymbol{\omega}) + 16\sqrt{u_0} \widehat{Q}_1(\boldsymbol{\omega}) \right). \quad (32)$$

Then it holds:

$$\frac{R_1(\widehat{\boldsymbol{\omega}})}{s_1^2} \leq \frac{1}{s_1^2} \min_{\boldsymbol{\omega} \in \mathcal{S}_B} \left[R_1(\boldsymbol{\omega})(1 + CB e^{-u_0/2}) + CQ_1(\boldsymbol{\omega})\sqrt{u_0} \right] + C \frac{u_0}{\sqrt{d_1^\bullet}}, \quad (33)$$

where $C > 0$ is an absolute constant, and (recalling $\Delta_k = \mu_k - \mu_1$)

$$Q_1(\boldsymbol{\omega}) := \sum_{k=2}^B \omega_k \sqrt{\frac{q_k}{N_1}}, \quad \text{with} \quad q_k := \Delta_k^T \Sigma_1 \Delta_k + \frac{\text{Tr} \Sigma_1 \Sigma_k}{N_k}. \quad (34)$$

The above bound (33) has the form of an ‘‘oracle inequality’’ relating the relative risk of the Q -aggregation approach to the minimum of the attainable relative risk of any aggregation estimator with fixed weight $\boldsymbol{\omega}$ but with a penalization term $Q_1(\boldsymbol{\omega})$. The extra additive term (outside the minimum) vanishes in high effective dimension, but indicates that the relative risk bound cannot be better than $O(\log B / \sqrt{d_1^\bullet})$. We also emphasise the requirement $\log B \lesssim N_1$ implicit in the condition on the calibration

parameter u_0 . The effect of the penalization term $Q_1(\boldsymbol{\omega})$ on the oracle bound (33) might appear obscure: depending on the weights $\boldsymbol{\omega}$, the penalization might outweigh the main risk term $R_1(\boldsymbol{\omega})$. It is noteworthy that this term penalises tasks with distant means (term $\Delta_k^T \Sigma_1 \Delta_k$) or with high variance (term $\text{Tr} \Sigma_1 \Sigma_k / N_k$). To provide further clarification, we present the following corollary which bounds the relative risk of the Q -aggregation method in terms of the relative risk of the oracle testing approach $\mathcal{B}(\tau, \nu)$:

Corollary 2. *Assume (GS) holds. Let $u_0 \in \mathbb{R}_+$ be fixed, such that $\log 17B \leq u_0 \leq (N_1 - 1)/2$, and $\widehat{\boldsymbol{\omega}}$ as defined in (32). Then it holds:*

$$\frac{R_1(\widehat{\boldsymbol{\omega}})}{s_1^2} \leq (1 + CB e^{-u_0/2}) \inf_{\substack{\tau \geq 0 \\ \varsigma \geq 1}} \left[\mathcal{B}(\tau, \nu(V_{\tau, \varsigma})) + C\varsigma \sqrt{\frac{u_0}{d_1^e}} \right]. \quad (35)$$

where $C > 0$ is an absolute constant, $\mathcal{B}(\cdot, \cdot)$, $\nu(\cdot)$ are as defined in (10), (9) and $V_{\tau, \varsigma}$ as in (21)-(22).

As a simple illustration, assume the tasks satisfy (ECSS) and have equal means ($\mu_k = \mu_1$ for $k \in \llbracket B \rrbracket$), but the estimator does not have this information. The oracle merges all tasks and has relative risk $\inf_{\tau, \varsigma} \mathcal{B}(\tau, \nu(V_{\tau, \varsigma})) = B^{-1}$ for $\tau \rightarrow 0, \varsigma = 1$. For $u_0 = \log 17B$, the relative risk of the Q -aggregation method (35) becomes

$$\frac{R_1(\widehat{\boldsymbol{\omega}})}{s_1^2} \leq C \max \left\{ \frac{1}{B}, \sqrt{\frac{\log B}{d_1^e}} \right\},$$

where $C \approx 1$ if d_1^e and B are large. We observe again a blessing of dimensionality; the best improvement is obtained when d_1^e is high ($d_1^e \geq B^2 \log B$ ensures a relative risk bound of order $1/B$, which is the best improvement even if the information of equal means had been known).

4.2. Comparison with the testing approach. Let us compare the bounds obtained for the test method (Proposition 4) to that for the Q -aggregation approach (Corollary 2), in high-dimensional asymptotics $d_1^\bullet, d_1^e \rightarrow \infty$. We start with an analysis of the conditions on the other parameters $\{\tau, \varsigma, B, (N_k)_{k \in \llbracket B \rrbracket}\}$ under which the obtained bounds guarantee that the relative risk of either method is bounded by the oracle bound $\mathcal{B}(\tau, \nu(V_{\tau, \varsigma}))$ up to a factor asymptotically converging to 1, a property which we call ‘‘oracle-consistency’’ for short.

Recall from Section 3.5 that the relative risk of the test method is oracle-consistent (as $d_1^\bullet \rightarrow \infty$), provided $\varsigma \log(B) / \sqrt{d_1^\bullet} = o(\tau)$ and $N_k \gtrsim \log^4 B$ hold. Aside from these conditions the parameters $\tau, \varsigma, B, (N_k)$ can vary with d_1^\bullet . On the other hand, (35) shows that the Q -aggregation method is oracle-consistent (as $d_1^e \rightarrow \infty$) with respect to *any* (τ, ς) provided that $N_1 \gtrsim \log(B d_1^e)$, and $\varsigma \sqrt{\log(B d_1^e) / d_1^e} = o(\tau)$ (taking $u_0 = 2 \log B d_1^e$). The additive terms in (35) are then negligible compared to $\mathcal{B}(\tau, \nu)$, due to $\mathcal{B}(\tau, \cdot) \geq \tau / (1 + \tau)$. Note also that it does not require any condition on N_k for $k \neq 1$.

If d_1^\bullet and d_1^e are of the same order (e.g. in the isotropic setting), the above parameter conditions for consistency of either method are very similar with only minor differences.

One such difference is that the test method is guaranteed to be oracle-consistent even if $B, \tau, \varsigma, (N_k)$ are fixed, i.e., must not change as $d_1^\bullet \rightarrow \infty$; while we require $N_1 \rightarrow \infty$ (though only at a logarithmic rate in B, d_1^\bullet) to warrant oracle consistency of the aggregation estimator. If d_1^e is of order $\sqrt{d_1^\bullet}$ (for example, for a slow power decrease of the eigenvalues λ_i of the covariance, $\lambda_i = i^{-\alpha}$ for $1 \leq i \leq d$ and $\alpha \in (1/2, 1)$), then the oracle consistency conditions for the Q -aggregation method are narrower.

Still, one has to keep in mind that oracle-consistency for the testing approach only holds for the specific parameters (τ, ς) that must be provided by the user, while the Q -aggregation method is oracle consistent with respect to any choice (τ, ς) satisfying the delineated conditions. In other words, the relative risk of the Q -aggregation method qualitatively enjoys the same asymptotic guarantees as the testing approach with *optimally selected* τ and ς subject to the above conditions. This and the fact that the Q -aggregation does not use data splitting is a strong argument in its favour. On the other hand, the testing method has the advantage of being more flexible and easily adapts to non-Gaussian distributions, e.g., bounded or heavy-tailed distributions (see Supplemental C.4). With a modification of the penalization term, the Q -aggregation method can also be applied to bounded distributions, as shown next, but it currently does not accommodate heavy-tailed data distributions.

4.3. Bounded setting. Our results for the Q -aggregation estimator can be extended to the bounded setting (**BS**) where the data lie in a ball of radius M centred in 0. A precise value for M is often known. For example, if the data lies in a reproducing kernel Hilbert space associated with a bounded kernel, M^2 will be the bound on the kernel. The methodology we propose for (**BS**) closely resembles the one outlined for the Gaussian setting. It utilises the same estimates, (29)-(30)-(31), for the risk estimation and its deviations. In order to compensate the lack of regularity of bounded compared to Gaussian data, an additional penalization term $\widehat{Q}^{\text{BS}}(\omega)$ is introduced, which depends on M .

Proposition 6. *Assume (**BS**). Let $u_0 \in \mathbb{R}_+$ with $2 \log N_1 + \log(B) \leq u_0 \leq N_1$, and*

$$\widehat{\omega} \in \arg \min_{\omega \in \mathcal{S}_B} \left(\widehat{L}_1(\omega) + 4\sqrt{2u_0} \widehat{Q}_1(\omega) + C_0 u_0 \widehat{Q}_1^{\text{BS}}(\omega) \right), \quad (36)$$

where $\widehat{L}_1, \widehat{Q}_1$ are defined in (29), (31) resp., $C_0 > 1424$ works, and

$$\widehat{Q}_1^{\text{BS}}(\omega) = \frac{M}{N_1} \sum_{i=2}^B \omega_i \|\widehat{\mu}_i^{NE} - \widehat{\mu}_1^{NE}\|. \quad (37)$$

Assume $N_k \geq (d_k^\bullet)^\beta$ for some $\beta > 0$ and all $k \in \llbracket B \rrbracket$, then:

$$\frac{R_1(\widehat{\omega})}{s_1^2} \leq \min_{\tau > 0, \varsigma \geq 1} (\mathcal{B}(\tau, \nu(V_{\tau, \varsigma})) + C \varsigma \varepsilon) + C \phi_1 \varepsilon, \quad \varepsilon := \max \left\{ \sqrt{\frac{u_0}{d_1^e}}, \frac{u_0}{(d_1^\bullet)^{\beta/2}} \right\}, \quad (38)$$

where $\mathcal{B}(\cdot, \cdot)$, $\nu(\cdot)$ are as defined in (10), (9), $V_{\tau, \varsigma}$ as in (21)-(22), C an absolute constant, and $\phi_1 := M^2 / \text{Tr} \Sigma_1$.

The quantity β reflects the trade-off between the requirement on the number of samples and the rate of convergence to the oracle bound. A bound similar to that in the Gaussian case will only be obtained if a stricter condition on the bag sizes is met ($N_k \gtrsim d_k^\bullet$ instead of $N_1 \gtrsim \log B$ as in Corollary 2). In contrast to (35), there is no multiplicative constant in front of the bound, however, the additive term now involves the quantity ϕ_1 (see Supplemental F for a discussion of this quantity in the framework of kernel mean embedding (KME) estimation with a bounded kernel, which is our primary motivation for analyzing the bounded setting).

5. MINIMAX RESULTS

This section explores if the oracle relative risk upper bound $\mathcal{B}(\tau, \nu(V_\tau))$ as defined in (12), which has been utilised as benchmark in previous sections, is optimal in a minimax sense. As before, we will first examine the estimation of a single mean. Subsequently, we extend the analysis to the compound relative risks averaged over tasks.

Our aim is to establish minimax bounds matching the upper bounds over distribution classes that are as restrictive as possible. Since a minimax lower bound on a distribution class also applies to every superclass containing it, bounds on restrictive classes are more insightful. To achieve this, we narrow down the distribution classes by fixing as many parameters as possible to arbitrary values. As employed throughout this manuscript, we will adopt a high-dimensional asymptotics viewpoint and focus on minimax statements as the effective dimension grows large.

5.1. Single task relative risk. We derive a lower minimax bound for a class of distributions that closely match the assumptions proposed to introduce the oracle bound (12): a known subset of τ -neighbours V in arbitrary position, all other parameters (sample sizes, covariances, ...) being fixed. We additionally assume that all task covariance matrices are proportional to each other ("aligned"), which appears to be the least favourable setting.

Definition 4. Let $\tau \in \mathbb{R}_+$; $B, V \in \mathbb{N}_{>0}$ with $B \geq V$, $\mathbf{s}^2 = (s_1^2, \dots, s_B^2) \in \mathbb{R}_+^B$, $(N_k)_{k \in \llbracket B \rrbracket} \in \mathbb{N}_{>0}^B$, and Σ a symmetric positive definite matrix of size d with $\text{Tr } \Sigma = 1$ be fixed. We denote by $\mathcal{P}_{\text{single}}(\tau, V, \Sigma, \mathbf{s}^2)$ the set of joint distributions for tasks following model (1) such that:

- (i) The total number of bags is B and the number of samples per bag is given by $(N_k)_{k \in \llbracket B \rrbracket}$. (Omitted from the distribution class notation for simplicity.)
- (ii) **(GS)** holds.
- (iii) The task covariances are given by $\Sigma_k = N_k s_k^2 \Sigma$ (i.e., all tasks have covariances proportional to Σ and the naive risks are specified by the vector \mathbf{s}^2).
- (iv) The mean vectors $(\mu_k)_{k \in \llbracket B \rrbracket}$ can vary freely subject to:

$$\|\mu_1 - \mu_k\|^2 \leq \tau s_1^2, k \in \llbracket V \rrbracket.$$

A minimax lower bound, as by Proposition 7 below, over that model holds over any larger model; for instance, the model where Σ_1 is arbitrarily fixed and the other covariances may vary freely provided that the naive risks still match the prescribed \mathbf{s}^2 .

Proposition 7. *It holds*

$$\inf_{\hat{\mu}_1} \sup_{\mathbb{Q} \in \mathcal{P}_{\text{single}}(\tau, V, \Sigma, \mathbf{s}^2)} \frac{R_1(\mathbb{Q}, \hat{\mu}_1)}{s_1^2} \geq \mathcal{B}(\tau, \nu(\llbracket V \rrbracket)) - \varepsilon(d^e(\Sigma)),$$

where \mathcal{B} is defined in (12), ν in (9), the infimum is over all estimators $\hat{\mu}_1$ for μ_1 , and $R_1(\mathbb{Q}, \hat{\mu}_1)$ indicates its risk (2) under distribution \mathbb{Q} . The function $\varepsilon(t)$ is independent of any parameters and satisfies $\varepsilon(t) = O((\log t)/t)$ as $t \rightarrow \infty$.

This minimax lower bound can be compared with the upper bounds obtained for the testing and Q -aggregation methods, Proposition 4 and Corollary 2, resp. In the case of (ECSS) (so that $V_{\tau, \varsigma} = V_\tau$ for any $\varsigma \geq 1$ and we can ignore the role of ς), the lower and upper bounds match. This shows that the oracle relative risk $\mathcal{B}(\tau, \nu(V_\tau))$ is indeed minimax in the sense of high-dimensional asymptotics, provided that $\log(B) = o(d_1^e)$. Furthermore, the Q -aggregation method is *asymptotically minimax adaptive* over the parameter $\tau > 0$. This can be seen as a generalization of classical results on the James-Stein estimator (see Supplemental B). Observe also that for the upper and lower bounds the dimension-dependent remainder terms do not depend on other parameters, which makes the dimensional asymptotics uniform with respect to those parameters.

If (ECSS) does not hold, there can be a discrepancy between the minimax lower bound and the obtained upper bounds due to the exclusion of high variance tasks in the latter (V_τ against $V_{\tau, \varsigma}$). An unfavourable regime illustrating this gap is the following: suppose there are many tasks that are τ -neighbours of the target (τ being fixed independently of the dimension but $V \approx d_1^\bullet$) with significantly higher variances though ($\varsigma = s_k^2/s_1^2 \approx V^{1/2}$ for all $2 \leq k \leq V$). In that scenario, the upper bounds of Proposition 4 and Corollary 2 do not guarantee convergence to $\mathcal{B}(\tau, \nu(V_\tau)) \approx \tau/(1+\tau)$, since the remainder terms $\varsigma/\sqrt{d_1^\bullet}$ (resp. $\varsigma/\sqrt{d_1^e}$) do not converge to zero for high-dimensional asymptotics. This gap can amount to an arbitrary large factor since τ can be arbitrarily small. However, the scenario where a target task is surrounded by numerous neighbours with significantly higher variance can only arise for a small proportion of the tasks. This implies that this concern is alleviated when evaluating the relative risk averaged across all tasks, as shown next.

5.2. Compound relative risk. We define the compound relative risk as the relative risk averaged over all tasks. As we only studied upper bounds for a single task so far, we first derive new upper bounds for the compound relative risk. We then proceed to derive minimax bounds on restrictive distribution classes under which the task means exhibit a certain clustering or covering structure.

Definition 5. Let $\boldsymbol{\mu} = (\mu_k)_{k \in \llbracket B \rrbracket}$ be a collection of vectors of \mathbb{R}^d , $J \in \mathbb{N}_{>0}$, and \mathcal{C} a J -partition of $\llbracket B \rrbracket$ (i.e., $\mathcal{C} = (\mathcal{C}_j)_{j \in \llbracket J \rrbracket}$ with $\mathcal{C}_1 \sqcup \dots \sqcup \mathcal{C}_J = \llbracket B \rrbracket$). The diameters of the

partition \mathbf{C} applied to $\boldsymbol{\mu}$ are defined as:

$$\text{diam}(\mathbf{C}, \boldsymbol{\mu}) = \left(\max_{k, \ell \in \mathcal{C}_j} \|\mu_k - \mu_\ell\| \right)_{j \in \llbracket J \rrbracket} \in \mathbb{R}_+^J. \quad (39)$$

We shall refer to parts as ‘‘groups’’ rather than clusters, because the partitioning can in principle be arbitrary. However, the intuition is that the set of vectors $\boldsymbol{\mu}$ exhibits more structure if it can be partitioned into a limited number of groups with small diameter. For instance, if it is strongly clustered, or supported on a set of small metric entropy such as a low-dimensional manifold. The compound relative risk of the Q -aggregation approach can then be upper bounded as follows:

Proposition 8. *Assume (GS) holds, and let $u_0 \in \mathbb{R}_+$ such that $\log 17B \leq u_0 \leq (\min_k N_k - 1)/2$. For $k \in \llbracket B \rrbracket$, define $\widehat{L}_k(\boldsymbol{\omega}), \widehat{Q}_k(\boldsymbol{\omega})$ analogously to (29),(31) and*

$$\widehat{\boldsymbol{\omega}}_k \in \arg \min_{\boldsymbol{\omega} \in \mathcal{S}_B} \left(\widehat{L}_k(\boldsymbol{\omega}) + 16\sqrt{u_0} \widehat{Q}_k(\boldsymbol{\omega}) \right). \quad (40)$$

Then it holds:

$$\frac{1}{B} \sum_{k=1}^B \frac{R_k(\widehat{\boldsymbol{\omega}}_k)}{s_k^2} \leq (1 + CB e^{-u_0/2}) \min_{\mathbf{C}} \left(\mathcal{L}^*(\mathbf{s}, \mathbf{C}, \text{diam}(\mathbf{C}, \boldsymbol{\mu})) + C \frac{u_0}{\min_{k \in \llbracket B \rrbracket} (d_k^e)^{1/2}} \right), \quad (41)$$

where the minimum is taken over all partitions \mathbf{C} of $\llbracket B \rrbracket$, C is an absolute constant, and for $\boldsymbol{\zeta} \in \mathbb{R}_+^J$:

$$\mathcal{L}^*(\mathbf{s}, \mathbf{C}, \boldsymbol{\zeta}) := \frac{1}{B} \sum_{j=1}^J \sum_{k \in \mathcal{C}_j} \mathcal{B}(\tau_{j,k}, \nu_{j,k}), \quad \tau_{j,k} := \frac{\zeta_j^2}{s_k^2}, \quad \nu_{j,k} := \frac{s^2(\mathcal{C}_j)}{s_k^2}, \quad (42)$$

and \mathcal{B} is defined in (12).

Similarly to the estimation of a single mean, the bound on the compound relative risk depends on the maximum distance between tasks of the same group relative to the naive risk of each task, and on the relative aggregated variances (9) in each group. Remarkably, the compound relative risk bound does not involve any ‘‘whittling down’’ of high-variance tasks as in the single task bound (22), and holds under arbitrary inhomogeneity of the tasks and sample sizes.

The quantity \mathcal{L}^* equates to an oracle compound relative risk and is minimax under high-dimensional asymptotics. To show this, we extend the single task model 6 to a joint distribution class such that the tasks are divided into inhomogeneous groups.

Definition 6. Let $B \in \mathbb{N}_{>0}$, $\mathbf{s}^2 = (s_1^2, \dots, s_B^2) \in \mathbb{R}_+^B$, $(N_k)_{k \in \llbracket B \rrbracket} \in \mathbb{N}_{>0}^B$, and Σ a symmetric positive definite matrix of size d with $\text{Tr} \Sigma = 1$ be fixed.

Let $J \in \mathbb{N}_{>0}$, \mathbf{C} be a J -partition of $\llbracket B \rrbracket$ and $\boldsymbol{\zeta} \in \mathbb{R}_+^J$. We define $\mathcal{P}_{\text{mult}}(\mathbf{C}, \boldsymbol{\zeta}, \Sigma, \mathbf{s})$ as the set of tasks according to model (1) with:

(i)-(iii) as in Definition 4;

(iv) The mean vectors $\boldsymbol{\mu} = (\mu_k)_{k \in [B]}$ can vary freely subject to

$$\boldsymbol{\mu} \in \{ \boldsymbol{\mu} \in \mathbb{R}^{d \times B} : \text{diam}(\mathbf{C}, \boldsymbol{\mu}) \leq \boldsymbol{\zeta} \text{ (coordinate-wise inequality)} \}.$$

In words, $\mathcal{P}_{\text{mult}}(\mathbf{C}, \boldsymbol{\zeta}, \Sigma, \mathbf{s})$ is the set of Gaussian tasks with fixed, aligned covariances and naive risks prescribed by the vector \mathbf{s} , such that the groups of mean vectors given by partition \mathbf{C} have diameters bounded by the respective entries of vector $\boldsymbol{\zeta}$.

Proposition 9. *Let $\mathbf{s} \in \mathbb{R}_+^B$, $J \in \mathbb{N}_{>0}$, \mathbf{C} a J -partition of $[B]$ and $\boldsymbol{\zeta} \in \mathbb{R}_+^J$ be fixed. It holds*

$$\lim_{d^e \rightarrow \infty} \sup_{\substack{\Sigma: \\ d^e(\Sigma) = d^e}} \inf_{\hat{\boldsymbol{\mu}}} \sup_{\mathbb{Q} \in \mathcal{P}_{\text{mult}}(\mathbf{C}, \boldsymbol{\zeta}, \Sigma, \mathbf{s})} \frac{1}{B} \sum_{k=1}^B \frac{R_k(\mathbb{Q}, \hat{\boldsymbol{\mu}}_k)}{s_k^2} \geq \mathcal{L}^*(\mathbf{s}, \mathbf{C}, \boldsymbol{\zeta}/2), \quad (43)$$

where the infimum is over all joint estimators $\hat{\boldsymbol{\mu}} = (\hat{\boldsymbol{\mu}}_1, \dots, \hat{\boldsymbol{\mu}}_B)$.

In particular, since it holds $\mathcal{L}^*(\mathbf{s}, \mathbf{C}, \boldsymbol{\zeta}/2) \geq \mathcal{L}^*(\mathbf{s}, \mathbf{C}, \boldsymbol{\zeta})/4$, the upper bound matches the lower minimax bound up to a fixed constant factor in a dimensional asymptotics sense (by choosing $u_0 = \log 17B$ and provided that $\log B / (\min_k (d_k^e)^{-1/2}) = o(\mathcal{L}^*)$). Moreover, (41) shows that the Q -aggregation estimator is (up to that constant factor) asymptotically minimax adaptive with respect to the choice of grouping \mathbf{C} of the task means, the corresponding group diameters, and the bag variances.

As in the single task case, the minimax bound \mathcal{L}^* only depends on the bag sizes through the naive risks \mathbf{s} : bags with large variance and many samples are statistically equivalent to bags with low variance and few samples. Similarly, the improvement only depends on the relative aggregated variance of each group, not on the number of bags. Proposition 10 gives an interpretable upper bound for \mathcal{L}^* :

Proposition 10. *Let $\mathbf{s} \in \mathbb{R}_+^B$, $J \in \mathbb{N}_{>0}$, \mathbf{C} a J -partition of $[B]$ and $\boldsymbol{\zeta} \in \mathbb{R}_+^J$, it holds:*

$$\mathcal{L}^*(\mathbf{s}, \mathbf{C}, \boldsymbol{\zeta}) \leq \sum_{j=1}^J \frac{|\mathcal{C}_j|}{B} \cdot \frac{\bar{\tau}_j + |\mathcal{C}_j|^{-1}}{\bar{\tau}_j + 1}, \quad \bar{\tau}_j := \frac{\zeta_j^2}{\bar{s}^2(\mathcal{C}_j)}, \quad \bar{s}^2(\mathcal{C}_j) := \left(\frac{1}{|\mathcal{C}_j|} \sum_{k \in \mathcal{C}_j} s_k^{-2} \right)^{-1}, \quad (44)$$

implying in particular:

$$\mathcal{L}^*(\mathbf{s}, \mathbf{C}, \boldsymbol{\zeta}) \leq \min \left(1, \frac{\bar{\tau}_*}{1 + \bar{\tau}_*} + \frac{J}{B} \right), \quad \bar{\tau}_* := \sum_{j=1}^J \frac{|\mathcal{C}_j|}{B} \bar{\tau}_j. \quad (45)$$

If all risks and diameters are equal, $s_k^2 = s^2$ and $\zeta_j^2 = \zeta^2$ for all $k \in [B]$ and $j \in [J]$, then the bound of (45) is sharp up to a factor at most 2.7.

Bound (45) elucidates that the compound oracle relative risk \mathcal{L}^* is small when (i) there are few groups relative to the number of bags (i.e., J/B small); and (ii) groups have on average a small squared diameter relative to the harmonic mean of the naive risks of its constituent tasks.

Eq. (41) implies that the compound risk is upper bounded by \mathcal{L}^* for any valid partitioning. As an illustrative example we consider the (**ECSS**) setting and \mathbf{C} as a $\sqrt{\tau}s$ -covering of $\boldsymbol{\mu}$ for a given τ . Then $\bar{\tau}_* = \tau$ and the number of groups J is the covering number $N(\boldsymbol{\mu}, \sqrt{\tau}s)$. This highlights that the Q -aggregation strategy will be very effective to reduce the compound risk if the set of true means can be covered by a relatively small number of balls, in comparison to the total number of tasks, with a radius significantly smaller than the standard deviation of the naive estimates.

This bound takes a form akin to the findings presented in Marienwald et al. (2021), who examined the (**ECSS**) setting only and used a testing strategy comparable to that of the previous section. The parameter of their (and our) test, though, has to be fixed by the user. In contrast, the Q -aggregation approach attains the oracle trade-off between the ‘‘bias’’ term $\tau/(1+\tau)$ and the ‘‘variance’’ term $N(\boldsymbol{\mu}, \sqrt{\tau}s)/B$ without the need to specify τ .

Finally, observe that the first term $\tau/(1+\tau)$ resembles the best potential improvement and is reminiscent of the oracle improvement factor of the James-Stein estimator, which can be conceived as a special case; see Supplemental B for additional details.

6. APPLICATION: ESTIMATION OF MULTIPLE KERNEL MEAN EMBEDDINGS

We emphasise that our discussion and theoretical results include the case when \mathcal{X} is a reproducing kernel Hilbert space (RKHS), in which case the mean corresponds to a kernel mean embedding (KME) (Muandet, Fukumizu, Sriperumbudur, and Schölkopf, 2017; Smola et al., 2007). Let \mathcal{Z} be a measurable space enriched with a reproducing kernel $\kappa : \mathcal{Z} \times \mathcal{Z} \rightarrow \mathbb{R}$ and its corresponding RKHS \mathcal{H} . The kernel mean embedding $\mu_{\mathbb{P}_Z} \in \mathcal{H}$ of distribution \mathbb{P}_Z on \mathcal{Z} and its empirical (naive) estimation $\hat{\mu}_{\mathbb{P}_Z}$, which is based on the samples $(Z_n)_{1 \leq n \leq N_Z} \sim \mathbb{P}_Z$, are defined as

$$\mu_{\mathbb{P}_Z} = \int_{\mathcal{Z}} \kappa(z, \cdot) d\mathbb{P}_Z(z) \quad , \quad \hat{\mu}_{\mathbb{P}_Z} = \frac{1}{N_Z} \sum_{n=1}^{N_Z} \kappa(Z_n, \cdot). \quad (46)$$

The estimation of multiple KMEs is an instance of model (1) once we identify $\mathcal{X} = \mathcal{H}$ and $X_k^{(i)} = \kappa(Z_k^{(i)}, \cdot)$ for a bounded reproducing kernel κ ; this allows a direct application of our theoretical results for the bounded setting.

For characteristic kernels the map from \mathbb{P} to $\mu_{\mathbb{P}}$ is injective and contains information about all moments of \mathbb{P} , so that $\mu_{\mathbb{P}}$ provides a unique representation of \mathbb{P} . Thus, KMEs can naturally be used to define a metric on probability distributions. Let \mathbb{P}, \mathbb{Q} denote distributions and their KMEs $\mu_{\mathbb{P}}, \mu_{\mathbb{Q}}$ respectively. The maximum mean discrepancy (MMD) expresses the distance between $\mu_{\mathbb{P}}$ and $\mu_{\mathbb{Q}}$ in \mathcal{H}

$$\begin{aligned} \text{MMD}^2(\mu_{\mathbb{P}}, \mu_{\mathbb{Q}}) &= \|\mu_{\mathbb{P}} - \mu_{\mathbb{Q}}\|_{\mathcal{H}}^2 \quad , \\ \widehat{\text{MMD}}^2(\mu_{\mathbb{P}}, \mu_{\mathbb{Q}}) &= \sum_{n \neq n'=1}^N \frac{\kappa(Z_n, Z_{n'})}{N(N-1)} - 2 \sum_{n=1}^N \sum_{m=1}^M \frac{\kappa(Z_n, Y_m)}{NM} + \sum_{m \neq m'}^M \frac{\kappa(Y_m, Y_{m'})}{M(M-1)}, \end{aligned}$$

where $\widehat{\text{MMD}}^2$ denotes an unbiased estimate based on the samples $(Z_n)_{1 \leq n \leq N} \sim \mathbb{P}$ and $(Y_m)_{1 \leq m \leq M} \sim \mathbb{Q}$. For characteristic kernels it holds that $\text{MMD}^2(\mu_{\mathbb{P}}, \mu_{\mathbb{Q}}) = 0$ iff $\mathbb{P} = \mathbb{Q}$ (Gretton, Borgwardt, et al., 2012), which enables a large range of possible applications.

6.1. Motivation and Related Work. KMEs are employed for a variety of statistical tests, e.g., two-sample tests (Gretton, Borgwardt, et al., 2012), goodness-of-fit tests (Chwialkowski et al., 2016), and tests on statistical independence based on the Hilbert Schmidt independence criterion (Gretton, Fukumizu, et al., 2007). It also finds application in machine learning, e.g., for unsupervised (Jegelka et al., 2009) or supervised distributional learning (Muandet, Fukumizu, Dinuzzo, et al., 2012; Szabó et al., 2016), density estimation (Muandet, Fukumizu, Sriperumbudur, Gretton, et al., 2014), as part of the optimization criterion of the learning (Brehmer and Cranmer, 2020; Fakoor et al., 2020), and so on. Due to the wide variety of kernel functions, kernel mean embeddings can in general be used on various data types and for structured data. See Muandet, Fukumizu, Sriperumbudur, and Schölkopf (2017) for an in-depth overview on KMEs and their applications.

The success of applying the KME or the MMD resp. relies heavily on the ability to accurately estimate the kernel mean based on sample data. The naive empirical estimator (46) was recently superseded by a James-Stein-like estimator (Muandet, Fukumizu, Sriperumbudur, Gretton, et al., 2014). They showed that this estimator is admissible and consistent for a suitable choice of shrinkage. Other single KME estimation strategies were proposed since then, e.g., non-linear shrinkage (Muandet, Sriperumbudur, et al., 2016), an empirical Bayesian approach (Filippi et al., 2016), and more robust estimations based on marginalised corrupted data (Xia et al., 2022), or a MOM approach (Lerasle et al., 2019). To the best of our knowledge, there is no prior work on the improved estimation of *multiple* kernel mean embeddings except for Marienwald et al. (2021).

6.2. Description of the Experiments. We evaluate the estimation of multiple kernel mean embeddings on artificial and real-world data.

6.2.1. Methods. We only sketch the best performing methods here. A complete list and detailed description of the tested methods can be found in Supplemental G, where we also provide pseudocode that demonstrates how the methods can be implemented in practice. More specifically, we found that methods based on Q-aggregation benefit from restricting the support of the weights from $\omega \in \mathcal{S}_B$ to $\omega \in \mathcal{S}_V$. However, the neighbouring test merely functions as a safeguard here with a much larger value for τ (cf. (8)) than that used for the testing approaches. We referred to methods, that are based on the testing procedure which finds neighbours for the construction of the convex combination as in Cor. 1, as similarity test-based (STB). The approaches differ in their weighting schemes for these neighbours. `STB opt` calculates the oracle weights (13) where the aggregated variances are replaced by their empirical estimations. `STB orth` performs constrained risk minimization and posits an orthogonality assumption,

$\langle \widehat{\mu}_j^{\text{NE}} - \widehat{\mu}_i^{\text{NE}}, \widehat{\mu}_{j'}^{\text{NE}} - \widehat{\mu}_i^{\text{NE}} \rangle = 0$ for all $j \neq j'$, which might be unrealistic in practice but yields a closed-form solution for the weights. Finally, `STB egd` minimises the Q-aggregation objective (36) and applies exponentiated gradient descent on the simplex (Kivinen and Warmuth, 1997) to approximate the solution.

We compare their performances to the naive estimation (NE), and we modify the multitask-averaging approach from Feldman et al. (2014) (`MTA const`) so that it is applicable to the estimation of KMEs. It assumes a constant similarity across tasks. In Supplemental G.4, we further report the results of our previously proposed approach, `STB weight` (Marienwald et al., 2021), which was not designed to handle inhomogenous data, and the regularised kernel mean shrinkage estimator `R-KMSE`, proposed in Muandet, Sriperumbudur, et al. (2016), that shrinks the estimation towards the origin and is performed separately on each bag. In Supplemental G.6 we discuss the computational complexity of all approaches.

The considered methods have data-dependent model parameters whose optimal values might be found by cross-validation. We also provide default parameter choices that we observed to perform well in most situations (see Supplemental G.5).

6.2.2. Experimental Metric. In the kernel case, the true KME μ is unknown even for synthesised data. We use a (naive) estimation based on an independent sample of the same distribution as approximation. Because this proxy is computed on a very large sample, it can be assumed to have low risk and to be more accurate than the estimation performed by any method on much smaller bags. The squared MMD between the (proxy) true KME μ_i of bag $i \in \llbracket B \rrbracket$ and its estimation $\widehat{\mu}_i^m$, of form (6), performed by method `m` with weights ω_i^m is then used as error measure

$$\begin{aligned} \widehat{\text{MMD}}^2(\mu_i, \widehat{\mu}_i^m) &= \sum_{j, j' \in \llbracket B \rrbracket} \omega_{ij}^m \omega_{ij'}^m \sum_{n=1}^{N_j} \sum_{n'=1}^{N_{j'}} \frac{\kappa(Z_n^{(j)}, Z_{n'}^{(j')})}{N_j N_{j'}} - \sum_{j \in \llbracket B \rrbracket} 2 \omega_{ij}^m \sum_{n=1}^{N_j} \sum_{m=1}^{M_i} \frac{\kappa(Z_n^{(j)}, Y_m^{(i)})}{N_j M_i} \\ &\quad + \sum_{m \neq m'} \frac{\kappa(Y_m^{(i)}, Y_{m'}^{(i)})}{M_i(M_i - 1)}, \end{aligned} \quad (47)$$

where $Y_i, Z_i \sim \mathbb{P}_i$ independent with $|Y_i| = M_i \gg N_i = |Z_i|$ for all $i \in \llbracket B \rrbracket$, so that Y_i can be used to calculate the proxy and Z_i for the estimation. Each method is validated on the same data to guarantee comparability. This estimation error is averaged over multiple trials $\overline{\text{MMD}}^2(\mu_i, \widehat{\mu}_i^m)$ and its decrease compared to the naive estimation $\widehat{\mu}_i^{\text{NE}}$ is reported for all experiments

$$\left(\left(\overline{\text{MMD}}^2(\mu_i, \widehat{\mu}_i^{\text{NE}}) - \overline{\text{MMD}}^2(\mu_i, \widehat{\mu}_i^m) \right) / \overline{\text{MMD}}^2(\mu_i, \widehat{\mu}_i^{\text{NE}}) \right) \cdot 100 \text{ [\%]}.$$

6.3. Artificial Gaussian Data. The toy data sets are Gaussian distributed in \mathbb{R}^2 with fixed means and randomly rotated covariance matrices. For $i \in \llbracket B \rrbracket$ and $B = 50$

$$Z_{\bullet}^{(i)}, Y_{\bullet}^{(i)} \sim \mathcal{N}\left(m_i, R(\theta_i) \Sigma R(\theta_i)^T\right) = \mathbb{P}_i, \quad \theta_i \sim \mathcal{U}\left(-\frac{\pi}{4}, \frac{\pi}{4}\right),$$




FIGURE 1. Decrease in average quadratic estimation error compared to NE in percent on Gaussian data settings (a) and (b) resp. Higher is better. The hashed histogram bars in (b) show the bag sizes for the bags 1 to 50, which vary between 10 and 300 (right axis).

where the rotation matrix $R(\theta_i)$ rotates the matrix $\Sigma = \text{diag}(1, 10)$ according to angle θ_i . We generate $|Y_{\bullet}^{(i)}| = 1000$ data for the “proxy truth”. A Gaussian RBF kernel, with a kernel width set to the average feature-wise standard deviation of the data, maps the data from the two-dimensional input space to the infinite dimensional RKHS. Two setups are tested:

- (a) Clustered Bags: $N_i = 50$ for all $i \in \llbracket B \rrbracket$. In the input space, each ten bags form a cluster where the cluster centres ($= m_i$) lie equally spaced on a circle. The radius of that circle varies between 0 and 3, which creates different amount of overlap between the clusters.
- (b) Imbalanced Bags: $m_i = 0$ for all $i \in \llbracket B \rrbracket$. The bags $Z_{\bullet}^{(i)}$ are highly imbalanced, i.e. $N_i \in [10, 300]$. Because the tasks only vary in the rotation of their covariance matrices, we know that their KMEs lie on a low dimensional manifold in the RKHS. Because of the different bag sizes, the individual KMEs have different estimation accuracies.

The experiments are repeated for 100 trials; the results of the methods with default parameter choices are shown in Fig. 1. We also report the performances with tuned parameters (optimised on i.i.d. training data) in Supplemental G.5, Figure 7.

All methods provide an improvement over NE, which is most significant for bags with few samples. This was already observed in other multi-task learning problems, e.g., see

Marienwald et al. (2021) or Feldman et al. (2014). The constant similarity assumption of `MTA const` leads to an inadequate estimation for large radii or large bags. Namely, a KME with large bag size is shrunk to the grand empirical mean of all bags even though it includes high-variance (low sample size) or distant bags. This impairs the improvement. This effect is alleviated by the proposed `STB` approaches, that define the shrinkage according to the variance of and the distances between the KMEs. They show high performance for the tested settings. For $0.5 < \text{radius} < 2$, the similarity test might mistake a bag of another cluster for a neighbour due to the strong overlap between the clusters, which explains the slight performance dip. All the proposed methods provide similarly accurate results. Despite its unrealistic orthogonality assumption, `STB orth` performs best on the artificial data.

6.4. Flow Cytometry Data. Flow cytometry is fundamental to biomedical research and clinical practice. It provides a multiparametric, single-cell analysis of a suspension or sample. The flow cytometer analyses the size, shape and internal complexity of cells and can detect the presence and amount of different fluorochromes (which in turn reveal insights about the presence of proteins or structures within the cell). These characteristics might then be used to classify the cells into different populations. Applications are vast, but well-known examples are differential blood count, or immunophenotyping of leukemia or in HIV infections (Adan et al., 2017; McKinnon, 2018).

The data set we use corresponds to the T-cell panel of the Human ImmunoPhenotyping Consortium (Finak et al., 2016). Seven laboratories were asked to perform a flow cytometry analysis of three replicates of blood samples of three patients. All laboratories were asked to follow the same experimental protocol and used the same seven markers to characterise the cells ($d = 7$). Based on the observed characteristics the cells were then classified into ten different populations or cell types. We use this structure (laboratory, replicate, patient, cell type) to divide the data into bags. We excluded bags with less than 1000 data points, which leads to 424 bags in total. Each data point $Z_n^{(i)} \in \mathbb{R}^7$ in a bag i corresponds to one cell. As the number of cells varies, the bags are highly imbalanced. We use a Gaussian RBF kernel with kernel width of 950 to map the cell features to a RKHS. The kernel choice and width are in accordance with Dussap et al. (2023). The (proxy) true KME is approximated by a naive estimation based on $Y_{\bullet}^{(i)}$ with $|Y_{\bullet}^{(i)}| = 1000$ (bags with more samples are capped). The sizes of the bags that are used for the estimation are chosen proportional to the bag sizes of the original input data, $N_i \in [7, 125]$, to mimic a realistic setting. In each one of the 100 trials, a subset of samples $Z_{\bullet}^{(i)}$ with $|Z_{\bullet}^{(i)}| = N_i$ is drawn randomly from $Y_{\bullet}^{(i)}$, on which the methods perform their estimation. We conducted experiments on each cell type separately so that $B \in [43, 62]$, and on all cell types jointly ($B = 424$). Cell types 5, 6 and 10, for which $B < 7$, are excluded for the separate but included in the joint analysis.

The results are depicted in Fig. 2. On average, all methods provide an improvement over NE. For some trials, `MTA const` gives worse estimations than NE (negative




FIGURE 2. Decrease in estimation error compared to NE in percent on the flow cytometry data. Higher is better. The number next to the boxplot quantifies the median, which is also depicted as a line. The mean is visualised as a circle. From left to right: results on individual cell types 1, 2, 3, 4, 7, 8, 9, and all cell types taken jointly.

improvement), see e.g., cell type 1. When all cell types are considered jointly, its performance drops significantly. The **STB** approaches give more accurate estimations than **MTA const** and provide an improvement of $\approx 50\%$ for all cell types. **STB egd** gives the most accurate and stable estimations across the different settings but also has high computational complexity.

In summary, our presented methods provide an improvement over the naive estimation and over other state-of-the-art methods. Although **R-KMSE** or **MTA const** give more accurate estimations than the sample average, the provided improvements vary whether a shrinkage towards a common reference point or the grand mean resp. complies with the underlying data. In contrast, our proposed methods identify inhomogeneous task similarities and are applicable to imbalanced data sets (which, therefore, surpass our previously introduced method **STB weight**). While **STB egd** provides in most cases the highest improvement with least variance, it also requires the most computational complexity. **STB orth** provides a good trade-off.

7. RELATION AND COMPARISON TO PREVIOUS WORK

We review related literature grouped along two axes: the first is rooted in statistics, compound decision rules and the empirical Bayes point of view, and secondly a more

recent one related to multitask learning. We first emphasise again the seminal importance of the James-Stein (JS) estimator (James and Stein, 1961) for a single vector mean, which can be seen as a particular setting of model (1). Historically important is the realization that the sample average $\hat{\mu}_i^{\text{NE}} := \frac{1}{N_i} \sum_{k=1}^{N_i} X_k^{(i)}$, despite being MLE (in the Gaussian model) and BLUE, is inadmissible and dominated by the shrinkage-based JS estimator. Pinsker (1980) should be credited for an early “dimensional asymptotics” point of view, analysing the minimax risk if the mean vector belongs to a ball of \mathbb{R}^d as a by-product of his celebrated minimax analysis of estimators in Sobolev ball models (see, e.g., Nussbaum, 1996 for a discussion). The risk of the JS estimator is asymptotically close to that minimax in the isotropic Gaussian model if $d \rightarrow \infty$, as well as adaptive to the radius of the ball (Beran, 1996); see more details in Supplemental B.

7.1. Empirical Bayes and compound decision point of view. The celebrated series of works by Efron and Morris (1972, 1973, 1976) advocated for an interpretation of the JS estimator as a compound decision problem and an empirical Bayes point of view (Robbins, 1951, 1964; Zhang, 2003): the problem of estimating a single mean vector in \mathbb{R}^B with standard Gaussian noise is better seen as B -many estimations of one-dimensional means observed with independent observation noise (which in model (1) corresponds to $B > 1$ means in dimension $d = 1$). The authors compare the performance of the JS estimator to that of a Bayesian model, i.e., the means are themselves drawn from a centred Gaussian prior. The Bayes rule under the fully Gaussian model (prior and observations) is solely determined by the prior variance, which is usually unknown, hence, called “oracle” in the present discussion. The JS estimator can then be interpreted as being empirically Bayes as it replaces the oracle (prior) variance with an empirically estimated counterpart. The compound risk is shown to converge to the oracle Bayes risk, as B grows.

Efron and Morris (1976) generalised this analysis to the multidimensional case which is an instance of model (1) for arbitrary d and Gaussian task distributions with identical covariances. They proposed a multidimensional version of the JS estimator. Similarly to the one-dimensional case, this is interpreted as an empirical Bayes procedure with a multidimensional Gaussian prior, whose unknown covariance is replaced by an empirically estimated counterpart. If $(d + 2)/B \rightarrow 0$, then the risk of the multidimensional JS estimator approaches that of the oracle Bayes rule.

The nonparametric empirical Bayes estimator developed by Brown and Greenshtein (2009) (see also Jiang and Zhang, 2009 for a closely related, independent work) is in the same line of thought, but considers a completely arbitrary prior on the means (in dimension $d = 1$). In that situation, the oracle Bayes procedure can be expressed in terms of the marginal, nonparametric mixture density of the observations across tasks and of its derivative (to establish this, Gaussian partial integration is used, thus, relying heavily on the assumption of isotropic Gaussian tasks). The proposed estimator replaces the true density with a kernel density estimate (while Jiang and Zhang, 2009 adopt a Generalised Maximum Likelihood Empirical Bayes estimator to estimate the

prior). For a Gaussian kernel and as $B \rightarrow \infty$, this estimator approaches the oracle Bayes rule.

Similar to our approach, George (1986) proposed a weighted combination of shrinkage estimators, e.g., multiple JS estimators. The weights are assumed to be known but can adapt to the data to some extent. He showed that an aggregation of Bayes rules is again Bayes on a mixture prior where the weights naturally translate to prior probabilities.

We emphasise the following key differences of this important line of work to the present one:

- (a) The above approaches focus on the *compound* risk, while we analyse the risk of each individual task. The compound relative risk, analysed in Section 5.2, is a different quantity from the ratio between compound risk and oracle Bayes risk.
- (b) In the empirical Bayes framework, the focus lies on asymptotics as the number of independent tasks B grows large, while ours is on the growing (effective) dimension. Consequently, the choice of “oracle” reference for analyzing risk ratios differs between the two perspectives. Within the empirical Bayes paradigm, the compound oracle Bayes risk serves as the reference. We adopt a task-specific oracle improvement relative to the naive estimator. Thus, the theoretical outcomes derived from these divergent approaches are not readily comparable.

Concerning the role of the dimension, consistency with the oracle Bayes reference requires $d/B \rightarrow 0$ for the parametric approach of Efron and Morris (1976) and presumably an even more stringent condition for the nonparametric approaches of Brown and Greenshtein (2009) or Jiang and Zhang (2009). In fact they only considered the case $d = 1$, but since both works rely on metric entropy estimates on appropriate function spaces, one would expect those to suffer of the curse of dimensionality.

Consistency with the oracle, as considered in our paper, requires roughly $\text{polylog}(B)/d \rightarrow 0$, thereby accommodating a broader spectrum of regimes. For instance, when $B = \Theta(d^\alpha)$ for arbitrary $\alpha > 0$, our approach ensures consistency with our oracle improvement, yet fails to achieve consistency with the oracle Bayes with a Gaussian prior if $\alpha \leq 1$. Conversely, the regime where $B \rightarrow \infty$ while d remains fixed, which is pertinent to empirical Bayes analyses, does not yield meaningful results in our framework (though, allowing the dimension to increase at an arbitrary small power of B remains viable).

In summary, our perspective is tailored towards high-dimensional scenarios, with possibly non-isotropic covariance structures, whereas the empirical Bayes methodology is not inherently designed for such settings. Moreover, we emphasise the minimax property of our oracle improvement across suitable models as the dimension grows.

- (c) We allow non-Gaussian data.
- (d) We allow strong task heterogeneity (e.g., the covariances are not shared across tasks).

7.2. Multitask learning point of view. Feldman et al. (2014) viewed the many means estimation problem (1) as a multi-task learning problem (Caruana, 1997; Zhang and Yang, 2021), which gave rise to the term multi-task averaging. Also inspired by the JS estimator, the proposed approach extends the empirical compound risk minimization with a regularization term that favours the alignment of mean estimations for “related” tasks. The notion of “task relatedness” is encoded as a similarity matrix considered as *a priori* information. In absence of specific information, the similarities are taken constant across tasks and the method reduces to shrinkage towards the grand mean. The theoretical analysis focused mainly on the low-dimensional setting and the oracle weights when $B = 2$. Their data-driven similarity estimation yielded inconclusive results. Martínez-Rego and Pontil (2013) mitigated the default constant similarities in the absence of information by first clustering the tasks into different groups and then applying the approach of Feldman et al. (2014) on each cluster separately; but a theoretical analysis of this approach was not conducted. In our work, we also propose to assimilate estimators of related tasks and thereby define an appropriate shrinkage direction. We eliminate the disadvantage of both approaches, i.e., constant or known similarities, by estimating them solely based on the available data. We also extend significantly our preliminary work (Marienwald et al., 2021) which was limited to the testing approach unfit for heterogeneous tasks, and with less precise theoretical results.

Recent work of Duan and Wang (2023) considers a general multi-task learning setting which includes the multiple mean estimation problem as a special case. Comparable to that of Feldman et al. (2014), their estimators are determined by compound empirical risk minimization with a regularization term measuring alignment to a predetermined model of task relatedness, e.g., the means form K clusters or are close to a linear subspace of dimension K . The proposed estimators depend on the considered task relatedness and on K . Once interpreted in terms of relative squared risk, the theoretical bounds obtained by Duan and Wang (2023) are not bounded by a constant but can grow as $\mathcal{O}(K^2)$ in the worst case where the fit to the posited task relatedness is poor. For the relative risk to be significantly less than $\mathcal{O}(1)$, the bounds require the condition $\delta \lesssim s_1/K$, where δ represents closeness to the model (cluster radius resp. distance to linear subspace). By contrast, in our analysis we do not posit a particular task relatedness or value of K to define the estimators; those are adaptive to the most advantageous grouping model, including cluster number and size, describing the structure of the true parameters (see Section 5). Our relative risk bounds are worst-case bounded (and even bounded close to 1), and show a significant improvement in favourable cases even for the number of groups K growing with the number of tasks B . On the other hand, our approach won’t result in a significant risk improvement if the task means belong to a low-dimensional subspace but are very far apart from each other. Still, using the covering complexity point of view discussed in Section 5, an improvement can be shown if the tasks increase in number and are drawn, say, from a fixed *a priori* distribution having a low-dimensional support while the ambient dimension grows.

8. CONCLUSION

Considering the estimation of multiple mean vectors in high dimensions from independent samples, we focused on estimators formed as convex combinations of empirical averages of each sample. We proposed a test-then-aggregate method generalizing the approach of Marienwald et al. (2021), and a direct Q -aggregation approach where the weights are found by minimization of an adequate objective. From a theoretical perspective, we established asymptotic convergence to an oracle risk in an appropriate “dimensional asymptotics” sense, as the effective dimensionality grows. This oracle risk was proved to be exactly minimax under certain homogeneity conditions for the single-task risk, and minimax up to a fixed factor for the compound relative risk (without homogeneity conditions). One advantage of the Q -aggregation method is its theoretical adaptivity with respect to parameters that have to be user-provided for the testing approach. We demonstrated the efficacy of the proposed methods on showcase experiments for estimating multiple kernel mean embeddings on controlled artificial datasets and real-world flow cytometry data.

Future investigations will aim to address the discrepancy between the lower and upper bounds for the single mean estimation in extremely inhomogeneous cases (we suspect the minimax lower bound could be too conservative in such a case because it does not take into account the problem of neighbour detection). Another important open direction is the integration in the multiple-mean estimation setting of recent advances on single-mean estimation in high dimension, achieving sub-Gaussian performance even under heavy-tailed distributions or samples that were adversarially corrupted, e.g., the median of means estimator (Lugosi and Mendelson, 2019b, 2020, see Fathi et al., 2022; Lugosi and Mendelson, 2019a for an overview), or efficiently computable estimators (e.g., Cheng et al., 2019; Depersin and Lecué, 2022). Finally, a significant future avenue is to extend our approach from mean estimation to more general high-dimensional multi-task learning problems such as those considered by Duan and Wang (2023).

REFERENCES

- Adan, A., Alizada, G., Kiraz, Y., Baran, Y., and Nalbant, A. (2017). Flow cytometry: basic principles and applications. *Critical Reviews in Biotechnology* 37 (2).
- Beran, R. (1996). Stein estimation in high dimensions: a retrospective. In: *Research Developments in Probability And Statistics: Festschrift in Honor of Madan L. Puri*. VSP International Science Publishers.
- Birgé, L. (2001). An alternative point of view on Lepski’s method. In: *State of the Art in Probability and Statistics*. Lecture Notes Monograph Series. Institute of Mathematical Statistics.
- Blanchard, G. and Fermanian, J.-B. (2023). Nonasymptotic one- and two-sample tests in high dimension with unknown covariance structure. *Foundations of Modern Statistics*.
- Boucheron, S., Lugosi, G., and Bousquet, O. (2004). Concentration inequalities. In: *Advanced Lectures on Machine Learning: Summer School on Machine Learning*. Springer.

- Brehmer, J. and Cranmer, K. (2020). Flows for simultaneous manifold learning and density estimation. *Advances in Neural Information Processing Systems*.
- Brown, L. D. and Greenshtein, E. (2009). Nonparametric empirical Bayes and compound decision approaches to estimation of a high-dimensional vector of normal means. *The Annals of Statistics* 37 (4).
- Caruana, R. (1997). Multitask learning. *Machine Learning* 28 (1).
- Cheng, Y., Diakonikolas, I., and Ge, R. (2019). High-dimensional robust mean estimation in nearly-linear time. *ACM-SIAM Symposium on Discrete Algorithms*.
- Chernoff, H. (1952). A measure of asymptotic efficiency for tests of a hypothesis based on the sum of observations. *The Annals of Mathematical Statistics* 23 (4).
- Chwialkowski, K., Strathmann, H., and Gretton, A. (2016). A kernel test of goodness of fit. *International Conference on Machine Learning*. Vol. 48.
- Collins, M., Globerson, A., Koo, T., Carreras Pérez, X., and Bartlett, P. (2008). Exponentiated gradient algorithms for conditional random fields and max-margin markov networks. *Journal of Machine Learning Research* 9.
- Depersin, J. and Lecué, G. (2022). Robust sub-Gaussian estimation of a mean vector in nearly linear time. *The Annals of Statistics* 50 (1).
- Devroye, L., Lerasle, M., Lugosi, G., and Oliveira, R. I. (2016). Sub-Gaussian mean estimators. *The Annals of Statistics* 44 (6).
- Duan, Y. and Wang, K. (2023). Adaptive and robust multi-task learning. *The Annals of Statistics* 51 (5).
- Dussap, B., Blanchard, G., and Chérif-Abdellatif, B.-E. (2023). Label shift quantification with robustness guarantees via distribution feature matching. *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*.
- Efron, B. and Morris, C. (1972). Empirical Bayes on vector observations: An extension of Stein’s method. *Biometrika* 59 (2).
- Efron, B. and Morris, C. (1973). Stein’s estimation rule and its competitors—an empirical Bayes approach. *Journal of the American Statistical Association* 68 (341).
- Efron, B. and Morris, C. (1976). Multivariate empirical Bayes and estimation of covariance matrices. *The Annals of Statistics* 4 (1).
- Fakoor, R., Chaudhari, P., Mueller, J., and Smola, A. J. (2020). Trade: Transformers for density estimation. *arXiv:2004.02441*.
- Fathi, M., Goldstein, L., Reinert, G., and Saumard, A. (2022). Relaxing the Gaussian assumption in shrinkage and SURE in high dimension. *The Annals of Statistics* 50 (5).
- Feldman, S., Gupta, M. R., and Frigiyik, B. A. (2014). Revisiting Stein’s paradox: multi-task averaging. *Journal of Machine Learning Research* 15.
- Filippi, S., Flaxman, S., Sejdinovic, D., and Cunningham, J. (2016). Bayesian learning of kernel embeddings. *Conference on Uncertainty in Artificial Intelligence*.
- Finak, G., Langweiler, M., Jaimes, M., Malek, M., Taghiyar, J., Korin, Y., Raddassi, K., Devine, L., Obermoser, G., Pekalski, M. L., et al. (2016). Standardizing flow cytometry immunophenotyping analysis from the Human Immunophenotyping Consortium. *Scientific Reports* 6.

- George, E. I. (1986). Minimax multiple shrinkage estimation. *The Annals of Statistics* 14 (1).
- Gretton, A., Borgwardt, K. M., Rasch, M. J., Schölkopf, B., and Smola, A. (2012). A kernel two-sample test. *Journal of Machine Learning Research* 13.
- Gretton, A., Fukumizu, K., Teo, C., Song, L., Schölkopf, B., and Smola, A. (2007). A kernel statistical test of independence. *Advances in Neural Information Processing Systems*.
- Houdré, C. and Reynaud-Bouret, P. (2003). Exponential inequalities, with constants, for U-statistics of order two. *Progress in Probability*. Vol. 56.
- Hsu, D., Kakade, S., and Zhang, T. (2012). Tail inequalities for sums of random matrices that depend on the intrinsic dimension. *Electronic Communications in Probability* 17.
- James, W. and Stein, C. (1961). Estimation with quadratic loss. *4th Berkeley Symposium on Mathematical Statistics and Probability*. Vol. 1.
- Jegelka, S., Gretton, A., Schölkopf, B., Sriperumbudur, B. K., and Von Luxburg, U. (2009). Generalized clustering via kernel embeddings. *Advances in Artificial Intelligence*. Vol. 32.
- Jiang, W. and Zhang, C.-H. (2009). General maximum likelihood empirical Bayes estimation of normal means. *The Annals of Statistics* 37 (4).
- Kivinen, J. and Warmuth, M. K. (1997). Exponentiated gradient versus gradient descent for linear predictors. *Information and Computation* 132 (1).
- Koltchinskii, V. and Lounici, K. (2016). Asymptotics and concentration bounds for bilinear forms of spectral projectors of sample covariance. *Annales de l'Institut Henri Poincaré, Probabilités et Statistiques* 52 (4).
- Laurent, B. and Massart, P. (2000). Adaptive estimation of a quadratic functional by model selection. *The Annals of Statistics* 28 (5).
- Lecué, G. and Rigollet, P. (2014). Optimal learning with Q-aggregation. *The Annals of Statistics* 42 (1).
- LeCun, Y., Bottou, L., Bengio, Y., and Haffner, P. (1998). Gradient-based learning applied to document recognition. *Proceedings of the IEEE* 86 (11).
- Lerasle, M., Szabo, Z., Mathieu, T., and Lecue, G. (2019). MONK outlier-robust mean embedding estimation by median-of-means. *International Conference on Machine Learning*.
- Lugosi, G. and Mendelson, S. (2019a). Mean estimation and regression under heavy-tailed distributions: A survey. *Foundations of Computational Mathematics* 19 (5).
- Lugosi, G. and Mendelson, S. (2019b). Sub-Gaussian estimators of the mean of a random vector. *The Annals of Statistics* 47 (2).
- Lugosi, G. and Mendelson, S. (2020). Multivariate mean estimation with direction-dependent accuracy. *Journal of the European Mathematical Society*.
- Marienwald, H., Fermanian, J.-B., and Blanchard, G. (2021). High-dimensional multi-task averaging and application to kernel mean embedding. *International Conference on Artificial Intelligence and Statistics*.

- Martínez-Rego, D. and Pontil, M. (2013). Multi-task averaging via task clustering. *Similarity-Based Pattern Recognition: 2nd International Workshop*.
- Massart, P. (2000). About the constants in Talagrand’s concentration inequalities for empirical processes. *The Annals of Probability* 28 (2).
- McAuley, J. (2022). *Personalized machine learning*. Cambridge University Press.
- McDiarmid, C. (1998). Concentration. In: *Probabilistic Methods for Algorithmic Discrete Mathematics*. Springer Berlin Heidelberg.
- McKinnon, K. M. (2018). Flow cytometry: an overview. *Current Protocols in Immunology* 120 (5).
- Minsker, S. and Strawn, N. (2019). Distributed statistical estimation and rates of convergence in normal approximation. *Electronic Journal of Statistics* 13 (2).
- Muandet, K., Fukumizu, K., Dinuzzo, F., and Schölkopf, B. (2012). Learning from distributions via support measure machines. *Advances in Neural Information Processing Systems*.
- Muandet, K., Fukumizu, K., Sriperumbudur, B., Gretton, A., and Schölkopf, B. (2014). Kernel mean estimation and Stein effect. *International Conference on Machine Learning*.
- Muandet, K., Fukumizu, K., Sriperumbudur, B., and Schölkopf, B. (2017). Kernel mean embedding of distributions: a review and beyond. *Foundations and Trends in Machine Learning* 10.
- Muandet, K., Sriperumbudur, B., Fukumizu, K., Gretton, A., and Schölkopf, B. (2016). Kernel mean shrinkage estimators. *Journal of Machine Learning Research* 17 (48).
- Nussbaum, M. (1996). The Pinsker bound: A review. In: *Encyclopedia of Statistical Sciences*. Wiley.
- Pinsker, M. S. (1980). Optimal filtering of square-integrable signals in Gaussian noise. *Problemy Peredachi Informatsii* 16 (2).
- Robbins, H. (1951). Asymptotically subminimax solutions of compound statistical decision problems. *Proceedings of the 2nd Berkeley Symposium on Mathematical Statistics and Probability*. Vol. 2.
- Robbins, H. (1964). The empirical Bayes approach to statistical decision problems. *The Annals of Mathematical Statistics* 35 (1).
- Rudelson, M. and Vershynin, R. (2007). Sampling from large matrices: an approach through geometric functional analysis. *Journal of the ACM* 54 (4).
- Smola, A., Gretton, A., Song, L., and Schölkopf, B. (2007). A Hilbert space embedding for distributions. *International Conference on Algorithmic Learning Theory*.
- Stein, C. (1956). Inadmissibility of the usual estimator for the mean of a multivariate normal distribution. *3rd Berkeley Symposium on Mathematical Statistics and Probability*.
- Szabó, Z., Sriperumbudur, B. K., Póczos, B., and Gretton, A. (2016). Learning theory for distribution regression. *Journal of Machine Learning Research* 17 (152).
- Tan, A. Z., Yu, H., Cui, L., and Yang, Q. (2022). Towards personalized federated learning. *IEEE Transactions on Neural Networks and Learning Systems* 34 (12).

- Tropp, J. A. et al. (2015). An introduction to matrix concentration inequalities. *Foundations and Trends in Machine Learning* 8 (1).
- Tsybakov, A. B. (2008). *Introduction to nonparametric estimation*. 1st. Springer Publishing Company.
- Vershynin, R. (2018). *High-dimensional probability: an introduction with applications to data science*. Cambridge University Press.
- Xia, X., Shan, S., Gong, M., Wang, N., Gao, F., Wei, H., and Liu, T. (2022). Sample-efficient kernel mean estimator with marginalized corrupted data. *28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*.
- Zhang, C.-H. (2003). Compound decision theory and empirical Bayes methods. *Annals of Statistics* 31 (2).
- Zhang, Y. and Yang, Q. (2021). A survey on multi-task learning. *IEEE Transactions on Knowledge and Data Engineering* 34 (12).

ACKNOWLEDGEMENTS AND FUNDING

HM gratefully receives funding from the German Federal Ministry of Education and Research under the grant BIFOLD24B. Part of this research was done when HM was hired at Universität Potsdam, Germany, and funded by the German Ministry for Education and Research as BIFOLD (01IS18025A and 01IS18037A). GB gratefully acknowledges funding from the grants ANR-21-CE23-0035 (ASCAI) and ANR-19-CHIA-0021-01 (BISCOTTE) of the French National Research Agency ANR. Part of this research was done as GB received support from DFG SFB1294 “Data Assimilation”, as “Mercator fellow” at the University of Potsdam.

APPENDIX A. NOMENCLATURE

B	number of tasks, Sec. 1
$\mathcal{B}(\tau, \nu)$	oracle risk, (10)
(BS)	bounded assumption, Sec. 2.3
\mathcal{C}	J -partition of $\llbracket B \rrbracket$, Def. 5
d	ambient dimension, Sec. 1
d_k^\bullet	effective dimension, (5)
d_k^e	effective dimension, (5)
$\text{diam}(\mathcal{C}, \boldsymbol{\mu})$	diameter of partition \mathcal{C} of $\boldsymbol{\mu}$, (39)
Δ_k	difference between μ_k and μ_1 , Sec. 2.5
(ECSS)	equal covariances and sample sizes, Sec. 2.3
η	relative estimation error of s_k^2 , (14)
(GS)	Gaussian assumption, Sec. 2.3
(HT)	heavy-tailed assumption, Sup. C.4
J	nr. of parts of the partition, Def. 5
k	index of task, Sec. 1
(KC)	known covariances, Sec. 2.3
$L_k(\hat{\boldsymbol{\mu}})$	loss of estimator $\hat{\boldsymbol{\mu}}$, (2)
$L_k(\boldsymbol{\omega})$	loss of aggregation estimator $\hat{\boldsymbol{\mu}}_{\boldsymbol{\omega}}$, (2.5)
$\hat{L}_k(\boldsymbol{\omega})$	estimator for cond. risk, (29), Sec. 1
$\mathcal{L}^*(\mathbf{s}, \mathcal{C}, \boldsymbol{\zeta})$	compound oracle risk, (42)
M	radius of ball in which the bounded data lies, Sec. 4.3
μ_k	expectation of distribution k
$\hat{\mu}_k$	estimator of μ_k
$\hat{\mu}_k^{\text{NE}}$	naive estimation (empirical average) of μ_k , Sec. 1
$\hat{\boldsymbol{\mu}}_{\boldsymbol{\omega}}$	aggregation estimator, (6)
$\llbracket n \rrbracket$	integers 1 to n , Sec. 1
N_k	number of samples (bag size) of task k , Sec. 1
$\ a\ $	canonical norm of vector a , Sec. 2
$\ \Sigma\ _p$	Schatten norm of matrix Σ , Sec. 2.2
$\ \Sigma\ _\infty$	operator norm of matrix Σ , Sec. 2.2
$\nu(U)$	relative aggregated variance, (9)
$\boldsymbol{\omega}$	aggregation weights, (6)
\mathbb{P}_k	k -th task (probability distribution), Sec. 1
$\mathcal{P}_{\text{mult}}(\mathcal{C}, \boldsymbol{\zeta}, \Sigma, \mathbf{s})$	class of distributions, Def. 6
$\mathcal{P}_{\text{single}}(\tau, V, \Sigma, \mathbf{s}^2)$	class of distributions, Def. 4
$\hat{Q}_1(\boldsymbol{\omega})$	prob. upper bound on $(\hat{L}_1(\boldsymbol{\omega}) - L_1(\boldsymbol{\omega}))$, (31)
$\hat{Q}_1^{\text{BS}}(\boldsymbol{\omega})$	additional penalization for (BS), (37)
$R_k(\hat{\boldsymbol{\mu}})$	risk of estimator $\hat{\boldsymbol{\mu}}$, (2)
$R_k(\boldsymbol{\omega})$	risk of aggregation estimator $\hat{\boldsymbol{\mu}}_{\boldsymbol{\omega}}$, (2.5)
$s^2(U)$	harmonic mean of the risks of the tasks in U (9)

\mathcal{S}_B	$(B - 1)$ -dimensional simplex, Sec. 2.5
s_k^2	naive risk, (3)
\mathcal{S}_V	set of convex weights of support incl. in V , Sec. 2.5
ς	threshold for $W_{(\varsigma)}$, (21)
Σ_k	covariance matrix of k -th task, Sec. 2
$\tilde{T}_k^{(\tau)}, \tilde{\tilde{T}}_k^{(\tau)}$	empirical similarity test on independent copy data, (17)-(63)
$\tau, \tau_{\min}^k, \tau_{\min}^\circ, \tau^\pm$	thresholds for similarity test, (8)-(18)-(23)-(23)
$\tau/1+\tau$	best potential improvement
\tilde{U}_k	unbiased estimator for $\ \Delta_k\ ^2$, (16)
\tilde{V}	estimation of V_τ , (14)
V_τ	τ -neighbouring tasks, (8)
$V_{\tau,\varsigma}$	trimmed V_τ , (22)
V^*	subset of \tilde{V} , (14)
$W_{(\varsigma)}$	set of tasks with bounded variance, (21)
$X_\bullet^{(k)}$	k -th bag, (1)
$\tilde{X}_\bullet^{(k)}$	independent copy of k -th bag, Sec. 3.2
ζ	bound on the diameter of the J -partition, Def. 6

APPENDIX B. PROPERTIES OF THE JAMES-STEIN ESTIMATOR IN THE LARGE DIMENSION REGIME

In this section, we provide a concise overview of the properties of the James-Stein estimator in high-dimensional settings and a comparison with the Q-aggregation approach. Let us first cast the standard James-Stein problem as a particular limiting case of our general setting (1) with only two bags, the second of which with known mean equal to 0 and serving as a reference point:

Assumption (JS, James-Stein setting). $B = 2$, $\mu_2 = 0$, formally $N_2 = \infty$ and $s_2^2 = 0$.

Since only μ_1 is of interest, we drop the index 1 everywhere from the notation in what follows. In that case, identifying ω with its first weight, renoted as $\omega \in [0, 1]$, $\hat{\mu}_\omega = \omega \hat{\mu}^{\text{NE}}$ (defined in (6)) is simply a shrinkage estimator towards 0.

This is the type of estimator that Stein (1956) used to demonstrate that the empirical mean is not admissible. Indeed, in an isotropic Gaussian setting ($\Sigma = \sigma^2 I_d$), a shrinkage estimator $\mu^{\text{JS}+}$ (James and Stein, 1961) outperforms the empirical mean by shrinking towards a chosen reference point $\mu_2 = 0$. Let us denote $\sigma_N^2 = \sigma^2/N$ and:

$$\hat{\mu}^{\text{JS}+} = \left(1 - \frac{\sigma^2 (d-2)}{N \|\hat{\mu}^{\text{NE}}\|_2^2}\right)_+ \hat{\mu}^{\text{NE}}. \quad (48)$$

The estimator is minimax for means inside a ball of radius $\tau \sigma_N^2 d$ but beats the empirical mean in general. For the model $\mathcal{G}_d(\tau, \sigma^2) = \{\mathcal{N}(\mu, \sigma^2 I_d)^{\otimes N}, \|\mu\|^2 \leq \tau \sigma_N^2 d\}$, the class of N -samples of an isotropic Gaussian distribution with bounded mean vector, Pinsker

(1980) shows that (see also Beran, 1996; Nussbaum, 1996; Tsybakov, 2008):

$$\liminf_{d \rightarrow \infty} \sup_{\hat{\mu} \in \mathbb{G}_d(\tau, \sigma^2)} \frac{\mathbb{E}[\|\hat{\mu} - \mu\|^2]}{d\sigma_N^2} = \frac{\tau}{1 + \tau}, \quad \frac{\mathbb{E}[\|\hat{\mu}^{\text{JS}^+} - \mu\|^2]}{d\sigma_N^2} \leq \frac{\tau}{1 + \tau} + \frac{4}{d}. \quad (49)$$

In a non isotropic setting (Σ is not necessary equal to $\sigma^2 I_d$ but remains known), a similar estimator achieving the same bounds can be constructed by replacing d by d^e and σ^2 by $\|\Sigma\|_\infty$ in (48) and (49).

We review three interpretations of the James-Stein estimator which relates it to our approaches in the general model: the oracle and testing ones in Section 3 and the Q-aggregation one in Section 4.

Oracle interpretation: the James-Stein shrinkage factor can be seen as an approximation of an oracle weight defined as the minimiser of the risk of $\hat{\mu}_\omega$:

$$R(\hat{\mu}_\omega) = (1 - \omega)^2 \|\mu\|^2 + \omega^2 \sigma_N^2 d,$$

which is minimised by $\omega^* = \frac{\|\mu\|^2}{d\sigma_N^2 + \|\mu\|^2}$. By remarking that $\omega^* = 1 - \frac{d\sigma_N^2}{d\sigma_N^2 + \|\mu\|^2}$, a natural estimation of the oracle weight is obtained in (48).

Test interpretation: As in Section 3, we could first want to detect if μ is close to $\mu_2 = 0$. Knowing the variance, a very simple test is $T = \mathbf{1}\{\|\hat{\mu}^{\text{NE}}\|_2^2 \leq (d - 2)\sigma_N^2\}$ for the hypothesis:

$$(H_0) : \|\mu\|^2 \geq d\sigma_N^2 \quad \text{against} \quad (H_1) : \mu = 0.$$

If (H_0) is rejected, then we choose 0 as an estimator of μ .

Regularization interpretation: Consider the following estimation by regularization, for $\lambda > 0$:

$$\hat{\mu}^\lambda \in \arg \min_{\mu \in \mathbb{R}^d} \frac{1}{n} \sum_{i=1}^n \|\mu - X_i\|^2 + 2\lambda \|\mu\|. \quad (50)$$

Then $\hat{\mu}^\lambda = \left(1 - \frac{\lambda}{\|\hat{\mu}^{\text{NE}}\|_2}\right)_+ \hat{\mu}^{\text{NE}}$ is a minimiser. Using Stein's Lemma, we recover (48) by choosing $\lambda = \frac{\sigma^2}{N} \frac{d-2}{\|\hat{\mu}^{\text{NE}}\|_2}$ (see for example Lemma 3.8 of Tsybakov, 2008).

In the **(JS)** setting, the Q-aggregation method exhibits the same asymptotic behaviour as the James-Stein estimator $\hat{\mu}^{\text{JS}^+}$ without knowing the covariance Σ . Corollary 3 is deduced from Proposition 5 and is proven in Supplemental D.

Corollary 3. *Assume **(JS)** and **(GS)**, let $N \geq 7$, $(N - 1)/2 \geq u_0 \geq 3$, and $\hat{\omega}$ as defined in (32). Then:*

$$\frac{R(\hat{\omega})}{s^2} \leq \frac{\|\mu\|^2}{s^2 + \|\mu\|^2} (1 + Ce^{-u_0/2}) + C \sqrt{\frac{u_0}{d^e}}, \quad (51)$$

where $C > 0$ is some absolute constant.

The first term is, up to the multiplicative factor, Stein's error $\tau/(1+\tau)$ with $\tau = \|\mu^2\|/s^2$. In the dimensional asymptotic $d^e \rightarrow \infty$, assume $u_0 \rightarrow \infty$ such that $u_0 = o(d^e)$ and suppose the mean satisfies $\|\mu\|^2 \leq \tau s^2$, then the estimator attains the Pinsker bound (49):

$$\lim_{d^e \rightarrow \infty} \sup_{\substack{\mu, s: \\ \|\mu\|^2 \leq \tau s^2}} \frac{R(\hat{\omega})}{s^2} \leq \frac{\tau}{1+\tau}.$$

We recover the same phenomenon in the bounded setting directly from Proposition 6 as $V_{\tau, \varsigma} = \{1, 2\}$ for all $\varsigma > 0$ and the relative aggregated variance is null: $\nu(V_{\tau, \varsigma}) = 0$.

APPENDIX C. PROOFS FOR SECTION 3

C.1. Proof of Lemma 1. The weights ω^* are obtained by minimizing the upper bound (11) using KKT conditions, for instance. However, to verify the bound (12), it suffices to substitute the weights (13) into (11). Let us denote $\nu = \frac{s^2(V)}{s_1^2}$, from (12):

$$\begin{aligned} \frac{R_1(\omega^*)}{s_1^2} &\leq \tau(1 - \omega_1^*)^2 + \sum_{k \in V} (\omega_k^*)^2 \frac{s_k^2}{s_1^2} \\ &= \tau\lambda^2(1 - \nu)^2 + (1 - \lambda)^2 + 2\lambda(1 - \lambda)\nu + \lambda^2\nu. \end{aligned}$$

By substituting λ with its value from Equation (13), we obtain:

$$\begin{aligned} \frac{R_1(\omega^*)}{s_1^2} &\leq \frac{\tau(1 - \nu)^2 + \tau^2(1 - \nu)^2 + 2\tau(1 - \nu)\nu + \nu}{(1 + \tau(1 - \nu))^2} \\ &= \frac{\tau(1 - \nu)((1 - \nu) + \tau(1 - \nu) + \nu) + \nu(\tau(1 - \nu) + 1)}{(1 + \tau(1 - \nu))^2} \\ &= \frac{\tau(1 - \nu) + \nu}{1 + \tau(1 - \nu)} = \mathcal{B}(\tau, \nu) = \mathcal{B}\left(\tau, \frac{s^2(V)}{s_1^2}\right). \end{aligned}$$

Thus, the inequality holds as claimed. \square

C.2. Proof of Proposition 1. Recall that we assume the following event holds:

$$\begin{cases} 1 \in V^* \subseteq \tilde{V} \subseteq V_\tau, \\ |\tilde{s}_k^2 - s_k^2| \leq \eta s_k^2, \text{ for all } k \in \tilde{V}, \end{cases} \quad (14)$$

for quantities \tilde{V}, \tilde{s} which are considered as nonrandom for this proof (e.g., they are computed from an independent sample and we argue conditionally to that sample). Denote

$$\bar{R}_1(\tilde{V}, \omega) := \tau s_1^2(1 - \omega_1)^2 + \sum_{k \in \tilde{V}} \omega_k^2 s_k^2$$

the risk upper bound from (11) wherein we used the index set \tilde{V} . Due to the first $\tilde{V} \subset V_\tau$ the same argument leading up to (11), it holds $R_1(\omega) \leq \bar{R}_1(\tilde{V}, \omega)$ for all

$\omega \in \mathcal{S}_{\tilde{V}}$. Denoting now

$$\tilde{R}_1(\tilde{V}, \omega) := \tau \tilde{s}_1^2 (1 - \omega_1)^2 + \sum_{k \in \tilde{V}} \omega_k^2 \tilde{s}_k^2$$

the plug-in version of $\bar{R}_1(\tilde{V}, \omega)$, we have, putting $\varepsilon_k := |s_k^2 - \tilde{s}_k^2|$:

$$\forall \omega \in \mathcal{S}_{\tilde{V}} : \left| \bar{R}_1(\tilde{V}, \omega) - \tilde{R}_1(\tilde{V}, \omega) \right| \leq \tau \varepsilon_1 (1 - \omega_1)^2 + \sum_{k \in \tilde{V}} \omega_k^2 \varepsilon_k \leq \left(\max_{k \in \tilde{V}} \frac{\varepsilon_k}{s_k^2} \right) \bar{R}_1(\tilde{V}, \omega),$$

which entails, from the second part of event (14):

$$(1 - \eta) \bar{R}_1(\tilde{V}, \omega) \leq \tilde{R}_1(\tilde{V}, \omega) \leq (1 + \eta) \bar{R}_1(\tilde{V}, \omega)$$

Since $\tilde{\omega}$ is a minimiser of $\tilde{R}_1(\tilde{V}, \omega)$, it holds for any other $\omega \in \mathcal{S}_{\tilde{V}}$:

$$R_1(\tilde{\omega}) \leq \bar{R}_1(\tilde{V}, \tilde{\omega}) \leq (1 - \eta)^{-1} \tilde{R}_1(\tilde{V}, \tilde{\omega}) \leq (1 - \eta)^{-1} \tilde{R}_1(\tilde{V}, \omega) \leq \left(\frac{1 + \eta}{1 - \eta} \right) \bar{R}_1(\tilde{V}, \omega).$$

Minimizing the latter inequality over ω yields (from Lemma 1):

$$\frac{R_1(\tilde{\omega})}{s_1^2} \leq \left(\frac{1 + \eta}{1 - \eta} \right) \mathcal{B}(\tau, \nu(\tilde{V})) \leq \left(\frac{1 + \eta}{1 - \eta} \right) \mathcal{B}(\tau, \nu(V^*)),$$

due to $V^* \subseteq \tilde{V}$ and the monotonicity properties of ν , \mathcal{B} .

C.3. Proofs for Section 3.3. We start with a generic result linking concentration of the test statistic to the properties of the associated test. It will allow to handle different distributional settings as particular cases.

We recall that \tilde{U}_k is the test U-statistic given by (16) using independent ‘‘tilde’’ data.

Assumption (TSC, Test Statistic Concentration). Assume that for all $k \in \llbracket B \rrbracket$ and $\alpha \in (0, 1)$, there exists $q_k(\alpha)$:

$$\mathbb{P} \left[\left| \tilde{U}_k - \|\Delta_k\|^2 \right| \geq \|\Delta_k\| q_k(\alpha) + c_0^2 q_k^2(\alpha) \right] \leq \alpha. \quad (52)$$

where $c_0 \geq 2$ is a numerical constant.

Put $u_\alpha := \log(8/\alpha)$, it is established that:

- The assumption is satisfied under **(GS)** for $q_k^2(\alpha) = 2 \left(\frac{\|\Sigma_1\|_2}{N_1} + \frac{\|\Sigma_k\|_2}{N_k} \right) u_\alpha$ and $c_0 = 4$. (Proposition 6 in Blanchard and Fermandian (2023))
- The assumption is satisfied under **(BS)** for $q_k^2(\alpha) = 16 \left(\frac{\|\Sigma_1\|_2}{N_1} + \frac{\|\Sigma_k\|_2}{N_k} \right) u_\alpha + 4 \frac{M^2 u_\alpha^2}{N_1^2 \wedge N_k^2}$ and $c_0 = 31$. (Proposition 9 in Blanchard and Fermandian (2023))
- The assumption is satisfied under **(HT)** for $q_k^2(\alpha) = 16 \left(\frac{\|\Sigma_1\|_2}{N_1} + \frac{\|\Sigma_i\|_2}{N_i} \right) u_\alpha$ and $c_0 = 2$ but for $\alpha \geq 8e^{-N_1 \wedge N_i}$. (Proposition 18).

Proposition 11. Grant assumption **(TSC)** and let $\alpha \in (0, 1)$, $\tau > 0$ be fixed. Let \tilde{T}_k be given by

$$\tilde{T}_k := \mathbf{1}\left\{\tilde{U}_k \leq \tau s_1^2\right\}. \quad (53)$$

Define $\tau_{\min}^k := 2c_0^2 s_1^{-2} q_k^2(\alpha)$, then it holds:

$$\text{if } \|\mu_1 - \mu_k\| > (\sqrt{\tau} + \sqrt{\tau_{\min}^k}) s_1 : \quad \mathbb{P}\left[\tilde{T}_k = 1\right] \leq \alpha; \quad (54)$$

$$\text{if } \|\mu_1 - \mu_k\| \leq (\sqrt{\tau} - \sqrt{\tau_{\min}^k}) s_1 : \quad \mathbb{P}\left[\tilde{T}_k = 0\right] \leq \alpha. \quad (55)$$

Proof of Prop. 11 We assume for the rest of the proof that

$$\left|\tilde{U}_k - \|\Delta_k\|^2\right| \leq \|\Delta_k\| q_k(\alpha) + c_0^2 q_k^2(\alpha)$$

holds, which according to Assumption **(TSC)** is the case with probability at least $1 - \alpha$. Using $q_k^2(\alpha) s_1^{-2} = \tau_{\min} c_0^{-2} / 2$ and putting $x := \frac{\|\Delta_k\|}{\sqrt{\tau} s_1}$, the above inequality entails

$$\left|\frac{\tilde{U}_k}{\tau s_1^2} - x^2\right| \leq x \sqrt{\frac{\tau_{\min}}{2\tau}} c_0^{-1} + \frac{\tau_{\min}}{2\tau} \leq x \frac{\varepsilon_\tau}{2\sqrt{2}} + \frac{\varepsilon_\tau^2}{2}, \quad (56)$$

where we have used $c_0 \geq 2$ and where $\varepsilon_\tau := \sqrt{\tau_{\min}/\tau}$. This entails

$$\tau^{-1} s_1^{-2} (\tilde{U}_k - \tau s_1^2) \leq x^2 + x \frac{\varepsilon_\tau}{2\sqrt{2}} + \frac{\varepsilon_\tau^2}{2} - 1.$$

Assuming $\varepsilon_\tau \leq 1$, the largest root of the quadratic polynomial on the right-hand-side above is lower bounded as

$$x_+ = -\frac{\varepsilon_\tau}{4\sqrt{2}} + \sqrt{1 - \frac{15}{32}\varepsilon_\tau^2} \geq 1 - \varepsilon_\tau,$$

using $\sqrt{1-a} \geq 1 - \sqrt{a}$ for $a \in [0, 1]$. Thus, $0 \leq x \leq 1 - \varepsilon_\tau$ is a sufficient condition ensuring $\tilde{T}_k = 1$, implying (55) since $(1 - \varepsilon_\tau)^2 \tau = (\sqrt{\tau} - \sqrt{\tau_{\min}})^2$. (The case $\varepsilon_\tau > 1$ is trivial since the statement is void in that configuration.)

Similarly, (56) entails

$$\tau^{-1} s_1^{-2} (\tilde{U}_k - \tau s_1^2) \geq x^2 - x \frac{\varepsilon_\tau}{2\sqrt{2}} - \frac{\varepsilon_\tau^2}{2} - 1;$$

the largest root of the quadratic polynomial on the right-hand-side above is upper bounded as

$$x'_+ = \frac{\varepsilon_\tau}{4\sqrt{2}} + \sqrt{1 + \frac{17}{32}\varepsilon_\tau^2} \leq 1 + \varepsilon_\tau,$$

using $\sqrt{1+a} \leq 1 + \sqrt{a}$. Thus, $x > 1 + \varepsilon_\tau$ is a sufficient condition ensuring $\tilde{T}_k = 0$, implying (54) since $(1 + \varepsilon_\tau)^2 \tau = (\sqrt{\tau} + \sqrt{\tau_{\min}})^2$.

Proof of Prop. 2. Proposition 6 of Blanchard and Fermanian (2023) states that under **(GS)** it holds with probability at least $1 - \alpha$ that

$$\left| \tilde{U}_k - \|\Delta_k\|^2 \right| \leq \|\Delta_k\| q'_k \sqrt{u_\alpha} + 16 q_k^2 u_\alpha, \quad (57)$$

where

$$q_k^2 = 2 \left(\frac{\|\Sigma_1\|_2}{N_1} + \frac{\|\Sigma_k\|_2}{N_k} \right) = 2 s_1^2 \left(\frac{1}{\sqrt{d_1^\bullet}} + \frac{s_k^2/s_1^2}{\sqrt{d_k^\bullet}} \right),$$

and

$$(q'_k)^2 := 2 \left(\frac{\|\Sigma_1\|_\infty}{N_1} + \frac{\|\Sigma_k\|_\infty}{N_k} \right);$$

since $\|\Sigma\|_\infty \leq \|\Sigma\|_2$, we have $q'_k \leq q_k$, so that assumption (52) is satisfied with $c_0 = 4$. The claim is then a consequence of Proposition 11.

Proof of Cor. 1. For any $k \in \tilde{V}$, we have $k \in W_{(c)}$, and since $\varsigma \geq 1$, it holds (with the notation used in Proposition 2, but using α/B in place of α)

$$\tau_{\min}^{(k)} = 64(u_\alpha + \log B) \left(\frac{1}{\sqrt{d_1^\bullet}} + \frac{s_k^2}{s_1^2 \sqrt{d_k^\bullet}} \right) \leq \frac{\tau_{\min}^\circ}{2} + \frac{\varsigma \tau_{\min}^\circ}{2} \leq \varsigma \tau_{\min}^\circ,$$

and the result is a direct consequence of Proposition 2 (combined with a union bound over $k \in \llbracket B \rrbracket$).

C.4. Proofs for Section 3.4: estimating Schatten norms and plug-in estimates. We will be concentrating on one bag at a time and for this reason omit the task index k in the next results. Thus, we assume $\tilde{X}_1, \dots, \tilde{X}_N$ (with $N \geq 4$) are i.i.d. data points in \mathbb{R}^d with expectation μ and known covariance matrix Σ . We start with estimators for the Schatten norms $\|\Sigma\|_p$, $p = 1, 2$.

We can use the natural unbiased estimator for any fixed $\|\Sigma\|_1 = \text{Tr } \Sigma$,

$$\tilde{Z}^{(1)} := \frac{1}{N-1} \sum_{i=1}^N \|\tilde{X}_i - \tilde{\mu}\|^2 = \frac{1}{2N(N-1)} \sum_{i \neq j} \|\tilde{X}_i - \tilde{X}_j\|^2, \quad (58)$$

where $\tilde{\mu} = N^{-1} \sum_{i=1}^N \tilde{X}_i$ is the empirical mean of the (sub-)sample.

C.4.1. *Gaussian setting.* We have the following error control in the Gaussian setting:

Proposition 12. *Assume **(GS)** holds. For $u \geq 1$, if $N \geq 2$*

$$\mathbb{P} \left[\left| \tilde{Z}^{(1)} - \text{Tr } \Sigma \right| \geq 4 \sqrt{\frac{2 \text{Tr } \Sigma^2}{N} u} \right] \leq 2e^{-u}.$$

Proof. Let $\mathbf{X} = (\tilde{X}_1 - \tilde{\mu}, \dots, \tilde{X}_N - \tilde{\mu}) \in \mathbb{R}^{dN}$. Then \mathbf{X} is a centred Gaussian vector with covariance matrix $\Sigma := \Gamma \otimes \Sigma$ where $\Gamma = I_N - \frac{1}{N} \mathbf{1}_N \mathbf{1}_N^T \in \mathbb{R}^{N \times N}$, $\mathbf{1}_N = (1, \dots, 1) \in \mathbb{R}^N$ and \otimes denotes the Kronecker product. Note that it holds $\text{Tr } \Gamma = (N-1)$, $\Sigma^2 = \Gamma^2 \otimes \Sigma^2 =$

$\Gamma \otimes \Sigma^2$, $\text{Tr } \Sigma = \text{Tr } \Gamma \text{Tr } \Sigma = (N-1) \text{Tr } \Sigma$, and $\text{Tr}(\Sigma^2) = (N-1) \text{Tr } \Sigma^2$. Then, according to Lemma 5, for $u \geq 1$, with probability greater than $1 - 2e^{-u}$:

$$\begin{aligned} \|\mathbf{X}\|_2^2 &\leq \text{Tr } \Sigma + 2\sqrt{\text{Tr } \Sigma^2 u} + 2\|\Sigma\|_\infty u \leq (N-1) \text{Tr } \Sigma + 4\sqrt{(N-1) \text{Tr } \Sigma^2 u}, \\ \|\mathbf{X}\|_2^2 &\geq \text{Tr } \Sigma - 2\sqrt{\text{Tr } \Sigma^2 u} \geq (N-1) \text{Tr } \Sigma - 2\sqrt{(N-1) \text{Tr } \Sigma^2 u}. \end{aligned}$$

We have used that $\sqrt{u} \leq u$ for $u \geq 1$. We conclude by remarking that $\|\mathbf{X}\|_2^2 = (N-1)\tilde{Z}^{(1)}$. \square

Following Blanchard and Fermanian (2023), we can estimate $\|\Sigma\|_2 = \sqrt{\text{Tr } \Sigma^2}$ using the following U-statistic, which is an unbiased estimator of $\text{Tr } \Sigma^2$:

$$(\tilde{Z}^{(2)})^2 := \frac{1}{4N(N-1)(N-2)(N-3)} \sum_{i \neq j \neq k \neq l} \langle X_i - X_k, X_j - X_l \rangle^2. \quad (59)$$

Proposition 13 (Blanchard and Fermanian, 2023, Prop. 12). *Assume (GS) holds and $N \geq 4$. Then for all $u \geq 0$:*

$$\mathbb{P} \left[\left| \tilde{Z}^{(2)} - \sqrt{\text{Tr } \Sigma^2} \right| \geq 30 \sqrt{\frac{\text{Tr } \Sigma^2}{N} u^2} \right] \leq e^4 e^{-u}. \quad (60)$$

Proof of Prop. 3 Proposition 3 is a consequence of the above Proposition 12, using the union bound over $k \in \llbracket B \rrbracket$. \square

Propositions 12 and 13 can now be used to handle the plug-in versions of the quantities considered in Section 3.3 when covariances are unknown:

Proposition 14. *Assume (GS) holds, let $c \geq 1$ be a fixed number and let $\alpha \in (0, 1)$. Assume that we have estimates $\tilde{Z}_1^{(1)}$ for $\|\Sigma_1\|_1$ and $\tilde{Z}_k^{(2)}$ for $\|\Sigma_k\|_2$, $k \in \llbracket B \rrbracket$, (depending on the independent ‘‘tilde’’ data only) such that with probability $1 - \alpha$ it holds simultaneously for some constants $\eta_1, \eta_2 \in (0, 1)$:*

$$\left| \tilde{Z}_1^{(1)} - \|\Sigma_1\|_1 \right| \leq \eta_1 \|\Sigma_1\|_1; \quad (61)$$

$$\left| \tilde{Z}_k^{(2)} - \|\Sigma_k\|_2 \right| \leq \eta_2 \|\Sigma_k\|_2, \text{ for all } k \in \llbracket B \rrbracket. \quad (62)$$

Consider the following plug-in versions of the quantities appearing in (17), (21):

$$\widetilde{W}_{(\varsigma)} := \left\{ k \in \llbracket B \rrbracket : \frac{\tilde{Z}_k^{(2)}}{N_k} \leq \varsigma \frac{\tilde{Z}_1^{(2)}}{N_1} \right\}, \quad \widetilde{T}_k^{(\tau)} := \mathbf{1} \left\{ \tilde{U}_k \leq \tau \frac{\tilde{Z}_1^{(1)}}{N_1} \right\}. \quad (63)$$

Then, defining

$$\widetilde{V}_{\tau, \varsigma} := \left\{ k \in \widetilde{W}_{(\varsigma)} : \widetilde{T}_k^{(\tau)} = 1 \right\},$$

with probability at least $1 - 3\alpha$ (with respect to the ‘‘tilde’’ data) it holds

$$V_{\tau-, \varsigma/\beta} \subseteq \widetilde{V}_{\tau, \varsigma} \subseteq V_{\tau+, \varsigma\beta}, \quad (64)$$

where $\tau_{\pm} := (1 \pm \eta_1)(\sqrt{\tau} \pm \sqrt{\tau_{\min}^{\circ}})^2_+$, (with $\tau_{\min}^{\circ} = 64(\log 8B\alpha^{-1})/\sqrt{d_1^{\bullet}}$), and $\beta := (1 + \eta_2)/(1 - \eta_2)$.

Proof. Assume that (61)-(62) are satisfied. Then $W_{(\varsigma/\beta)} \subseteq \widetilde{W}_{(\varsigma)} \subseteq W_{(\beta\varsigma)}$, with $\beta := (1 + \eta_2)/(1 - \eta_2)$. Furthermore, recalling $\widetilde{T}_k^{(\tau)} := \mathbf{1}\{\widetilde{U}_k \leq \tau s_1^2\}$, then we have $\widetilde{T}_k^{((1-\eta_1)\tau)} \leq \widetilde{T}_k^{(\tau)} \leq \widetilde{T}_k^{((1+\eta_1)\tau)}$; therefore

$$\left\{k \in W_{(\varsigma/\beta)} : \widetilde{T}_k^{((1-\eta_1)\tau)} = 1\right\} =: \widetilde{V}_- \subseteq \widetilde{V} \subseteq \widetilde{V}_+ := \left\{k \in W_{(\beta\varsigma)} : \widetilde{T}_k^{((1+\eta_1)\tau)} = 1\right\}.$$

We can apply Corollary 1 separately to $\widetilde{V}_-, \widetilde{V}_+$ and get that with probability $1 - 3\alpha$ (accounting for the union bound together with event (61)-(62)), (64) holds.. \square

Proof of Prop. 4. From Proposition 12 with $u = \log(4B\alpha^{-1})$ and a union bound over tasks, with probability at least $1 - \alpha/2$ it holds

$$\forall k \in \llbracket B \rrbracket : \quad \left| \widetilde{Z}_k^{(1)} - \|\Sigma_k\|_1 \right| \leq \|\Sigma_k\|_2 \frac{\sqrt{32} \log(4B\alpha^{-1})}{\sqrt{N_k}} \leq \frac{1}{\sqrt{a}} \|\Sigma_k\|_2, \quad (65)$$

where for the last inequality we used the assumption $N_k \geq a(4 + \log(2B\alpha^{-1}))^4 \geq a(4 + \log 6)^2(4 + \log(2B\alpha^{-1}))^2 \geq 32a(\log(4B\alpha^{-1}))^2$ (also using $\alpha \leq 1/3$ in that estimate).

Similarly, from Proposition 13 with $u = (4 + \log(2B\alpha^{-1}))$, with probability at least $1 - \alpha/2$ it holds

$$\forall k \in \llbracket B \rrbracket : \quad \left| \widetilde{Z}_k^{(2)} - \|\Sigma_k\|_2 \right| \leq 30\|\Sigma_k\|_2 \frac{(4 + \log(2B\alpha^{-1}))^2}{\sqrt{N_k}} \leq \frac{30}{\sqrt{a}} \|\Sigma_k\|_2. \quad (66)$$

Therefore, conditions (61)-(62) are satisfied simultaneously with probability $1 - \alpha$, with $\eta_1 = \frac{1}{\sqrt{a}} \frac{1}{\sqrt{d_1^{\bullet}}}$ and $\eta_2 = \frac{30}{\sqrt{a}}$ (with $a \geq 4400$).

We apply Proposition 14, but using the the values $(\widetilde{\tau}, 3c)$ given by (27) in place of (τ, ς) . As a result we get with high probability the sandwiching property (64),

$$V_{\widetilde{\tau}, c} \subseteq \widetilde{V}_{\widetilde{\tau}, 3c} \subseteq V_{\widetilde{\tau}+}, \quad (67)$$

denoting $\widetilde{\tau}_{\pm}$ the formula for τ_{\pm} of Proposition 14 where we replace (τ, ς) by $(\widetilde{\tau}, 3\varsigma)$. We proceed to get bounds for $\widetilde{\tau}_{\pm} = (1 \pm \eta_1)(\sqrt{\widetilde{\tau}} \pm \sqrt{3\varsigma\tau_{\min}^{\circ}})^2$.

Let us start with bounding the estimation error of d_1^{\bullet} by $\widetilde{d}_1^{\bullet}$: it holds

$$\sqrt{\widetilde{d}_1^{\bullet}} = \frac{N_1 \widetilde{s}_1^2}{\widetilde{Z}_1^{(2)}} = \frac{\widetilde{Z}_1^{(1)}}{\widetilde{Z}_1^{(2)}} \leq \frac{1 + \eta_1}{1 - \eta_2} \sqrt{d_1^{\bullet}} \leq 2\sqrt{d_1^{\bullet}},$$

where the last inequality holds if $a \geq 4400$. We deduce

$$\widetilde{\tau}_{\min}^{\circ} = \frac{32 \log(8B\alpha^{-1})}{\sqrt{\widetilde{d}_1^{\bullet}}} \geq \frac{1}{2} \cdot \frac{32 \log(8B\alpha^{-1})}{\sqrt{d_1^{\bullet}}} = \tau_{\min}^{\circ}/2,$$

as defined in Proposition 14. Furthermore, we have for $\eta_1 = \frac{1}{\sqrt{ad_1^\bullet}} \leq \frac{1}{\sqrt{a}}$ and $a \geq 4400$:

$$\frac{1}{1 - \eta_1} = 1 + \frac{\eta_1}{1 - \eta_1} \leq 1 + \frac{1/\sqrt{a}}{1 - 1/\sqrt{a}} \frac{1}{\sqrt{d_1^\bullet}} \leq 1 + \frac{1}{60\sqrt{d_1^\bullet}}.$$

Using the previous estimates we obtain

$$\tilde{\tau} := \left(1 + \frac{1}{60\sqrt{d_1^\bullet}}\right) \left(\sqrt{\tau} + \sqrt{6\zeta\tilde{\tau}_{\min}^\circ}\right)^2 \geq \frac{1}{1 - \eta_1} \left(\sqrt{\tau} + \sqrt{3\zeta\tau_{\min}^\circ}\right)^2.$$

It follows :

$$\tilde{\tau}_- = (1 - \eta_1) \left(\sqrt{\tilde{\tau}} - \sqrt{3\zeta\tau_{\min}^\circ}\right)^2 \geq \tau.$$

Now to get an upper bound on $\tilde{\tau}_+$, similarly to above we have

$$\sqrt{\tilde{d}_1^\bullet} \geq \frac{1 - \eta_1}{1 + \eta_2} d_1^\bullet \geq \frac{\sqrt{d_1^\bullet}}{2},$$

and thus $\tilde{\tau}_{\min}^\circ \leq 2\tau_{\min}^\circ$; it follows

$$\begin{aligned} \tilde{\tau}_+ &= (1 + \eta_1) \left(\sqrt{\tilde{\tau}} + \sqrt{3\zeta\tau_{\min}^\circ}\right)^2 \leq \left(1 + \frac{1}{66\sqrt{d_1^\bullet}}\right) \left(1 + \frac{1}{30\sqrt{d_1^\bullet}}\right) \left(\sqrt{\tau} + 3\sqrt{3\zeta\tau_{\min}^\circ}\right)^2 \\ &= \xi\tau, \end{aligned}$$

where $\xi := (1 + 1/(30\sqrt{d_1^\bullet}))(1 + 1/(66\sqrt{d_1^\bullet}))(1 + 3\sqrt{3\zeta\tau_{\min}^\circ}/\tau)^2$.

With these estimates in hand the sandwiching property (67) implies

$$V_{\tau,\zeta} \subseteq \tilde{V}_{\tilde{\tau},3\zeta} \subseteq V_{\xi\tau}.$$

We use this property to apply Proposition (1) as earlier, and obtain

$$\frac{R_1(\tilde{\omega})}{s_1^2} \leq \left(\frac{1 + \eta}{1 - \eta}\right) \mathcal{B}(\xi\tau, \nu(V_{\tau,\zeta})) \leq \left(1 + \frac{1}{25\sqrt{\min_k d_k^\bullet}}\right) \xi \mathcal{B}(\tau, \nu(V_{\tau,\zeta})).$$

Elementary estimates lead to

$$\left(1 + \frac{1}{25\sqrt{\min_k d_k^\bullet}}\right) \xi \leq \left(1 + \frac{1}{10\sqrt{\min_k d_k^\bullet}}\right) \left(1 + \frac{30\sqrt{\zeta \log(8B\alpha^{-1})}}{(d_1^\bullet)^{\frac{1}{4}}\sqrt{\tau}}\right)^2.$$

□

C.4.2. *Bounded setting.* Proposition 15 and Proposition 16 give concentration bounds for $\tilde{Z}^{(1)}$ and $\tilde{Z}^{(2)}$ in bounded setting.

Proposition 15. *Assume (BS) holds. For $u \geq 1$, if $N \geq 2$*

$$\mathbb{P} \left[\left| \tilde{Z}^{(1)} - \text{Tr } \Sigma \right| \geq 2\sqrt{2\frac{\text{Var}[\|X_1 - \mu\|^2]}{N}u} + 32\frac{M^2u}{N} \right] \leq 4e^{-u}.$$

Proof. Let us first remark that:

$$\tilde{Z}^{(1)} = \frac{1}{N-1} \sum_{i=1}^N \left\| \tilde{X}_i - \mu \right\|^2 - \frac{N \|\tilde{\mu} - \mu\|^2}{N-1}$$

Using Bernstein's inequality (Lemma 7), with probability greater than $1 - 2e^{-u}$:

$$\left| \sum_{i=1}^N \left\| \tilde{X}_i - \mu \right\|^2 - N \operatorname{Tr} \Sigma \right| \leq \sqrt{2N \operatorname{Var}[\|X_1 - \mu\|^2] u} + 8M^2 u.$$

Using McDiarmid's inequality (Boucheron et al., 2004; McDiarmid, 1998), for $f(x_1, \dots, x_N) = \|N^{-1} \sum_{i=1}^N (x_i - \mu)\|$, with probability greater than $1 - 2e^{-u}$:

$$\begin{aligned} -\frac{4M^2}{N} \leq \|\tilde{\mu} - \mu\|^2 - \frac{\operatorname{Tr} \Sigma}{N} &\leq \left(\mathbb{E}[\|\tilde{\mu} - \mu\|] + \sqrt{\frac{2M^2 u}{N}} \right)^2 - \frac{\operatorname{Tr} \Sigma}{N} \\ &\leq \left(\mathbb{E}[\|\tilde{\mu} - \mu\|^2] - \frac{\operatorname{Tr} \Sigma}{N} \right) + 2\mathbb{E}[\|\tilde{\mu} - \mu\|] \sqrt{\frac{2M^2 u}{N}} + \frac{2M^2 u}{N} \leq 8 \frac{M^2 u}{N}, \end{aligned}$$

where we have used successively Jensen's inequality, that $\operatorname{Tr} \Sigma \leq 4M^2$ and $u \geq 1$. It only stays to use that $(N-1)^{-1} \leq 2N^{-1}$ for $N \geq 2$ and a triangle inequality to conclude the proof, with probability at least $1 - 4e^{-u}$:

$$\begin{aligned} \left| \tilde{Z}^{(1)} - \operatorname{Tr} \Sigma \right| &\leq \frac{\sqrt{2N \operatorname{Var}[\|X_1 - \mu\|^2] u}}{N-1} + \frac{8M^2 u}{N-1} + \frac{8M^2 u}{N-1} \\ &\leq 2\sqrt{2 \frac{\operatorname{Var}[\|X_1 - \mu\|^2]}{N} u} + 32 \frac{M^2 u}{N}. \end{aligned}$$

□

Similarly as in the Gaussian setting, we can estimate $\|\Sigma\|_2$ using the U-statistic (59):

Proposition 16 (Blanchard and Fermanian, 2023, Prop. 13). *Assume (BS) holds and $N \geq 4$. Then for all $u \geq 0$:*

$$\mathbb{P} \left[\left| \tilde{Z}^{(2)} - \sqrt{\operatorname{Tr} \Sigma^2} \right| \geq 12M^2 \sqrt{\frac{u}{N}} \right] \leq 2e^{-u}. \quad (68)$$

Thanks to these concentration results, we are able to give a bound on the estimation error of the test method for bounded data on the model of Proposition 4.

Proposition 17. *Assume (BS) holds. Let $\alpha \in (0, 1/3)$. Consider the set of estimated τ -neighbours $\tilde{V}_{\tau, \varsigma}$ defined in (26), assume $N_k \geq a\phi_k^2 d_k^\bullet \log(8B\alpha^{-1})$ for all $k \in \llbracket B \rrbracket$, for a big enough constant a ($a = 576$ works), and where $\phi_k := M^2 / (\operatorname{Tr} \Sigma_k)$.*

For fixed $\tau > 0$, $\varsigma \geq 1$, consider the weights $\tilde{\omega}^\sharp$ obtained by the modified plug-in $(\tilde{V}_{\tilde{\tau}, 3\varsigma}, \tilde{\mathbf{s}}^2)$ for (V, \mathbf{s}^2) in (13), where

$$\tilde{\tau} := \left(1 + \frac{1}{2\sqrt{d_1^\bullet}}\right) \left(\sqrt{\tau} + \sqrt{6\tilde{\tau}_{\min}^\circ}\right)^2; \quad \tilde{\tau}_{\min}^\circ := \frac{80c_0^2\varsigma(\log(8B\alpha^{-1}))}{\sqrt{d_1^\bullet}}; \quad \sqrt{d_1^\bullet} := \frac{N_1\tilde{s}_1^2}{\tilde{Z}_1^{(2)}}. \quad (69)$$

and $c_0 = 31$. Then with probability at least $1 - 3\alpha$ over the draw of the ‘‘tilde’’ sample $(\tilde{X}_{\bullet}^{(k)})_{k \in \llbracket B \rrbracket}$, it holds

$$\frac{R_1(\tilde{\omega}^\sharp)}{s_1^2} \leq \left(1 + \frac{4}{\sqrt{\min_k d_k^\bullet}}\right) \left(1 + \frac{900\sqrt{\varsigma \log(8B\alpha^{-1})}}{(d_1^\bullet)^{\frac{1}{4}}\sqrt{\tau}}\right)^2 \mathcal{B}(\tau, \nu(V_{\tau, \varsigma})),$$

where the expected risk is with respect to the main sample $(X_{\bullet}^{(k)})_{k \in \llbracket B \rrbracket}$.

Proof of Prop. 17. From Proposition 15 with $u = \log(8B\alpha^{-1})$ and a union bound over tasks, with probability at least $1 - \alpha/2$ it holds

$$\forall k \in \llbracket B \rrbracket : \quad \left| \tilde{Z}_k^{(1)} - \|\Sigma_k\|_1 \right| \leq 2\sqrt{2\frac{\|\Sigma_k\|_1 M^2 u}{N_k}} + 32\frac{M^2 u}{N_k} \leq \frac{1}{3}\|\Sigma_k\|_2 \quad (70)$$

where for the last inequality we used the assumption $N_k \geq 64a\phi_k^2 d_k^\bullet \log(8B\alpha^{-1})$. Similarly, from Proposition 16 with $u = \log(4B\alpha^{-1})$, with probability at least $1 - \alpha/2$ it holds

$$\forall k \in \llbracket B \rrbracket : \quad \left| \tilde{Z}_k^{(2)} - \|\Sigma_k\|_2 \right| \leq 12M^2\sqrt{\frac{u}{N}} \leq \frac{1}{6}\|\Sigma_k\|_2. \quad (71)$$

Therefore, as in the Gaussian case (see proof of Proposition 4), with $\eta_2 = 1/6$ and $\beta = (1 + \eta_2)/(1 - \eta_2) \leq 3$:

$$\tilde{W}_\varsigma \subseteq W_{3\varsigma}$$

Let $\tau_{\min}^\circ = 80c_0^2\varsigma u (d_1^\bullet)^{-1/2}$, one can check that $\tau_{\min}^\circ \geq \tau_{\min}^k$ for all $k \in V_{\tau, \varsigma}$. Indeed, in bounded setting:

$$\tau_{\min}^k \leq 2c_0^2 \left(16u\frac{1+\varsigma}{\sqrt{d_1^\bullet}} + 4u^2 s_1^{-2} \left(\frac{\Phi_1 s_1^2}{N_1} + \frac{\Phi_k s_k^2}{N_k}\right)\right) \leq 2c_0^2 \left(\frac{32\varsigma u}{\sqrt{d_1^\bullet}} + 4u\frac{1+\varsigma}{d_1^\bullet}\right) \leq \frac{80c_0^2\varsigma u}{\sqrt{d_1^\bullet}}$$

where we have used that $\varsigma \geq 1$, the assumption on N_k and the expression of τ_{\min}^k given by Proposition 11. We apply Proposition 11 to $\tilde{\tau}$ defined in (69), then with high probability:

$$V_{\tilde{\tau}, \varsigma} \subseteq \tilde{V}_{\tilde{\tau}, 3\varsigma} \subseteq V_{\tilde{\tau}_\pm}, \quad (72)$$

where $\tilde{\tau}_\pm = (1 \pm \eta_1)\left(\sqrt{\tilde{\tau}} \pm \sqrt{\tau_{\min}^\circ}\right)$. We proceed to get bounds for $\tilde{\tau}_\pm$.

Let us start with bounding the estimation error of d_1^\bullet by \tilde{d}_1^\bullet : it holds

$$\sqrt{\tilde{d}_1^\bullet} = \frac{N_1 \tilde{s}_1^2}{\tilde{Z}_1^{(2)}} = \frac{\tilde{Z}_1^{(1)}}{\tilde{Z}_1^{(2)}} \leq \frac{1 + \eta_1}{1 - \eta_2} \sqrt{d_1^\bullet} \leq 2\sqrt{d_1^\bullet},$$

where $\eta_1 = (d_1^\bullet)^{-1/2}/3 \leq 1/3$. We deduce

$$\tilde{\tau}_{\min}^\circ = \frac{80c_0^2 \zeta u}{\sqrt{\tilde{d}_1^\bullet}} \geq \frac{1}{2} \cdot \frac{80c_0^2 \zeta u}{\sqrt{d_1^\bullet}} = \tau_{\min}^\circ/2,$$

Furthermore, as $\eta_1 \leq 1/3$:

$$\frac{1}{1 - \eta_1} = 1 + \frac{\eta_1}{1 - \eta_1} \leq 1 + \frac{1}{2\sqrt{d_1^\bullet}}.$$

Using the previous estimates we obtain

$$\tilde{\tau} := \left(1 + \frac{1}{2\sqrt{\tilde{d}_1^\bullet}}\right) \left(\sqrt{\tau} + \sqrt{6\zeta \tilde{\tau}_{\min}^\circ}\right)^2 \geq \frac{1}{1 - \eta_1} \left(\sqrt{\tau} + \sqrt{3\zeta \tau_{\min}^\circ}\right)^2.$$

It follows :

$$\tilde{\tau}_- = (1 - \eta_1) \left(\sqrt{\tilde{\tau}} - \sqrt{3\zeta \tau_{\min}^\circ}\right)^2 \geq \tau.$$

Now to get an upper bound on $\tilde{\tau}_+$, similarly to above we have

$$\sqrt{\tilde{d}_1^\bullet} \geq \frac{1 - \eta_1}{1 + \eta_2} d_1^\bullet \geq \frac{\sqrt{d_1^\bullet}}{2},$$

and thus $\tilde{\tau}_{\min}^\circ \leq 2\tau_{\min}^\circ$; it follows

$$\begin{aligned} \tilde{\tau}_+ &= (1 + \eta_1) \left(\sqrt{\tilde{\tau}} + \sqrt{3\zeta \tau_{\min}^\circ}\right)^2 \leq \left(1 + \frac{1}{3\sqrt{d_1^\bullet}}\right) \left(1 + \frac{1}{\sqrt{d_1^\bullet}}\right) \left(\sqrt{\tau} + 3\sqrt{3\zeta \tau_{\min}^\circ}\right)^2 \\ &= \xi \tau, \end{aligned}$$

where $\xi := (1 + 1/(3\sqrt{d_1^\bullet}))(1 + 1/\sqrt{d_1^\bullet})(1 + 3\sqrt{3\zeta \tau_{\min}^\circ}/\tau)^2$.

With these estimates in hand the sandwiching property (72) implies

$$V_{\tau, \zeta} \subseteq \tilde{V}_{\tilde{\tau}, 3\zeta} \subseteq V_{\xi \tau}.$$

We use this property to apply Proposition (1) as earlier, and obtain

$$\frac{R_1(\tilde{\omega})}{s_1^2} \leq \left(\frac{1 + \eta}{1 - \eta}\right) \mathcal{B}(\xi \tau, \nu(V_{\tau, \zeta})) \leq \left(1 + \frac{1}{2\sqrt{\min_k d_k^\bullet}}\right) \xi \mathcal{B}(\tau, \nu(V_{\tau, \zeta})).$$

Elementary estimates lead to

$$\left(1 + \frac{1}{2\sqrt{\min_k d_k^\bullet}}\right) \xi \leq \left(1 + \frac{4}{\sqrt{\min_k d_k^\bullet}}\right) \left(1 + \frac{900\sqrt{\zeta \log(8B\alpha^{-1})}}{(d_1^\bullet)^{1/4} \sqrt{\tau}}\right)^2.$$

□

C.4.3. *Heavy-tailed setting.* Similarly as in Sup. C.4.1 and C.4.2, we provide in this section estimators of $\|\Delta_k\|^2$, $\|\Sigma_k\|_1$ and $\|\Sigma_k\|_2$ but for heavy-tailed data. These estimators can be directly used to estimate the neighbours $V_{\tau, \varsigma}$ and the oracle weights to then apply the testing approach in this setting.

Assumption (HT, Heavy-tailed setting). For all $k \in \llbracket B \rrbracket$, \mathbb{P}_k has a finite fourth moment.

Consider a statistic $T(N; x_1, \dots, x_N)$ in \mathbb{R} , the Median of Blocks statistics $\text{MOB}_b(T)$ for b a divisor of N is defined by the median of the statistics T^a , $1 \leq a \leq b$ built from a b -partition of x_1, \dots, x_N :

$$\text{MOB}_k(T) := \text{Median}(T^a, 1 \leq a \leq b)$$

where $T^a = T(N/b; x_{aN/b+1}, \dots, x_{(a+1)N/b})$. If b does not divide N , it suffices to partition the sample into sub-samples of size $\lfloor N/b \rfloor$ and $\lceil N/b \rceil$. If the original estimator is constructed from different samples (e.g., (16)), each sample is partitioned into b subsamples.

Proposition 18. *Assume (HT) holds, let $0 \leq u \leq N$ and $b = \lceil u \rceil$, let $U(X_{\bullet}^{(1)}, X_{\bullet}^{(k)})$ the estimator of $\|\Delta_k\|^2$ defined in (16), then, with probability greater than $1 - e^{-u/8}$:*

$$|\text{MOB}_b(U(X_{\bullet}^{(1)}, X_{\bullet}^{(k)})) - \|\Delta_k\|^2| \leq 4\sqrt{\Delta_k^T \left(\frac{\Sigma_1}{N_1} + \frac{\Sigma_k}{N_k} \right) \Delta_k} u + 4 \left(\frac{\|\Sigma_1\|_2}{N_1} + \frac{\|\Sigma_k\|_2}{N_k} \right) u. \quad (73)$$

In the kernel setting, the statistic $U(X_{\bullet}^{(1)}, X_{\bullet}^{(k)})$ is an estimator of the MMD distance between \mathbb{P}_1 and \mathbb{P}_k . (Lerasle et al., 2019 proposed a different robust estimator of this quantity called MONK, but we focus here on the MOB estimator, which has the advantage to be easier to compute and to study.)

Proposition 19. *Assume (HT) holds, let $0 \leq u \leq N/4$ and $b = \lceil u \rceil$:*

$$\mathbb{P} \left[\left| \text{MOB}_b(Z^{(1)}) - \text{Tr} \Sigma \right| \geq C \sqrt{\frac{\text{Var}[\|X_1 - \mu\|^2]u}{N}} + C \frac{\sqrt{\text{Tr} \Sigma^2 u}}{N} \right] \leq e^{-u/8},$$

$$\mathbb{P} \left[\left| \sqrt{\text{MOB}_b(Z^{(2)})} - \sqrt{\text{Tr} \Sigma^2} \right| \geq C \sqrt{\frac{M_X u}{N}} \right] \leq e^{-u/8},$$

where $Z^{(1)}$ is defined in (58), $Z^{(2)}$ in (59), $C > 0$ is an absolute constant and $M_X = \mathbb{E}[\|X_1 - \mu\|^4]$.

Proposition 18 and Proposition 19 are different consequences of Lemma 2 below. Some more refined concentration bounds can be derived for MOB-type statistics (see, e.g., Devroye et al., 2016; Minsker and Strawn, 2019), but the present results are sufficient to show that in the (HT) setting suitable statistics satisfy Assumption (TSC)

and (61)-(62).

Proof of Proposition 18. According to Lemma 2, we only need compute the variances of the statistics \tilde{U}_a ,

$$\begin{aligned} \text{Var}[\tilde{U}_k] &= 4\sqrt{\Delta_k^T \left(\frac{\Sigma_1}{N_1} + \frac{\Sigma_k}{N_k} \right) \Delta_k} + 2 \text{Tr} \left(\frac{\Sigma_1}{N_1} + \frac{\Sigma_k}{N_k} \right)^2 + 2 \left(\frac{\|\Sigma_1\|_2}{N_1^2(N_1-1)} + \frac{\|\Sigma_k\|_2}{N_k^2(N_k-1)} \right) \\ &\leq 4\sqrt{\Delta_k^T \left(\frac{\Sigma_1}{N_1} + \frac{\Sigma_i}{N_i} \right) \Delta_k} + 4 \left(\frac{\|\Sigma_1\|_2}{N_1} + \frac{\|\Sigma_k\|_2}{N_k} \right) =: \tilde{v}(N_1, N_i) \end{aligned}$$

We apply Lemma 2 with $N = N_1 + N_i$ and $v(N/u) := \tilde{v}(N_1/u, N_k/u)$. \square

Proof of Proposition 19.

For $Z^{(1)}$ the concentration bound is deduced directly from the variance:

$$\text{Var}[Z^{(1)}] = \frac{\text{Var}[\|X - \mu\|^2]}{N} + \frac{2\|\Sigma\|_2^2}{N(N-1)}.$$

For $Z^{(2)}$ we can first assume w.l.g. that X is centred. Then $Z^{(2)}$ can be developed as:

$$(Z^{(2)})^2 = \frac{1}{N^{(2)}} \sum_{i \neq j} \langle X_i, X_j \rangle^2 - \frac{2}{N^{(3)}} \sum_{i \neq j \neq k} \langle X_i, X_j \rangle \langle X_i, X_k \rangle - \frac{1}{N^{(4)}} \sum_{i \neq j \neq k \neq q} \langle X_i, X_j \rangle \langle X_k, X_q \rangle.$$

where $n^{(p)} = n(n-1)\dots(n-p+1)$ for $n \geq p \in \mathbb{N}$. Let us first compute $\text{Var}[(Z^{(2)})^2]$:

$$\begin{aligned} \text{Var}[(Z^{(2)})^2] &\leq \frac{2}{N^{(2)}} \mathbb{E}[\langle X, X' \rangle^4] + \frac{4(N-2)}{N^{(2)}} \mathbb{E}[(X^T \Sigma X)^2] \\ &\quad + \frac{4}{N^{(3)}} ((3!)M_X^2 + 2(N-3) \text{Tr} \Sigma^4) + \frac{4!}{N^{(4)}} M_X^2 \\ &\leq C \frac{\|\Sigma\|_\infty^2 M_X}{N} + C \frac{M_X^2}{N^2} \end{aligned}$$

where $C > 0$ is some absolute constant. Then according to Lemma 2, for $u \leq N/4$, with probability grater than $1 - e^{-u/8}$:

$$|\text{MOB}_b((Z^{(2)})^2) - \text{Tr} \Sigma^2| \leq C \|\Sigma\|_\infty \sqrt{\frac{M_X u}{N}} + C \frac{M_X u}{N}, \quad (74)$$

Using that $|\sqrt{(a^2 + b)_+} - a| \leq \min(\sqrt{|b|}, \frac{b}{a})$ for $a \in \mathbb{R}_+$ and $b \in \mathbb{R}$, (see, e.g., Lemma 15 of Blanchard and Fermanian, 2023), assuming (74), then

$$\begin{aligned} \left| \text{MOB}_b(Z^{(2)}) - \sqrt{\text{Tr} \Sigma^2} \right| &\leq \max_{\varepsilon \in \{-1, 1\}} \left| \sqrt{\text{Tr} \Sigma^2 + \varepsilon C \|\Sigma\|_\infty \sqrt{\frac{M_X u}{N}}} - \sqrt{\text{Tr} \Sigma^2} \right| + C \sqrt{\frac{M_X u}{N}} \\ &\leq C \frac{\|\Sigma\|_\infty}{\sqrt{\text{Tr} \Sigma^2}} \sqrt{\frac{M_X u}{N}} + C \sqrt{\frac{M_X u}{N}} \leq C \sqrt{\frac{M_X u}{N}}. \end{aligned}$$

□

Lemma 2. Let $T(N; x_1, \dots, x_N)$ a statistic build from N i.i.d. random variables such that for all $N \geq N_0$:

$$\mathbb{E}[T(N; X_1, \dots, X_N)] = \mathbb{E}[T], \quad \text{Var}[T(N; X_1, \dots, X_N)] \leq v(N),$$

where $v : \mathbb{R}_+ \rightarrow \mathbb{R}_+$ is nonincreasing. Let $1 \leq u \leq N/(N_0 + 1)$ and $b = \lceil u \rceil$, then

$$\mathbb{P} \left[|\text{MOB}_b(T) - \mathbb{E}[T]| \geq \sqrt{4v \left(\frac{N}{4u} \right)} \right] \leq e^{-u/8}.$$

Proof of Lemma 2.

First assume that $b|N$. Let us denote $T_a := T(N/b; x_{(a-1)N/b+1}, \dots, x_{aN/b})$ for $a \in \llbracket b \rrbracket$. Then for all $a \in \llbracket b \rrbracket$, by Markov's inequality:

$$\mathbb{P} \left[|T_a - \mathbb{E}[T]| \geq \sqrt{4v(N/b)} \right] \leq \frac{1}{4}. \quad (75)$$

Then, $|\text{MOB}_b(T) - \mathbb{E}[T]| \geq \sqrt{4v(N/b)}$ implies that at least $b/2$ of T_a satisfies

$$|T_a - \mathbb{E}[T]| \geq \sqrt{4v(N/b)}.$$

By independence of the T_a and Hoeffding's inequality:

$$\mathbb{P} \left[|\text{MOB}_b(T) - \mathbb{E}[T]| > \sqrt{4v(N/b)} \right] \leq \mathbb{P} \left[\text{Bin} \left(b, \frac{1}{4} \right) \geq \frac{b}{2} \right] \leq e^{-b/8},$$

where Bin denotes the Binomial distribution. Because $u \leq b \leq u + 1$ and v is a nonincreasing function, we can conclude:

$$e^{-b/8} \leq e^{-u/8}, \quad v \left(\frac{N}{b} \right) \leq v \left(\frac{N}{u+1} \right) \leq v \left(\frac{N}{4u} \right).$$

If $b \nmid N$, equation (75) is still verified with $v(\lfloor \frac{N}{b} \rfloor)$ instead of $v(\frac{N}{b})$ and:

$$\begin{aligned} \left\lfloor \frac{N}{\lceil u \rceil} \right\rfloor &\geq \frac{N}{\lceil u \rceil} - 1 \geq \frac{N}{2\lceil u \rceil} && \text{if } \lceil u \rceil \leq N/2 \\ \left\lfloor \frac{N}{\lceil u \rceil} \right\rfloor &= 1 \geq \frac{N}{2\lceil u \rceil} && \text{if } N \geq \lceil u \rceil > N/2. \end{aligned}$$

We conclude using that $\lceil u \rceil \leq (u + 1) \leq 2u$ for $u \geq 1$. □

APPENDIX D. PROOFS FOR SECTION 4

D.1. Proof of Proposition 5. Let $\hat{\omega} \in \arg \min_{\omega \in S_B} (\widehat{L}_1(\omega) + 16\sqrt{u_0}\widehat{Q}_1(\omega))$. Denote $\mathcal{X}^{-1} = (X_{\bullet}^{(k)})_{k \neq 1}$ the observed bag data except for the first bag, which corresponds to the target task.

First step : bound in conditional probability. As a first step, we obtain a high-probability bound for $L_1(\widehat{\omega})$. For $x \geq 1$, define the event $A(x)$:

$$A(x) := \left\{ \begin{array}{ll} \sqrt{\mathbf{q}_k} \leq c_1(x) \sqrt{\widehat{q}_k} + C \frac{s_1^2}{d_1^e} \sqrt{N_1 x}, & 2 \leq k \leq B, \quad (\text{a}) \\ \sqrt{\widehat{q}_k} \leq \left(1 + \sqrt{\frac{2x}{N_1 - 1}}\right) \left(\sqrt{\mathbf{q}_k} + \frac{s_1^4}{d_1^e} N_1 + \frac{s_1^2}{d_1^e} \sqrt{2N_1 x}\right), & 2 \leq k \leq B, \quad (\text{b}) \\ |\widehat{s}_1^2 - s_1^2| \leq C \frac{s_1^2}{\sqrt{d_1^e N_1}} x, & (\text{c}) \\ \|\widehat{\mu}_1^{\text{NE}} - \mu_1\|^2 \leq s_1^2 + C \frac{s_1^2}{\sqrt{d_1^e}} x, & (\text{d}) \\ |\langle \widehat{\mu}_k^{\text{NE}} - \mu_1, \widehat{\mu}_1^{\text{NE}} - \mu_1 \rangle| \leq \sqrt{2 \frac{\mathbf{q}_k}{N_1} x}, & 2 \leq k \leq B, \quad (\text{e}) \end{array} \right.$$

where $\mathbf{q}_k = (\widehat{\mu}_k^{\text{NE}} - \mu_1)^T \Sigma_1 (\widehat{\mu}_k^{\text{NE}} - \mu_1)$ and $c_1(x) = \sqrt{e} \exp(x/(N_1 - 1))$. For the whole proof, the notation C will denote an absolute numeric constant whose value can change between lines. The probability of the event A conditionally to $\mathcal{X}^{(-1)}$ is bounded as:

$$\mathbb{P}[A^c(x, y) | \mathcal{X}^{(-1)}] \leq (6B + 4)e^{-x}. \quad (76)$$

We combine a union bound with estimates for each individual bound: bounds (a) and (b) are consequences of Proposition 20 with $\nu = \widehat{\mu}_k^{\text{NE}}$. For (a), we have used that $\sqrt{\mathbf{q}_k} \leq \sqrt{\mathbf{q}_k + \text{Tr} \Sigma_1^2 / N_1}$. Bound (c) is a rewriting of Proposition 12. Bound (d) is a consequence of Lemma 5 with $X = \widehat{\mu}_1^{\text{NE}} - \mu_1$, $\mu = 0$, $\Sigma = \Sigma_1 / N_1$; bounding \sqrt{x} by x , and $\|\Sigma_1\|_\infty$ by $\sqrt{\text{Tr} \Sigma_1^2}$. Finally (e) is deduced from Lemma 4 with $X = \langle \widehat{\mu}_k^{\text{NE}} - \mu_1, \widehat{\mu}_1^{\text{NE}} - \mu_1 \rangle$, $m = 0$ and $\sigma^2 = \mathbf{q}_k$.

From now on, assume that event $A(x)$ holds. Then,

$$\begin{aligned} L_1(\widehat{\omega}) &= \left\| \sum_{k=1}^B \widehat{\omega}_k (\widehat{\mu}_k^{\text{NE}} - \widehat{\mu}_1^{\text{NE}}) + (\widehat{\mu}_1^{\text{NE}} - \mu_1) \right\|^2 \\ &= \left\| \sum_{k=2}^B \widehat{\omega}_k (\widehat{\mu}_k^{\text{NE}} - \widehat{\mu}_1^{\text{NE}}) \right\|^2 + 2 \sum_{k=2}^B \widehat{\omega}_k \langle \widehat{\mu}_k^{\text{NE}} - \widehat{\mu}_1^{\text{NE}}, \widehat{\mu}_1^{\text{NE}} - \mu_1 \rangle + \|\widehat{\mu}_1^{\text{NE}} - \mu_1\|^2 \\ &= \left\| \sum_{k=2}^B \widehat{\omega}_k (\widehat{\mu}_k^{\text{NE}} - \widehat{\mu}_1^{\text{NE}}) \right\|^2 + 2 \sum_{k=2}^B \widehat{\omega}_k \langle \widehat{\mu}_k^{\text{NE}} - \mu_1, \widehat{\mu}_1^{\text{NE}} - \mu_1 \rangle + (2\widehat{\omega}_1 - 1) \|\widehat{\mu}_1^{\text{NE}} - \mu_1\|^2 \\ &= \widehat{L}_1(\widehat{\omega}) + 2 \sum_{k=2}^B \widehat{\omega}_k \langle \widehat{\mu}_k^{\text{NE}} - \mu_1, \widehat{\mu}_1^{\text{NE}} - \mu_1 \rangle + (2\widehat{\omega}_1 - 1) (\|\widehat{\mu}_1^{\text{NE}} - \mu_1\|^2 - s_1^2) \\ &\quad + (2\widehat{\omega}_1 - 1) (s_1^2 - \widehat{s}_1^2) \end{aligned}$$

Using (e) and then (a) for the second term, (d) for the third and (c) for the last, we get:

$$\begin{aligned} L_1(\widehat{\boldsymbol{\omega}}) &\leq \widehat{L}_1(\widehat{\boldsymbol{\omega}}) + 2c_1(x)\sqrt{2x} \sum_{k=2}^B \widehat{\omega}_k \sqrt{\frac{\widehat{q}_k}{N_1}} + Cs_1^2 \left(\frac{x}{\sqrt{d_1^\bullet}} + \frac{x}{d_1^e} \right) \\ &\leq \left(1 \vee \frac{c_1(x)\sqrt{2x}}{8\sqrt{u_0}} \right) \min_{\boldsymbol{\omega} \in \mathcal{S}_B} \left(\widehat{L}_1(\boldsymbol{\omega}) + 16\sqrt{u_0} \sum_{k=2}^B \omega_k \sqrt{\frac{\widehat{q}_k}{N_1}} \right) + Cs_1^2 \frac{x}{\sqrt{d_1^\bullet}}. \end{aligned} \quad (77)$$

The appearance of the minimum is a consequence of the definition of $\widehat{\boldsymbol{\omega}}$.

Second step : conditional bound in expectation. We can now deduce, from the previous step, a bound in expectation conditionally to all samples except the first one. For any fixed $\boldsymbol{\omega} \in \mathcal{S}_B$, we first want to compare $\widehat{L}_1(\boldsymbol{\omega})$ to its conditional expectation $\mathbb{E}[\widehat{L}_1(\boldsymbol{\omega})|\mathcal{X}^{(-1)}]$ which is equal to the conditional expectation of the loss L_1 :

$$\mathbb{E}[\widehat{L}_1(\boldsymbol{\omega})|\mathcal{X}^{(-1)}] = \left\| \sum_{k=2}^B \omega_k (\widehat{\mu}_k^{\text{NE}} - \mu_1) \right\|^2 + \omega_1^2 s_1^2 = \mathbb{E}[L_1(\boldsymbol{\omega})|\mathcal{X}^{(-1)}].$$

For any fixed $\boldsymbol{\omega} \in \mathcal{S}_B$, as $x \geq 1$:

$$\begin{aligned} \widehat{L}_1(\boldsymbol{\omega}) &= \left\| \sum_{k=2}^B \omega_k (\widehat{\mu}_k^{\text{NE}} - \widehat{\mu}_1^{\text{NE}}) \right\|^2 + (2\omega_1 - 1)\widehat{s}_1^2 \\ &= \left\| \sum_{k=2}^B \omega_k (\widehat{\mu}_k^{\text{NE}} - \mu_1) + (1 - \omega_1)(\mu_1 - \widehat{\mu}_1^{\text{NE}}) \right\|^2 + (2\omega_1 - 1)\widehat{s}_1^2 \\ &= \mathbb{E}[L_1(\boldsymbol{\omega})|\mathcal{X}^{(-1)}] + 2(1 - \omega_1) \sum_{k=2}^B \omega_k \langle \widehat{\mu}_k^{\text{NE}} - \mu_1, \mu_1 - \widehat{\mu}_1^{\text{NE}} \rangle \\ &\quad + (1 - \omega_1)^2 (\|\widehat{\mu}_1^{\text{NE}} - \mu_1\|^2 - s_1^2) + (2\omega_1 - 1)(\widehat{s}_1^2 - s_1^2) \\ &\leq \mathbb{E}[L_1(\boldsymbol{\omega})|\mathcal{X}^{(-1)}] + 2\sqrt{2x} \sum_{k=2}^B \omega_k \sqrt{\frac{\mathbf{q}_k}{N_1}} + C \frac{s_1^2 x}{\sqrt{d_1^\bullet}}, \end{aligned} \quad (78)$$

using (c), (d), (e) again. From (b), for all $k \in \llbracket B \rrbracket$, and using again $x \geq 1$:

$$\begin{aligned} \sqrt{\widehat{q}_k} &\leq \left(1 + \sqrt{\frac{2x}{N_1 - 1}} \right) \left(\sqrt{\mathbf{q}_k + \frac{s_1^4}{d_1^\bullet} N_1} + \frac{s_1^2}{d_1^e} \sqrt{2N_1 x} \right) \\ &\leq \left(1 + \sqrt{\frac{2x}{N_1 - 1}} \right) \sqrt{\mathbf{q}_k} + C \left(\sqrt{x} + \frac{x}{\sqrt{N_1 - 1}} \right) \sqrt{\frac{N_1}{d_1^\bullet}} s_1^2. \end{aligned} \quad (79)$$

Then, plugging (78) and (79) into (77), for all $\omega \in \mathcal{S}_B$, as $x \geq 1$:

$$\begin{aligned} L_1(\widehat{\omega}) &\leq \left(1 \vee \frac{c_1(x)\sqrt{2x}}{8\sqrt{u_0}}\right) \left[\mathbb{E}[L_1(\omega)|\mathcal{X}^{(-1)}] \right. \\ &\quad + \left(2\sqrt{2x} + 16\sqrt{u_0}\left(1 + \sqrt{\frac{2x}{N_1-1}}\right)\right) \sum_{k=2}^B \omega_k \sqrt{\frac{\mathbf{q}_k}{N_1}} \\ &\quad \left. + \frac{Cs_1^2}{\sqrt{d_1^\bullet}} \left(x + C\sqrt{u_0}\left(\sqrt{x} + \frac{x}{\sqrt{N_1-1}}\right)\right) \right]. \end{aligned}$$

By rearranging the terms and using that $u_0 \leq N_1 - 1$ and $x \geq 1$:

$$\begin{aligned} L_1(\widehat{\omega}) &\leq \left(1 \vee \frac{c_1(x)\sqrt{2x}}{8\sqrt{u_0}}\right) \left[\mathbb{E}[L_1(\omega)|\mathcal{X}^{(-1)}] \right. \\ &\quad \left. + C(\sqrt{u_0} + \sqrt{x}) \sum_{k=2}^B \omega_k \sqrt{\frac{\mathbf{q}_k}{N_1}} + \frac{Cs_1^2}{\sqrt{d_1^\bullet}}(\sqrt{u_0x} + x) \right] =: \psi(x)P(x) \end{aligned}$$

where $\psi(x) := 1 \vee \frac{c_1(x)\sqrt{2x}}{8\sqrt{u_0}}$ and P is a degree 2 polynomial in \sqrt{x} with coefficients that are constant conditionally to $\mathcal{X}^{(-1)}$. We will denote the shifted version of ψ and P by ψ_s and P_s , for $v \geq 0$:

$$\psi_s(v) := \psi(v + \log(6B + 4)), \quad P_s(v) = P(v + \log(6B + 4)). \quad (80)$$

Both notations will be used depending on the case for the sake of readability. So for all $v \geq 0$

$$\mathbb{P}[L_1(\widehat{\omega}) \geq \psi_s(v)P_s(v)|\mathcal{X}^{(-1)}] \leq e^{-v}.$$

thanks to (76) after taking $x = v + \log(6B + 4) \geq 1$. Then there exists a random variable ξ following an exponential distribution of parameter 1 conditionally to $\mathcal{X}^{(-1)}$, such that $L_1(\widehat{\omega}) \leq \psi_s(\xi)P_s(\xi)$ almost surely. Let us first simplify the expression of ψ , recalling that by assumption $(N_1 - 1)/2 \geq u_0 \geq \log(17B) \geq 1/2 + \log(6B + 4) \geq 1/2 + \log(10) \geq 5/2$, then for $x \leq u_0$:

$$\sqrt{2}c_1(x) \leq \sqrt{2} \exp\left(\frac{1}{2} + \frac{u_0}{N_1 - 1}\right) \leq \sqrt{2} \exp\left(\frac{1}{2} + \frac{1}{2}\right) \leq \sqrt{2}e \leq 8.$$

Thus, for $x \leq u_0$, $\psi(x) = 1$. For $x \geq u_0$, it holds $c_1(x) \geq \sqrt{e} \geq 1$, so that:

$$\begin{aligned} \psi(x) &\leq \frac{c_1(x)\sqrt{x}}{\sqrt{u_0}} \leq \exp\left(\frac{1}{2} + \frac{x - u_0}{N_1 - 1} + \frac{u_0}{N_1 - 1}\right) \sqrt{\frac{x}{u_0}} \\ &\leq e \exp\left(\frac{x - \log(6B + 4)}{5}\right) \sqrt{\frac{x}{u_0}}. \end{aligned} \quad (81)$$

We can now bound the conditional expectation $\mathbb{E}[L_1(\widehat{\omega})|\mathcal{X}^{(-1)}]$ separating the values before and after u_0 :

$$\begin{aligned} \mathbb{E}[L_1(\widehat{\omega})|\mathcal{X}^{(-1)}] &\leq \mathbb{E}[\psi_s(\xi)P_s(\xi)|\mathcal{X}^{(-1)}] \\ &= \mathbb{E}[\psi_s(\xi)P_s(\xi)(\mathbf{1}_{\xi+\log(6B+4)\leq u_0} + \mathbf{1}_{\xi+\log(6B+4)>u_0})|\mathcal{X}^{(-1)}] \\ &\leq P_s(u_0 - \log(6B + 4)) + \mathbb{E}[\psi_s(\xi)P_s(\xi)\mathbf{1}_{\xi+\log(6B+4)>u_0}|\mathcal{X}^{(-1)}] \\ &\leq P(u_0) + \mathbb{E}\left[e \exp(\xi/5) \sqrt{\frac{\xi + \log(6B + 4)}{u_0}} P_s(\xi) \mathbf{1}_{\xi+\log(6B+4)>u_0} |\mathcal{X}^{(-1)}\right]. \end{aligned} \quad (82)$$

We have used that P (and P_s) is increasing on \mathbb{R}_+ (P is a polynomial with positive coefficients) and the bound (81). The second term in (82) can be upper bounded using Lemma 3, as $\sqrt{\xi + \log(6B + 4)}P_s(\xi)$ can be seen as a polynomial of degree 3 evaluated in $\sqrt{\xi + \log(6B + 4)}$. We apply (85) to this polynomial with $a = \log(6B + 4)$, $\delta = u_0 - \log(6B + 4)$, $\rho = 1/5$, $d = 3$ and $\gamma = 1/2$. As $a \geq \log(10) \geq 2$ and $\delta \geq 1/2$, the condition required by Lemma 85 is satisfied: $(\delta + a)(1 - \rho) \geq 2 \geq 3/2 = \gamma d$. Then it holds:

$$\begin{aligned} &\mathbb{E}\left[\exp(\xi/5) \sqrt{\frac{\xi + \log(6B + 4)}{u_0}} P_s(\xi) \mathbf{1}_{\xi+\log(6B+4)>u_0} |\mathcal{X}^{(-1)}\right] \\ &\leq C \sqrt{\frac{u_0}{u_0}} P_s(u_0 - \log(6B + 4)) e^{-(4/5)(u_0 - \log(6B+4))} \leq CP(u_0) B e^{-u_0/2}. \end{aligned} \quad (83)$$

Combining (82) and (83) and replacing $P(u_0)$ by its value, we obtain:

$$\mathbb{E}[L_1(\widehat{\omega})|\mathcal{X}^{(-1)}] \leq \mathbb{E}[L_1(\omega)|\mathcal{X}^{(-1)}](1 + CB e^{-u_0/2}) + C\sqrt{u_0} \sum_{k=2}^B \omega_k \sqrt{\frac{\mathbf{q}_k}{N_1}} + Cs_1^2 \frac{u_0}{\sqrt{d_1^\bullet}}.$$

Third step : unconditional bound. We now simply take the expectation with respect to $\mathcal{X}^{(-1)}$. From the previous bound, using Jensen's inequality, for all $\omega \in \mathcal{S}_B$:

$$\mathbb{E}[L_1(\widehat{\omega})] \leq \mathbb{E}[L_1(\omega)](1 + CB e^{-u_0/2}) + C\sqrt{u_0} \sum_{k=2}^B \omega_k \sqrt{\frac{\mathbb{E}[\mathbf{q}_k]}{N_1}} + Cs_1^2 \frac{u_0}{\sqrt{d_1^\bullet}}.$$

We obtain (33) as $\mathbb{E}[\mathbf{q}_k] = q_k$. □

Lemma 3. Let $\xi \sim \mathcal{E}(1)$ be an exponential random variable, and ρ, a, δ be positive real numbers. Then for all $p \geq 0$ such that $p < (\delta + a)(1 - \rho)$, it holds:

$$\mathbb{E}[(\xi + a)^p e^{\rho\xi} \mathbf{1}_{\xi \geq \delta}] \leq \left(1 - \rho - \frac{p}{a + \delta}\right)^{-1} (\delta + a)^p e^{-\delta(1-\rho)}. \quad (84)$$

Let P a polynomial of degree d and $\gamma > 0$ such that $\gamma d < (\delta + a)(1 - \rho)$, then:

$$\mathbb{E}[P((\xi + a)^\gamma)e^{\rho\xi}1_{\xi \geq \delta}] \leq \left(1 - \rho - \frac{d\gamma}{a + \delta}\right)^{-1} P((\delta + a)^\gamma)e^{-\delta(1-\rho)}. \quad (85)$$

Proof. As $p < (\delta + a)(1 - \rho)$, then $p < (\delta + a)(1 - \rho - \varepsilon)$ for all $\varepsilon < 1 - \rho - p/(a + \delta)$. The function $x \mapsto F(x) := (x + a)^p e^{(\rho - (1-\varepsilon))x}$ on \mathbb{R}_+ attains its maximum in $x_* := p(1 - \rho - \varepsilon)^{-1} - a$ and then decreases to 0. As $x_* < \delta$, we have $F(x) \leq F(\delta)$ for all $x \geq \delta$, thus:

$$\mathbb{E}[(\xi + a)^p e^{\rho\xi} 1_{\xi \geq \delta}] = \mathbb{E}[F(\xi)e^{(1-\varepsilon)\xi} 1_{\xi \geq \delta}] \leq F(\delta)\mathbb{E}[e^{(1-\varepsilon)\xi} 1_{\xi \geq \delta}] = (\delta + a)^p e^{-(1-\rho)\delta} \varepsilon^{-1}.$$

As the inequality is true for all $\varepsilon < 1 - \rho - p/(a + \delta)$ we get (84). Equation (85) is obtained by applying (84) to each of the monomials of degree $k \leq d$ as $k\gamma \leq d\gamma < (\delta + a)(1 - \rho)$, upper bounding the first factor and summing. \square

D.2. Proofs of Corollary 3 and Corollary 2.

Proof of Corollary 3. According to Proposition 5, for $B = 2$, $\mu_2 = 0$ and $\Sigma_2 = 0$; for all $\omega_1 \in (0, 1)$:

$$R_1(\hat{\omega}) \leq ((1 - \omega_1)^2 \|\mu_1\|^2 + \omega_1 s_1^2 + 2(1 - \omega_1)\eta)(1 + Ce^{-u_0/2}) + Cs_1^2 \sqrt{\frac{u_0}{d_1^e}},$$

where $\eta = C \frac{\|\mu_1\|s_1}{\sqrt{d_1^e}} \sqrt{u_0}$. Let us choose $\omega_1 = \min\left(\frac{\|\mu_1\|^2 + \eta}{\|\mu_1\|^2 + s_1^2}, 1\right)$. Then if $\eta \leq s_1^2$:

$$\begin{aligned} R_1(\hat{\omega}) &\leq (1 + Ce^{-u_0/2}) \frac{\|\mu_1\|^2 s_1^2 + 2s_1^2 \eta - \eta^2}{\|\mu_1\|^2 + s_1^2} + Cs_1^2 \sqrt{\frac{u_0}{d_1^e}} \\ &\leq (1 + Ce^{-u_0/2}) \frac{\|\mu_1\|^2 s_1^2}{\|\mu_1\|^2 + s_1^2} + Cs_1^2 \sqrt{\frac{u_0}{d_1^e}} \frac{2\|\mu_1\|s_1}{\|\mu_1\|^2 + s_1^2} + Cs_1^2 \sqrt{\frac{u_0}{d_1^e}} \\ &\leq (1 + Ce^{-u_0/2}) \frac{\|\mu_1\|^2 s_1^2}{\|\mu_1\|^2 + s_1^2} + Cs_1^2 \sqrt{\frac{u_0}{d_1^e}}, \end{aligned}$$

where we have used that $2ab \leq a^2 + b^2$. Otherwise, if $\eta \geq s_1^2$:

$$\begin{aligned} R_1(\hat{\omega}) &\leq s_1^2(1 + Ce^{-u_0/2}) + Cs_1^2 \sqrt{\frac{u_0}{d_1^e}} \\ &\leq (1 + Ce^{-u_0/2}) \frac{\|\mu_1\|^2 s_1^2}{\|\mu_1\|^2 + s_1^2} + (1 + Ce^{-u_0/2}) \frac{s_1^4}{\|\mu_1\|^2 + s_1^2} + Cs_1^2 \sqrt{\frac{u_0}{d_1^e}}. \end{aligned}$$

We conclude using that $s_1^2 \leq C \frac{\|\mu_1\|^2}{d_1^e} u_0$ in this case.

Proof of Corollary 2. Let $\tau \geq 0, \varsigma \geq 1$ be fixed. Let k be an element of $V_{\tau, \varsigma} = W_{(\varsigma)} \cap V_\tau$ with $k \neq 1$. We start by upper bounding q_k , with q_k defined in (34). Since

$k \in W_{(\varsigma)}$, it holds $\text{Tr } \Sigma_k^2 \leq \varsigma^2 \frac{N_k^2}{N_1^2} \text{Tr } \Sigma_1^2$, so that

$$\begin{aligned} \text{Tr } \Sigma_1 \Sigma_k &\leq \frac{1}{2} \left(\frac{N_k}{N_1} \text{Tr } \Sigma_1^2 + \frac{N_1}{N_k} \text{Tr } \Sigma_k^2 \right) \leq \frac{1 + \varsigma^2}{2} \frac{N_k}{N_1} \text{Tr } \Sigma_1^2 \\ &\leq \frac{N_k (1 + \varsigma^2) (\text{Tr } \Sigma_1)^2}{N_1 2d_1^\bullet} \\ &= N_k N_1 \frac{\varsigma^2 s_1^4}{d_1^\bullet}. \end{aligned}$$

Since $k \in V_\tau$, it holds

$$\frac{\Delta_k^T \Sigma_1 \Delta_k}{N_1} \leq \frac{\|\Sigma_1\|_\infty}{N_1} \|\Delta_k\|^2 \leq \frac{\text{Tr } \Sigma_1}{N_1} \frac{1}{d_1^e} \tau s_1^2 = \frac{\tau s_1^4}{d_1^e}.$$

Joining these estimates, we get

$$\frac{q_k}{N_1} \leq \frac{\Delta_k^T \Sigma_1 \Delta_k}{N_1} + \frac{\text{Tr } \Sigma_1 \Sigma_k}{N_1 N_k} \leq s_1^4 \left(\frac{\tau}{d_1^e} + \frac{\varsigma^2}{d_1^\bullet} \right).$$

Therefore, for $\boldsymbol{\omega}$ a vector of the simplex \mathcal{S}_B having support in $W^{(\varsigma)} \cap V_\tau$, using $d_1^e \leq d_1^\bullet$ it holds

$$Q_1(\boldsymbol{\omega}) = \sum_{k \geq 2} \omega_k \sqrt{\frac{q_k}{N_1}} \leq (1 - \omega_1) \sqrt{\tau + \varsigma^2} \frac{s_1^2}{\sqrt{d_1^e}}. \quad (86)$$

We now choose the weight vector $\boldsymbol{\omega}^* = \boldsymbol{\omega}_{V_{\tau, \varsigma}}^*$ given by the oracle weights of (13), for the set $V = V_{\tau, \varsigma}$. From Lemma 1, this gives rise to $R_1(\boldsymbol{\omega}^*) \leq \mathcal{B}(\tau, \nu)$, where $\nu = \nu(V_{\tau, \varsigma})$; furthermore we have the explicit expression

$$(1 - \omega_1^*) = \lambda(1 - \nu), \quad \text{where } \lambda = \frac{1}{1 + \tau(1 - \nu)},$$

so that it holds (since $\nu \in [0, 1]$)

$$(1 - \omega_1^*) \sqrt{\tau} = \frac{(1 - \nu) \sqrt{\tau}}{1 + \tau(1 - \nu)} \leq \max \left(\frac{\tau(1 - \nu)}{1 + \tau(1 - \nu)}, \frac{\sqrt{\tau}(1 - \nu)}{1 + \sqrt{\tau}(1 - \nu)} \right) \leq 1.$$

Plugging this into (86), we get $Q_1(\boldsymbol{\omega}^*) \leq 2\varsigma s_1^2 / \sqrt{d_1^e}$, then (35) since the obtained estimate holds for any $\tau \geq 0, \varsigma \geq 1$.

D.3. Proof of Proposition 6. We follow the same general canvas as in the proof of Proposition 5.

First step : bound in conditional probability. Let us recall the definitions of $Q^{\text{BS}}(\boldsymbol{\omega})$ and \hat{q}_k :

$$\hat{Q}^{\text{BS}}(\boldsymbol{\omega}) := \frac{M}{N_1} \sum_{k=2}^B \omega_k \|\hat{\mu}_k^{\text{NE}} - \hat{\mu}_1^{\text{NE}}\|, \quad \hat{q}_k = \frac{1}{N_1 - 1} \sum_{p=1}^{N_1} \langle \hat{\mu}_k^{\text{NE}} - \hat{\mu}_1^{\text{NE}}, X_p^{(1)} - \hat{\mu}_1^{\text{NE}} \rangle^2.$$

We will need the following quantity \widehat{q}'_k which is close to \widehat{q}_k but easier to control:

$$\widehat{q}'_k = \frac{1}{N_1 - 1} \sum_{p=1}^{N_1} \langle \widehat{\mu}_k^{\text{NE}} - \mu_1, X_p^{(1)} - \widehat{\mu}_1^{\text{NE}} \rangle^2.$$

The estimated weight vector $\widehat{\omega}$ for the estimation of μ_1 is chosen as

$$\widehat{\omega} \in \underset{\omega \in \mathcal{S}_B}{\text{Arg Min}} \left(\widehat{L}_1(\omega) + 4\sqrt{2u_0} \widehat{Q}_1(\omega) + 1424u_0 \widehat{Q}(\omega)^{\text{BS}} \right).$$

Let $u := u_0 - \log B$, and define the events:

$$A_1 = \left\{ \|\widehat{\mu}_k^{\text{NE}} - \mu_1\|_{\Sigma_1} \leq 2\sqrt{\widehat{q}'_k} + 711 \frac{\|\widehat{\mu}_k^{\text{NE}} - \mu_1\| M}{\sqrt{N_1}} (u + \log B), 2 \leq k \leq B \right\},$$

$$A_2 = \left\{ \left| \|\widehat{\mu}_1^{\text{NE}} - \mu_1\|^2 - \widehat{s}_1^2 \right| \leq C \frac{s_1^2}{\sqrt{d_1^\bullet}} u + C \frac{M^2}{N_1^2} u^2 \right\},$$

and

$$A_3 = \left\{ \begin{aligned} \langle \widehat{\mu}_k^{\text{NE}} - \mu_1, \widehat{\mu}_1^{\text{NE}} - \mu_1 \rangle &\leq \sqrt{2 \frac{u + \log B}{N_1}} \|\widehat{\mu}_k^{\text{NE}} - \mu_1\|_{\Sigma_1} \\ &+ \frac{2\|\widehat{\mu}_k^{\text{NE}} - \mu_1\| M}{3N_1} (u + \log B), 2 \leq k \leq B \end{aligned} \right\},$$

where we recall that for ν a vector and Σ an operator, $\|\nu\|_{\Sigma}^2 := \langle \nu, \Sigma \nu \rangle$. For $i \in \{1, 3\}$, $\mathbb{P}[A_i | \mathcal{X}^{(-1)}] \geq 1 - e^{-u}$ and $\mathbb{P}[A_2 | \mathcal{X}^{(-1)}] \geq 1 - 2e^{-u}$ because of Proposition 21 for A_1 , Lemma 7 for A_3 and for A_2 , because $\|\widehat{\mu}_1^{\text{NE}} - \mu_1\|^2 - \widehat{s}_1^2$ is a U-statistic:

$$\|\widehat{\mu}_1^{\text{NE}} - \mu_1\|^2 - \widehat{s}_1^2 = \frac{1}{N_1(N_1 - 1)} \sum_{\ell \neq p=1}^{N_1} \langle X_\ell^{(1)} - \mu_1, X_p^{(1)} - \mu_1 \rangle, \quad (87)$$

the concentration is a direct consequence of Houdré and Reynaud-Bouret (2003) (or see Proposition 9 in Blanchard and Fermanian (2023) for this specific statistic). Then the event $A = A_1 \cap A_2 \cap A_3$ conditionally to $\mathcal{X}^{(-1)}$ is of probability greater than $1 - 4e^{-u}$. The differences between respectively \widehat{q}_k and \widehat{q}'_k for $k \in \llbracket B \rrbracket$ can be bounded independently of k :

$$\left| \sqrt{\widehat{q}_k} - \sqrt{\widehat{q}'_k} \right| \leq \sqrt{\frac{1}{N_1(N_1 - 1)} \sum_{p=1}^{N_1} \langle \widehat{\mu}_1^{\text{NE}} - \mu_1, X_p - \widehat{\mu}_1^{\text{NE}} \rangle^2} =: \Delta q. \quad (88)$$

Assume A , then:

$$\begin{aligned}
L_1(\widehat{\omega}) &= \widehat{L}_1(\widehat{\omega}) + 2 \sum_{k=2}^B \widehat{\omega}_k \langle \widehat{\mu}_k^{\text{NE}} - \mu_1, \widehat{\mu}_1^{\text{NE}} - \mu_1 \rangle + (2\widehat{\omega}_1 - 1) (\|\widehat{\mu}_1^{\text{NE}} - \mu_1\|^2 - \widehat{s}_1^2) \\
&\leq \widehat{L}_1(\widehat{\omega}) + 2 \sum_{k=2}^B \widehat{\omega}_k \left(\sqrt{2 \frac{u + \log B}{N_1}} \|\widehat{\mu}_k^{\text{NE}} - \mu_1\|_{\Sigma_1} + \frac{2 \|\widehat{\mu}_k^{\text{NE}} - \mu_1\| M}{3N_1} (u + \log B) \right) \\
&\quad + \frac{Cs_1^2}{\sqrt{d_1^\bullet}} u + \frac{CM^2}{N_1^2} u^2,
\end{aligned}$$

where we have used the events A_2 and A_3 . Then using the event A_1 , the bound (88) and a triangle inequality we get:

$$\begin{aligned}
L_1(\widehat{\omega}) &\leq \widehat{L}_1(\widehat{\omega}) + 4\sqrt{\log B + u} \sum_{k=2}^B \widehat{\omega}_k \sqrt{\frac{2\widehat{q}_k}{N_1}} + 1424(\log B + u) \sum_{k=2}^B \widehat{\omega}_k \frac{M \|\widehat{\mu}_k^{\text{NE}} - \widehat{\mu}_1^{\text{NE}}\|}{N_1} \\
&\quad + C \frac{\Delta q}{\sqrt{N_1}} \sqrt{\log B + u} + C \frac{\|\widehat{\mu}_1^{\text{NE}} - \mu_1\| M}{N_1} (\log B + u) + \frac{Cs_1^2}{\sqrt{d_1^\bullet}} u + \frac{CM^2}{N_1^2} u^2.
\end{aligned}$$

Using the choice of $\widehat{\omega}$, conditionally to A :

$$\begin{aligned}
L_1(\widehat{\omega}) &\leq \min_{\omega \in \mathcal{S}_B} \left(\widehat{L}_1(\omega) + 4\sqrt{2u_0} \widehat{Q}(\omega) + 1424u_0 \widehat{Q}^{\text{BS}}(\omega) \right) \\
&\quad + C \frac{\Delta q}{\sqrt{N_1}} \sqrt{\log B + u} + C \frac{\|\widehat{\mu}_1^{\text{NE}} - \mu_1\| M}{N_1} (\log B + u) + \frac{Cs_1^2}{\sqrt{d_1^\bullet}} u + \frac{CM^2}{N_1^2} u^2.
\end{aligned}$$

Second and third steps: bound in expectation. Let us bound some expectation using Jensen's inequality:

$$\mathbb{E} \left[\sqrt{\widehat{q}_k} \right] \leq \sqrt{\|\mu_k - \mu_1\|_{\Sigma_1}^2 + \frac{\text{Tr}(\Sigma_1 \Sigma_k)}{N_k}}, \quad \mathbb{E}[\Delta q] \leq \frac{M \sqrt{\text{Tr} \Sigma_1}}{N_1} + \frac{\sqrt{\text{Tr} \Sigma_1^2}}{\sqrt{N_1}}. \quad (89)$$

The expectation of $\sqrt{\widehat{q}_k}$ can be bounded using that $\sqrt{\widehat{q}_k} \leq \sqrt{\widehat{q}_k} + \Delta q$. We can now bound the risk. Let $\omega \in \mathcal{S}_B$:

$$\begin{aligned}
R_1(\widehat{\omega}) &\leq \mathbb{E}[L_1(\widehat{\omega}) 1_A] + M^2 \mathbb{P}[A^c] \\
&\leq L_1(\omega) + 4\sqrt{2u_0} \sum_{k=2}^B \omega_k \frac{\mathbb{E} \left[\sqrt{\widehat{q}_k} \right]}{\sqrt{N_1}} + 1424u_0 \sum_{k=2}^B \omega_k \frac{M (\|\mu_k - \mu_1\| + s_1 + s_k)}{N_1} \\
&\quad + C \frac{\mathbb{E}[\Delta q]}{\sqrt{N_1}} \sqrt{\log B + u} + C \frac{s_1 M}{N_1} (\log B + u) + C \frac{s_1^2}{\sqrt{d_1^\bullet}} u + C \frac{M^2}{N_1^2} u^2 + 3M^2 e^{-u}
\end{aligned}$$

Because $u \geq 2 \log N_1$, the last term is upper bounded by the previous one. Using (89) and by bringing together the terms:

$$\begin{aligned} R_1(\widehat{\boldsymbol{\omega}}) &\leq R_1(\boldsymbol{\omega}) + 4\sqrt{2(\log B + u)}Q(\boldsymbol{\omega}) + 1424(\log B + u) \sum_{k=2}^B \omega_k \frac{M(\|\mu_k - \mu_1\| + s_k)}{N_1} \\ &\quad + C \frac{s_1^2}{\sqrt{d_1^\bullet}}(u + \sqrt{\log B + u}) + C \frac{Ms_1}{N_1}(\log B + u) + C \frac{M^2}{N_1^2}u^2, \end{aligned} \quad (90)$$

where Q is defined in (34). Let $\tau, \varsigma > 0$ and $\boldsymbol{\omega}^* = \boldsymbol{\omega}_{V_{\tau, \varsigma}}^*$ be defined as in (13). Then as in the proof of Corollary 2:

$$R_1(\boldsymbol{\omega}^*) = s_1^2 \mathcal{B}(\tau, \nu(V_{\tau, \varsigma})), \quad Q(\boldsymbol{\omega}^*) \leq C \sqrt{\frac{1 + \varsigma^2}{d_1^e}} s_1^2. \quad (91)$$

Up to bound the third term in the upper bound (90), let us bound s_k^2 for $k \in V_{\tau, \varsigma}$. On the one hand:

$$s_k^2 = \frac{\text{Tr} \Sigma_k}{N_k} \leq \frac{4M^2}{N_k} = 4 \text{Tr} \Sigma_1 \frac{\Phi_1}{N_k} = 4s_1^2 \frac{\Phi_1 N_1}{N_k}.$$

On the other hand, as $k \in V_{\tau, \varsigma} \subset W_{(\varsigma)}$:

$$s_k^2 = \frac{\text{Tr} \Sigma_k}{N_k} = \sqrt{d_k^\bullet} \frac{\sqrt{\text{Tr} \Sigma_k^2}}{N_k} \leq \sqrt{d_k^\bullet} \varsigma \frac{\sqrt{\text{Tr} \Sigma_1^2}}{N_1} = s_1^2 \varsigma \sqrt{\frac{d_k^\bullet}{d_1^\bullet}}.$$

Combining these two bounds:

$$s_k^2 \leq 4s_1^2 \min\left(\frac{\Phi_1 N_1}{N_k}, \varsigma \sqrt{\frac{d_k^\bullet}{d_1^\bullet}}\right).$$

As we assume $N_k \geq (d_k^\bullet)^\beta$, for $k \in V_{\tau, \varsigma}$:

$$s_k^2 \leq 4s_1^2 \min\left(\frac{\Phi_1 N_1}{(d_k^\bullet)^\beta}, \varsigma \sqrt{\frac{d_k^\bullet}{d_1^\bullet}}\right) \leq 4s_1^2 \max_{d \geq 1} \min\left(\frac{\Phi_1 N_1}{d^\beta}, \varsigma \sqrt{\frac{d}{d_1^\bullet}}\right) = 4s_1^2 (\Phi_1 N_1)^{\frac{1}{1+2\beta}} \left(\frac{\varsigma}{\sqrt{d_1^\bullet}}\right)^{\frac{2\beta}{1+2\beta}}.$$

We can now bound the third term in (90). As $\omega_k^* = 0$ for $k \notin V_{\tau, \varsigma}$:

$$\begin{aligned} \sum_{k=2}^B \omega_k^* \frac{M(\|\mu_k - \mu_1\| + s_k)}{N_1} &\leq \frac{M}{N_1} (1 - \omega_1^*) \left(\sqrt{\tau} s_1 + 2s_1 (\Phi_1 N_1)^{\frac{1}{2(1+2\beta)}} \left(\frac{\varsigma}{\sqrt{d_1^\bullet}}\right)^{\frac{\beta}{1+2\beta}} \right) \\ &\leq s_1^2 \left((1 - \omega_1^*) \sqrt{\frac{\tau \Phi_1}{N_1}} + 2\Phi_1^{\frac{1+\beta}{1+2\beta}} \left(\frac{\varsigma}{N_1 \sqrt{d_1^\bullet}}\right)^{\frac{\beta}{1+2\beta}} \right). \end{aligned}$$

As $N_1 \geq (d_1^\bullet)^\beta$ and $(1 - \omega_1^*)\sqrt{\tau} \leq 1$ (by definition of ω_1^*), we get:

$$\sum_{k=2}^B \omega_k^* \frac{M(\|\mu_k - \mu_1\| + s_k)}{N_1} \leq 2s_1^2 \left(\frac{\sqrt{\Phi_1}}{(d_1^\bullet)^{\beta/2}} + \frac{\Phi_1^{\frac{1+\beta}{1+2\beta}} \varsigma^{\frac{\beta}{1+2\beta}}}{(d_1^\bullet)^{\beta/2}} \right). \quad (92)$$

Injecting the bounds (91) and (92) into (90) leads to:

$$\begin{aligned} \frac{R_1(\widehat{\omega})}{s_1^2} &\leq \min_{\tau>0, \varsigma>0} \left(\mathcal{B}(\tau, \nu(V_{\tau, \varsigma})) + C\varsigma \sqrt{\frac{u_0}{d_1^e}} + C u_0 \frac{\Phi_1^{\frac{1+\beta}{1+2\beta}} \varsigma^{\frac{\beta}{1+2\beta}}}{(d_1^\bullet)^{\beta/2}} \right) \\ &\quad + C u_0 \frac{\sqrt{\Phi_1}}{(d_1^\bullet)^{\beta/2}} + C \sqrt{\frac{u_0}{d_1^e}} + C \frac{u_0}{\sqrt{d_1^\bullet}} + C \frac{u_0 \sqrt{\Phi_1}}{\sqrt{N_1}} + C \frac{\Phi_1 u^2}{N_1} \\ &\leq \min_{\tau>0, \varsigma>0} \left(\mathcal{B}(\tau, \nu(V_{\tau, \varsigma})) + C\varsigma \sqrt{\frac{u_0}{d_1^e}} + C u_0 \frac{\Phi_1^{\frac{1+\beta}{1+2\beta}} \varsigma^{\frac{\beta}{1+2\beta}}}{(d_1^\bullet)^{\beta/2}} \right) + C \sqrt{\frac{u_0}{d_1^e}} + C \frac{u_0 \Phi_1}{(d_1^\bullet)^{\beta/2}}. \end{aligned}$$

As $\Phi_1^{\frac{1+\beta}{1+2\beta}} \varsigma^{\frac{\beta}{1+2\beta}} \leq \max(\Phi_1, \varsigma) \leq \Phi_1 + \varsigma$, we obtain:

$$\frac{R_1(\widehat{\omega})}{s_1^2} \leq \min_{\tau>0, \varsigma>0} \left(\mathcal{B}(\tau, \nu(V_{\tau, \varsigma})) + C\varsigma \max\left(\sqrt{\frac{u_0}{d_1^e}}, \frac{u_0}{(d_1^\bullet)^{\beta/2}}\right) \right) + C \sqrt{\frac{u_0}{d_1^e}} + C \frac{u_0 \Phi_1}{(d_1^\bullet)^{\beta/2}}.$$

□

D.4. Concentration inequalities.

D.4.1. *Concentration for \widehat{q} .* Consider first the Gaussian setting (**GS**).

Proposition 20. *Let X_1, \dots, X_N i.i.d. Gaussian random vectors of distribution $\mathcal{N}(\mu_1, \Sigma_1)$ and $\nu \in \mathbb{R}^d$. Let $\widehat{q} = \frac{1}{N-1} \sum_{k=1}^N \langle \widehat{\mu}_1^{\text{NE}} - \nu, X_k - \widehat{\mu}_1^{\text{NE}} \rangle^2$, then for all $x \geq 0$:*

$$\mathbb{P} \left[\sqrt{\widehat{q}} \geq \left(1 + \sqrt{\frac{2x}{N-1}} \right) \left(\sqrt{\|\mu_1 - \nu\|_{\Sigma_1}^2} + \frac{\text{Tr} \Sigma_1^2}{N} + \|\Sigma_1\|_\infty \sqrt{\frac{2x}{N}} \right) \right] \leq 2e^{-x}, \quad (93)$$

and

$$\mathbb{P} \left[\sqrt{\widehat{q}} \leq e^{-1/2-x/(N-1)} \left(\sqrt{\|\mu_1 - \nu\|_{\Sigma_1}^2} + \frac{\text{Tr} \Sigma_1^2}{N} - 2\|\Sigma_1\|_\infty \sqrt{\frac{2x}{N}} \right) \right] \leq 2e^{-x}, \quad (94)$$

where $\|\mu_1 - \nu\|_{\Sigma_1}^2 = (\mu_1 - \nu)^T \Sigma_1 (\mu_1 - \nu)$.

Proof. Let us consider the random vector $Z \in \mathbb{R}^N$ with $Z_k = \langle \widehat{\mu}_1^{\text{NE}} - \nu, X_k - \widehat{\mu}_1^{\text{NE}} \rangle$, then $\widehat{q} = \|Z\|_N^2 / (N-1)$, where $\|\cdot\|_N$ is the Euclidian norm in \mathbb{R}^N . Conditionally to $\widehat{\mu}_1^{\text{NE}}$, Z is a Gaussian vector of distribution $\mathcal{N}(0, e(\widehat{\mu}_1^{\text{NE}})\Gamma)$, where $e(\widehat{\mu}_1^{\text{NE}}) = (\widehat{\mu}_1^{\text{NE}} - \nu)^T \Sigma_1 (\widehat{\mu}_1^{\text{NE}} - \nu)$ and $\Gamma = I_N - \mathbf{1}_N \mathbf{1}_N^T / N$ with $\mathbf{1}_N = (1, \dots, 1) \in \mathbb{R}^N$. The eigenvalues of Γ are 1

with multiplicity $N - 1$ and 0. So $\|Z\|^2/e(\widehat{\mu}_1^{\text{NE}})$ has a $\chi^2(N - 1)$ distribution. Then conditionally to $\widehat{\mu}_1^{\text{NE}}$:

$$\widehat{q} = \frac{\|Z\|^2}{N-1} \sim \frac{e(\widehat{\mu}_1^{\text{NE}})}{N-1} \chi^2(N-1).$$

Then according to Lemma 5 and Lemma 6, for all $x \geq 0$:

$$\mathbb{P} \left[\sqrt{\frac{\widehat{q}}{e(\widehat{\mu}_1^{\text{NE}})}} \geq 1 + \sqrt{\frac{2x}{N-1}} \middle| \widehat{\mu}_1^{\text{NE}} \right] \leq e^{-x}, \quad \mathbb{P} \left[\sqrt{\frac{\widehat{q}}{e(\widehat{\mu}_1^{\text{NE}})}} \leq e^{-1/2} e^{-x/(N-1)} \middle| \widehat{\mu}_1^{\text{NE}} \right] \leq e^{-x}.$$

Let $g = \Sigma_1^{1/2}(\widehat{\mu}^{\text{NE}} - \nu) \sim \mathcal{N}(\Sigma_1^{1/2}(\mu_1 - \nu), \Sigma_1^2/N)$, as $\|g\|^2 = e(\widehat{\mu}_1^{\text{NE}})$, from Lemma 5 with $\Sigma_1^{1/2}(\mu_1 - \nu) \rightarrow \mu$ and $\Sigma_1^2/N \rightarrow \Sigma$, we get that for all $x \geq 0$:

$$\begin{aligned} \mathbb{P} \left[\sqrt{e(\widehat{\mu}_1^{\text{NE}})} \geq \sqrt{(\mu_1 - \nu)^T \Sigma_1 (\mu_1 - \nu) + \frac{\text{Tr} \Sigma_1^2}{N} + \|\Sigma_1\|_\infty \sqrt{\frac{2x}{N}}} \right] &\leq e^{-x}, \\ \mathbb{P} \left[\sqrt{e(\widehat{\mu}_1^{\text{NE}})} \leq \sqrt{(\mu_1 - \nu)^T \Sigma_1 (\mu_1 - \nu) + \frac{\text{Tr} \Sigma_1^2}{N} - 2\|\Sigma_1\|_\infty \sqrt{\frac{2x}{N}}} \right] &\leq e^{-x}. \end{aligned}$$

We have used that for all $\mu \in \mathbb{R}^d$, $\Sigma \in \mathbb{R}^{d \times d}$ and $x \geq 0$:

$$\begin{aligned} \left(\sqrt{\|\mu\|^2 + \text{Tr} \Sigma} + \sqrt{2\|\Sigma\|_\infty x} \right)^2 &\geq (\|\mu\|^2 + \text{Tr} \Sigma) + 2\sqrt{(\text{Tr} \Sigma^2 + 2\mu^T \Sigma \mu)x} + 2\|\Sigma\|_\infty x, \\ \left(\sqrt{\|\mu\|^2 + \text{Tr} \Sigma} - \sqrt{2\|\Sigma\|_\infty x} \right)^2 &\leq \left((\|\mu\|^2 + \text{Tr} \Sigma) - 2\sqrt{(\text{Tr} \Sigma^2 + 2\mu^T \Sigma \mu)x} \right)_+, \end{aligned}$$

as $(a - b)_+^2 \leq (a^2 - ab)_+$ for $a, b > 0$.

Equations (93) and (94) are obtained by combining these concentration inequalities. \square

In the bounded setting **(BS)**, Proposition 21 gives a concentration bound for \widehat{q}' , which is a slightly different statistic from \widehat{q} because we consider $\mu_1 - \nu$ known for \widehat{q}' .

Proposition 21. *Assume **(BS)**, let $\nu \in \mathbb{R}^d$ and $\widehat{q}' = \frac{1}{N-1} \sum_{k=1}^N \langle \mu_1 - \nu, X_k - \widehat{\mu}_1^{\text{NE}} \rangle^2$. Then for all $u \geq 1$:*

$$\mathbb{P} \left[2\sqrt{\widehat{q}'} \leq \sqrt{(\mu_1 - \nu)^T \Sigma_1 (\mu_1 - \nu)} - 711 \frac{\|\mu_1 - \nu\| M}{\sqrt{N-1}} u \right] \leq e^{-u}.$$

Proof. Let us first denote $\delta := \mu_1 - \nu$ and $Z' := \sqrt{\widehat{q}'}$. We are going to use Talagrand's inequality (Theorem 1). So let us first rewrite Z' :

$$\begin{aligned} Z' &= \sup_{\|v\|_N=1} \frac{1}{\sqrt{N-1}} \sum_{k=1}^N v_k \langle \delta, X_k - \widehat{\mu}_1^{\text{NE}} \rangle \\ &= \sup_{\|v\|_N=1} \frac{1}{\sqrt{N-1}} \sum_{k=1}^N \langle \delta, X_k - \mu_1 \rangle \left(v_k - \frac{1}{N} \sum_{q=1}^N v_q \right). \end{aligned}$$

Let $T = \{v \in \mathbb{R}^N, \|v\|_N = 1\}$ (or a countable dense subset) and define for $v \in T$:

$$X_k^v := \frac{1}{\sqrt{N-1}} \langle \delta, X_k - \mu_1 \rangle \left(v_k - \frac{1}{N} \sum_{q=1}^N v_q \right),$$

then:

$$|X_k^v| \leq \frac{2\|\delta\|M}{\sqrt{N-1}}, \quad \sup_{v \in T} \sum_{k=1}^N \mathbb{E}[(X_k^v)^2] \leq \frac{\delta^T \Sigma \delta}{N-1} \leq \frac{4\|\delta\|^2 M^2}{N-1}.$$

Using Theorem 1, with probability greater than $1 - e^{-u}$, $u \geq 1$:

$$Z' \geq \mathbb{E}[Z'](1 - \varepsilon) - C(\varepsilon) \frac{\|\delta\|M}{\sqrt{N-1}} u,$$

where $C(\varepsilon) = 8(2 + \varepsilon^{-1})$ for some $\varepsilon > 0$. We just need to lower bound $\mathbb{E}[Z']$ by $\sqrt{\mathbb{E}[(Z')^2]} = \sqrt{\delta^T \Sigma_1 \delta}$. For that, using again Talagrand's inequality, it exists an exponential random variable $\xi \sim \mathcal{E}(1)$ such that:

$$Z' \leq \mathbb{E}[Z'](1 + \varepsilon) + C(\varepsilon) \frac{\|\delta\|M}{\sqrt{N-1}} \xi$$

Then:

$$\begin{aligned} \mathbb{E}[(Z')^2] &\leq \mathbb{E} \left[\left(\mathbb{E}[Z'](1 + \varepsilon) + C(\varepsilon) \frac{\|\delta\|M}{\sqrt{N-1}} \xi \right)^2 \right] \\ &\leq \left(\mathbb{E}[Z'](1 + \varepsilon) + \sqrt{2} C(\varepsilon) \frac{\|\delta\|M}{\sqrt{N-1}} \right)^2, \end{aligned}$$

and we get that $(1 + \varepsilon)\mathbb{E}[Z'] \geq \sqrt{\mathbb{E}[(Z')^2]} - \sqrt{2} C(\varepsilon) \frac{\|\delta\|M}{\sqrt{N-1}}$. Putting together the two bounds, we get a first lower bound for Z' : for $u \geq 1$ and probability greater than $1 - e^{-u}$:

$$Z' \geq \sqrt{\delta^T \Sigma_1 \delta} \frac{1 - \varepsilon}{1 + \varepsilon} - C(\varepsilon) \left(\sqrt{2} \frac{1 - \varepsilon}{1 + \varepsilon} + 1 \right) \frac{\|\delta\|M}{\sqrt{N-1}} u. \quad (95)$$

Let us choose $\varepsilon = 1/3$ to conclude. \square

D.4.2. Classical concentration inequalities.

Concentration inequalities for Gaussian random variables.

Lemma 4. *Let $X \sim \mathcal{N}(m, \sigma^2)$, then for all $x \geq 0$:*

$$\mathbb{P} \left[|X - m| \geq \sqrt{2\sigma^2 x} \right] \leq 2e^{-x}$$

Proof. It is a direct consequence of the Chernoff bound (Chernoff, 1952). \square

Lemma 5. [*Concentration of Gaussian vectors*] Let $X \sim \mathcal{N}(\mu, \Sigma)$, then for all $x \geq 0$:

$$\begin{aligned} \mathbb{P}\left[\|X\|^2 \geq (\|\mu\|^2 + \text{Tr } \Sigma) + 2\sqrt{(\text{Tr } \Sigma^2 + 2\mu^T \Sigma \mu)x} + 2\|\Sigma\|_\infty x\right] &\leq e^{-x}, \\ \mathbb{P}\left[\|X\|^2 \leq (\|\mu\|^2 + \text{Tr } \Sigma) - 2\sqrt{(\text{Tr } \Sigma^2 + 2\mu^T \Sigma \mu)x}\right] &\leq e^{-x}, \end{aligned}$$

Proof. The Lemma is a consequence of Lemma 1 of Laurent and Massart (2000) and Lemma 8.1 of Birgé (2001). \square

Lemma 6. [*Lower bound for χ^2*] Let $Z \sim \chi^2(n)$, then for all $x \geq 0$:

$$\mathbb{P}[Z \leq ne^{-(1+2x/n)}] \leq e^{-x}.$$

Proof. Let $\delta \in (0, 1)$, $\lambda \in \mathbb{R}^+$:

$$\mathbb{P}[Z \leq n\delta] = \mathbb{P}[e^{-\lambda Z} \geq e^{-n\lambda\delta}] \leq \mathbb{E}[e^{-\lambda Z}] e^{n\lambda\delta} = \exp\left(-\frac{n}{2}(\log(1+2\lambda) - 2\lambda\delta)\right)$$

where the inequality is due to Markov. Fix $\lambda = (-1 + \delta^{-1})/2 > 0$, then:

$$\mathbb{P}[Z \leq n\delta] \leq \exp\left(-\frac{n}{2}(-\log(\delta) + \delta - 1)\right) \leq \exp\left(-\frac{n}{2}(-\log(\delta) - 1)\right)$$

Let us choose $\delta = \exp(-1 - 2x/n)$ to conclude the proof. \square

Concentration inequalities for bounded random variables.

Lemma 7. [*Bernstein's concentration inequality*] Let X_1, \dots, X_N i.i.d. real centred random variables bounded by M such that $\mathbb{E}[X_1^2] \leq \sigma^2$, then for all $x \geq 0$:

$$\mathbb{P}\left[\sum_{i=1}^N X_i \geq \sqrt{2N\sigma^2 x} + \frac{2Mx}{3}\right] \leq e^{-x}$$

Proof. See for instance Vershynin (2018), Exercise 2.8.5. \square

Theorem 1. [*Talagrand's inequality*] Let X_1^t, \dots, X_n^t independant random variables indexed by $t \in T$ (T countable) in \mathbb{R} and $L > 0$ such that for all $t \in T$, $i \leq n$,

$$\mathbb{E}[X_i^t] = 0, \quad |X_i^t| \leq L \tag{96}$$

Let

$$Z := \sup_{t \in T} \sum_{i=1}^n X_i^t, \quad \sigma^2 = \sup_{t \in T} \sum_{i=1}^n \mathbb{E}[(X_i^t)^2]$$

then for all $x \geq 0$ and $\varepsilon \in (0, 1)$:

$$\mathbb{P}\left[Z \geq \mathbb{E}[Z](1 + \varepsilon) + 2\sqrt{2\sigma^2 x} + 2Lx(1 + 8\varepsilon^{-1})\right] \leq e^{-x} \tag{97}$$

$$\mathbb{P}\left[Z \leq \mathbb{E}[Z](1 - \varepsilon) - 2\sqrt{4\sigma^2 x} - 4Lx(1 + 8\varepsilon^{-1})\right] \leq e^{-x} \tag{98}$$

Proof. See for instance Massart (2000). \square

APPENDIX E. PROOFS FOR SECTION 5

E.1. Proof of Proposition 7. This proof follows the same scheme as the Pinsker's bound (Pinsker, 1980 or see Tsybakov, 2008 for a recent version).

The proof is provided for $V = B$ but can be directly adapted for $V < B$ by assuming μ_k independent of μ_1 for $k > V$ when constructing the distribution \mathbb{Q} (100).

Let us first restrict ourselves to the case where μ_1 is in a ball around 0:

$$\inf_{\hat{\mu}_1} \sup_{\mu_i \in B(\mu_1, \sqrt{\tau} s_1)} R_1(\hat{\mu}_1) \geq \inf_{\hat{\mu}_1} \sup_{\substack{\mu_1 \in B(0, \sqrt{\beta} s_1) \\ \mu_i \in B(\mu_1, \sqrt{\tau} s_1)}} R_1(\hat{\mu}_1).$$

Then the infimum over the estimators is now attained for an estimator $\hat{\mu}_1$ bounded by $2\sqrt{\beta} s_1$. Indeed, any estimator $\hat{\mu}$ further perform less well than the deterministic estimator $\hat{\mu} = 0$. If $\|\hat{\mu}\| > 2\sqrt{\beta} s_1$:

$$\|\hat{\mu} - \mu_1\| \geq \|\hat{\mu}\| - \|\mu_1\| > \sqrt{\beta} s_1 > \|0 - \mu_1\|. \quad (99)$$

We introduce now the probability measure \mathbb{Q} :

$$\mu_1 \stackrel{\mathbb{Q}}{\sim} \mathcal{N}(0, \alpha\beta s_1^2 \Sigma), \quad \mu_2 = \dots = \mu_B = \mu_\circ \stackrel{\mathbb{Q}}{\sim} \mathcal{N}(\mu_1, \alpha\tau s_1^2 \Sigma), \quad (100)$$

where $\beta > 0$ and $\alpha \in (0, 1)$. Let A be the event $\{\|\mu_1\|^2 \leq \beta s_1^2, \|\mu_\circ - \mu_1\|^2 \leq \tau s_1^2\}$ and $\mathbb{E}_{\mathbb{Q}}$ denote the expectation over the distribution \mathbb{Q} , then:

$$\begin{aligned} \inf_{\hat{\mu}_1} \sup_{\mu_i \in B(\mu_1, \tau s_1)} R_1(\hat{\mu}_1) &\geq \inf_{\hat{\mu}_1: \|\hat{\mu}_1\| \leq 2\sqrt{\beta} s_1} \sup_{\substack{\mu_1 \in B(0, \sqrt{\beta} s_1) \\ \mu_i \in B(\mu_1, \sqrt{\tau} s_1)}} R_1(\hat{\mu}_1) \\ &\geq \inf_{\hat{\mu}_1: \|\hat{\mu}_1\| \leq 2\sqrt{\beta} s_1} \frac{1}{\mathbb{Q}(A)} \int_A R_1(\hat{\mu}_1) d\mathbb{Q}(\nu, \mu_1, \dots, \mu_B) \\ &\geq \inf_{\hat{\mu}_1} \mathbb{E}_{\mathbb{Q}}[R_1(\hat{\mu}_1)] - \sup_{\hat{\mu}_1: \|\hat{\mu}_1\| \leq 2\sqrt{\beta} s_1} \mathbb{E}_{\mathbb{Q}}[R_1(\hat{\mu}_1) 1_{A^c}] \\ &=: I - r, \end{aligned} \quad (101)$$

Let us now bound I and r .

Lower bound for I : The first infimum (term I) is attained for $\hat{\mu}_1 = \mathbb{E}[\mu_1 | X_{\bullet}^{(1)}, \dots, X_{\bullet}^{(B)}]$.

Let us calculate $\hat{\mu}_1$.

$$\mathbb{E}[\mu_1 | \mu_\circ, X_{\bullet}^{(1)}, \dots, X_{\bullet}^{(B)}] = \mathbb{E}[\mu_1 | \mu_\circ, X_{\bullet}^{(1)}] = ((\alpha\beta)^{-1} + 1 + (\alpha\tau)^{-1})^{-1} \left(\hat{\mu}_1^{\text{NE}} + \frac{1}{\alpha\tau} \mu_\circ \right),$$

$$\mathbb{E}[\mu_\circ | \mu_1, X_{\bullet}^{(1)}, \dots, X_{\bullet}^{(B)}] = ((\alpha\tau)^{-1} + \|\rho\|^2)^{-1} \left(\frac{1}{\alpha\tau} \mu_1 + \sum_{k=2}^B \rho_k^2 \hat{\mu}_k^{\text{NE}} \right)$$

where $\rho = (s_1/s_k)_{k \neq 1}$ and $\|\rho\|^2 = \sum_{k=2}^B \rho_k^2$. Combining these two expressions we get:

$$\begin{aligned} \mathbb{E}[\mu_1 | X_{\bullet}^{(1)}, \dots, X_{\bullet}^{(B)}] &= ((\alpha\beta)^{-1} + 1 + (\alpha\tau)^{-1})^{-1} \\ &\times \left(\widehat{\mu}_1^{\text{NE}} + \frac{1}{\alpha\tau} ((\alpha\tau)^{-1} + \|\rho\|^2)^{-1} \left(\frac{1}{\alpha\tau} \mathbb{E}[\mu_1 | X_{\bullet}^{(1)}, \dots, X_{\bullet}^{(B)}] + \sum_{k=2}^B \rho_k^2 \widehat{\mu}_k^{\text{NE}} \right) \right), \end{aligned}$$

and then:

$$\mathbb{E}[\mu_1 | X_{\bullet}^{(1)}, \dots, X_{\bullet}^{(B)}] = \left((\alpha\beta)^{-1} + 1 + \frac{\|\rho\|^2}{1 + \alpha\tau\|\rho\|^2} \right)^{-1} \left(\widehat{\mu}_1^{\text{NE}} + \frac{1}{1 + \alpha\tau\|\rho\|^2} \sum_{k=2}^B \rho_k^2 \widehat{\mu}_k^{\text{NE}} \right),$$

Let us first notice that:

$$\begin{aligned} \mathbb{E}[\mu_1 | X_{\bullet}^{(\cdot)}] - \mu_1 &= \left((\alpha\beta)^{-1} + 1 + \frac{\|\rho\|^2}{1 + \alpha\tau\|\rho\|^2} \right)^{-1} \\ &\times \left[\left(\widehat{\mu}_1^{\text{NE}} - \mu_1 \right) + \frac{1}{1 + \alpha\tau\|\rho\|^2} \sum_{k=2}^B \rho_k^2 (\widehat{\mu}_k^{\text{NE}} - \mu_{\circ}) + \frac{\|\rho\|^2}{1 + \alpha\tau\|\rho\|^2} (\mu_{\circ} - \mu_1) - \frac{1}{\alpha\beta} \mu_1 \right] \end{aligned}$$

Using that $\widehat{\mu}_1^{\text{NE}} - \mu_1$, $\widehat{\mu}_k^{\text{NE}} - \mu_{\circ}$ (for $k \neq 1$), $\mu_{\circ} - \mu_1$ and μ_1 are pairwise independent we get that:

$$\begin{aligned} \frac{\mathbb{E}[\|\widehat{\mu}_1 - \mu_1\|^2]}{s_1^2} &= \left((\alpha\beta)^{-1} + 1 + \frac{\|\rho\|^2}{1 + \alpha\tau\|\rho\|^2} \right)^{-2} \\ &\times \left[1 + \frac{1}{(1 + \alpha\tau\|\rho\|^2)^2} \sum_{k=2}^B \rho_k^4 \rho_k^{-2} + \frac{\alpha\tau\|\rho\|^4}{(1 + \alpha\tau\|\rho\|^2)^2} + \frac{1}{\alpha\beta} \right] \end{aligned}$$

After simplification:

$$I = s_1^2 \left((\alpha\beta)^{-1} + 1 + \frac{\|\rho\|^2}{1 + \alpha\tau\|\rho\|^2} \right)^{-1} \quad (102)$$

Upper bound for r : Using the triangle and Cauchy-Schwartz inequalities we have:

$$\begin{aligned} r &= \sup_{\widehat{\mu}_1: \|\widehat{\mu}_1\| \leq 2\sqrt{\beta}s_1} \mathbb{E}[\|\widehat{\mu}_1(X_{\bullet}^{(k)}, k \in \llbracket B \rrbracket) - \mu_1\|^2 1_{A^c}] \quad (103) \\ &\leq \mathbb{E}[2(4\beta s_1^2 + \|\mu_1\|^2) 1_{A^c}] \\ &\leq 8\beta s_1^2 \mathbb{P}[A^c] + 2\sqrt{\mathbb{E}[\|\mu_1\|^4] \mathbb{P}[A^c]} \\ &\leq 2s_1^2 (4\beta + \sqrt{3}\alpha\beta) \sqrt{\mathbb{P}[A^c]} \leq 20\beta s_1^2 \sqrt{\mathbb{P}[A^c]} \end{aligned}$$

It stays to show the exponential decrease of $\mathbb{P}[A^c]$. Let $\xi \sim \mathcal{N}(0, \Sigma)$:

$$\begin{aligned} \mathbb{P}[\|\mu_1\|^2 \geq \beta s_1^2] &= \mathbb{P}[\|\mu_{\circ} - \mu_1\|^2 \geq \tau s_1^2] = \mathbb{P}[\|\xi\|^2 \geq \alpha^{-1}] \\ &\leq \exp\left(-\frac{d_1^e}{2} \left(\sqrt{\frac{2}{\alpha}} - 1 - 1\right)\right). \end{aligned}$$

This follows from the concentration of the norm of Gaussian vectors (Lemma 5). By union bound we get that:

$$r \leq 30s_1^2\beta \exp\left(-\frac{d_1^e}{4}\left(\sqrt{\frac{2}{\alpha}} - 1 - 1\right)\right).$$

Conclusion : The lower bound finally obtained is :

$$\inf_{\hat{\mu}_1} \sup_{\mu_i \in B(\mu_1, \tau s_1)} \frac{R_1(\hat{\mu}_1)}{s_1^2} \geq \left((\alpha\beta)^{-1} + 1 + \frac{\|\rho\|^2}{1 + \alpha\tau\|\rho\|^2} \right)^{-1} - 30\beta \exp\left(-\frac{d_1^e}{4}\left(\sqrt{\frac{2}{\alpha}} - 1 - 1\right)\right),$$

where $\alpha \in (0, 1)$ and $\beta \in \mathbb{R}_+$ are two free parameters. We can choose $\beta = d_1^e / \log d_1^e$ and $\alpha = \frac{2}{1+(1+8\beta^{-1})^2}$, then:

$$\beta \exp\left(-\frac{d_1^e}{4}\left(\sqrt{\frac{2}{\alpha}} - 1 - 1\right)\right) = \beta \exp\left(-\frac{2d_1^e}{\beta}\right) = \frac{1}{d_1^e \log d_1^e}$$

$$\left((\alpha\beta)^{-1} + 1 + \frac{\|\rho\|^2}{1 + \alpha\tau\|\rho\|^2} \right)^{-1} - \left(1 + \frac{\|\rho\|^2}{1 + \alpha\tau\|\rho\|^2} \right)^{-1} \geq -(\alpha\beta)^{-1} \geq -41 \frac{\log d_1^e}{d_1^e}.$$

and

$$\left(1 + \frac{\|\rho\|^2}{1 + \alpha\tau\|\rho\|^2} \right)^{-1} - \left(1 + \frac{\|\rho\|^2}{1 + \tau\|\rho\|^2} \right)^{-1} = -(1 - \alpha) \frac{\tau\|\rho\|^2}{1 + \tau\|\rho\|^2} \frac{\frac{\|\rho\|^2}{1 + \alpha\tau\|\rho\|^2}}{1 + \frac{\|\rho\|^2}{1 + \alpha\tau\|\rho\|^2}} \frac{1}{1 + \frac{\|\rho\|^2}{1 + \tau\|\rho\|^2}}$$

$$\geq -(1 - \alpha) \geq -40 \frac{\log d_1^e}{d_1^e}$$

where we recall $\|\rho\|^2 = \sum_{i=1}^B \frac{s_i^2}{s_i^2} - 1 = (\nu(V_\tau))^{-1} - 1$. Hence:

$$\left(1 + \frac{\|\rho\|^2}{1 + \tau\|\rho\|^2} \right)^{-1} = \mathcal{B}(\tau, \nu(V_\tau))$$

By combining these three inequalities, we get that:

$$\inf_{\hat{\mu}_1} \sup_{\mu_i \in B(\mu_1, \tau s_1)} \frac{R_1(\hat{\mu}_1)}{s_1^2} \geq \mathcal{B}(\tau, \nu(V_\tau)) - 111 \frac{\log d_1^e}{d_1^e}$$

E.2. Proof of Proposition 8. Let \mathcal{C} be a fixed J -partition of the means $(\mu_k)_{k \in \llbracket B \rrbracket}$ and denote $\zeta = \text{diam}(\mathcal{C})$. Let us focus first on a specific group $j \in \llbracket J \rrbracket$ and task $k \in \mathcal{C}_j$. Denote $\tau_{j,k} = \zeta_j^2 / s_k^2$ and $\nu_{j,k} = s^2(\mathcal{C}_j) / s_k^2$. Consider the vector of oracle weights ω_k^* given by (13), wherein the target task 1 is replaced by k everywhere, and the subset of

neighbouring tasks is taken as $\mathcal{C}_j \subseteq V_{\tau_{j,k}}$. Lemma 1 then states $R_k(\boldsymbol{\omega}_k^*)/s_k^2 \leq \mathcal{B}(\tau_{j,k}, \nu_{j,k})$. As a consequence, according to Proposition 5, it holds

$$\frac{R_k(\widehat{\boldsymbol{\omega}}_k)}{s_k^2} \leq (1 + CB e^{-u_0}) \left(\mathcal{B}(\tau_{j,k}, \nu_{j,k}) + C \sqrt{u_0} \frac{Q_k(\boldsymbol{\omega}_k^*)}{s_k^2} \right) + C \frac{u_0}{\sqrt{d_1^\bullet}}.$$

The rest of the proof is dedicated to bounding the terms $Q_k(\boldsymbol{\omega}_k^*)s_k^{-1}$ (and their sum over $k \in \mathcal{C}_j$). Denote $\omega_{k,\ell}^*$ the ℓ -th component of $\boldsymbol{\omega}_k^*$. It holds

$$\begin{aligned} \frac{Q_k(\boldsymbol{\omega}_k^*)}{s_k^2} &= s_k^{-2} \sum_{\ell \in \mathcal{C}_j \setminus \{k\}} \omega_{k,\ell}^* \sqrt{\frac{(\mu_\ell - \mu_k)^T \Sigma_k (\mu_\ell - \mu_k)}{N_k} + \frac{\text{Tr} \Sigma_\ell \Sigma_k}{N_\ell N_k}} \\ &\leq s_k^{-2} \sum_{\ell \in \mathcal{C}_j \setminus \{k\}} \omega_{k,\ell}^* \frac{\|\Sigma_k\|_\infty^{1/2}}{\sqrt{N_k}} \sqrt{\zeta_j^2 + s_\ell^2} \\ &\leq \frac{1}{\sqrt{d_k^e}} \left((1 - \omega_{k,k}^*) \sqrt{\tau_{j,k}} + \frac{\nu_{j,k} s_k}{1 + \tau_{j,k} (1 - \nu_{j,k})} \sum_{\ell \in \mathcal{C}_j \setminus \{k\}} s_\ell^{-1} \right) \\ &\leq \frac{1}{\sqrt{d_k^e}} \left((1 - \omega_{k,k}^*) \sqrt{\tau_{j,k}} + \nu_{j,k} s_k \sum_{\ell \in \mathcal{C}_j} s_\ell^{-1} \right), \end{aligned} \quad (104)$$

where we have used: $\|\mu_\ell - \mu_k\| \leq \zeta_j$ as tasks k and ℓ are in the group \mathcal{C}_j ; $(\|\Sigma_k\|_\infty/N_k)^{1/2} = s_k/\sqrt{d_k^e}$; and the explicit expression (13) for the oracle weights $\omega_{k,\ell}^*$ for group \mathcal{C}_j . For the first term of (104), for all $k \in \mathcal{C}_j$ we have:

$$(1 - \omega_{k,k}^*) \sqrt{\tau_{j,k}} = \frac{1 - \nu_{j,k}}{1 + \tau_{j,k} (1 - \nu_{j,k})} \sqrt{\tau_{j,k}} \leq \frac{\sqrt{\tau_{j,k}}}{1 + \tau_{j,k}} \leq 1.$$

For the second term of (104), introduce the vector $\boldsymbol{\rho} := (s_\ell^{-1})_{\ell \in \mathcal{C}_j}$ and observe that $\nu_{j,k} = \rho_k^2 / \|\boldsymbol{\rho}\|_2^2$, thus, when summing over $k \in \mathcal{C}_j$:

$$\sum_{k \in \mathcal{C}_j} \left(\nu_{j,k} s_k \sum_{\ell \in \mathcal{C}_j} s_\ell^{-1} \right) = \sum_{k \in \mathcal{C}_j} \rho_k \frac{\|\boldsymbol{\rho}\|_1}{\|\boldsymbol{\rho}\|_2^2} = \frac{\|\boldsymbol{\rho}\|_1^2}{\|\boldsymbol{\rho}\|_2^2} \leq |\mathcal{C}_j|.$$

We deduce from the above estimates:

$$\sum_{k \in \mathcal{C}_j} \frac{Q_k(\boldsymbol{\omega}_k^*)}{s_k^2} \leq \frac{2|\mathcal{C}_j|}{\min_k (d_k^e)^{1/2}},$$

implying

$$\frac{1}{B} \sum_{k=1}^B \frac{Q(\boldsymbol{\omega}_k^*)}{s_k^2} \leq \frac{2}{\min_k (d_k^e)^{1/2}}.$$

Therefore for any J -partition \mathcal{C} , since $d_k^\bullet \geq d_k^e$:

$$\frac{1}{B} \sum_{k=1}^B \frac{R_k(\widehat{\omega}_k)}{s_k^2} \leq \left(1 + CB e^{-u_0}\right) \left(\frac{1}{B} \sum_{j=1}^J \sum_{k \in \mathcal{C}_j} \mathcal{B}(\tau_{j,k}, \nu_{j,k}) + C' \frac{u_0}{\min_{k \in \llbracket B \rrbracket} (d_k^e)^{1/2}}\right).$$

□

E.3. Proof of Proposition 9. The proof follows the same steps as the proof of Proposition 7. Let \mathcal{C} a J -partition of $\llbracket B \rrbracket$, $\zeta \in \mathbb{R}_+^J$ and Σ a definite positive matrix in $\mathbb{R}^{d \times d}$. W.l.g. we can assume that $\text{Tr} \Sigma = 1$. In a first time, we are going to lower bound the minimax risk for the estimation of μ_1 that we can assume to be in the cluster 1 ($1 \in \mathcal{C}_1$).

If for $j \in \llbracket J \rrbracket$ the means of \mathcal{C}_j are in a ball of radius $\zeta_j/2$, then two means are at a distance at most ζ_j :

$$\inf_{\widehat{\mu}_1} \sup_{\mathbb{P} \in \mathcal{P}_{\text{mult}}(\mathcal{C}, \zeta, \Sigma, \mathbf{s}^2)} R_1(\widehat{\mu}_1) \geq \inf_{\widehat{\mu}_1} \sup_{\substack{\exists \nu_1, \dots, \nu_J \in \mathbb{R}^d \\ \mu_k \in B(\nu_j, \zeta_j/2), \forall k \in \mathcal{C}_j}} R_1(\widehat{\mu}_1).$$

For simplicity, the supremum over the vectors means μ_k is used to denote the supremum over the Gaussian distributions $\mathbb{P}_k = \mathcal{N}(\mu_k, s_k^2 \Sigma)$.

We can restrict ourself in the case where the centres ν_j are in a ball around 0 of radius $\sqrt{\beta}$:

$$\inf_{\widehat{\mu}_1} \sup_{\substack{\exists \nu_1, \dots, \nu_J \in \mathbb{R}^d \\ \mu_k \in B(\nu_j, \zeta_j/2), \forall k \in \mathcal{C}_j}} R_1(\widehat{\mu}_1) \geq \inf_{\widehat{\mu}_1} \sup_{\substack{\exists \nu_1, \dots, \nu_J \in B(0, \sqrt{\beta}) \\ \mu_k \in B(\nu_j, \zeta_j/2), \forall k \in \mathcal{C}_j}} R_1(\widehat{\mu}_1)$$

Let $\alpha \in (0, 1)$, $\beta > 0$, we introduce now the probability measure $\mathbb{Q} = \mathbb{Q}(\alpha, \beta)$ on $(\mathbb{R}^d)^{B+J}$ such that a random vector $(\nu_1, \dots, \nu_J, \mu_1, \dots, \mu_B) \in (\mathbb{R}^d)^{B+J}$ follows the distribution \mathbb{Q} if:

$$\nu_j \stackrel{\mathbb{Q}}{\sim} \mathcal{N}(0, \alpha\beta\Sigma) \text{ for } k \in \llbracket \mathcal{M} \rrbracket, \quad \mu_k \stackrel{\mathbb{Q}}{\sim} \mathcal{N}(\nu_j, \alpha \frac{\zeta_j^2}{4} \Sigma) \text{ for } k \in \mathcal{C}_j.$$

Hence, considering the events $H_j := \{\|\nu_j\|^2 \leq \beta, \|\mu_k - \nu_j\|^2 \leq \zeta_j^2/4, k \in \mathcal{C}_j\}$, $H := \cap_{j=1}^J H_j$, as in the equations (101):

$$\inf_{\widehat{\mu}_1} \sup_{\mathbb{P} \in \mathcal{P}_{\text{mult}}(\mathcal{C}, \zeta, \Sigma, \mathbf{s}^2)} R_1(\widehat{\mu}_1) \geq \inf_{\widehat{\mu}_1} \mathbb{E}_{\mathbb{Q}}[R_1(\widehat{\mu}_1) | H].$$

The distribution \mathbb{Q} can be decomposed into a product of J probability measure: $\mathbb{Q} = \otimes_{j=1}^J \mathbb{Q}_j$ where \mathbb{Q}_j is the distribution of $(\nu_j, (\mu_k)_{k \in \mathcal{C}_j})$. By independence, the Bayes estimator of μ_1 only consider the means of \mathcal{C}_1 and following equations (101) we get:

$$\inf_{\widehat{\mu}_1} \sup_{\mathbb{P} \in \mathcal{P}_{\text{mult}}(\mathcal{C}, \zeta, \Sigma, \mathbf{s}^2)} R_1(\widehat{\mu}_1) \geq \inf_{\widehat{\mu}_1} \mathbb{E}_{\mathbb{Q}_1}[R_1(\widehat{\mu}_1) | H_1] \geq \frac{1}{\mathbb{Q}(H_1)} (I_1 - r_1),$$

where

$$I_1 := \inf_{\widehat{\mu}_1} \mathbb{E}_{\mathbb{Q}_1}[R_1(\widehat{\mu}_1)], \quad r_1 := \sup_{\widehat{\mu}_1: \|\widehat{\mu}_1\| \leq 2\sqrt{\beta} + \zeta_1} \mathbb{E}_{\mathbb{Q}}[R_1(\widehat{\mu}_1) 1_{H_1^c}] \quad (105)$$

We have used that the infimum is attained for an estimator $\widehat{\mu}_1$ bounded by $2\sqrt{\beta} + \zeta_1$, because the estimator $\widehat{\mu} = 0$ beats the estimators outside that ball (as in (99)).

Lower bound for I_1 : The infimum is attained for $\widehat{\mu}_1 = \mathbb{E}[\mu_1 | X_{\bullet}^{(k)} \text{ } k \in \mathcal{C}_1]$. Let us calculate $\widehat{\mu}_1$. We will denote in the rest of the proof $\tilde{\zeta}_j := \zeta_j/2$:

$$\begin{aligned} \mathbb{E}[\mu_1 | \nu_1, X_{\bullet}^{(k)} \text{ } k \in \mathcal{C}_1] &= \mathbb{E}[\mu_1 | \nu_1, X_{\bullet}^{(1)}] = \frac{\alpha \tilde{\zeta}_1^2}{s_1^2 + \alpha \tilde{\zeta}_1^2} \widehat{\mu}_1^{\text{NE}} + \frac{s_1^2}{s_1^2 + \alpha \tilde{\zeta}_1^2} \nu_1, \\ \mathbb{E}[\nu_1 | X_{\bullet}^{(k)} \text{ } k \in \mathcal{C}_1] &= \left((\alpha\beta)^{-1} + \sum_{i \in \mathcal{C}_1} (\alpha \tilde{\zeta}_1^2 + s_k^2)^{-1} \right)^{-1} \sum_{k \in \mathcal{C}_1} \frac{1}{\alpha \tilde{\zeta}_1^2 + s_k^2} \widehat{\mu}_k^{\text{NE}} \end{aligned}$$

Combining these two expressions:

$$\begin{aligned} \mathbb{E}[\mu_1 | X_{\bullet}^{(k)} \text{ } k \in \mathcal{C}_1] &= \\ &= \frac{\alpha \tilde{\zeta}_1^2}{s_1^2 + \alpha \tilde{\zeta}_1^2} \widehat{\mu}_1^{\text{NE}} + \frac{s_1^2}{s_1^2 + \alpha \tilde{\zeta}_1^2} \left((\alpha\beta)^{-1} + \sum_{k \in \mathcal{C}_1} (\alpha \tilde{\zeta}_1^2 + s_k^2)^{-1} \right)^{-1} \sum_{k \in \mathcal{C}_1} \frac{1}{\alpha \tilde{\zeta}_1^2 + s_k^2} \widehat{\mu}_k^{\text{NE}} \end{aligned}$$

Let $\kappa_1 := \left((\alpha\beta)^{-1} + \sum_{k \in \mathcal{C}_1} (\alpha \tilde{\zeta}_1^2 + s_k^2)^{-1} \right)^{-1}$, we can first notice that:

$$\begin{aligned} \mathbb{E}[\mu_1 | X_{\bullet}^{(\cdot)}] - \mu_1 &= \left[\frac{\alpha \tilde{\zeta}_1^2}{s_1^2 + \alpha \tilde{\zeta}_1^2} + \frac{\kappa_1 s_1^2}{(s_1^2 + \alpha \tilde{\zeta}_1^2)^2} \right] (\widehat{\mu}_1^{\text{NE}} - \mu_1) \\ &+ \frac{\kappa_1 s_1^2}{s_1^2 + \alpha \tilde{\zeta}_1^2} \sum_{k \in \mathcal{C}_1 \setminus \{1\}} \frac{1}{\alpha \tilde{\zeta}_1^2 + s_k^2} (\widehat{\mu}_k^{\text{NE}} - \nu_1) \\ &- \frac{s_1^2}{s_1^2 + \alpha \tilde{\zeta}_1^2} \left(1 - \frac{\kappa_1 s_1^2}{s_1^2 + \alpha \tilde{\zeta}_1^2} \right) (\mu_1 - \nu_1) - \frac{\kappa_1 s_1^2}{s_1^2 + \alpha \tilde{\zeta}_1^2} \frac{1}{\alpha\beta} \nu_1. \end{aligned}$$

Using that $\widehat{\mu}_k^{\text{NE}} - \nu_1$ for $k \in \mathcal{C}_1 \setminus \{1\}$, $\widehat{\mu}_1^{\text{NE}} - \mu_1$, $\mu_1 - \nu_1$ and ν_1 are pairwise independent we get that:

$$\begin{aligned} \mathbb{E}[\|\widehat{\mu}_1 - \mu_1\|^2] &= \left[\frac{\alpha \tilde{\zeta}_1^2}{s_1^2 + \alpha \tilde{\zeta}_1^2} + \frac{\kappa_1 s_1^2}{(s_1^2 + \alpha \tilde{\zeta}_1^2)^2} \right]^2 s_1^2 + \frac{\kappa_1^2 s_1^4}{(s_1^2 + \alpha \tilde{\zeta}_1^2)^2} \sum_{k \in \mathcal{C}_1 \setminus \{1\}} \frac{1}{\alpha \tilde{\zeta}_1^2 + s_k^2} \\ &+ \frac{s_1^4}{(s_1^2 + \alpha \tilde{\zeta}_1^2)^2} \left(1 - \frac{\kappa_1 s_1^2}{s_1^2 + \alpha \tilde{\zeta}_1^2} \right)^2 \alpha \tilde{\zeta}_1^2 + \frac{\kappa_1^2 s_1^4}{(s_1^2 + \alpha \tilde{\zeta}_1^2)^2} \frac{1}{\alpha\beta}. \end{aligned}$$

After simplification:

$$\frac{I_1}{s_1^2} = \frac{\alpha \tilde{\zeta}_1^2}{s_1^2 + \alpha \tilde{\zeta}_1^2} + \frac{\kappa_1 s_1^2}{(s_1^2 + \alpha \tilde{\zeta}_1^2)^2} \quad (106)$$

Upper bound for r_1 : By the same arguments of equations (103):

$$\sup_{\hat{\mu}_1: \|\hat{\mu}_1\| \leq 2\sqrt{\beta} + \zeta_1} \mathbb{E} \left[\|\hat{\mu}_1(X_{\bullet}^{(k)}, k \in \mathcal{C}_1) - \mu_1\|^2 \mathbf{1}_{H_1^c} \right] \leq 20(\beta + \zeta_1^2) \sqrt{\mathbb{P}[H_1^c]}$$

From Lemma 5, for all $k \in \mathcal{C}_1$:

$$\mathbb{P}[\|\nu_1\|^2 \geq \beta] = \mathbb{P}[\|\mu_k - \nu_1\|^2 \geq \zeta_1^2/2] \leq \exp\left(-\frac{d^e}{2} \left(\sqrt{\frac{2}{\alpha}} - 1 - 1\right)\right),$$

and by union bound we get that :

$$r_1 \leq 20(\beta + \zeta_1^2) \sqrt{|\mathcal{C}_1| + 1} \exp\left(-\frac{d^e}{4} \left(\sqrt{\frac{2}{\alpha}} - 1 - 1\right)\right).$$

where $d^e = \text{Tr } \Sigma / \|\Sigma\|_\infty$.

Compound bound. We recall that $\mathbb{Q} = \bigotimes_{j=1}^J \mathbb{Q}_j$ where \mathbb{Q}_j is the distribution of $(\nu_j, \mu_k$ for $k \in \mathcal{C}_j$). Then let $\hat{\mu} = (\hat{\mu}_k)_{k \in [B]} \in (\mathbb{R}^d)^B$ be an estimator of the vectors $(\mu_k)_{k \in [B]}$:

$$\begin{aligned} \inf_{\hat{\mu}} \sup_{\mathbb{P} \in \mathcal{P}_{\text{mult}}(\mathcal{C}, \zeta, \Sigma, \mathbf{s}^2)} \frac{1}{B} \sum_{k=1}^B \frac{R_k(\hat{\mu}_k)}{s_k^2} &\geq \inf_{\hat{\mu}} \frac{1}{\mathbb{Q}(H)} \int_H \frac{1}{B} \sum_{k=1}^B \frac{R_k(\hat{\mu}_k)}{s_k^2} d\mathbb{Q}(\nu_1, \dots, \nu_N, \mu_1, \dots, \mu_B) \\ &= \inf_{\hat{\mu}} \frac{1}{B} \sum_{j=1}^J \sum_{k \in \mathcal{C}_j} \frac{\mathbb{Q}(H_{-j})}{\mathbb{Q}(H)} \int_{H_j} \frac{R_k(\hat{\mu}_k)}{s_k^2} d\mathbb{Q}_j(\nu_j, (\mu_\ell)_{\ell \in \mathcal{C}_j}) \end{aligned}$$

where we recall $H_j = \{\|\nu_j\|^2 \leq \beta, \|\mu_k - \nu_j\|^2 \leq \tilde{\zeta}_j^2, \forall k \in \mathcal{C}_j\}$, $H = \bigcap_{j=1}^J H_j$ and $H_{-j} = \bigcap_{\ell \neq j} H_\ell$. Using that $\mathbb{Q}(H_{-j})/\mathbb{Q}(H) = \mathbb{Q}_j(H_j)^{-1} \geq 1$ and that the infimum over estimators $\hat{\mu}$ of the sum is the sum of the infimum over estimators $\hat{\mu}_k$, we get that:

$$\begin{aligned} \inf_{\hat{\mu}} \sup_{\mathbb{P} \in \mathcal{P}_{\text{mult}}(\mathcal{C}, \zeta, \Sigma, \mathbf{s}^2)} \frac{1}{B} \sum_{k=1}^B \frac{R_k(\hat{\mu}_k)}{s_k^2} &\geq \frac{1}{B} \sum_{j=1}^J \sum_{k \in \mathcal{C}_j} (I_k - r_k) \\ &\geq \frac{1}{B} \sum_{j=1}^J \sum_{k \in \mathcal{C}_j} \frac{\alpha \tilde{\zeta}_j^2}{s_k^2 + \alpha \tilde{\zeta}_j^2} + \frac{\kappa_j s_k^2}{(s_k^2 + \alpha \tilde{\zeta}_j^2)^2} - \frac{20}{B} \left(\sum_{j=1}^J |\mathcal{C}_j|^{3/2} \frac{\beta + \tilde{\zeta}_j^2}{s^2(\mathcal{C}_j)} \right) \exp(-d^e c(\alpha)) \quad (107) \end{aligned}$$

where $\kappa_j = \left((\alpha\beta)^{-1} + \sum_{k \in \mathcal{C}_j} (\alpha \tilde{\zeta}_j^2 + s_k^2)^{-1} \right)^{-1}$ and $c(\alpha) = \left(\sqrt{\frac{2}{\alpha}} - 1 - 1 \right) / 4$.

Conclusion : Let $d^e \rightarrow \infty$ in (107), then:

$$\lim_{d^e \rightarrow \infty} \inf_{\hat{\mu}} \sup_{\mathbb{P} \in \mathcal{P}_{\text{mult}}(\mathcal{C}, \zeta, \Sigma, \mathbf{s}^2)} \frac{1}{B} \sum_{k=1}^B \frac{R_k(\hat{\mu}_k)}{s_k^2} \geq \frac{1}{B} \sum_{j=1}^J \sum_{k \in \mathcal{C}_j} \frac{\alpha \tilde{\zeta}_j^2}{s_k^2 + \alpha \tilde{\zeta}_j^2} + \frac{\kappa_j s_k^2}{(s_k^2 + \alpha \tilde{\zeta}_j^2)^2}$$

Let $\alpha \rightarrow 1$ and $\beta \rightarrow \infty$, then:

$$\liminf_{d^e \rightarrow \infty} \inf_{\hat{\boldsymbol{\mu}}} \sup_{\mathbb{P} \in \mathcal{P}_{\text{mult}}(\mathcal{C}, \zeta, \Sigma, \mathbf{s}^2)} \frac{1}{B} \sum_{k=1}^B \frac{R_k(\hat{\boldsymbol{\mu}}_k)}{s_k^2} \geq \frac{1}{B} \sum_{j=1}^J \sum_{k \in \mathcal{C}_j} \frac{\tilde{\zeta}_j^2}{s_k^2 + \tilde{\zeta}_j^2} + \frac{s_k^2}{s_k^2 + \tilde{\zeta}_j^2} \frac{1}{\sum_{\ell \in \mathcal{C}_j} \frac{s_k^2 + \tilde{\zeta}_j^2}{s_\ell^2 + \tilde{\zeta}_j^2}} \quad (108)$$

We conclude by remarking that for all $j \in \llbracket J \rrbracket$:

$$\sum_{\ell \in \mathcal{C}_j} \frac{s_k^2 + \tilde{\zeta}_j^2}{s_\ell^2 + \tilde{\zeta}_j^2} = 1 + \sum_{\ell \in \mathcal{C}_j \setminus \{k\}} \frac{s_k^2 + \tilde{\zeta}_j^2}{s_\ell^2 + \tilde{\zeta}_j^2} \leq 1 + \sum_{\ell \in \mathcal{C}_j \setminus \{k\}} \frac{s_k^2 + \tilde{\zeta}_j^2}{s_\ell^2}$$

Then:

$$\begin{aligned} & \frac{1}{B} \sum_{j=1}^J \sum_{k \in \mathcal{C}_j} \frac{\tilde{\zeta}_j^2}{s_k^2 + \tilde{\zeta}_j^2} + \frac{s_k^2}{s_k^2 + \tilde{\zeta}_j^2} \frac{1}{\sum_{\ell \in \mathcal{C}_j} \frac{s_k^2 + \tilde{\zeta}_j^2}{s_\ell^2 + \tilde{\zeta}_j^2}} \\ & \geq \frac{1}{B} \sum_{j=1}^J \sum_{k \in \mathcal{C}_j} \frac{\tilde{\zeta}_j^2}{s_k^2 + \tilde{\zeta}_j^2} + \frac{s_k^2}{s_k^2 + \tilde{\zeta}_j^2} \frac{1}{1 + \sum_{\ell \in \mathcal{C}_j \setminus \{k\}} \frac{s_k^2 + \tilde{\zeta}_j^2}{s_\ell^2}} = \mathcal{L}^*(\mathbf{s}, \mathbf{C}, \zeta/2). \end{aligned}$$

□

E.4. Proof of Proposition 10. We start with the following elementary bounds on the function \mathcal{B} (for $\tau \geq 0, \nu \in [0, 1]$)

$$\mathcal{B}(\tau, \nu) \leq \frac{\tau + \nu}{1 + \tau} \leq \max\left(1, \frac{\tau}{1 + \tau} + \nu\right). \quad (109)$$

Now consider the quantity $A_j := |\mathcal{C}_j|^{-1} \sum_{k \in \mathcal{C}_j} \mathcal{B}(\tau_{j,k}, \nu_{j,k})$. Observe that $\sum_{k \in \mathcal{C}_j} \nu_{j,k} = 1$ and $\tau_{j,k} = \nu_{j,k} B_j$, where $B_j := \zeta_j^2 / s^2(\mathcal{C}_j)$. Thus

$$A_j := |\mathcal{C}_j|^{-1} \sum_{k \in \mathcal{C}_j} \mathcal{B}(B_j \nu_{j,k}, \nu_{j,k}) \leq (B_j + 1) |\mathcal{C}_j|^{-1} \sum_{k \in \mathcal{C}_j} \frac{\nu_{j,k}}{1 + B_j \nu_{j,k}}.$$

where we have used the first inequality in (109). By concavity of $t \mapsto t/(1+t)$ we conclude to

$$A_j \leq \frac{B_j |\mathcal{C}_j|^{-1} + |\mathcal{C}_j|^{-1}}{1 + B_j |\mathcal{C}_j|^{-1}} = \frac{\bar{\tau}_j + |\mathcal{C}_j|^{-1}}{1 + \bar{\tau}_j},$$

and thus to (44) by summation over $j \in \llbracket J \rrbracket$. Now using the second inequality in (109), we obtain

$$\sum_{j \in \llbracket J \rrbracket} \frac{|\mathcal{C}_j|}{B} \frac{\bar{\tau}_j + |\mathcal{C}_j|^{-1}}{1 + \bar{\tau}_j} \leq \sum_{j \in \llbracket J \rrbracket} \frac{|\mathcal{C}_j|}{B} \min\left(1, \frac{\bar{\tau}_j}{1 + \bar{\tau}_j} + |\mathcal{C}_j|^{-1}\right) \leq \min\left(1, \frac{\bar{\tau}_*}{1 + \bar{\tau}_*} + \frac{J}{B}\right),$$

where we have used the second inequality in (109) and the biconcave character of the function $(x, y) \mapsto \min(1, y + x/(1+x))$; thus establishing (45). Assume now that all risks and the diameters are equal, i.e. $s_k^2 = s^2$ and $\zeta_j = \zeta$ for $k \in \llbracket B \rrbracket$ and $j \in \llbracket J \rrbracket$.

Then for all $j \in \llbracket J \rrbracket$ and $k \in \llbracket B \rrbracket$, $\bar{s}^2(\mathcal{C}_j) = s^2$, $\bar{\tau}_{j,k} = \zeta^2/s^2 = \bar{\tau}$ and $\nu_{j,k} = |\mathcal{C}_j|^{-1}$. Using the elementary bound

$$\mathcal{B}(\tau, \nu) \geq \frac{\tau}{1 + \tau} + \frac{\nu}{(1 + \tau)^2},$$

we thus have in this case

$$\begin{aligned} \mathcal{L}^*(\mathbf{s}, \mathcal{C}, \zeta) &= \frac{1}{B} \sum_{j=1}^J \sum_{k \in \mathcal{C}_j} \mathcal{B}(\tau_{j,k}, \nu_{j,k}) = \frac{1}{B} \sum_{j=1}^J \sum_{k \in \mathcal{C}_j} \mathcal{B}(\bar{\tau}, |\mathcal{C}_j|^{-1}) \\ &\geq \frac{1}{B} \sum_{j=1}^J \sum_{k \in \mathcal{C}_j} \left(\frac{\bar{\tau}}{1 + \bar{\tau}} + \frac{|\mathcal{C}_j|^{-1}}{(1 + \bar{\tau})^2} \right) \\ &= \frac{\bar{\tau}}{1 + \bar{\tau}} + \frac{J}{B} \frac{1}{(1 + \bar{\tau})^2}, \end{aligned} \tag{110}$$

Finally, since for $\tau \geq 0, \nu \in [0, 1]$:

$$\begin{aligned} \frac{\tau}{1 + \tau} + \frac{\nu}{(1 + \tau)^2} &\geq \max \left(\frac{\tau}{1 + \tau}, \frac{1}{(1 + \tau)^2} \left(\frac{\tau}{1 + \tau} + \nu \right) \right) \\ &\geq \max \left(\frac{\tau}{1 + \tau}, \frac{1}{(1 + \tau)^2} \right) \min \left(1, \frac{\tau}{1 + \tau} + \nu \right) \\ &\geq 0.38 \min \left(1, \frac{\tau}{1 + \tau} + \nu \right), \end{aligned}$$

we conclude that in the case of equal risks and diameters the upper bound (45) and the lower bound (110) differ by a factor at most $1/0.38 \leq 2.7$. \square

APPENDIX F. ABOUT THE CONSTANT ϕ IN THE TRANSLATION-INVARIANT KERNEL SETTING

In this section, we investigate the distribution-dependent constant $\phi = M^2/(\text{Tr } \Sigma)$ in the **(BS)** setting (i.e., for data bounded in norm by the constant M). This constant comes into play in the risk bounds for our methods, in relation to sufficient sample sizes, see e.g. Props. 17, 6. Rewriting $\text{Tr } \Sigma = \mathbb{E}[\|X - \mathbb{E}[X]\|^2]$ yields a direct interpretation of ϕ , namely it is the ratio between the known bound on $\|X\|$ and the ‘‘variance’’ of X ; in other words, ϕ is all the bigger as the variable X is more concentrated in relation to the size of its support.

We are interested in an understanding more detailed than this simple observation in the situation of kernel mean embedding (KME), which was our primary motivation for investigating the **(BS)** setting. Namely, in that situation the user might choose between different kernels and their associated Hilbert space mappings, in particular choosing or tuning the ‘‘kernel bandwidth’’. Even if kernels under consideration are all bounded by the same constant, different kernels may give rise to different constants ϕ for the same underlying data distribution.

We look into this issue under the following general conditions:

- (K1) the original data takes values in $\mathcal{Z} = \mathbb{R}^\ell$, and the data whose means we wish to estimate have been obtained via a Hilbert space mapping $X = \Phi_\kappa(Z)$, $\Phi_\kappa : \mathbb{R}^\ell \rightarrow \mathcal{H}$, associated to the kernel $\kappa(z, z') = \langle \Phi_\kappa(z), \Phi_\kappa(z') \rangle$.
- (K2) κ is a translation-invariant kernel on \mathbb{R}^ℓ , of the form $\kappa(z, z') = F(z - z')$, where $F : \mathbb{R}^\ell \rightarrow \mathbb{R}$, with $M^2 := F(0)$.
- (K3) For any $u \in \mathbb{R}^\ell$, the function $\lambda \mapsto F(\lambda u)$ is nonincreasing on \mathbb{R}_+ . Furthermore, there exist constants $h > 0, c \leq 1$ such that

$$F(u) \leq M^2 \left(1 - c \frac{\|u\|^2}{h^2} \right), \text{ for all } u \in \mathbb{R}^\ell \text{ s.t. } 0 \leq \|u\| \leq h. \quad (111)$$

Observe that (K1)-(K2) imply that the mapped data X satisfies **(BS)**; as for (K3), it means that the kernel is locally upper bounded by a strongly concave function in a neighbourhood of 0 of size h . The latter quantity can therefore be interpreted as a proxy bandwidth for the kernel; and if F_1 satisfies (111) for $h = 1$ then the rescaled kernel function $F_h(u) := F_1(u/h)$ satisfies (111) for the bandwidth parameter $h > 0$. The classical Gaussian, exponential, and Matérn kernels, for example, satisfy such conditions.

Proposition 22. *Assume (K1)-(K2)-(K3) hold, and that the distribution P of the original data Z in \mathbb{R}^ℓ satisfies the following norm moment condition for some $p \geq 1, C > 0$:*

$$\frac{\mathbb{E}[\xi^{2p}]}{\mathbb{E}[\xi^p]^2} \leq C, \quad \text{where } \xi := \|Z - \mathbb{E}[Z]\|. \quad (112)$$

Then it holds

$$\phi = \frac{M^2}{\mathbb{E}[\|X - \mathbb{E}[X]\|^2]} \leq \frac{4 \cdot 2^{\frac{2}{p} + 2p} C}{c} \max \left(1, \frac{h}{2\mathbb{E}[\|Z - \mathbb{E}[Z]\|^p]^{\frac{1}{p}}} \right)^2.$$

Assume $p = 2$ to simplify (we allowed for other values of p in the moment condition (112) mainly with the possible value $p = 1$ in mind, which makes the condition weaker; the discussion below can be readily adapted to other values of p). This result shows that, provided the bandwidth parameter h is chosen of the order of $\sigma_Z := \mathbb{E}[\|Z - \mathbb{E}[Z]\|^2]^{\frac{1}{2}}$ or smaller, the constant ϕ for the mapped data is bounded independently of h . The bound depends on (1) the strong concavity parameter c of the upper bound on the (unit scaled) kernel function in a neighbourhood of the origin, and (2) the norm moment ratio (112) of the original data distribution. Since $\mathbb{E}[\xi^4] \leq \mathbb{E}[\xi^2] \|\xi\|_{L^\infty}^2$, in the case where the original data is itself bounded in norm by a constant R , (112) holds with $C = (R/\sigma_Z)^2$. Thus, if the original X data is bounded, the distribution of the mapped data Z under the above conditions ‘‘inherits’’ the constant ϕ from that of the original data, up to factors. However, the norm moment condition

is much milder than a boundedness condition and can also accommodate unbounded distributions with heavy tails of the original data.

Proof of Proposition 22.

For $Z, Z' \sim \mathbb{P}$ independent, denote $D := \|Z - Z'\|$, $\theta := \min\left(\frac{h^p}{\mathbb{E}[D^p]}, \frac{1}{2}\right)$, and $t^p := \theta \mathbb{E}[D^p] = \min\left(h^p, \frac{\mathbb{E}[D^p]}{2}\right)$, it holds

$$\begin{aligned} \|\mathbb{E}[\Phi_\kappa(Z)]\|^2 / M^2 &= M^{-2} \mathbb{E}[\langle \Phi_\kappa(Z), \Phi_\kappa(Z') \rangle] \\ &= M^{-2} \mathbb{E}[F(Z - Z')] \\ &\leq 1 - c \frac{t^2}{h^2} \mathbb{P}[D^p > t^p] \\ &\leq 1 - c \frac{\mathbb{E}[D^p]^{\frac{2}{p}}}{h^2} \theta^{\frac{2}{p}} (1 - \theta)^2 \frac{\mathbb{E}[D^p]^2}{\mathbb{E}[D^{2p}]} \\ &\leq \frac{\mathbb{E}[\|\Phi_\kappa(Z)\|^2]}{M^2} - \frac{c}{4} \min\left(1, \frac{\mathbb{E}[D^p]^{\frac{2}{p}}}{2^{\frac{2}{p}} h^2}\right) \frac{\mathbb{E}[D^p]^2}{\mathbb{E}[D^{2p}]}, \end{aligned}$$

where the first inequality stems from (K3); the second comes from the Paley-Zygmund inequality; and we used $\theta \leq \frac{1}{2}$ for the third. Since $\mathbb{E}[\|\Phi_\kappa(Z)\|^2] - \|\mathbb{E}[\Phi_\kappa(Z)]\|^2 = \mathbb{E}[\|X\|^2] - \|\mathbb{E}[X]\|^2 = \mathbb{E}[\|X - \mathbb{E}[X]\|^2]$, we deduce

$$\frac{M^2}{\mathbb{E}[\|X - \mathbb{E}[X]\|^2]} \leq \frac{4.2^{\frac{2}{p}}}{c} \max\left(1, \frac{h}{\|D\|_{L^p(P)}}\right)^2 \left(\frac{\mathbb{E}[D^{2p}]}{\mathbb{E}[D^p]^2}\right).$$

Finally, note that

$$\mathbb{E}[D^{2p}] = \mathbb{E}[\|Z - Z'\|^{2p}] \leq \mathbb{E}[(\|Z - \mathbb{E}[Z]\| + \|Z' - \mathbb{E}[Z']\|)^{2p}] \leq 2^{2p} \mathbb{E}[\|Z - \mathbb{E}[Z]\|^{2p}],$$

and by Jensen's inequality

$$\mathbb{E}[\|Z - \mathbb{E}[Z]\|^p] = \mathbb{E}[\|Z - \mathbb{E}[Z']\|^p] \leq \mathbb{E}[\|Z - Z'\|^p] = \mathbb{E}[D^p].$$

(Observe that the equality $\mathbb{E}[\|Z - Z'\|^2] = 2\mathbb{E}[\|Z - \mathbb{E}[Z]\|^2]$ holds, so the constants in the first, resp. second inequality above can be improved for the special cases $p = 1$, resp. $p = 2$.) \square

APPENDIX G. DESCRIPTION OF THE TESTED METHODS

The tested methods propose KME estimations of the form

$$\hat{\mu}_i^{\mathbf{m}} := \sum_{j \in \llbracket B \rrbracket} \omega_{ij}^{\mathbf{m}} \cdot \hat{\mu}_j^{\text{NE}},$$

where the definition of the weighting $\omega_{ij}^{\mathbf{m}}$ depends on the applied method \mathbf{m} .

G.1. State-of-the-Art Approaches.

(i) NE considers each bag individually.

$$\omega_{ij}^{\text{NE}} = \begin{cases} 1, & \text{for } i = j \\ 0, & \text{otherwise.} \end{cases}$$

(ii) R-KMSE (Muandet, Sriperumbudur, et al., 2016) estimates each KME individually but shrinks it towards $\mathbf{0}$. The amount of shrinkage is data dependent

$$\omega_{ij}^{\text{R-KMSE}} = \begin{cases} 1 - \frac{\lambda_i}{1+\lambda_i}, & \text{for } i = j \\ 0, & \text{otherwise} \end{cases}$$

where

$$\lambda_i = \frac{\varrho_i - \rho_i}{(1/N_i - 1)\varrho_i + (N_i - 1)\rho_i}$$

with $\varrho_i = 1/N_i \sum_{n=1}^{N_i} \kappa(Z_n^{(i)}, Z_n^{(i)})$ and $\rho_i = 1/N_i^2 \sum_{n,n'=1}^{N_i} \kappa(Z_n^{(i)}, Z_{n'}^{(i)})$.

(iii) MTA `const` (Feldman et al., 2014) was initially proposed for the estimation for multiple real means. We adapted the approach such that it can be applied to the estimation of multiple kernel means

$$\omega_{ij}^{\text{MTA const}} = \left(\left(I + \frac{\gamma}{B} \hat{S} \cdot L(A) \right)^{-1} \right)_{ij}. \quad (113)$$

Here, $\hat{S} = \text{diag}((\hat{s}_i^2)_{i \in [B]})$, as defined in (3), can be estimated as

$$\hat{s}_i^2 = \frac{1}{2N_i^2(N_i - 1)} \sum_{n \neq n'}^{N_i} \kappa(Z_n^{(i)}, Z_n^{(i)}) - 2\kappa(Z_n^{(i)}, Z_{n'}^{(i)}) + \kappa(Z_{n'}^{(i)}, Z_{n'}^{(i)}), \quad (114)$$

which corresponds to (30), and $L(A)$ denotes the graph Laplacian of task-similarity matrix A . For MTA `const` the similarity is assumed to be constant, i.e., $A = a \cdot (\mathbf{1}\mathbf{1}^T)$ with $a = \frac{1}{B(B-1)} \sum_{i,j \in [B]} \|\hat{\mu}_i^{\text{NE}} - \hat{\mu}_j^{\text{NE}}\|_{\mathcal{H}}^2$. The optimal value for model parameter γ may be found using model optimization. As default value we propose $\gamma = 1$.

G.2. AGG Approaches. The aggregation approaches form a convex combination of possibly all bags whose weights are found directly by minimization of quantities related to the squared risk.

(iv) AGG `orth` is based on the constraint optimization problem

$$\omega_i = \underset{w_i}{\text{argmin}} \left\{ \mathbb{E} \left\| \sum_{j \in [B]} w_{ij} \hat{\mu}_j^{\text{NE}} - \mu_i \right\|_{\mathcal{H}}^2 \right\} \text{ s.t. } \sum_{j \in [B]} \omega_{ij} = 1, \quad \forall i, j \in [B] : \omega_{ij} \geq 0.$$

Using Lagrangian multipliers the optimal solution can be derived as

$$\omega_i \simeq (S + \Lambda^{(i)})^{(-1)} \mathbf{1} \quad (115)$$

where $S = \text{diag}((s_i^2)_{i \in \llbracket B \rrbracket})$ and $\Lambda^{(i)} \in \mathbb{R}^{B \times B}$ with $\Lambda_{j,j'}^{(i)} = \langle \mu_j - \mu_i, \mu_{j'} - \mu_i \rangle_{\mathcal{H}}$. Central assumption of **AGG orth** is $\Lambda_{j,j'}^{(i)} = \langle \mu_j - \mu_i, \mu_{j'} - \mu_i \rangle_{\mathcal{H}} = 0$ for all $j \neq j'$ such that $\Lambda^{(i)}$ becomes a diagonal matrix with $\Lambda^{(i)} = \text{diag}((\|\mu_j - \mu_i\|_{\mathcal{H}}^2)_{j \in \llbracket B \rrbracket})$. An unbiased estimation of $\|\mu_j - \mu_i\|_{\mathcal{H}}^2$ is given by (16) which in the kernel setting translates to

$$\begin{aligned} \widehat{\text{MMD}}^2(\mu_i, \mu_j) &= \sum_{n \neq n'} \frac{N_i}{N_i(N_i - 1)} \kappa(Z_n^{(i)}, Z_{n'}^{(i)}) + \sum_{m \neq m'} \frac{N_j}{N_j(N_j - 1)} \kappa(Z_m^{(j)}, Z_{m'}^{(j)}) \\ &\quad - 2 \sum_{n=1}^{N_i} \sum_{m=1}^{N_j} \frac{\kappa(Z_n^{(i)}, Z_m^{(j)})}{N_i N_j}. \end{aligned} \quad (116)$$

Eq. (115) reduces to

$$\omega_{ij}^{\text{AGG orth}} = \frac{1}{\hat{s}_j^2 + \gamma \cdot \widehat{\text{MMD}}^2(\mu_i, \mu_j)}.$$

We add a multiplicative constant γ for more flexibility, whose value is either found by model optimization or $\gamma = 13$ taken as default. If the distances between bags is inhomogeneous, e.g., the data set contains close but also far distant unrelated bags, higher values of γ might be advisable. Finally the weights are normalised such that they sum to one.

(v) **AGG egd** is based on Q-Aggregation and resembles (36)

$$\begin{aligned} \omega_{i \cdot}^{\text{AGG egd}} &= \underset{w_i}{\text{argmin}} \left\{ \hat{L}_i + c_q \hat{Q}_i + c_1 \sum_{j \in \llbracket B \rrbracket} w_{ij} \frac{\sqrt{\text{Tr} \Sigma_j^2}}{N_j} + c_2 \sum_{j \in \llbracket B \rrbracket} w_{ij}^2 \frac{\sqrt{\text{Tr} \Sigma_j^2}}{N_j} \right\}, \\ \hat{L}_i &= \left\| \sum_{j \in \llbracket B \rrbracket} w_{ij} (\hat{\mu}_j^{\text{NE}} - \hat{\mu}_i^{\text{NE}}) \right\|_{\mathcal{H}}^2 + s_i^2 (2w_{ii} - 1), \quad \hat{Q}_i = \sum_{j \in \llbracket B \rrbracket} w_{ij} \sqrt{\frac{\hat{\Delta}_j^T \Sigma_i \hat{\Delta}_j}{N_i}}, \end{aligned}$$

such that $\sum_{j \in \llbracket B \rrbracket} \omega_{ij} = 1$ and $\forall i, j \in \llbracket B \rrbracket : \omega_{ij} \geq 0$. There is no instantiation of \hat{Q}^{BS} . It is required for the theoretical results to hold on bounded data which is less regularised than Gaussian data. In practice, we add two regularization terms instead. The c_1 term favours sparse results whereas the c_2 regularization leads to diffuse, small weights. Their effect can be compared to that of ℓ_1 - and ℓ_2 -regularization respectively. Distant means are penalised by the c_q term.

The optimization over the probability simplex is done by exponentiated gradient descent (Kivinen and Warmuth, 1997) with gradient

$$\nabla \omega_{i \cdot}^{\text{AGG egd}} = 2(\Lambda^{(i)} + c_2 \text{diag}(\vartheta)) \omega_i + 2S_i + c_q \varrho^{(i)} + c_1 \vartheta,$$

where S_i denotes the i -th column of matrix $S = \text{diag}((s_i^2)_{i \in [B]})$, ϑ and $\varrho^{(i)}$ are B -dimensional vectors and defined as $\vartheta_j = \sqrt{\text{Tr}\Sigma_j^2/N_j}$ and $\varrho_j^{(i)} = \sqrt{\Delta_j^T \Sigma_i \Delta_j / N_i}$. We propose the following estimators for these terms: s_i^2 is estimated as shown in (114). Matrix $\check{\Lambda}^{(i)}$ is a biased estimator of $\Lambda^{(i)}$ with $\check{\Lambda}_{j,j'}^{(i)} = \langle \hat{\mu}_j^{\text{NE}} - \hat{\mu}_i^{\text{NE}}, \hat{\mu}_{j'}^{\text{NE}} - \hat{\mu}_i^{\text{NE}} \rangle$ that can be computed as

$$\check{\Lambda}_{j,j'}^{(i)} = \begin{cases} 0, & \text{for } i = j, \text{ or } i = j', \text{ or } i = j = j' \\ \frac{1}{N_j N_{j'}} \sum_{m^j}^{N_j} \sum_{m^{j'}}^{N_{j'}} \kappa(Z_m^{(j)}, Z_{m'}^{(j')}) & \\ - \frac{1}{N_j N_i} \sum_{m^j}^{N_j} \sum_{n^i}^{N_i} \kappa(Z_m^{(j)}, Z_n^{(i)}) & \\ - \frac{1}{N_i N_{j'}} \sum_{n^i}^{N_i} \sum_{m^{j'}}^{N_{j'}} \kappa(Z_n^{(i)}, Z_{m'}^{(j')}) & \\ + \frac{1}{N_i N_i} \sum_{n^i}^{N_i} \sum_{n'^i}^{N_i} \kappa(Z_n^{(i)}, Z_{n'}^{(i)}) & \end{cases} \quad \text{otherwise.} \quad (117)$$

Vector ϑ is based on $\text{Tr}\Sigma_j^2$. Let X_1, X_2, X_3, X_4 denote independent copies, then

$$\Sigma = \mathbb{E}[(X - \mathbb{E}(X))(X - \mathbb{E}(X))^T] = \frac{1}{2} \mathbb{E}[(X_1 - X_2)(X_1 - X_2)^T]$$

such that

$$\Sigma^2 = \frac{1}{4} \mathbb{E}[(X_1 - X_2)(X_1 - X_2)^T (X_3 - X_4)(X_3 - X_4)^T].$$

By linearity of the trace, we then have

$$\begin{aligned} \text{Tr}(\Sigma^2) &= \frac{1}{4} \mathbb{E} \left[\text{Tr} \left((X_1 - X_2)(X_1 - X_2)^T (X_3 - X_4)(X_3 - X_4)^T \right) \right] \\ &= \frac{1}{4} \mathbb{E} \left[(X_1 - X_2)^T (X_3 - X_4) \cdot (X_3 - X_4)^T (X_1 - X_2) \right] \\ &= \mathbb{E} \left[\langle X_1, X_3 \rangle^2 - \langle X_1, X_3 \rangle \langle X_2, X_3 \rangle - \langle X_1, X_3 \rangle \langle X_1, X_4 \rangle + \langle X_1, X_3 \rangle \langle X_2, X_4 \rangle \right]. \end{aligned}$$

For $N_i \geq 4$ an unbiased estimation for $\text{Tr}(\Sigma_i^2)$ is then given by

$$\begin{aligned} & \frac{1}{N_i(N_i - 1)} \sum_{n_1 \neq n_2}^{N_i} \kappa(Z_{n_1}^{(i)}, Z_{n_2}^{(i)})^2 \\ & - \frac{2}{N_i(N_i - 1)(N_i - 2)} \sum_{n_1 \neq n_2 \neq n_3}^{N_i} \kappa(Z_{n_1}^{(i)}, Z_{n_2}^{(i)}) \kappa(Z_{n_1}^{(i)}, Z_{n_3}^{(i)}) \\ & + \frac{1}{N_i(N_i - 1)(N_i - 2)(N_i - 3)} \sum_{n_1 \neq n_2 \neq n_3 \neq n_4}^{N_i} \kappa(Z_{n_1}^{(i)}, Z_{n_2}^{(i)}) \kappa(Z_{n_3}^{(i)}, Z_{n_4}^{(i)}), \end{aligned}$$

and we recover (59). However this estimator has computational complexity $\mathcal{O}(N_i^4)$ and is infeasible in practice. Instead, we propose in Algorithm 1 a subsampling strategy that gives an approximation which operates in $\mathcal{O}(N_i)$.

Algorithm 1 Approximation of estimation of $\text{Tr}(\Sigma_i^2)$

Require: data $Z_{\bullet}^{(i)}$, bag size N_i , number of repetitions r

- 1: # initialise
- 2: $t_1 \leftarrow 0$
- 3: $t_2 \leftarrow 0$
- 4: $t_3 \leftarrow 0$
- 5: # first term can be calculated directly in linear time
- 6: $t_1 \leftarrow \sum_{n,n'}^{N_i} k(Z_n^{(i)}, Z_{n'}^{(i)})^2 - \sum_n^{N_i} k(Z_n^{(i)}, Z_n^{(i)})^2$
- 7: # other terms are approximated in r iterations
- 8: **for** 1 to r **do**
- 9: # select four distinct samples
- 10: $n_1, n_2, n_3, n_4 \leftarrow \text{randint}(1, N_i, 4)$
- 11: # approximate second and third term
- 12: $t_2 \leftarrow t_2 + \kappa(Z_{n_1}^{(i)}, Z_{n_2}^{(i)}) \cdot \kappa(Z_{n_1}^{(i)}, Z_{n_3}^{(i)})$
- 13: $t_3 \leftarrow t_3 + \kappa(Z_{n_1}^{(i)}, Z_{n_2}^{(i)}) \cdot \kappa(Z_{n_3}^{(i)}, Z_{n_4}^{(i)})$
- 14: **end for**
- 15: # normalise and add
- 16: $\text{trS}_i \leftarrow t_1 / (N_i(N_i - 1)) - 2t_2 / r + t_3 / r$
- 17: **return** trS_i

For the vector $\varrho^{(i)}$ we need an estimation of $\hat{\Delta}_j^T \Sigma_i \hat{\Delta}_j$ for which we propose a biased estimate

$$\begin{aligned} \text{dSd}_j^{(i)} &= \frac{1}{N_i - 1} \sum_{n=1}^{N_i} \left(\frac{1}{N_j} \sum_{m=1}^{N_j} \kappa(Z_n^{(i)}, Z_m^{(j)}) - \frac{1}{N_i} \sum_{n'=1}^{N_i} \kappa(Z_n^{(i)}, Z_{n'}^{(i)}) \right)^2 \\ &\quad - \frac{N_i}{N_i - 1} \left(\frac{1}{N_i N_j} \sum_{n=1}^{N_i} \sum_{m=1}^{N_j} \kappa(Z_n^{(i)}, Z_m^{(j)}) - \frac{1}{N_i N_i} \sum_{n=1}^{N_i} \sum_{n'=1}^{N_i} \kappa(Z_n^{(i)}, Z_{n'}^{(i)}) \right)^2 \end{aligned} \quad (118)$$

Note that estimator $\text{dSd}_j^{(1)}$ is a rewriting of \hat{q}_j (31) in the kernel setting. For translation invariant kernels we obtain a less biased estimate

$$\begin{aligned} \text{dSd}_j^{(i)} &= \frac{1}{N_i - 1} \sum_{n=1}^{N_i} \left(\frac{1}{N_j} \sum_{m=1}^{N_j} \kappa(Z_n^{(i)}, Z_m^{(j)}) - \frac{1}{N_i - 2} \sum_{n'=1}^{N_i} \kappa(Z_n^{(i)}, Z_{n'}^{(i)}) \right)^2 \\ &\quad - \frac{N_i}{N_i - 1} \left(\frac{1}{N_i N_j} \sum_{n=1}^{N_i} \sum_{m=1}^{N_j} \kappa(Z_n^{(i)}, Z_m^{(j)}) - \frac{1}{N_i(N_i - 2)} \sum_{n=1}^{N_i} \sum_{n'=1}^{N_i} \kappa(Z_n^{(i)}, Z_{n'}^{(i)}) \right)^2. \end{aligned}$$

Its computational complexity is in $\mathcal{O}(N_i^2)$.

The final procedure of AGG egd is shown in Algorithm 2. We suggest $c_q =$

Algorithm 2 AGG egd

Require: matrix $\check{\Lambda}^{(i)}$ (Eq. (117)), vectors trS (Alg. 1), $\text{dSd}^{(i)}$ (Eq. (118)), $(\hat{s}_j)_{j \in \llbracket B \rrbracket}$ (Eq. (114)), model parameters c_q, c_1, c_2 , learning rate η , maximum nr. of iterations t_{\max}

- 1: # initialise
- 2: $\tilde{\vartheta}_j \leftarrow (\text{trS}_j)^{1/2}/N_j$, $\forall j \in \llbracket B \rrbracket$
- 3: $\check{\varrho}_j^{(i)} \leftarrow (\text{dSd}_j^{(i)}/N_i)^{1/2}$, $\forall j \in \llbracket B \rrbracket$
- 4: $\omega_i^{(0)} \leftarrow \mathbf{1}$
- 5: # until maximum nr. of iterations or convergence
- 6: **while** $t \leq t_{\max}$ and $(\omega_i^{(t-1)} - \omega_i^{(t)})^2 > 10^{-8}$ **do**
- 7: # compute gradient
- 8: $\nabla \omega_i^{(t-1)} \leftarrow 2(\check{\Lambda}^{(i)} + c_2 \text{diag}(\tilde{\vartheta}))\omega_i^{(t-1)} + 2\hat{S}_i + c_q \check{\varrho}^{(i)} + c_1 \tilde{\vartheta}$
- 9: # perform exponentiated gradient descent
- 10: $\omega_i^{(t)} \leftarrow \omega_i^{(t-1)} \cdot \exp\{-\eta^{(t)} \cdot \nabla \omega_i^{(t-1)}\}$
- 11: # normalise
- 12: $\omega_i^{(t)} \leftarrow \frac{\omega_i^{(t)}}{\mathbf{1}^T \omega_i^{(t)}}$
- 13: **end while**
- 14: # estimated optimal weight vector for bag i
- 15: **return** $\omega_i^{(t)}$

1.4, $c_1 = 1, c_2 = 4$ and $r = 100, t_{\max} = 500, \eta^{(t)} = 50/(1 + (t/B))$ as default parameter values.

G.3. STB Approaches. The similarity test based approaches shrink the estimation only towards neighbouring means. Neighbors are found as described in Cor. 1,

$$\begin{aligned}
 W_i &= \left\{ j \in \llbracket B \rrbracket : \sqrt{\text{Tr}(\Sigma_j^2)}/N_j \leq 5 \cdot \sqrt{\text{Tr}(\Sigma_i^2)}/N_i \right\} \\
 V_i &= \left\{ j \in W_i : \|\mu_i - \mu_j\|_{\mathcal{H}}^2 \leq \tau \cdot s_i^2 \right\}.
 \end{aligned} \tag{119}$$

In practice the quantities are estimated. Alg. 1 provides an approximation of $\text{Tr}(\Sigma_i^2)$. Eq. (114) shows an unbiased estimate for s_i^2 and (116) for $\|\mu_i - \mu_j\|_{\mathcal{H}}^2$.

(vi) **STB weight** (Marienwald et al., 2021) assigns a uniform weight to all neighbours except for ω_{ii} which is higher

$$\omega_{ij}^{\text{STB weight}} = \begin{cases} \gamma + \frac{1-\gamma}{|V_i|}, & \text{for } i = j \\ \frac{1-\gamma}{|V_i|}, & \text{for } i \neq j, j \in V_i \\ 0, & \text{otherwise.} \end{cases}$$

STB weight was proposed for balanced bags and under independence of test and data, for which it has strong theoretical results (cf. Theorem 3.1 of Marienwald

et al. (2021)). The optimal values of τ and γ are found by model optimization or $\tau = 2.2, \gamma = 0.2$ taken as default. Larger values of τ allow higher distances between μ_i and its neighbours, thus, potentially increase the number of neighbours and the bias of the estimation. Higher γ values put emphasis on μ_i , i.e., $\omega_{ii} > \omega_{ij}$ for $i \neq j$, and the solution reduces to NE for $\gamma = 1$.

- (vii) **STB opt** corresponds to Lemma 1 and minimises an upper bound on the risk

$$\omega_{i.}^{\text{STB opt}} = \underset{w_{i.}}{\operatorname{argmin}} \left\{ \tau s_i^2 (1 - w_{ii})^2 + \sum_{j \in V_i} w_{ij}^2 s_j^2 \right\},$$

such that $\sum_{j \in [B]} \omega_{ij} = 1$ and $\forall i, j \in [B]. \omega_{ij} \geq 0$. Using Lagrangian multipliers the optimal solution is (cf. (13))

$$\omega_{ij}^{\text{STB opt}} = \begin{cases} \lambda_i \nu_i + (1 - \lambda_i), & \text{for } i = j \\ \lambda_i \nu_j, & \text{for } i \neq j, j \in V_i \\ 0, & \text{otherwise.} \end{cases}$$

where $\nu_j := s_j^{-2} / \sum_{j' \in V_i} s_{j'}^{-2}$ and $\lambda_i := (1 + \gamma\tau(1 - \nu_i))^{(-1)}$. An unbiased estimator for s_i^2 is given in (114). The additional multiplicative constant γ allows for more flexibility and tends to put emphasis on ω_{ii} . Model optimization can be used to find suitable values for τ and γ . Otherwise, we recommend $\tau = 2.2, \gamma = 0.2$ as default values.

- (viii) **STB orth** performs the similarity test and applies **AGG orth** on neighbouring means

$$\omega_{ij}^{\text{STB orth}} = \begin{cases} \omega_{ij}^{\text{AGG orth}}, & \text{for } j \in V_i \\ 0, & \text{otherwise.} \end{cases}$$

The similarity test merely functions as a safeguard here and excludes high distant neighbours and does not play such a central role as for the other **STB** methods. Therefore, τ can be fixed to a large value, e.g., $\tau := 5$. Even though $\omega_{ij}^{\text{AGG orth}}$ is reduced when $\|\mu_i - \mu_j\|_{\mathcal{H}}^2$ is high, **AGG orth** does not perform well when there are many high distant neighbours. Their weights accumulate and reduce the weights of important bags because of the normalization step. The similarity test alleviates this problem.

Either model optimization can be used to find suitable values for τ and γ , or their default values $\tau := 5, \gamma = 3$ can be chosen. Note that, compared to **STB weight** and **STB opt**, τ is larger which highlights the fact that here the similarity test only excludes distant bags. Because of this safeguard, γ , which penalises large distances, can be reduced ($\gamma = 2.2$ vs $\gamma = 13$ for **AGG orth**).

- (ix) **STB egd** performs the similarity test and applies **AGG egd** on neighbouring means

$$\omega_{ij}^{\text{STB egd}} = \begin{cases} \omega_{ij}^{\text{AGG egd}}, & \text{for } j \in V_i \\ 0, & \text{otherwise.} \end{cases}$$

Analogous to the discussion of **STB orth** the similarity test functions as a safeguard to exclude high distant neighbours. It can also be seen as another instrument to replace $\widehat{Q}^{\mathbf{BS}}$ (see also discussion of **AGG egd**). **STB egd** relies on several model parameters. We recommend to set $r = 100, t_{\max} = 500, \eta^{(t)} = 50/(1 + (t/B))$ and $\tau := 5, c_q = 1, c_1 = 1, c_2 = 5$ as default. Compared to the default values of **AGG egd**, diffuse weights should be favoured whereas regularization based on the distances (c_q) or sparse weights (c_1) become less important because of the preselection of neighbouring means.

G.4. Experimental Results for Additional Methods. The results of **R-KMSE** (Muan-det, Sriperumbudur, et al., 2016), **STB weight** (Marienwald et al., 2021), **AGG orth**, and **AGG egd** can be found in Figure 3 and Figure 4. **R-KMSE** estimates each KME




FIGURE 3. Decrease in estimation error compared to NE in percent on Gaussian data settings (a) and (b) resp. Higher is better. The bars (right axis) in (b) show the bag sizes for the bags 1 to 50 which vary between 10 and 300. Compare with Fig. 1

individually and only provides marginal improvement over NE. The performance of **STB weight** is comparable to that of **STB opt**, however, it gives less accurate estimations of large bags in setting (b). **STB weight** assumes equal variances of the estimations, thus, can not handle imbalanced bags. Its weights estimation of large bags might be corrupted by small bags with high variance. Comparing the performances of the aggregation methods **AGG orth** and **AGG egd** with **STB orth** and **STB egd** resp. shows that the similarity test is beneficial and functions as a safeguard to discard distant bags.




FIGURE 4. Decrease in estimation error compared to NE in percent on the flow cytometry data. Higher is better. The number next to the boxplot quantifies the median which is also depicted as a line. The mean is visualised as a circle. From left to right: results on individual cell types 1, 2, 3, 4, 7, 8, 9, and all cell types jointly. The performance of **AGG orth** is partly occluded which allows a more detailed display of the remaining results. Its mean is ≈ -30 , $Q_1 \approx -60$, and its lower whisker ends at ≈ -140 . Compare with Fig. 2

The same observations can be made on the cytometry data. Especially on 'All' cell types, the importance of the similarity test can be noted. **AGG orth** considers the distance between the bags in the definition of the weights. However, due to the normalization $\sum_j \omega_{ij}^{\text{AGG orth}} = 1$ small weights $\omega_{ij}^{\text{AGG orth}}$ of many distant bags j can accumulate and thereby reduce the impact of important weights. The similarity test preselects only close means and eliminates this problem.

G.5. Default Model Parameter Values. The presented models have up to four data dependent model parameters and three hyperparameters. Parameter tuning is possible, whenever the user wishes to estimate the (kernel) means of bags of size N but also has bags of much larger size $N' \gg N$. For the optimization, e.g., in form of cross-validation, subsets of size N are sampled from the N' bags. A method-specific parameter combination is then tested on all bags of size N , while the test error is (only) computed wrt. the N' bags (again, a proxy true mean can be estimated using the complete N' samples).

In most practical applications this scheme is not possible. For this reason we propose default values that we observed to perform well in various settings. Table G.5 shows

an overview. To determine these values we considered the Clustered setting (1). We

TABLE 2. Summary of the default parameter values of each method.

Method	Parameter(s)	Method	Parameter(s)
SOTA		STB	
NE	\emptyset	STB weight	$\tau = 2.2$
R-KMSE	\emptyset		$\gamma = 0.2$
MTA const	$\gamma = 1.$	STB opt	$\tau = 2.2$
			$\gamma = 0.2$
AGG		STB orth	$\tau := 5.$
AGG orth	$\gamma = 13.$		$\gamma = 3.$
AGG egd	$c_q = 1.4$	STB egd	$\tau := 5.$
	$c_1 := 1.$		$c_q := 1.$
	$c_2 = 4.$		$c_1 := 1.$
	$r := 100$		$c_2 = 5.$
	$t_{\max} := 500$		$r := 100$
	$\eta^{(t)} := 50/(1 + (t/B))$		$t_{\max} := 500$
			$\eta^{(t)} := 50/(1 + (t/B))$

generated additional independent training samples on which we ran cross-validation to determine for each radius an optimal set of parameter values (25 repetitions). Most default values correspond to their optimal choices for radius= 1.5. We chose a radius of 1.5 because it presents a balance between overlapping and yet distinct clusters. It also corresponds to a setting which is close to practice and for which multi-task averaging approaches generally promise an improvement over the naive estimation.

The methods based on egd form a peculiarity as the selection procedure did not uncover suitable parameter values. For **STB egd** we fixed $c_q = c_1 = 1$ because the values selected by CV were too low (c_q) or even zero (c_1) at radius 1.5. Because the similarity test selects only close neighbours, the distance controlled by parameter c_q , i.e., $\hat{\Delta}_j^T \Sigma_i \hat{\Delta}_j$, is small. Still, we found the term to be important in practice and, hence, fix $c_q = 1$. Regularization by the c_1 -term leads to sparse weights. For the artificial data, the found neighbours are relatively homogeneous, therefore, no sparse solution must be acquired. However, we found $c_1 = 1$ to be better in more general settings. The same argument holds for **AGG egd** for which we had to fix $c_1 = 1$ as well.

We fixed $\tau := 5$ for **STB orth** and **STB egd**, because their similarity test merely functions as a precaution to discard high-distant bags. As shown by the theoretical discussion, the **AGG** methods are equivalent to the **STB** methods with optimally selected τ .

We did not optimise for the hyperparameters r, t_{\max} and $\eta^{(t)}$. For larger values of r the approximation in Alg. 1 becomes more accurate. It also leads to a higher computational complexity. We found $r = 100$ to be a good trade-off. The same trade-off can be

observed for t_{\max} . However, we found that for the artificial data `egd` converged usually before $t_{\max} = 500$ was reached. For the selection of $\eta^{(t)}$ we follow the recommendation given in Collins et al. (2008) and LeCun et al. (1998). Figures 5-6 show the effect of the model parameter choice on the performance for each method. `STB weight` is highly affected by improper parameter choices, whereas the new `STB` methods are more stable. Specifically, the effect of τ plateaus. Even for small τ , an improvement over `NE` can be observed. When the `AGG` methods are complemented with a preselection of bags (similarity test as safeguard), the methods become more robust. Compared to the c_2 - and c_q -terms of `AGG egd` and `STB egd`, the c_1 -term has only little effect. This lets us conclude that ℓ_1 -regularization is not as important as ℓ_2 -regularization of the weights. Furthermore, it seems that an improper value of one parameter can be alleviated if the remaining parameters are chosen correctly.

For completeness, we also report the optimised (again by cross-validation on iid training data) model parameter values for the Imbalanced setting in Table G.5.

TABLE 3. Summary of the optimised parameter values of each method for the imbalanced setting.

Method	Parameter(s)	Method	Parameter(s)
SOTA		STB	
NE	\emptyset	STB weight	$\tau = 1.4$
R-KMSE	\emptyset		$\gamma = 0.05$
MTA const	$\gamma = 1.1$	STB opt	$\tau = 1.4$
			$\gamma = 0.1$
AGG		STB orth	$\tau = 3.9$
AGG orth	$\gamma = 5.$		$\gamma = 2.$
AGG egd	$c_q = 0.$	STB egd	$\tau = 2.7$
	$c_1 = 3.5$		$c_q = 0.$
	$c_2 = 14.$		$c_1 = 2.9$
	$r := 100$		$c_2 = 34.$
	$t_{\max} := 500$		$r := 100$
	$\eta^{(t)} := 50/(1 + (t/B))$		$t_{\max} := 500$
			$\eta^{(t)} := 50/(1 + (t/B))$

Finally, Figure 7 shows the results on the artificial Gaussian data with optimised model parameters which can be compared to the results with the default parameters shown in Figures 1 and 3. The performance of `MTA const` is unchanged because the optimised and default values of γ are similar. For `STB weight`, `STB opt` and `STB orth` their optimised model parameters differ from their default counterparts for small radii, so that we observe a performance difference respectively. Even though the cross-validation values of `STB orth` vary with the radius, the performance difference is only small which suggests that `STB orth` is stable for its parameter choice. The range

of suitable values of `AGG orth` is large as can be seen in the marginal performance difference even though the value of γ fluctuates between 7 and 14. However, when we also consider the poor performance on the cytometry data, it can be seen that it is sensitive to improper choices of γ (very larger values for γ , e.g., $\gamma = 85$, improve the performance on cytometry). Unsurprisingly, the performance difference is most noticeably for `AGG egd` and `STB egd` for which we manually forced the values of c_q or c_1 . Note, however, that c_1 is of clear importance for the Imbalanced data (optimal values are non-zero), and we observe an improvement on the cytometry data if c_q and c_1 are non-zero. This justifies the decision to choose non-zero values for c_q and c_1 in more general settings.

G.6. Computational Complexity. Based on the previous Section, we can give estimations of the computational complexity of each method. We first identify the complexities of prerequisites on which the methods are based on (items (a)-(g)). Then we analyse each approach individually, where the stated complexity relates to the calculation of the weighting matrix, i.e., all pairwise weights (items (i)-(ix)). Table G.6.4 summarises the total complexities as the combination of all required operations.

G.6.1. *Prerequisites.*

- (a) Intra-task kernel matrices $\kappa(Z_i, Z_i)$ and inter-task kernel matrices $\kappa(Z_i, Z_j)$:

Requires: -

Complexity: $\mathcal{O}(B^2 N^2 D)$

Before we analyse the computations of the weights, we note that the kernel mean embedding is not computed explicitly. Instead, it occurs only in terms of inner products with other KMEs (kernel trick). A kernelised multiple instance problem is usually also not completed with the computation of the KMEs. The computations are required as intermediate step as part of, e.g., statistical testing or distributional learning. We conclude that the computation of the intra-bag kernel matrices $\kappa(Z_i, Z_i)$ and that of the inter-bag kernel matrices $\kappa(Z_i, Z_j)$ for all pairs $i, j \in \llbracket B \rrbracket$ is done anyway, and not just required for the computation of the weights of a multi-task averaging approach.

For simplicity we assume that $N = \max(N_i)_{i \in \llbracket B \rrbracket}$. The complexity of the computation of all kernel matrices depends on the kernel choice. If it is mostly defined by the complexity of matrix multiplications of the data matrices, e.g., linear, Gaussian kernel, etc., it can be assumed to be in $\mathcal{O}(B^2 N^2 D)$.

Except for the approximation of $\text{Tr}(\Sigma^2)$ (Algorithm 1) and the estimations of the naive variance s^2 and $\hat{\Delta}_j^T \Sigma_i \hat{\Delta}_j$ (Eq (118)), the methods rely on the kernel matrices only in term of their sums, $\sum_{n=1}^{N_i} \sum_{m=1}^{N_j} \kappa(Z_n^{(i)}, Z_m^{(j)})$ and $\sum_{n \neq n'}^{N_i} \kappa(Z_n^{(i)}, Z_{n'}^{(i)})$. This reduces the memory consumption from $B^2 N^2$ to $B^2 + B$ because the kernel matrices do not have to be stored but only their sums.

- (b) Distance matrix $\text{MMD}^2(\mu_i, \mu_j)$, Eq. (116):

Requires: (a)

Complexity: $\mathcal{O}(B^2)$

Once (a) is precomputed, the distances $\widehat{\text{MMD}}^2(\mu_i, \mu_j)$, Eq. (116), can be computed in linear time for all pairs $i, j \in \llbracket B \rrbracket$. The complexity is then in $\mathcal{O}(B^2)$.

(c) Naive risk \hat{s}_i^2 , Eq. (114):

Requires: (a) intra-task kernel matrices

Complexity: $\mathcal{O}(BN^2)$

Performs linear operations, e.g., sum and elementwise multiplication of $N \times N$ kernel matrices for every bag individually.

(d) trS_i , Algorithm 1:

Requires: (a) intra-task kernel matrices

Complexity: $\mathcal{O}(BN^2r)$

As seen in Algorithm 1, the first term can be calculated explicitly (line 6) in $\mathcal{O}(N^2)$ and the other terms must be approximated in r iterations.

(e) $\text{dSd}_j^{(i)}$, Eq. (118):

Requires: (a)

Complexity: $\mathcal{O}(B^2N^2)$

The first sum of (118) is computed explicitly on $N \times N$ kernel matrices for every pair $i, j \in \llbracket B \rrbracket$ individually, $\mathcal{O}(B^2N^2)$. The second sum can be computed more efficiently, as it operates only on the sums of the kernels, which requires $\mathcal{O}(B^2)$.

G.6.2. *State-of-the-Art Approaches.*

(i) NE:

Requires: -

Complexity: -

(ii) R-KMSE:

Requires: (a) intra-task kernel matrices

Complexity: $\mathcal{O}(B)$

Performs linear operations on the precomputed sums of the intra-task kernel matrices.

(iii) MTA `const`:

Requires: (b), (c)

Complexity: $\mathcal{O}(B^3)$

The similarity matrix is constant but data dependent and computes the sum of the naive distance matrix, i.e., is in $\mathcal{O}(B^2)$. The calculation of the Laplacian also operates linearly on $B \times B$ matrices. Finally, the computation of the weighting matrix performs a matrix multiplication and a matrix inversion which both require $\mathcal{O}(B^3)$.

G.6.3. *AGG Approaches.*

(iv) AGG `orth`:

Requires: (b), (c)

Complexity: $\mathcal{O}(B^2)$

Only elementwise multiplications and sums are required for the computation of the weights.

(v) **AGG egd**:

Requires: (a), (c), (d), (e)

Complexity: $\mathcal{O}(B^4 t_{\max})$

Algorithm 2 is performed for each bag individually. The biased estimation of $\hat{\Lambda}$, Eq. (117), can be computed in $\mathcal{O}(B^2)$. Line 2 to 4 are linear operations on vectors. The weights are then iteratively updated at most t_{\max} times. In each iteration the gradient is computed in $\mathcal{O}(B^2)$ (matrix-vector multiplication), and egd and normalization are performed in $\mathcal{O}(B)$. In total, the computational complexity of the weighting matrix is then in $\mathcal{O}(B \cdot (B^2 + B + t_{\max} \cdot (B^2 + B))) = \mathcal{O}(B^3 t_{\max})$.

G.6.4. *STB Approaches.*

(f) **STB safeguard** W_i , Eq. (119):

Requires: (d)

Complexity: $\mathcal{O}(B^2)$

Before neighbours are found, bags with a very large variance are excluded as a safeguard. Because the inequality is checked for every pair of bags, the complexity is in $\mathcal{O}(B^2)$.

(g) **STB neighbours** V_i , Eq. (119):

Requires: (b), (c)

Complexity: $\mathcal{O}(B^2)$

The similarity test is performed for every pair of bags such that the complexity is in $\mathcal{O}(B^2)$.

(vi) **STB weight**:

Requires: (f), (g)

Complexity: $\mathcal{O}(B^2)$

The computation of the weighting matrix only requires elementwise operations, i.e., sums and multiplications.

(vii) **STB opt**:

Requires: (c), (f), (g)

Complexity: $\mathcal{O}(B^2)$

The computation of the weighting matrix only requires elementwise operations, i.e., sums and multiplications.

(viii) **STB orth**:

Requires: (f), (g), (iv)

Complexity: $\mathcal{O}(B^2)$

STB orth combines the similarity test with **AGG orth**. Its computational complexity is the sum of both approaches accordingly.

(ix) STB egd:

Requires: (f), (g), (v)

Complexity: $\mathcal{O}(B^3 t_{\max})$

The computational complexity is composed of finding the neighbours and of AGG egd.

TABLE 4. Summary of the individual and total computational complexities. The total complexity is the sum of the complexities of all required computations. Task (a) does not affect the total complexity because its computation is required not only for the estimation of the KMEs.

Task	Individual	Required Computations	Total
PREREQUISITES			
(a)	$\mathcal{O}(B^2 N^2 D)$	(a)	$\mathcal{O}(B^2 N^2 D)$
(b)	$\mathcal{O}(B^2)$	(a), (b)	$\mathcal{O}(B^2)$
(c)	$\mathcal{O}(BN^2)$	(a), (c)	$\mathcal{O}(BN^2)$
(d)	$\mathcal{O}(BN^2 r)$	(a), (d)	$\mathcal{O}(BN^2 r)$
(e)	$\mathcal{O}(B^2 N^2)$	(a), (e)	$\mathcal{O}(B^2 N^2)$
SOTA			
(i)	-	(i)	-
(ii)	$\mathcal{O}(B)$	(a), (ii)	$\mathcal{O}(B)$
(iii)	$\mathcal{O}(B^3)$	(b), (c), (iii)	$\mathcal{O}(BN^2 + B^3)$
AGG			
(iv)	$\mathcal{O}(B^2)$	(b), (c), (iv)	$\mathcal{O}(BN^2 + B^2)$
(v)	$\mathcal{O}(B^3 t_{\max})$	(a), (c), (d), (e), (v)	$\mathcal{O}(BN^2 r + B^2 N^2 + B^3 t_{\max})$
STB			
(f)	$\mathcal{O}(B^2)$	(d), (f)	$\mathcal{O}(BN^2 r + B^2)$
(g)	$\mathcal{O}(B^2)$	(b), (c), (g)	$\mathcal{O}(BN^2 + B^2)$
(vi)	$\mathcal{O}(B^2)$	(f), (g), (vi)	$\mathcal{O}(BN^2 r + B^2)$
(vii)	$\mathcal{O}(B^2)$	(c), (f), (g), (vii)	$\mathcal{O}(BN^2 r + B^2)$
(viii)	$\mathcal{O}(B^2)$	(f), (g), (iv), (viii)	$\mathcal{O}(BN^2 r + B^2)$
(ix)	$\mathcal{O}(B^3 t_{\max})$	(f), (g), (v), (ix)	$\mathcal{O}(BN^2 r + B^2 N^2 + B^3 t_{\max})$




FIGURE 5. See Figure 6.

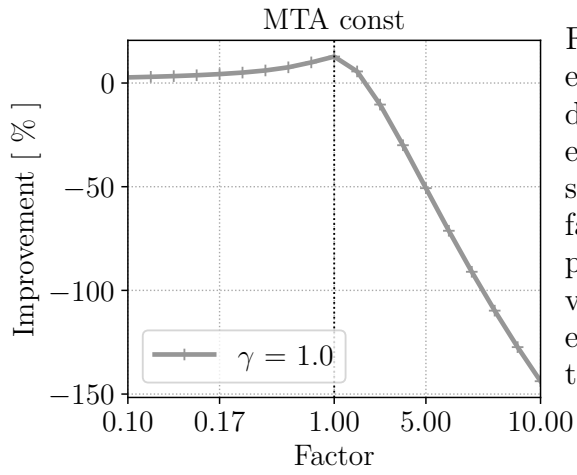


FIGURE 6. Decrease in estimation error compared to NE on Clusters, radius = 1.5, for different model parameter values. Each parameter is considered individually and changed by factors $\times 0.1 - \times 10$. The remaining parameters are kept at their optimal value. (Note that the y-axis is different for MTA const. Some settings lead to worse results than NE.)




FIGURE 7. Decrease in estimation error compared to NE in percent on Gaussian data settings (a) and (b) resp. Higher is better. The model parameters of the methods were optimised on iid training samples using cross-validation. The bars (right axis) in (b) show the bag sizes for the bags 1 to 50 which vary between 10 and 300. Compare with Fig. 1, 3.