

Submitted to *Operations Research*
manuscript (Please, provide the manuscript number!)

Nonparametric Empirical Bayes Estimation On Heterogeneous Data

Trambak Banerjee

Analytics, Information and Operations Management, University of Kansas, Lawrence, KS 66045, trambak@ku.edu

Luella Fu

Department of Mathematics, San Francisco State University, San Francisco, CA 94132, luella@sfsu.edu

Gareth M. James

Information Systems & Operations Management, Emory University, Atlanta, GA 30322, gareth@emory.edu

Wenguang Sun

Center for Data Science, Zhejiang University, Hangzhou, Zhejiang 310058, China, wgsun@zju.edu.cn

The simultaneous estimation of many parameters based on data collected from corresponding studies is a key research problem that has received renewed attention in the high-dimensional setting. Many practical situations involve heterogeneous data where heterogeneity is captured by a nuisance parameter. Effectively pooling information across samples while correctly accounting for heterogeneity presents a significant challenge in large-scale estimation problems. We address this issue by introducing the “Nonparametric Empirical Bayes Structural Tweedie” (NEST) estimator, which efficiently estimates the unknown effect sizes and properly adjusts for heterogeneity via a generalized version of Tweedie’s formula. For the normal means problem, NEST simultaneously handles the two main selection biases introduced by heterogeneity: one, the selection bias in the mean, which cannot be effectively corrected without also correcting for, two, selection bias in the variance. We develop theory to show that NEST is asymptotically as good as the optimal Bayes rule that uniquely minimizes a weighted squared error loss. In our simulation studies NEST outperforms competing methods, with much efficiency gains in many settings. The proposed method is demonstrated on estimating the batting averages of baseball players and Sharpe ratios of mutual fund returns. Extensions to other members of the two-parameter exponential family are discussed.

Key words: compound decision, double shrinkage estimation, kernelized Stein’s discrepancy, non-parametric empirical Bayes, Tweedie’s formula.

1. Introduction

Suppose that we are interested in estimating a vector of parameters $\boldsymbol{\mu} = (\mu_1, \dots, \mu_n)$ based on the summary statistics Y_1, \dots, Y_n from n study units. The setting where $Y_i | \mu_i \sim N(\mu_i, \sigma^2)$ is the most well-known example, but the broader scope includes the compound estimation of Poisson parameters λ_i , Binomial parameters p_i , and other members of the exponential family.

In modern large-scale applications it is often of interest to perform simultaneous and selective inference ([Benjamini and Yekutieli 2011](#), [Berk et al. 2013](#), [Weinstein et al. 2013](#)), which has called for solving the compound estimation problem in new ways and for new purposes. For example, there has been recent work on how to construct valid simultaneous confidence intervals of μ_i 's after a selection procedure is applied ([Lee et al. 2016](#)). In multiple testing, as well as related ranking and selection problems, it is often desirable to incorporate estimates of the effect sizes μ_i in the decision process to prioritize the selection of more scientifically meaningful hypotheses ([Benjamini and Hochberg 1997](#), [Sun and McLain 2012](#), [He et al. 2015](#), [Henderson and Newton 2016](#), [Basu et al. 2017](#)).

However, the simultaneous inference of thousands of means, or other parameters, is challenging because, as described in [Efron \(2011\)](#), the large scale of the problem introduces selection bias, wherein some data points are large merely by chance, causing traditional estimators to overestimate the corresponding means. Shrinkage estimation, exemplified by the seminal work of [James and Stein \(1961\)](#), has been widely used in simultaneous inference. There are several popular classes of methods, including linear shrinkage estimators ([James and Stein 1961](#), [Efron and Morris 1975](#), [Berger 1976](#)), non-linear thresholding-based estimators motivated by sparse priors ([Donoho and Johnstone 1994](#), [Johnstone and Silverman 2004](#), [Abramovich et al. 2006](#)), and both Bayes or empirical Bayes estimators with unspecified priors ([Brown and Greenshtein 2009](#), [Jiang and Zhang 2009](#), [Castillo and van der Vaart 2012](#)). This article focuses on a class of estimators based on Tweedie's formula ([Dyson 1926](#), [Eddington 1940](#), [Robbins 1956](#)). The formula is an elegant shrinkage estimator, for distributions from the exponential family, that has recently received renewed interest

(Brown and Greenshtein 2009, Efron 2011, Koenker and Mizera 2014). Tweedie’s formula is simple and intuitive, and its implementation only requires estimating the marginal distribution of Y_i . This property is particularly appealing for large-scale estimation problems where such estimates can be easily constructed from the observed data. The resultant empirical Bayes estimator enjoys optimality properties (Brown and Greenshtein 2009) and the work of Efron (2011) further convincingly demonstrates that Tweedie’s formula provides an effective bias correction tool when estimating thousands of parameters simultaneously.

1.1. Issues with heterogeneous data

Most of the research in this area has been restricted to models where the distribution of Y_i is solely a function of the parameter of interest μ_i . In situations involving a nuisance parameter τ_i it is generally assumed to be known and identical for all Y_i . For example, homoskedastic Gaussian models of the form $Y_i | \mu_i, \sigma \stackrel{ind}{\sim} N(\mu_i, \sigma^2)$ involve a common nuisance parameter $\tau_i = 1/\sigma^2$ for all i . However, in large-scale studies when the data are collected from heterogeneous sources, the nuisance parameters may vary over the n study units. Perhaps the most common example, and the setting we concentrate most on, involves heteroskedastic errors, where σ^2 varies over Y_i . Microarray data (Erickson and Sabatti 2005, Chiaretti et al. 2004), returns on mutual funds (Brown et al. 1992), and the state-wide school performance gaps (Sun and McLain 2012) are all instances of large-scale data where genes, funds, or schools have heterogeneous variances. Heteroskedastic errors also arise in analysis of variance (Weinstein et al. 2018) and linear regression settings (Kou and Yang 2017). Moreover, in compound binomial problems, heterogeneity arises through unequal sample sizes across different study units. Unfortunately, the conventional Tweedie’s formula assumes identical nuisance parameters across study units and so cannot eliminate selection bias for heterogeneous data. Moreover, various works show that failing to account for heterogeneity leads to inefficient shrinkage estimators (Weinstein et al. 2018), methods with invalid false discovery rates (Efron 2008, Cai and Sun 2009), unstable multiple testing procedures (Tusher et al. 2001) and suboptimal

ranking and selection algorithms (Henderson and Newton 2016), exacerbating the replicability crisis in large-scale studies. Few methodologies are available to address this issue.

For Gaussian data, a common goal is to find the estimator of the means μ_i that minimizes the expected squared error loss. A plausible-seeming solution might be to scale each sample mean Y_i by its estimated standard deviation $S_i/\sqrt{m_i}$ so that a homoskedastic method could be applied to $X_i = \sqrt{m_i}Y_i/S_i$, before undoing the scaling on the final estimate of μ_i . Indeed, this is essentially the approach taken whenever we compute standardized test statistics, such as t -values and z -values. However, this approach, which disregards important structural information, can be highly inefficient. More advanced methods have been developed, but all suffer from various limitations. For instance, the methods proposed by Xie et al. (2012), Tan (2015), Jing et al. (2016a), Kou and Yang (2017), and Zhang and Bhattacharya (2017) are designed for heteroscedastic data but assume a parametric Gaussian prior or semi-parametric Gaussian mixture prior, which leads to loss of efficiency when the prior is misspecified. Moreover, existing methods, such as Xie et al. (2012) and Weinstein et al. (2018), often assume that the nuisance parameters, the variances $1/\tau_i$, are known and use a consistent estimator for implementation. However, when a large number of units are investigated simultaneously, traditional sample variance estimators may similarly suffer from selection bias, which often leads to severe deterioration in the MSE for estimating the means.

1.2. The proposed approach and main contributions

In the homogeneous setting, Tweedie’s formula for Gaussian data estimates μ_i using the score function at Y_i , which is the gradient of the log marginal density of the sufficient statistic Y_i . However, this approach does not immediately extend to heterogeneous data. A significant challenge in the heterogeneous setting is how to pool information from different study units effectively while accounting for the heterogeneity captured by the possibly unknown nuisance parameter τ_i . In this article, we address this issue by proposing a double shrinkage method that simultaneously incorporates the structural information encoded in both the primary (e.g. Y_i) and auxiliary (e.g. S_i) data. We develop a two-step approach, “Nonparametric Empirical Bayes Structural Tweedie” (NEST),

which first estimates the bivariate score function at the primary and auxiliary data, and second predicts μ_i under a weighted squared error loss function using a generalized version of Tweedie's formula that effectively incorporates the structural information encoded in the nuisance parameter. Unlike the squared error loss, which has been widely used for large-scale shrinkage estimation, the weighted squared error loss provides a more suitable performance metric for analyzing heterogeneous data by utilizing a weighting scheme that captures the underlying heterogeneity across the study units. The corresponding shrinkage estimator that minimizes the expected weighted squared error loss naturally encompasses the ability to differentiate between the study units both with respect to μ_i and the heterogeneity captured by the unknown nuisance parameter τ_i .

NEST has several clear advantages. First, it simultaneously handles the two main selection biases introduced by heterogeneity: one, selection bias in the primary data (e.g. sample means), which cannot be effectively corrected without also correcting for, two, selection bias in the auxiliary data (e.g. sample variances). By producing more accurate estimates for the nuisance parameters, NEST in general renders improved shrinkage factors for estimating the primary parameters. Second, NEST makes no parametric assumptions about the prior since it uses a nonparametric method to directly estimate the bivariate score function. Third, NEST exploits the structure of the entire sample and avoids the information loss that occurs in the discretization step used in grouping methods (Weinstein et al. 2018). Fourth, we establish that NEST is asymptotically as good as the optimal Bayes rule that uniquely minimizes the weighted squared error loss. Finally, NEST provides a general estimation framework for members of the two-parameter exponential family and is robust against model mis-specification. We demonstrate numerically that NEST can provide high levels of estimation accuracy relative to a host of benchmark methods.

1.3. Connection to existing works

For empirical Bayes (EB) estimation of the means, recent works such as Xie et al. (2012), Tan (2015), Jing et al. (2016a), Kou and Yang (2017) assume a parametric prior distribution for the means. This is in contrast to the setting where such a prior is unknown and in those settings there

are two main modeling strategies for EB estimation. They are known as the g -modeling and f -modeling strategies in the terminology of Efron (2014). The idea of g -modeling is to first obtain a deconvolution estimate of the unknown prior distribution of μ_i , and then predict μ_i by plugging this estimate into Bayes rule for various loss functions. The deconvolution estimate can be constructed via the nonparametric maximum likelihood estimate (NPMLE; Kiefer and Wolfowitz 1956, Laird 1978), or by modeling the unknown prior as a low-dimensional exponential family distribution (Efron 2016). Some notable works along this line include Jiang and Zhang (2009), Koenker and Mizera (2014), Gu and Koenker (2017a), Saha and Guntuboyina (2020), and Soloff et al. (2021). In contrast, the f -modeling strategy usually assumes that the nuisance parameter, such as the variance, is known and sidesteps the need of deconvolution estimation by directly predicting μ_i based on Tweedie’s formula (or its generalized version), which only depends on the score function at the sufficient statistic Y_i and the known nuisance parameter. Notable works along this line include Brown and Greenshtein (2009) and Efron (2011), both of which rely on fixed and known variances. In particular, their approach involves kernel density estimation techniques for separately estimating the marginal density and its gradient, and then taking their ratio to construct an estimate of the score function. NEST adopts the f -modeling strategy. However, in contrast to existing f -modeling methods, we allow the variances to be unknown and develop a convex optimization approach that directly provides consistent estimate of the score function and is capable of incorporating various structural constraints in the data-driven NEST estimator.

The g -modeling approach via the NPMLE provides an excellent tool for EB estimation of heteroskedastic means. However, to the best of our knowledge, the asymptotic properties of the NPMLE are highly nontrivial to establish and often require strong assumptions. For instance, the analysis in Saha and Guntuboyina (2020) only works for a limited class of covariance structures, and the theory on the rate of convergence is applicable only when the degree of heteroskedasticity is “mild”; alternatively the analysis in Soloff et al. (2021) assumes that μ_i are independent of σ_i , which is often violated in practice (Weinstein et al. 2018). In contrast, we establish the asymptotic properties of NEST without assumptions on the degree of heteroskedasticity or independence

between μ_i and σ_i . A key advantage of g -modeling is its capability to deal with a wider range of problems, particularly those in which direct use of the marginal density of the sufficient statistic itself cannot yield a solution. Meanwhile, the f -modeling approach, which often has a simple and intuitive form (e.g. Tweedie's formula), is attractive when only the information on the marginal distribution is needed for solving the problem of interest.

1.4. Organization

The rest of the paper is structured as follows. In Section 2 we present our hierarchical Gaussian model where both the mean and variance parameters are unknown. In Section 2.1 we discuss the weighted squared error loss function and in Section 2.2 we rely on the natural-parameter Tweedie's formula for our hierarchical model to introduce the oracle NEST estimator under the weighted squared error loss in Definition 1. We then develop a convex optimization approach in Section 3 for estimating the unknown shrinkage factors in the oracle NEST formula. In Section 4, we describe the theoretical setup, justify the optimization criterion, and finally establish asymptotic theories for the proposed NEST estimator. Simulation studies are carried out in Section 5 to compare NEST to competing methods. The article concludes with a discussion in Section 6. Two data applications are presented in Section EC.4 of the Electronic Companion (EC). The proofs, as well as additional theoretical and numerical results, are provided in the EC.

2. Double shrinkage estimation on heteroskedastic Normal data

In the main text of this article, we focus on the normal means problem. Extensions of the methodology to other members of the two-parameter exponential family are discussed in Section EC.5 of the EC.

Suppose we collect m_i observations for the i^{th} study unit, $i = 1, \dots, n$. The data are normally distributed obeying the following hierarchical model:

$$\begin{aligned}
 Y_{ij} \mid \mu_i, \tau_i &\stackrel{i.i.d.}{\sim} N(\mu_i, 1/\tau_i), \quad j = 1, \dots, m_i, \\
 \mu_i \mid \tau_i &\stackrel{ind.}{\sim} G_\mu(\cdot \mid \tau_i), \quad \tau_i \stackrel{i.i.d.}{\sim} H_\tau(\cdot), \quad i = 1, \dots, n.
 \end{aligned}
 \tag{1}$$

Let $Y_i = m_i^{-1} \sum_{j=1}^{m_i} Y_{ij}$ and $S_i^2 = (m_i - 1)^{-1} \sum_{j=1}^{m_i} (Y_{ij} - Y_i)^2$ respectively denote the sample mean and sample variance under Model (1). Further, let $\mathbf{Y} = (Y_1, \dots, Y_n)$ and $\mathbf{S} = (S_1^2, \dots, S_n^2)$ be the vectors of summary statistics, and $\mathbf{y} = (y_1, \dots, y_n)$ and $\mathbf{s} = (s_1^2, \dots, s_n^2)$ the observed values. We view μ_i and τ_i , both of which are unknown, as the primary and nuisance parameters, respectively. The prior distributions $G_\mu(\cdot | \tau_i)$ and $H_\tau(\cdot)$ are unspecified. When the precisions τ_i are known in Model (1), compound estimation of the means $\boldsymbol{\mu} = (\mu_1, \dots, \mu_n)^T$ under the squared error loss has received significant attention in recent years (see for example [Weinstein et al. \(2018\)](#), [Xie et al. \(2012\)](#), [Cai et al. \(2021\)](#) and the references therein) and the estimator minimizing the expected squared error loss is,

$$\mathbb{E}(\mu_i | Y_i = y_i, \tau_i, m_i) = y_i + \frac{1}{m_i \tau_i} w_1(y_i; m_i, \tau_i), \quad (2)$$

where $w_1(y; m, \tau) := \frac{\partial}{\partial y} \log f_{m, \tau}(y)$ and $f_{m, \tau}(\cdot)$ is the pdf of the marginal distribution of Y . Equation (2) is the celebrated Tweedie's formula with known variances ([Efron 2011](#)) which forms the basis for f -modeling strategies for estimating μ_i , and only requires estimation of the score functions $w_1(y_i; m_i, \tau_i)$ in order to compute the estimator. This is particularly appealing in large-scale studies where one observes thousands of (Y_i, τ_i) , making it possible to obtain an accurate estimate of $w_1(y_i; m_i, \tau_i)$ (see Section EC.5 of the EC for more details). Here, we develop a nonparametric empirical Bayes method for estimating $\boldsymbol{\mu}$ under the weighted squared error loss (Section 2.1) when τ_i are unknown.

2.1. The weighted squared error loss

Let $\boldsymbol{\delta} = (\delta_1, \dots, \delta_n)^T$ be an estimator for $\boldsymbol{\mu}$ based on (\mathbf{Y}, \mathbf{S}) . Consider a class of loss functions

$$\ell^{(p)}(\mu_i, \delta_i; \tau_i) = \tau_i^p (\mu_i - \delta_i)^2, \quad (3)$$

for $p \in \{0, 1\}$, where $\ell^{(0)}(\mu_i, \delta_i; \tau_i)$ represents the usual squared error loss and $\ell^{(1)}(\mu_i, \delta_i; \tau_i) = \tau_i (\mu_i - \delta_i)^2$ is the weighted squared error loss with weight given by the precision τ_i . While the squared error loss treats any two study units equally with respect to their precisions τ_i , the weighted squared error loss $\ell^{(1)}(\mu_i, \delta_i; \tau_i)$ is a particularly natural loss function which effectively captures

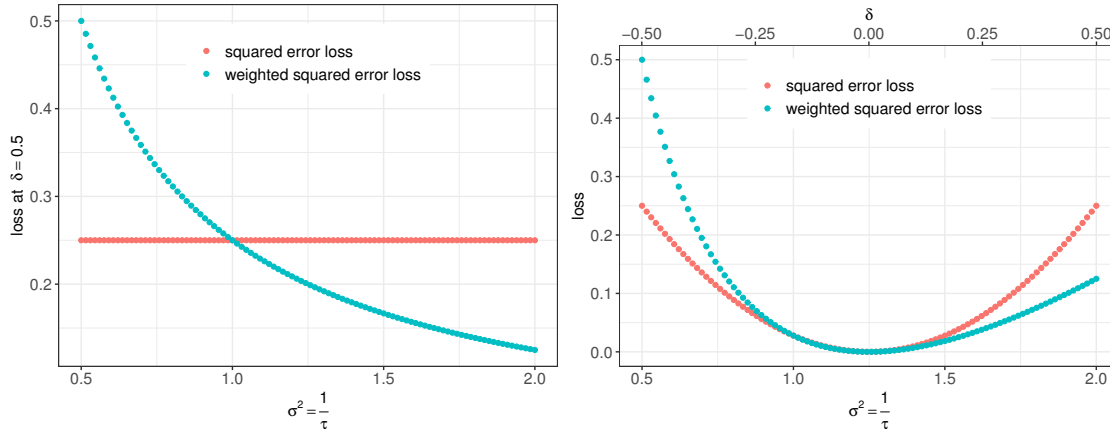


Figure 1 A toy illustration of the squared error and weighted squared error losses. Here $n = 100$ and $\mu_i = 0$ for all $i = 1, \dots, n$. The variances $\sigma_i^2 = 1/\tau_i$ are n equispaced numbers between 0.5 and 2. In the right panel δ_i are n equispaced numbers between -0.5 and 0.5 . The red dots represent the squared error losses while the green dots represent the corresponding weighted squared error losses for estimating μ_i using δ_i for $i = 1, \dots, n$.

the underlying heterogeneity across the study units and appropriately adjusts their contributions to the total estimation error. Under Model (1), and unlike the squared error loss, $\ell^{(1)}(\mu_i, \delta_i; \tau_i)$ is also invariant under a linear transformation of the data Y_{ij} to \tilde{Y}_{ij} i.e. $\ell^{(1)}(\mu_i, \delta_i; \tau_i) = \ell^{(1)}(\tilde{\mu}_i, \tilde{\delta}_i; \tilde{\tau}_i)$ where $\tilde{Y}_{ij} = a_i + b_i Y_{ij}$, $\tilde{\mu}_i = a_i + b_i \mu_i$, $\tilde{\tau}_i = \tau_i / b_i^2$ and $\tilde{\delta}_i = a_i + b_i \delta_i$ for $(a_i, b_i) \in \mathbb{R}^2$ with $b_i \neq 0$. Moreover, the weighting scheme alleviates potential concerns on the fairness issue in terms of the estimation accuracy that each study unit achieves since estimating μ_i as accurately when τ_i is small is difficult compared to when τ_i is large. Intuitively, the weighted loss down-weights the impacts of observations with extremely large variances and allows a relatively larger estimation error when the precisions are small while penalizing the estimation error more when the precisions are large. In Figure 1 we present a toy illustration of this phenomenon. We fix $n = 100$ and $\mu_i = 0$ for all $i = 1, \dots, n$. The variances $\sigma_i^2 = 1/\tau_i$ are n equispaced numbers between 0.5 and 2. The red dots in Figure 1 represent the squared error losses while the green dots represent the corresponding weighted squared error losses for estimating μ_i using δ_i for $i = 1, \dots, n$. In the left panel of Figure 1 we see that when the precisions are large, the same δ_i , say $\delta_i = 0.5$, incurs a relatively larger loss under $\ell^{(1)}(\mu_i, \delta_i; \tau_i)$ than when the precisions are small. In the right panel, we pair each σ_i^2 with a

δ_i , where δ_i are n equispaced numbers between -0.5 and 0.5 , and calculate the two losses. Both the left and the right panels reveal that $\ell^{(0)}(\mu_i, \delta_i; \tau_i)$ does not adjust the estimation error in response to the underlying heteroskedasticity represented by the precisions.

In the following section, we present the oracle NEST estimator of $\boldsymbol{\mu}$ under this weighted squared error loss. Thereafter in Section EC.1 of the EC we present an analysis under the usual squared error loss.

2.2. The oracle NEST estimator under the weighted squared error loss

Denote the average loss $l_n^{(p)}(\boldsymbol{\mu}, \boldsymbol{\delta}; \boldsymbol{\tau}) = n^{-1} \sum_{i=1}^n \ell^{(p)}(\mu_i, \delta_i; \tau_i)$. The compound Bayes risk is

$$r_p(\boldsymbol{\delta}, \mathcal{G}) = \mathbb{E}\{l_n^{(p)}(\boldsymbol{\mu}, \boldsymbol{\delta}; \boldsymbol{\tau})\} = \frac{1}{n} \sum_{i=1}^n \int \int \int \ell^{(p)}(\mu_i, \delta_i; \tau_i) f_{m_i}(y, s^2 | \boldsymbol{\psi}_i) dy ds^2 d\mathcal{G}(\boldsymbol{\psi}_i), \quad (4)$$

where $\boldsymbol{\psi}_i = (\mu_i, \tau_i)$, $\mathcal{G}(\boldsymbol{\psi}_i) = G_\mu(\mu_i | \tau_i) H_\tau(\tau_i)$, and $f_{m_i}(y, s^2 | \boldsymbol{\psi}_i)$ is the likelihood function of (Y_i, S_i^2) .

The Bayes estimator that uniquely minimizes Equation (4) is given by $\boldsymbol{\delta}_{(p)}^\pi = (\delta_{1,(p)}^\pi, \dots, \delta_{n,(p)}^\pi)$,

where

$$\delta_{i,(p)}^\pi := \delta_{(p)}^\pi(y_i, s_i^2, m_i) = \frac{\mathbb{E}(\tau_i^p \mu_i | y_i, s_i^2, m_i)}{\mathbb{E}(\tau_i^p | y_i, s_i^2, m_i)}. \quad (5)$$

Denote $\zeta_i := \tau_i \mu_i$. In Definition 1 we present the oracle NEST estimator of $\boldsymbol{\mu}$ under the weighted squared error loss which is the Bayes estimator $\boldsymbol{\delta}_{(1)}^\pi$. The key idea is to exploit the exponential family representation of the posterior distribution of (μ_i, τ_i) to construct a Tweedie-type formulae for $\delta_{i,(1)}^\pi$.

DEFINITION 1. Consider hierarchical Model (1) with $m_i > 3$. Then the oracle NEST estimator of μ_i under the weighted squared error loss is $\delta_{i,(1)}^\pi$ where

$$\delta_{i,(1)}^\pi := \delta_{(1)}^\pi(y_i, s_i^2, m_i) = \frac{\mathbb{E}(\zeta_i | y_i, s_i^2, m_i)}{\mathbb{E}(\tau_i | y_i, s_i^2, m_i)} = y_i + \frac{s_i^2}{m_i} \gamma(y_i, s_i^2, m_i) w_1(y_i, s_i^2, m_i), \quad (6)$$

$$\text{and } \gamma(y_i, s_i^2, m_i) = \frac{m_i - 1}{m_i - 3 - 2s_i^2 w_2(y_i, s_i^2; m_i)}.$$

We first discuss the oracle NEST estimator in Equation (6) and then explain its genesis.

In Equation (6) $\delta_{i,(1)}^\pi$ is a ratio of $\hat{\zeta}_i^\pi := \mathbb{E}(\zeta_i | y_i, s_i^2, m_i)$ and $\hat{\tau}_i^\pi := \mathbb{E}(\tau_i | y_i, s_i^2, m_i)$, and it involves the shrinkage factor $\gamma_i := \gamma(y_i, s_i^2, m_i)$ that should be applied to s_i^2 when τ_i is unknown. This shrinkage

factor is, by construction, positive and depends on $w_{2,i} := w_2(y_i, s_i^2; m_i)$ that controls the magnitude of shrinkage that is applied to the sample variance s_i^2 . In practical applications, however, the score functions $w_{1,i}$ and $w_{2,i}$ are unknown. In Section 3, we develop a data-driven NEST estimator with estimated scores.

The particularly useful representation of $\delta_{i,(1)}^\pi$ in Equation (6) relies on the construction of Tweedie's formulae for the natural parameters (ζ_i, τ_i) under Model (1). Specifically, we represent the posterior distribution of (μ_i, τ_i) as a two-parameter exponential family with natural parameters (ζ_i, τ_i) which allows one to explicitly compute the posterior means of ζ_i and τ_i as follows:

$$\hat{\zeta}_i^\pi := \hat{\zeta}^\pi(y_i, s_i^2, m_i) = \mathbb{E}(\zeta_i | y_i, s_i^2, m_i) = y_i \mathbb{E}(\tau_i | y_i, s_i^2, m_i) + m_i^{-1} w_1(y_i, s_i^2; m_i), \quad (7)$$

$$\hat{\tau}_i^\pi := \hat{\tau}^\pi(y_i, s_i^2, m_i) = \mathbb{E}(\tau_i | y_i, s_i^2, m_i) = \frac{m_i - 3}{(m_i - 1)s_i^2} - \frac{2}{m_i - 1} w_2(y_i, s_i^2; m_i). \quad (8)$$

Equations (7) and (8) represent the natural parameter Tweedie's formulae for the parameters (ζ_i, τ_i) under Model (1) and are significant for several reasons. First, they generalize Tweedie's formula with known variances (Equation (2)), which can be recovered from Equation (7) by treating τ_i as a known constant and dividing both sides by τ_i . Second, $\hat{\tau}_i^\pi$ in Equation (8) has an interesting interpretation. Apart from being the Bayes estimate of τ_i under the squared error loss, $1/\hat{\tau}_i^\pi$ is the Bayes estimate of $1/\tau_i$ under Stein's loss (James and Stein 1961). Third, Equations (7) and (8) can be used for compound estimation of the parameters $\theta_i = T(\zeta_i, \tau_i, m_i)$ where $T(\cdot)$ is a known function of (ζ_i, τ_i, m_i) . For instance, in finance applications the Sharpe ratio for a portfolio i is given by $T(\zeta_i, \tau_i, m_i) = \zeta_i / \sqrt{\tau_i}$ and one can construct an estimator $\hat{\theta}_i := T(\hat{\zeta}_i^\pi, \hat{\tau}_i^\pi, m_i)$ using Equations (7) and (8) to estimate θ_i . In Section EC.4.2 of the EC we present a real data application that is dedicated to the compound estimation of Sharpe ratios for a large number of mutual fund portfolios. However we note that $\hat{\theta}_i$ constructed in the aforementioned fashion may not, in general, be the optimal estimator of θ_i under the squared error or the weighted squared error loss considered here. In Section EC.2.7 of the EC we provide a proof of Equations (7) and (8).

3. The data-driven NEST estimator

In this section we discuss the estimation of the shrinkage factors and introduce the data-driven NEST estimator in Definition 2. We begin by introducing some notation. Let $\mathbf{x} = (y, s^2)$ be a generic pair of observations from the distribution with marginal density $f_m(\mathbf{x})$, which we assume is continuously differentiable with support on $\mathcal{X} \subseteq \mathbb{R} \times \mathbb{R}^+$. Denote the score function

$$\mathbf{w}(\mathbf{x}; m) = \nabla_{\mathbf{x}} \log f_m(\mathbf{x}) := \{w_1(\mathbf{x}; m), w_2(\mathbf{x}; m)\}. \quad (9)$$

Next, for $i = 1, \dots, n$, let $\mathbf{z}^i = (\mathbf{x}^i, m_i)$ where \mathbf{x}^i is an observation from a distribution with density $f_{m_i}(\mathbf{x})$.

3.1. Convex optimization

We first describe the methodology and then provide explanations. Let,

$$\mathcal{K}_\lambda(\mathbf{z}^i, \mathbf{z}^j) = \exp\left\{-\frac{1}{2\lambda^2}(\mathbf{z}^i - \mathbf{z}^j)^T \Omega (\mathbf{z}^i - \mathbf{z}^j)\right\}, \quad (10)$$

be the Radial Basis Function (RBF) kernel with bandwidth parameter $\lambda > 0$ and $\Omega^{3 \times 3}$ being the inverse of the sample covariance matrix of $(\mathbf{z}^1, \dots, \mathbf{z}^n)$. Denote

$$\mathcal{W}_0^{n \times 2} = \{\mathbf{w}(\mathbf{x}^1; m_1), \dots, \mathbf{w}(\mathbf{x}^n; m_n)\}^T, \text{ where} \quad (11)$$

$$\mathbf{w}(\mathbf{x}^i; m_i) = \nabla_{\mathbf{x}} \log f_{m_i}(\mathbf{x}) \Big|_{\mathbf{x}=\mathbf{x}^i} := \{w_1(\mathbf{x}^i; m_i), w_2(\mathbf{x}^i; m_i)\}.$$

We denote $w_k(\mathbf{x}^i; m_i)$ by $w_{k,i}$, $k = 1, 2$, for the remainder of this article.

Let $\nabla_{z_k^j} \mathcal{K}_\lambda(\mathbf{z}^i, \mathbf{z}^j)$ be the partial derivative of $\mathcal{K}_\lambda(\mathbf{z}^i, \mathbf{z}^j)$ with respect to the k^{th} component of \mathbf{z}^j . The following matrices are needed in our proposed estimator:

$$\mathbf{K}_\lambda^{n \times n} = [K_{ij,\lambda}]_{1 \leq i \leq n, 1 \leq j \leq n}, \quad \nabla \mathbf{K}_\lambda^{n \times 2} = [\nabla K_{ik,\lambda}]_{1 \leq i \leq n, 1 \leq k \leq 2},$$

where $K_{ij,\lambda} = \mathcal{K}_\lambda(\mathbf{z}^i, \mathbf{z}^j)$ and $\nabla K_{ik,\lambda} = \sum_{j=1}^n \nabla_{z_k^j} \mathcal{K}_\lambda(\mathbf{z}^i, \mathbf{z}^j)$. Next we formally define our proposed Nonparametric Empirical-Bayes Structural Tweedie (NEST) estimator.

DEFINITION 2. Consider hierarchical Model (1) with $m_i > 3$. For a fixed bandwidth parameter $\lambda > 0$, let $\hat{\mathcal{W}}_n(\lambda) = \{\hat{\mathbf{w}}_{\lambda,n}(1), \dots, \hat{\mathbf{w}}_{\lambda,n}(n)\}^T$, where $\hat{\mathbf{w}}_{\lambda,n}(i) = \{\hat{w}_{\lambda,n}^{(1)}(i), \hat{w}_{\lambda,n}^{(2)}(i)\}^T$, be the solution to the following quadratic optimization problem:

$$\min_{\mathcal{W} \in \mathbb{R}^{n \times 2}} \frac{1}{n^2} \text{trace} \left(\mathcal{W}^T \mathbf{K}_\lambda \mathcal{W} + 2\mathcal{W}^T \nabla \mathbf{K}_\lambda \right). \quad (12)$$

Then the NEST estimator for μ_i is

$$\delta_{i,n}^{\text{ds}}(\lambda) = y_i + \frac{s_i^2}{m_i} \hat{\gamma}_{i,n}(\lambda) \hat{w}_{\lambda,n}^{(1)}(i), \quad (13)$$

$$\text{where } \hat{\gamma}_{i,n}(\lambda) = \frac{m_i - 1}{m_i - 3 - 2s_i^2 \hat{w}_{\lambda,n}^{(2)}(i)}, \quad (14)$$

with the superscript *ds* denoting “double shrinkage”.

Although not immediately obvious, we show in Section 4 that, under the compound estimation setting, minimizing the objective function (12) is asymptotically equivalent to minimizing the kernelized Stein’s discrepancy (KSD; Liu et al. 2016, Chwialkowski et al. 2016). Roughly speaking, the KSD measures how far a given $n \times 2$ matrix \mathcal{W} is from the true score matrix \mathcal{W}_0 . A key property of the KSD is that it is always non-negative and is equal to 0 if and only if \mathcal{W} and \mathcal{W}_0 are equal. Hence, solving the convex program (12) is equivalent to finding a $\hat{\mathcal{W}}$ that is as close as possible to \mathcal{W}_0 . Since the oracle NEST estimator in Definition 1 is constructed based on \mathcal{W}_0 , we can expect that the data-driven NEST estimator based on $\hat{\mathcal{W}}_n(\lambda)$ would be asymptotically close to its oracle counterpart. Theory underpinning this intuition are established in Section 4.

The bandwidth parameter λ in the RBF kernel \mathcal{K}_λ (Equation (10)) controls the bias-variance trade-off in the score function estimate. For instance, a small value of λ allows unbiasedness but the resulting n -dimensional score function estimator has more variance compared to when λ is large, which forces the estimated shrinkage factors towards 0, and in the limit the NEST estimate is simply the unbiased estimate y_i . In Section 3.2 we discuss a data-driven approach to choose λ for practical implementation of the NEST estimator.

A key characteristic of $\delta_{i,n}^{\text{ds}}(\lambda)$ in Equation (13) is that it exploits the joint structural information available in both Y_i and S_i^2 through $\hat{\mathcal{W}}_n(\lambda)$. Since the weighted loss function involves both the

mean and the variance, we perform shrinkage on both these dimensions. Inspecting Equations (13) and (14), we expect that the improved accuracy achieved by $\hat{\gamma}_{i,n}(\lambda)$ will lead to better shrinkage factors for $\delta_{i,n}^{\text{ds}}(\lambda)$ and hence additional reduction in the estimation risk. Our numerical results in Sections 5 and EC.4 reveal that this is indeed true and the proposed NEST estimator dominates other shrinkage estimators across many settings. In Section EC.5 of the EC we adopt a similar strategy to extend the estimation framework presented in Definition 2 to other distributions in the two-parameter exponential family where the nuisance parameter is known.

We end this section with a discussion of the simpler case of equal sample sizes, i.e. $m_i = m$ for all i . For instance, the leukemia dataset analyzed in Jing et al. (2016a) consists of the expression levels of $n = 5,000$ genes for $m = 27$ acute lymphoblastic leukemia patients. The heterogeneity across the n units in this case is due to the intrinsic variability instead of the varied number of replicates. When the m_i 's are equal, the RBF kernel $\mathcal{K}_\lambda(\cdot, \cdot)$ needs to be modified to avoid a singular sample covariance matrix. Denote $\mathcal{K}_\lambda(\mathbf{x}^i, \mathbf{x}^j) = \exp\{-0.5\lambda^{-2}(\mathbf{x}^i - \mathbf{x}^j)^T \Omega (\mathbf{x}^i - \mathbf{x}^j)\}$ the modified RBF kernel with $\Omega^{2 \times 2}$ being the inverse of the sample covariance matrix of $(\mathbf{x}^1, \dots, \mathbf{x}^n)$. Correspondingly in Definition 2, $\hat{\mathcal{W}}_n(\lambda)$ are the estimates of the shrinkage factors $\mathcal{W}_0^{n \times 2} = \{\mathbf{w}(\mathbf{x}^1; m), \dots, \mathbf{w}(\mathbf{x}^n; m)\}^T$, where $\mathbf{w}(\mathbf{x}^i; m) = \{w_1(\mathbf{x}^i; m), w_2(\mathbf{x}^i; m)\}^T := (w_{1,i}, w_{2,i})^T$.

3.2. Details around implementation

In this section we discuss details around the implementation of NEST. First note that Equation (12) can be solved separately for the two columns of \mathcal{W} , which respectively yield the estimates for $w_{1,i}$ and $w_{2,i}$. Next, the solution to Equation (12) is available in the closed form of $\hat{\mathcal{W}}_n(\lambda) = -\mathbf{K}_\lambda^{-1} \nabla \mathbf{K}_\lambda$. However, in our implementation the closed form solution is replaced by a convex program that directly solves (12) with the following two constraints: (1) $\mathcal{W} \mathbf{a} \preceq \mathbf{b}$, where $\mathbf{a} = (0, 1)^T$ and $\mathbf{b} = (b_1, \dots, b_n)$ with $b_i = \frac{1}{2}(m_i - 3)/s_i^2 - \kappa$ for some $\kappa > 0$, and (2) $\mathbf{1}^T \mathcal{W} \mathbf{e}_k = 0$, $k = 1, 2$, where \mathbf{e}_k is the canonical basis vector with 1 at coordinate k . Inspecting Equation (14) shows that adding the first constraint guarantees that $\hat{\gamma}_{i,n}(\lambda) < \infty$. This is desirable in both numerical and theoretical analyses. Similar ideas have been used in the seminal work of Koenker and Mizera (2014). The

second constraint ensures that the sample mean of the k^{th} estimated shrinkage factor is zero since the expectation of the gradient of the log marginal density is indeed zero.

The practical implementation requires a data-driven scheme for choosing λ . We propose to use a variation of the modified cross validation scheme of [Brown et al. \(2013\)](#), which involves splitting Y_{ij} into two parts: $U_{ij} = Y_{ij} - (1/\alpha)\epsilon_{ij}$ and $V_{ij} = Y_{ij} + \alpha\epsilon_{ij}$, where $\epsilon_{ij} \sim N(0, \tau_i^{-1})$, and U_{ij} and V_{ij} are used to construct the estimator and to choose the tuning parameter, respectively. However in our setup τ_i is unknown; hence we define $U_{ij} = Y_{ij} - (1/\alpha)\sqrt{S_i^2}\epsilon_{ij}$, $V_{ij} = Y_{ij} + \alpha\sqrt{S_i^2}\epsilon_{ij}$ and sample ϵ_{ij} independently from $N(0, 1)$. Let $\bar{V}_i = m_i^{-1} \sum_{j=1}^{m_i} V_{ij}$, $\bar{U}_i = m_i^{-1} \sum_{j=1}^{m_i} U_{ij}$, $\mathcal{U} = \{U_{ij} : 1 \leq i \leq n, 1 \leq j \leq m_i\}$ and $\mathcal{V} = \{V_{ij} : 1 \leq i \leq n, 1 \leq j \leq m_i\}$. Then conditional on (μ_i, τ_i) , \bar{U}_i and \bar{V}_i are uncorrelated with mean μ_i and variances $(1 + \alpha^{-2})/(m_i\tau_i)$ and $(1 + \alpha^2)/(m_i\tau_i)$, respectively. Define

$$\vartheta_n(\lambda; \mathcal{U}, \mathcal{V}) = \frac{1}{n} \sum_{i=1}^n \frac{\{\bar{V}_i - \delta_{i,n}^{\text{ds}}(\bar{U}_i; \mathcal{U}, \lambda)\}^2}{S_{V,i}^2},$$

where $\delta_{i,n}^{\text{ds}}(\bar{U}_i; \mathcal{U}, \lambda)$ is the NEST estimator of μ_i based on \mathcal{U} and $S_{V,i}^2$ is the sample variance of $\{V_{ij} : 1 \leq j \leq m_i\}$. The tuning parameter will be chosen as $\hat{\lambda} := \arg \min_{\lambda \in \Lambda} \vartheta_n(\lambda)$. In our numerical studies of Section 5, we set $\alpha = 1/2$, $\Lambda = [0.5, 10^2]$. The tuning parameter $\hat{\lambda}$ is obtained from this scheme and then used to estimate μ_i via Equation (13).

REMARK 1. Note that for large m , \bar{V}_i and $S_{V,i}^2$ are independent and $\mathbb{E}[\vartheta_n(\lambda; \mathcal{U}, \mathcal{V})]$ can be well approximated by the sum $n^{-1} \sum_{i=1}^n \mathbb{E}[\ell^{(1)}\{\mu_i, \delta_{i,n}^{\text{ds}}(\bar{U}_i; \mathcal{U}, \lambda); \tau_i/(1 + \alpha^2)\}] + n^{-1} \sum_{i=1}^n \mathbb{E}[(\bar{V}_i - \mu_i)^2/S_{V,i}^2]$ for any fixed $\lambda > 0$ and n . The first term depends on λ while the second term is independent of it. We choose the value of λ that minimizes $\vartheta_n(\lambda; \mathcal{U}, \mathcal{V})$ to construct the data-driven NEST estimator of $\boldsymbol{\mu}$ based on the original sample (\mathbf{Y}, \mathbf{S}) . For $\alpha \rightarrow 0$, this scheme is equivalent to choosing a value of λ that minimizes the true weighted squared error loss of $\delta_{i,n}^{\text{ds}}(\bar{U}_i; \mathcal{U}, \lambda)$, $i = 1, \dots, n$. We note in Section 5 that when m is small, the risk of the data-driven NEST estimator is generally higher than the risk of the NEST estimator that chooses λ by minimizing the true weighted squared error loss of $\boldsymbol{\delta}_n^{\text{ds}}(\lambda)$ based on (\mathbf{Y}, \mathbf{S}) (see for example the left panels in Figures 2 and 3). This is not unexpected since $\vartheta_n(\lambda; \mathcal{U}, \mathcal{V})$ is not an unbiased estimator of the true risk. However, $\vartheta_n(\lambda; \mathcal{U}, \mathcal{V})$

provides a practical criterion for choosing λ and works well empirically in our numerical studies for larger m . The development of an unbiased risk estimate, such as a SURE-type criterion, for this setting is a challenging topic requiring further research. See also Ignatiadis and Wager (2019) for related discussions.

We are developing an R package, `nest`, to implement the NEST estimator in Definition 2. The R code that reproduces the numerical results in this paper can be downloaded from the link: <https://www.dropbox.com/sh/5yptcj4epxdgbqs/AADAYHZNDcv4Hqagsi97vC2Da?dl=0>.

4. Theory

In this section we introduce the Kernelized Stein's Discrepancy (KSD) measure (Liu et al. 2016, Chwialkowski et al. 2016) and discuss its connection to the quadratic program (12). While the KSD has been used in various contexts including goodness of fit tests (Liu et al. 2016, Yang et al. 2018), variational inference (Liu and Wang 2016) and Monte Carlo integration (Oates et al. 2017), its connections to compound estimation and empirical Bayes methodology was established only recently (Banerjee et al. 2021, Luo et al. 2022). The analysis in this and following sections is geared towards the case $m_i = m$ for $i = 1, \dots, n$. Under this setting, $(\mathbf{x}^1, \dots, \mathbf{x}^n)$ constitute an i.i.d random sample from $f_m(\mathbf{x})$. The case of unequal m_i 's can be analyzed in a similar fashion by assuming that m_i 's are a random sample from a distribution with mass function $q(\cdot)$. Then $\mathbf{z} = (\mathbf{x}, m)$ has distribution with density $p(\mathbf{z}) := q(m)f_m(\mathbf{x})$, where $\mathbf{z}^i = (\mathbf{x}^i, m_i)$, $(\mathbf{z}^1, \dots, \mathbf{z}^n)$ are realizations of an i.i.d random sample from $p(\mathbf{z})$.

4.1. Kernelized Stein's Discrepancy

Suppose \mathbf{X} and \mathbf{X}' are i.i.d. copies from the marginal distribution of (Y, S^2) that has density f wherein the dependence on m is implicit. Denote $\mathbf{w}(\mathbf{X})$ and $\mathbf{w}(\mathbf{X}')$, defined in Equation (9), to be the score functions at \mathbf{X} and \mathbf{X}' respectively. Suppose \tilde{f} is an arbitrary density function on the support of (Y, S^2) , for which we similarly define $\tilde{\mathbf{w}}(\mathbf{X})$. The KSD, formally defined as

$$\mathcal{S}_\lambda(f, \tilde{f}) = \mathbb{E}_{\mathbf{X}, \mathbf{X}' \sim f} \left[\left\{ \tilde{\mathbf{w}}(\mathbf{X}) - \mathbf{w}(\mathbf{X}) \right\}^T \mathcal{K}_\lambda(\mathbf{X}, \mathbf{X}') \left\{ \tilde{\mathbf{w}}(\mathbf{X}') - \mathbf{w}(\mathbf{X}') \right\} \right],$$

provides a discrepancy measure between f and \tilde{f} in the sense that $\mathcal{S}_\lambda(f, \tilde{f})$ tends to increase when there is a bigger disparity between \mathbf{w} and $\tilde{\mathbf{w}}$ (or equivalently, between f and \tilde{f}), and

$$\mathcal{S}_\lambda(f, \tilde{f}) \geq 0 \text{ and } \mathcal{S}_\lambda(f, \tilde{f}) = 0 \text{ if and only if } f = \tilde{f}.$$

The direct evaluation of $\mathcal{S}_\lambda(f, \tilde{f})$ is difficult because \mathbf{w} is unknown. Liu et al. (2016) introduced an alternative representation of the KSD that does not directly involve \mathbf{w} :

$$\begin{aligned} \mathcal{S}_\lambda(f, \tilde{f}) &= \mathbb{E}_f \kappa_\lambda[\tilde{\mathbf{w}}(\mathbf{X}), \tilde{\mathbf{w}}(\mathbf{X}')](\mathbf{X}, \mathbf{X}') \\ &= \mathbb{E}_f \left\{ \frac{1}{n(n-1)} \sum_{i=1}^n \sum_{j=1}^n \kappa_\lambda[\tilde{\mathbf{w}}(\mathbf{X}^i), \tilde{\mathbf{w}}(\mathbf{X}^j)](\mathbf{X}^i, \mathbf{X}^j) \mathbb{I}(i \neq j) \right\} \\ &= \mathbb{E}_f [\bar{\mathbb{M}}_{\lambda, n}(\tilde{\mathcal{W}})] := \mathbb{M}_\lambda(\tilde{\mathcal{W}}), \end{aligned} \quad (15)$$

where $\{\mathbf{X}^1, \dots, \mathbf{X}^n\}$ is a random sample from f , \mathbb{E}_f denotes expectation under f and

$$\begin{aligned} \kappa_\lambda[\tilde{\mathbf{w}}(\mathbf{x}), \tilde{\mathbf{w}}(\mathbf{x}')](\mathbf{x}, \mathbf{x}') &= \tilde{\mathbf{w}}(\mathbf{x})^T \tilde{\mathbf{w}}(\mathbf{x}') \mathcal{K}_\lambda(\mathbf{x}, \mathbf{x}') + \tilde{\mathbf{w}}(\mathbf{x})^T \nabla_{\mathbf{x}'} \mathcal{K}_\lambda(\mathbf{x}, \mathbf{x}') + \nabla_{\mathbf{x}} \mathcal{K}_\lambda(\mathbf{x}, \mathbf{x}')^T \tilde{\mathbf{w}}(\mathbf{x}) \\ &\quad + \text{trace}(\nabla_{\mathbf{x}} \nabla_{\mathbf{x}'} \mathcal{K}_\lambda(\mathbf{x}, \mathbf{x}')), \end{aligned}$$

is a smooth and symmetric positive definite kernel function associated with the U-statistic $\bar{\mathbb{M}}_{\lambda, n}(\tilde{\mathcal{W}})$.

The implementation of the NEST estimator in Definition 2 boils down to the estimation of \mathcal{W}_0 via the convex program (12), which corresponds to minimizing

$$\hat{\mathbb{M}}_{\lambda, n}(\tilde{\mathcal{W}}) = \frac{1}{n^2} \sum_{i=1}^n \sum_{j=1}^n \kappa_\lambda[\tilde{\mathbf{w}}(\mathbf{X}^i), \tilde{\mathbf{w}}(\mathbf{X}^j)](\mathbf{X}^i, \mathbf{X}^j) \quad (16)$$

w.r.t. $\tilde{\mathcal{W}}$. A key observation is that if the empirical criterion $\hat{\mathbb{M}}_{\lambda, n}(\tilde{\mathcal{W}})$ is asymptotically equal to the population KSD criterion $\mathbb{M}_\lambda(\tilde{\mathcal{W}})$, then minimizing $\hat{\mathbb{M}}_{\lambda, n}(\tilde{\mathcal{W}})$ with respect to $\tilde{\mathcal{W}}$ is effectively the process of finding a $\tilde{\mathcal{W}}$ that is as close as possible to \mathcal{W}_0 in Equation (11). This intuitively justifies the NEST estimator in Definition 2.

In Theorem 1, we formally establish the asymptotic consistency of the sample criterion $\hat{\mathbb{M}}_{\lambda, n}(\tilde{\mathcal{W}})$ around the population criterion $\mathbb{M}_\lambda(\tilde{\mathcal{W}})$ for any fixed $\lambda > 0$. We impose the following regularity condition where f denotes the density function of the joint marginal distribution of (Y, S^2) and \mathbb{E}_f denotes expectation under f .

ASSUMPTION 1. $\mathbb{E}_f |\kappa_\lambda[\tilde{\mathbf{w}}(\mathbf{X}^i), \tilde{\mathbf{w}}(\mathbf{X}^j)](\mathbf{X}^i, \mathbf{X}^j)|^2 < \infty$ for any $(i, j) \in \{1, \dots, n\}$.

Assumption 1 is a standard moment condition on $\kappa_\lambda[\tilde{\mathbf{w}}(\mathbf{X}^i), \tilde{\mathbf{w}}(\mathbf{X}^j)](\mathbf{X}^i, \mathbf{X}^j)$ [see, for example, Section 5.5 in Serfling (2009)], which is needed for establishing the Central Limit Theorem for the U-statistic $\bar{\mathbb{M}}_{\lambda,n}(\tilde{\mathcal{W}})$ in Equation (15).

THEOREM 1. *Under Assumption 1 and for any fixed $\lambda > 0$, we have*

$$|\hat{\mathbb{M}}_{\lambda,n}(\tilde{\mathcal{W}}) - \mathbb{M}_\lambda(\tilde{\mathcal{W}})| = O_p(n^{-1/2}).$$

Moreover, along with the fact that $\mathbb{M}_\lambda(\mathcal{W}_0) = 0$, Theorem 1 justifies $\hat{\mathbb{M}}_{\lambda,n}(\tilde{\mathcal{W}})$ as an appropriate optimization criterion. Next we show that the NEST estimator in Definition 2 is asymptotically close to its oracle counterpart.

4.2. Asymptotic Properties of NEST

This section studies the asymptotic properties of the NEST estimator under the setting where $\lambda := \lambda(n)$ varies with n . We begin by recalling the oracle NEST estimator $\delta_{(1)}^\pi = (\delta_{1,(1)}^\pi, \dots, \delta_{n,(1)}^\pi)$ for $\boldsymbol{\mu}$ in Equation (6), where

$$\delta_{i,(1)}^\pi := \delta_{(1)}^\pi(y_i, s_i^2, m) = \frac{\hat{\zeta}^\pi(y_i, s_i^2, m)}{\hat{\tau}^\pi(y_i, s_i^2, m)} = y_i + \frac{s_i^2}{m} \gamma(y_i, s_i^2; m_i) w_1(y_i, s_i^2, m),$$

and $\hat{\zeta}^\pi$ and $\hat{\tau}^\pi$ are respectively the Bayes estimators of $\tau\mu$ and τ as defined in Equations (7) and (8). Viewing the proposed NEST estimator $\delta_{i,n}^{\text{ds}}(\lambda)$ as a data-driven approximation to $\delta_{i,(1)}^\pi$, we study the quality of this approximation for large n and fixed m .

In Theorem 2 we first establish the consistency of the estimated score functions $\hat{\mathcal{W}}_n(\lambda)$ obtained as the solution to the quadratic optimization problem in Equation (12). In what follows, we denote the Frobenius norm of a matrix \mathbf{A} by $\|\mathbf{A}\|_F$ and for two sequences a_n and b_n , we denote $a_n \asymp b_n$ to mean $d_1 a_n \leq b_n \leq d_2 a_n$ for large n and some constants $d_2 \geq d_1 > 0$. The following regularity condition is needed in our technical analysis.

ASSUMPTION 2. *The joint density f and its score functions are Lipschitz continuous.*

THEOREM 2. *If $\lim_{n \rightarrow \infty} c_n (\log n)^4 n^{-1/3} = 0$ then under Assumption 2 and with $\lambda \asymp n^{-1/3}$, we have*

$$\lim_{n \rightarrow \infty} \mathbb{P} \left\{ \frac{1}{n} \left\| \hat{\mathcal{W}}_n(\lambda) - \mathcal{W}_0 \right\|_F^2 \geq c_n^{-1} \epsilon \right\} = 0, \text{ for any } \epsilon > 0.$$

We prove Theorem 2 for sub-exponential random variables \mathbf{X} whose density f and the corresponding score functions are Lipschitz continuous (Assumption 2). The assumption of Lipschitz continuity is mild and allows a cleaner statement of our theoretical result. Theorem 2 reveals a slower than parametric \sqrt{n} rate of convergence, barring the logarithmic terms. This is related to the observation that the KSD approach for estimating the score functions essentially works by first mapping the data into a reproducing kernel Hilbert space (RKHS) associated with the integrally strictly positive definite kernel \mathcal{K}_λ in Equation (10) and then estimating the scores by minimizing $\hat{\mathbb{M}}_{\lambda,n}(\tilde{\mathcal{W}})$ with respect to $\tilde{\mathcal{W}}$, where $\hat{\mathbb{M}}_{\lambda,n}(\tilde{\mathcal{W}})$ is a data-driven approximation to the squared RKHS norm $\mathbb{M}_\lambda(\tilde{\mathcal{W}})$. The statement of Theorem 2, however, is based on the metric $L_2(\tilde{\mathcal{W}}, \mathcal{W}_0) = n^{-1} \|\tilde{\mathcal{W}} - \mathcal{W}_0\|_F^2$ and not the RKHS norm. As such, it is not necessarily true that if $\hat{\mathcal{W}}_n(\lambda)$ is a good estimate of \mathcal{W}_0 in the RKHS norm then it is a good estimate of \mathcal{W}_0 in the L_2 metric too for every $\lambda > 0$. Theorem 2 precisely tabulates how λ must behave as $n \rightarrow \infty$ to establish consistency of the estimated scores in the L_2 metric and the slower than \sqrt{n} rate of convergence is a consequence of the cost of inversion from the RKHS norm to this L_2 metric.

Next we provide decision theoretic guarantees on $\delta_n^{\text{ds}}(\lambda)$ in relation to $\delta_{(1)}^\pi$. We impose the following moment condition on the prior distributions of μ and τ .

ASSUMPTION 3. *For some $\epsilon_i \in (0, 1), i = 1, 2, 3$, $\mathbb{E}_G\{\exp(\epsilon_1|\mu|)\} < \infty$, $\mathbb{E}_H\{\exp(\epsilon_2/\tau)\} < \infty$ and $\mathbb{E}_H\{\exp(\epsilon_3\tau)\} < \infty$.*

The moment conditions in Assumption 3 ensure that with high probability $|\mu| \leq \log n$ and $1/\log n \leq \tau \leq \log n$ as $n \rightarrow \infty$. This is formalized in Lemma EC.4 in Section EC.2.6 of the EC. It is likely that Assumption 3 can be further relaxed but we do not seek the full generality here.



THEOREM 3. *Suppose Assumptions 2 – 3 hold. If $\lim_{n \rightarrow \infty} c_n n^{-1/3} (\log n)^8 = 0$, then as $n \rightarrow \infty$ with $\lambda \asymp n^{-1/3}$, we have $\frac{c_n}{n} \left\| \delta_n^{\text{ds}}(\lambda) - \delta_{(1)}^\pi \right\|_2^2 = o_p(1)$. Furthermore, under the same conditions, $c_n^{1/2} |l_n^{(1)}\{\boldsymbol{\mu}, \delta_n^{\text{ds}}(\lambda); \boldsymbol{\tau}\} - l_n^{(1)}(\boldsymbol{\mu}, \delta_{(1)}^\pi; \boldsymbol{\tau})| = o_p(1)$ as $n \rightarrow \infty$.*

Theorem 3 establishes the optimality theory of $\delta_n^{\text{ds}}(\lambda)$ by showing that (a) the average squared error between $\delta_n^{\text{ds}}(\lambda)$ and $\delta_{(1)}^\pi$ is asymptotically small, and (b) the estimation loss of NEST converges in probability to that of its oracle counterpart as $n \rightarrow \infty$.

5. Numerical experiments

In this section we assess the performance of the NEST estimation framework for compound estimation of Normal means under the weighted squared error loss. We focus on the hierarchical Model of Equation (1) and compare the following four approaches for estimating $\boldsymbol{\mu}$ when the precisions $\boldsymbol{\tau}$ are unknown: (1) the proposed NEST method where the tuning parameter λ is chosen using a modified cross-validation approach described in Section 3.2, (2) the NEST method which estimates λ by minimizing the true loss (NEST Orc. λ), (3) the g-modelling approach of Gu and Koenker (2017a,b) which first estimates the joint prior distribution of (μ_i, τ_i) using nonparametric maximum likelihood estimation (NPMLE) techniques (Kiefer and Wolfowitz 1956, Laird 1978) and then plugs them into the Bayes rule for estimating μ_i under the weighted squared error loss, and (4) the oracle Bayes estimator $\delta_{(1)}^\pi$ in Equation (6) that has full knowledge of the prior distributions of (μ_i, τ_i) . For NPMLE we use the R-function `WGLVmix`, available within the R-package `REBayes` (Koenker and Gu 2017), to estimate the joint prior distribution of (μ_i, σ_i^2) . To compute $\delta_{(1)}^\pi$ we rely on the R-package `RStan` (Stan Development Team 2022, Carpenter et al. 2017).

The aforementioned four approaches are evaluated on nine different simulation settings, with the goal of assessing the relative performance of the competing estimators as the heterogeneity in the variances σ_i^2 is varied while keeping the sample sizes m_i fixed at m . The nine simulation settings can be categorized into three types: two settings where mean and variances are independent, five settings where mean and variance are correlated and two settings that represent departures from the Normal data-generating model. For each setting we set $n = 1,000$ and compute the average Bayes risk for each competing estimator of $\boldsymbol{\mu}$ across 50 Monte Carlo repetitions. Figures 2 to 5 plot the relative risk of each competing estimator, which is the ratio of the average Bayes risk of the competing estimator to that of the oracle Bayes estimator $\delta_{(1)}^\pi$.

(a) Setting 1: $\mu_i \stackrel{i.i.d}{\sim} N(0,3)$ and $\sigma_i^2 \stackrel{i.i.d}{\sim} U(0.25, u)$.(b) Setting 2: Here $\sigma_i^2 \stackrel{i.i.d}{\sim} N(u, 1)$ truncated below 0.1 and $\mu_i \stackrel{i.i.d}{\sim} \text{Laplace}(\text{location} = 0, \text{scale} = 1)$ **Figure 2** Comparison of relative risks when (μ_i, σ_i^2) are independent. Plots show $m = 10, 15, 20$ left to right.

The first two settings, Figures 2a and 2b, correspond to the independent case. We describe the two settings below.


- Setting 1: $\mu_i \stackrel{i.i.d}{\sim} N(0,3)$ and $\sigma_i^2 \stackrel{i.i.d}{\sim} U(0.25, u)$ where we let u to vary across five levels, $\{1, 2, 3, 4, 5\}$.
- Setting 2: $\mu_i \stackrel{i.i.d}{\sim} \text{Laplace}(\text{location} = 0, \text{scale} = 1)$ and $\sigma_i^2 \stackrel{i.i.d}{\sim} N(u, 1)$ truncated below 0.1 where u varies across $\{0.25, 0.5, 0.75, 1, 1.25\}$.

The three plots in Figures 2a and 2b show the relative risks as u varies for $m = 10, 15$ and 20 (left to right). We see that under Setting 1, NEST and NEST Orc. λ have relative risks substantially closer to 1 than NPMLE. As m increases, the performance of NEST and NEST Orc. λ improves as their relative risks almost coincide with 1. The NPMLE, on the other hand, exhibits a slightly higher relative risk at $m = 20$ than at $m = 10$ and 15.


The next five settings, Figures 3 and 4, correspond to the correlated case. We describe these five settings below.

- Setting 3: $\sigma_i^2 \stackrel{i.i.d.}{\sim} U(0.25, 3)$ and $\mu_i \mid \sigma_i^2 \stackrel{ind.}{\sim} 0.5 N(-\sigma_i^2, u^2) + 0.5 N(\sigma_i^2, u^2)$ where u varies across $\{0.25, 0.35, 0.45, 0.55, 0.65, 0.75\}$.
- Setting 4: $\sigma_i^2 \stackrel{i.i.d.}{\sim} U(0.25, u)$ and $\mu_i \mid \sigma_i^2 \stackrel{ind.}{\sim} 0.5 N(\sigma_i^2, \sigma_i^2) + 0.5 N(-\sigma_i^2, \sigma_i^2)$ where we let $u \in \{1, 2, 3, 4, 5\}$.
- Setting 5: we sample σ_i^2 from an Inverse Gamma (IG) distribution with shape parameter fixed at 5 and rate at u where $u \in \{1, 1.5, 2, 2.5, 3\}$. Conditional on the variances, the means μ_i are sampled from the design described in Setting 4.
- Setting 6: $\sigma_i^2 \stackrel{i.i.d.}{\sim} U(0.1, u)$ and $\mu_i \mid \sigma_i^2 \stackrel{ind.}{\sim} \text{Laplace}(\text{location} = 0, \text{scale} = \sqrt{\sigma_i^2/2})$ where $u \in \{1, 1.5, 2, 2.5, 3\}$.
- Setting 7: σ_i^2 are independently sampled from a two component mixture of truncated Normal distributions that are truncated below 0.1 with means 0.5 and u , standard deviation 0.1 and mixing weights 0.5. Conditional on the variances, the means μ_i are sampled from a Logistic distribution with location σ_i^2 and scale $\sqrt{\sigma_i^2}$. We let u to vary across five levels, $\{1, 1.5, 2, 2.5, 3\}$.


Setting 3 is presented in Figure 3a where u controls the spread of the prior distribution of μ_i given σ_i^2 and reveals that when u is small, NPMLE has a relatively better risk performance than NEST. This is expected because when u is small, the prior distribution of μ_i can be well approximated by a discrete distribution with just two mass points and the NEST estimator finds it particularly difficult to estimate the means well in these scenarios because such a discrete prior on μ_i introduces potential multimodality of the underlying marginal distribution of (Y_i, S_i^2) in which case the true score functions may no longer be Lipschitz continuous, an assumption that is needed for the KSD approach to produce consistent estimates of the true scores in Section 3. The NPMLE, on the other hand, has a near parametric rate of convergence in Gaussian denoising problems when the underlying prior on the means is discrete (Saha and Guntuboyina 2020), which explains its improved performance in Setting 3, especially when $u \rightarrow 0$.



(a) Setting 3: $\sigma_i^2 \stackrel{i.i.d.}{\sim} U(0.25, 3)$ and $\mu_i \mid \sigma_i^2 \stackrel{ind.}{\sim} 0.5 N(-\sigma_i^2, u) + 0.5 N(\sigma_i^2, u)$.




(b) Setting 4: $\sigma_i^2 \stackrel{i.i.d.}{\sim} U(0.25, u)$ and $\mu_i \mid \sigma_i^2 \stackrel{ind.}{\sim} 0.5 N(\sigma_i^2, \sigma_i^2) + 0.5 N(-\sigma_i^2, \sigma_i^2)$.




(c) Setting 5: $\sigma_i^2 \stackrel{i.i.d.}{\sim} \text{IG}(\text{shape} = 5, \text{rate} = u)$ and $\mu_i \mid \sigma_i^2 \stackrel{ind.}{\sim} 0.5 N(\sigma_i^2, \sigma_i^2) + 0.5 N(-\sigma_i^2, \sigma_i^2)$.

Figure 3 Comparison of relative risks when (μ_i, σ_i^2) are dependent - Settings 3, 4 and 5. Plots show $m = 10, 15, 20$ left to right.



(a) Setting 6: $\sigma_i^2 \stackrel{i.i.d.}{\sim} U(0.1, u)$ and $\mu_i | \sigma_i^2 \stackrel{ind.}{\sim} \text{Laplace}(\text{location} = 0, \text{scale} = \sqrt{\sigma_i^2/2})$




(b) Setting 7: $\sigma_i^2 \stackrel{i.i.d.}{\sim} 0.5 \text{ Trunc.}N(0.5, 0.1^2; 0.1) + 0.5 \text{ Trunc.}N(u, 0.1^2; 0.1)$ where $\text{Trunc.}N(a, b; c)$ represents a truncated Normal distribution with mean a , variance b and truncated below c . Here $\mu_i | \sigma_i^2 \stackrel{ind.}{\sim} \text{Logistic}(\text{location} = \sigma_i^2, \text{scale} = \sqrt{\sigma_i^2})$.


Figure 4 Comparison of relative risks when (μ_i, σ_i^2) are dependent - Settings 6 and 7. Plots show $m = 10, 15, 20$ left to right.

Settings 4 to 7, Figures 3b, 3c and 4, represent scenarios where NEST dominates NPMLE uniformly for all values of u and m . The relative risk of NPMLE in these scenarios is sometimes as high as 1.6 while NEST and NEST Oracle λ exhibit relative risk profiles substantially closer to 1.

The eight and ninth settings, Figures 5a and 5b, correspond to the scenarios where the data $Y_{ij} | (\mu_i, \sigma_i^2)$ are not normally distributed. In Figure 5a, (μ_i, σ_i^2) are sampled according to Setting 5 but $Y_{ij} | (\mu_i, \sigma_i^2)$ are generated independently from a Laplace distribution with location μ_i and scale $\sqrt{\sigma_i^2/2}$. Similarly, in Figure 5b, (μ_i, σ_i^2) are sampled according to Setting 4 but $Y_{ij} | (\mu_i, \sigma_i^2) = \mu_i + \xi_i$



(a) Setting 8: (μ_i, σ_i^2) generated according to Setting 5 and $Y_{ij} | (\mu_i, \sigma_i^2) \stackrel{i.i.d.}{\sim} \text{Laplace}(\text{location} = \mu_i, \text{scale} = \sqrt{\sigma_i^2/2})$.



(b) Setting 9: (μ_i, σ_i^2) generated according to Setting 4 and $Y_{ij} | (\mu_i, \sigma_i^2) = \mu_i + \xi_i$ where ξ_i are independently sampled from a central t -distribution with 4 degrees of freedom.

Figure 5 Comparison of relative risks when the data $Y_{ij} | (\mu_i, \sigma_i^2)$ are not normally distributed. Plots show $m = 10, 15, 20$ left to right.

where ξ_i are independently sampled from a central t -distribution with 4 degrees of freedom. Across both these settings the proposed NEST method demonstrates robustness to departures from the Normal model. Proposition 7 in Barp et al. (2019) guarantees that, in general, the influence function of minimum KSD estimators, such as the NEST estimator, is bounded under data corruption and the behavior of the NEST estimator in Settings 8 and 9 is potentially an example of such a robustness property. The scenario in Setting 9 is particularly interesting because we note from Figure 5b that the NEST estimator exhibits a smaller estimation risk than the Bayes oracle across all values of u . This is not surprising because in Settings 8 and 9, the Bayes oracle estimator $\delta_{(1)}^\pi$

in Equation (6), while having full knowledge of the underlying prior distributions, continues to rely on the incorrect Gaussian data generation process for Y_{ij} to estimate the means. The NEST estimator, on the other hand, exhibits potential robustness to such mis-specification for estimating the shrinkage factors from the data.

Overall, the results of the preceding nine simulation settings reveal that when the variances are unknown, the NEST estimation framework, in general, enjoys a relatively better risk performance for estimating μ under the weighted squared error loss than the NPMLE based g -modeling approach. However, as Setting 3 reveals, the NPMLE may exhibit a substantially better risk performance when the underlying prior distribution on the means can be well approximated by a discrete mixture with a finite number of mass points. In this scenario, the NEST estimator is unable to effectively account for the potential multimodality of the marginal distribution while estimating the shrinkage factors and therefore may return a poorer estimate of the means.

In Section EC.3 of the EC, we present additional numerical experiments for the following cases: compound estimation of Normal means under the squared error loss (Section EC.3.1), under unequal and small sample sizes m_i (Section EC.3.2) and compound estimation of ratios under the squared error loss (Section EC.3.3). Two real data examples are discussed in Section EC.4.

6. Discussion

In this article we propose a nonparametric empirical Bayes framework, NEST, for large scale estimation of normal means when the corresponding precisions are unknown. We show that the data-driven NEST estimator is asymptotically as good as the optimal Bayes rule for estimating the means under the weighted squared error loss. Under the usual squared error loss our analysis in Section EC.1 reveals that the oracle NEST estimator dominates Tweedie's formula that relies on sample variances for shrinkage estimation of the means. Furthermore, in our numerical experiments and real data examples NEST delivers a better risk performance, often substantially, than competing shrinkage estimators under both the weighted squared error loss and the squared error loss.

A key component of the NEST framework is the consistent estimation of the bivariate shrinkage factors that appear in the oracle NEST estimator. Our approach for estimating these shrinkage factors relies on the KSD measure and we solve a convex optimization problem to derive the shrinkage factors. However, the advantage of the KSD approach for empirical Bayes methods extends far beyond the hierarchical Model (1) considered in this paper. For instance, a natural extension of Model (1) is in large-scale regression problems where the goal is to estimate a large number of low-dimensional regression coefficients when some potentially useful side information on these unknown regression coefficients are available. In this setting the KSD approach is a natural candidate for estimating the multivariate shrinkage factors that arise in this context. Our future research will be directed towards studying this problem and developing empirical Bayes methods for the same.

Acknowledgments

The authors thank the Editor, the Associate Editor and the anonymous referees for their valuable comments and feedback that have significantly improved the quality and presentation of this paper. T. Banerjee was partially supported by the University of Kansas General Research Fund allocation #2302216. W. Sun was supported in part by NSF grant DMS-2015339.

References

- Abramovich F, Benjamini Y, Donoho DL, Johnstone IM (2006) Adapting to unknown sparsity by controlling the false discovery rate. *Ann. Statist.* 34:584–653, ISSN 0090-5364. 2
- Banerjee T, Liu Q, Mukherjee G, Sun W (2021) A general framework for empirical bayes estimation in discrete linear exponential family. *The Journal of Machine Learning Research* 22(1):3108–3153. 16
- Barp A, Briol FX, Duncan A, Girolami M, Mackey L (2019) Minimum stein discrepancy estimators. *Advances in Neural Information Processing Systems*, 12964–12976. 25
- Basu P, Cai TT, Das K, Sun W (2017) Weighted false discovery rate control in large-scale multiple testing. *Journal of the American Statistical Association* 0(ja):0–0, URL <http://dx.doi.org/10.1080/01621459.2017.1336443>. 2

- Benjamini Y, Hochberg Y (1997) Multiple hypotheses testing with weights. *Scandinavian Journal of Statistics* **24**:407–418. [2](#)
- Benjamini Y, Yekutieli D (2011) False discovery rate-adjusted multiple confidence intervals for selected parameters. *Journal of the American Statistical Association* 100(469):71–81, URL <http://dx.doi.org/10.1198/016214504000001907>. [2](#)
- Berger JO (1976) Admissible minimax estimation of a multivariate normal mean with arbitrary quadratic loss. *Ann. Statist.* 4(1):223–226, URL <http://dx.doi.org/10.1214/aos/1176343356>. [2](#)
- Berk R, Brown L, Buja A, Zhang K, Zhao L (2013) Valid post-selection inference. *Ann. Statist.* 41(2):802–837, URL <http://dx.doi.org/10.1214/12-AOS1077>. [2](#)
- Brown LD (2008) In-season prediction of batting averages: A field test of empirical bayes and bayes methodologies. *The Annals of Applied Statistics* 113–152. [ec32](#), [ec33](#)
- Brown LD, Greenshtein E (2009) Nonparametric empirical Bayes and compound decision approaches to estimation of a high-dimensional vector of normal means. *The Annals of Statistics* 37:1685–1704. [2](#), [3](#), [6](#)
- Brown LD, Greenshtein E, Ritov Y (2013) The poisson compound decision problem revisited. *Journal of the American Statistical Association* 108(502):741–749. [15](#)
- Brown SJ, Goetzmann W, Ibbotson RG, Ross SA (1992) Survivorship bias in performance studies. *The Review of Financial Studies* 5(4):553–580, URL <http://dx.doi.org/10.1093/rfs/5.4.553>. [3](#)
- Cai J, Han X, Ritov Y, Zhao L (2021) Nonparametric empirical bayes estimation and testing for sparse and heteroscedastic signals. *arXiv preprint arXiv:2106.08881* . [8](#)
- Cai TT, Sun W (2009) Simultaneous testing of grouped hypotheses: Finding needles in multiple haystacks. *J. Amer. Statist. Assoc.* 104:1467–1481. [3](#)
- Carpenter B, Gelman A, Hoffman MD, Lee D, Goodrich B, Betancourt M, Brubaker M, Guo J, Li P, Riddell A (2017) Stan: A probabilistic programming language. *Journal of statistical software* 76(1). [20](#)
- Castillo I, van der Vaart A (2012) Needles and straw in a haystack: Posterior concentration for possibly sparse sequences. *Ann. Statist.* 40(4):2069–2101, URL <http://dx.doi.org/10.1214/12-AOS1029>. [2](#)

- Chiaretti S, Li X, Gentleman R, Vitale A, Vignetti M, Mandelli F, Ritz J, Foa R (2004) Gene expression profile of adult t-cell acute lymphocytic leukemia identifies distinct subsets of patients with different response to therapy and survival. *Blood* 103(7):2771–2778, ISSN 0006-4971, URL <http://dx.doi.org/10.1182/blood-2003-09-3243>. 3
- Chwialkowski K, Strathmann H, Gretton A (2016) A kernel test of goodness of fit (JMLR: Workshop and Conference Proceedings). 13, 16
- Diaconis P, Ylvisaker D (1979) Conjugate priors for exponential families. *The Annals of statistics* 269–281. ec3
- Donoho DL, Jonhstone JM (1994) Ideal spatial adaptation by wavelet shrinkage. *Biometrika* 81(3):425, URL <http://dx.doi.org/10.1093/biomet/81.3.425>. 2
- Dyson F (1926) A Method for Correcting Series of Parallax Observations. *Monthly Notices of the Royal Astronomical Society* 86(9):686–706, ISSN 0035-8711, URL <http://dx.doi.org/10.1093/mnras/86.9.686>. 2
- Eddington AS (1940) The Correction of Statistics for Accidental Error. *Monthly Notices of the Royal Astronomical Society* 100(5):354–361, ISSN 0035-8711, URL <http://dx.doi.org/10.1093/mnras/100.5.354>. 2
- Efron B (2008) Microarrays, empirical Bayes and the two-groups model. *Statist. Sci.* **23**:1–22, ISSN 0883-4237. 3
- Efron B (2011) Tweedie’s formula and selection bias. *Journal of the American Statistical Association* 106(496):1602–1614. 2, 3, 6, 8
- Efron B (2014) Two modeling strategies for empirical bayes estimation. *Statistical science* 29(2):285–301. 6
- Efron B (2016) Empirical bayes deconvolution estimates. *Biometrika* 103(1):1–20. 6
- Efron B, Morris CN (1975) Data analysis using stein’s estimator and its generalizations. *Journal of the American Statistical Association* 70(350):311–319. 2
- Erickson S, Sabatti C (2005) Empirical Bayes estimation of a sparse vector of gene expression change. *Statistical applications in genetics and molecular biology* 4(1):1132. 3

- Gu J, Koenker R (2017a) Empirical bayesball remixed: Empirical bayes methods for longitudinal data. *Journal of Applied Econometrics* 32(3):575–599. 6, 20, ec23, ec32, ec33
- Gu J, Koenker R (2017b) Unobserved heterogeneity in income dynamics: An empirical bayes perspective. *Journal of Business & Economic Statistics* 35(1):1–16. 20, ec23
- He L, Sarkar SK, Zhao Z (2015) Capturing the severity of type ii errors in high-dimensional multiple testing. *Journal of Multivariate Analysis* 142:106 – 116, ISSN 0047-259X, URL <http://dx.doi.org/https://doi.org/10.1016/j.jmva.2015.08.005>. 2
- Henderson NC, Newton MA (2016) Making the cut: improved ranking and selection for large-scale inference. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 78:1467–9868. 2, 4
- Ignatiadis N, Wager S (2019) Covariate-powered empirical bayes estimation. *Advances in Neural Information Processing Systems*, 9620–9632. 16
- James W, Stein C (1961) Estimation with quadratic loss. *Proceedings of the Fourth Berkeley Symposium on Mathematical Statistics and Probability, Volume 1: Contributions to the Theory of Statistics*, 361–379 (Berkeley, Calif.: University of California Press), URL <http://projecteuclid.org:443/euclid.bsmmsp/1200512173>. 2, 11
- Jiang W, Zhang CH (2009) General maximum likelihood empirical bayes estimation of normal means. *Ann. Statist.* 37(4):1647–1684, URL <http://dx.doi.org/10.1214/08-AOS638>. 2, 6
- Jiang W, Zhang CH, et al. (2010) Empirical bayes in-season prediction of baseball batting averages. *Borrowing Strength: Theory Powering Applications—A Festschrift for Lawrence D. Brown*, 263–273 (Institute of Mathematical Statistics). ec33
- Jing BY, Li Z, Pan G, Zhou W (2016a) On sure-type double shrinkage estimation. *Journal of the American Statistical Association* 111(516):1696–1704. 4, 5, 14, ec2
- Jing BY, Li Z, Pan G, Zhou W (2016b) On sure-type double shrinkage estimation. *Journal of the American Statistical Association* 111(516):1696–1704. ec4, ec23, ec25, ec31, ec38
- Johnstone IM, Silverman BW (2004) Needles and straw in haystacks: empirical Bayes estimates to possibly sparse sequences. *Annals of Statistics* 32(4):1594–1649. 2

- Kiefer J, Wolfowitz J (1956) Consistency of the maximum likelihood estimator in the presence of infinitely many incidental parameters. *The Annals of Mathematical Statistics* 887–906. 6, 20
- Kim Y, Carbonetto P, Stephens M, Anitescu M (2020) A fast algorithm for maximum likelihood estimation of mixture proportions using sequential quadratic programming. *Journal of Computational and Graphical Statistics* 29(2):261–273. ec39
- Koenker R, Gu J (2017) Rebayes: Empirical bayes mixture methods in r. *Journal of Statistical Software* 82(8):1–26. 20, ec23, ec38
- Koenker R, Mizera I (2014) Convex optimization, shape constraints, compound decisions, and empirical Bayes rules. *Journal of the American Statistical Association* 109(506):674–685. 3, 6, 14, ec38
- Kou SC, Yang JJ (2017) *Optimal Shrinkage Estimation in Heteroscedastic Hierarchical Linear Models*, chapter 25, 249–284 (Cham: Springer International Publishing), URL http://dx.doi.org/10.1007/978-3-319-41573-4_13. 3, 4, 5
- Kwon Y, Zhao Z (2022) On F-modelling-based empirical Bayes estimation of variances. *Biometrika* ISSN 1464-3510, URL <http://dx.doi.org/10.1093/biomet/asac019>, asac019. ec2
- Laird N (1978) Nonparametric maximum likelihood estimation of a mixing distribution. *Journal of the American Statistical Association* 73(364):805–811. 6, 20
- Laurent B, Massart P (2000) Adaptive estimation of a quadratic functional by model selection. *Annals of Statistics* 1302–1338. ec19, ec21
- Lee JD, Sun DL, Sun Y, Taylor JE (2016) Exact post-selection inference, with application to the lasso. *Ann. Statist.* 44(3):907–927, URL <http://dx.doi.org/10.1214/15-AOS1371>. 2
- Liu Q, Lee JD, Jordan MI (2016) A kernelized stein discrepancy for goodness-of-fit tests. *Proceedings of the International Conference on Machine Learning (ICML)*. 13, 16, 17, ec12
- Liu Q, Wang D (2016) Stein variational gradient descent: A general purpose bayesian inference algorithm. *Advances In Neural Information Processing Systems*, 2378–2386. 16
- Luo J, MUukherjee G, Sun W (2022) Transfer learning for empirical bayes estimation: A nonparametric integrative tweedie approach URL <https://gmukherjee.github.io/pdfs/nit.pdf>. 16, ec14

- Oates CJ, Girolami M, Chopin N (2017) Control functionals for monte carlo integration. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 79(3):695–718. [16](#)
- Robbins H (1956) An empirical Bayes approach to statistics. *Proc. Third Berkeley Symp. on Math. Statistic. and Prob.* 1:157–163. [2](#)
- Saha S, Guntuboyina A (2020) On the nonparametric maximum likelihood estimator for gaussian location mixture densities with application to gaussian denoising. *The Annals of Statistics* 48(2):738–762. [6](#), [22](#)
- Serfling RJ (2009) *Approximation theorems of mathematical statistics*, volume 162 (John Wiley & Sons). [18](#), [ec12](#)
- Soloff JA, Guntuboyina A, Sen B (2021) Multivariate, heteroscedastic empirical bayes via nonparametric maximum likelihood. *arXiv preprint arXiv:2109.03466* . [6](#)
- Stan Development Team (2022) RStan: the R interface to Stan. URL <https://mc-stan.org/>, r package version 2.21.5. [20](#)
- Sun W, McLain AC (2012) Multiple testing of composite null hypotheses in heteroscedastic models. *Journal of the American Statistical Association* 107(498):673–687. [2](#), [3](#)
- Tan Z (2015) Improved minimax estimation of a multivariate normal mean under heteroscedasticity. *Bernoulli* 21:574–603. [4](#), [5](#)
- Tusher VG, Tibshirani R, Chu G (2001) Significance analysis of microarrays applied to the ionizing radiation response. *Proceedings of the National Academy of Sciences* 98(9):5116–5121. [3](#)
- Weinstein A, Fithian W, Benjamini Y (2013) Selection adjusted confidence intervals with more power to determine the sign. *Journal of the American Statistical Association* 108(501):165–176, URL <http://dx.doi.org/10.1080/01621459.2012.737740>. [2](#)
- Weinstein A, Ma Z, Brown LD, Zhang CH (2018) Group-linear empirical bayes estimates for a heteroscedastic normal mean. *Journal of the American Statistical Association* 0(0):1–13, URL <http://dx.doi.org/10.1080/01621459.2017.1280406>. [3](#), [4](#), [5](#), [6](#), [8](#), [ec1](#), [ec4](#), [ec23](#), [ec36](#), [ec38](#)
- Wharton School (1993) Wharton research data services. <https://wrds-web.wharton.upenn.edu/wrds/> . [ec34](#)
- Xie X, Kou S, Brown LD (2012) Sure estimates for a heteroscedastic hierarchical model. *Journal of the American Statistical Association* 107(500):1465–1479. [4](#), [5](#), [8](#), [ec1](#), [ec29](#), [ec34](#), [ec38](#)

Yang J, Liu Q, Rao V, Neville J (2018) Goodness-of-fit testing for discrete distributions via stein discrepancy.

International Conference on Machine Learning, 5561–5570. [16](#)

Zhang X, Bhattacharya A (2017) Empirical Bayes, sure, and sparse normal mean models, preprint. [4](#)

Electronic Companion for “Nonparametric Empirical Bayes Estimation On Heterogeneous Data”

This Electronic Companion (EC) is organized as follows: in Section [EC.1](#) we present an analysis of the oracle NEST estimator (Definition [1](#)) under the squared error loss. The proofs of all theoretical results in the paper are presented in Section [EC.2](#). In Section [EC.3](#), we provide additional numerical experiments for the following cases: compound estimation of Normal means under the squared error loss (Section [EC.3.1](#)), under unequal and small sample sizes m_i (Section [EC.3.2](#)) and compound estimation of ratios under the squared error loss (Section [EC.3.3](#)). Two real data examples are presented in Section [EC.4](#). In Section [EC.5](#) we discuss extensions of our methodology to several well known members in the two-parameter exponential family when the nuisance parameter is known. We conclude with a discussion on the computational complexity of NEST in Section [EC.6](#).

EC.1. Analysis under the squared error loss

In this section we present an analysis of the oracle NEST estimator $\delta_{(1)}^\pi$ with respect to the squared error loss. We first introduce a counterpart to Tweedie’s formula for μ_i (Equation [\(2\)](#)) that relies on the sample variances S_i^2 instead of the unknown variances. Thereafter, we compare these two shrinkage estimators under the squared error loss.

When the precisions are unknown existing shrinkage methods, such as [Xie et al. \(2012\)](#), [Weinstein et al. \(2018\)](#), usually rely on the sample variance S_i^2 , a consistent estimator of the unknown population variance $1/\tau_i$, for practical implementation. For instance, in Definition [EC.1](#) we present the oracle Pseudo-Tweedie’s formula with sample variance which is a natural counterpart to Equation [\(2\)](#) when the variance is unknown.

DEFINITION EC.1. (Oracle Pseudo-Tweedie’s formula with sample variances) Consider the hierarchical Model [\(1\)](#). Then an estimator for μ_i is δ_i^{TF} where,

$$\delta_i^{\text{TF}} := \delta^{\text{TF}}(y_i, s_i^2, m_i) = y_i + \frac{s_i^2}{m_i} w_1(y_i, s_i^2; m_i),$$

and $w_1(y, s^2; m) := \frac{\partial}{\partial y} \log f_m(y, s^2)$.

The oracle Pseudo-Tweedie's formula in Definition EC.1 has a striking similarity to the oracle NEST estimator $\delta_{i,(1)}^\pi$ (Definition 1) in the sense that both δ_i^{TF} and $\delta_{i,(1)}^\pi$ involve an unbiased estimate y_i of μ_i plus a shrinkage factor. The key difference, however, is that (1) while the shrinkage factor in $\delta_{i,(1)}^\pi$ relies on γ_i , δ_i^{TF} directly uses sample variances s_i^2 , and (2) while it is not immediately clear what is the underlying loss function that δ_i^{TF} is minimizing for estimating μ_i , $\delta_{i,(1)}^\pi$ uniquely minimizes the expected weighted squared error loss.

Notwithstanding, Proposition EC.1 establishes that the oracle Pseudo-Tweedie's estimator $\boldsymbol{\delta}^{\text{TF}} = (\delta_1^{\text{TF}}, \dots, \delta_n^{\text{TF}})$ dominates the sample mean estimator under a squared error loss function.

PROPOSITION EC.1. *Denote $w_1'(y, s^2; m)$ as the partial derivative of $w_1(y, s^2; m)$ with respect to y . Suppose $f_m(y, s^2)$ is a log-concave density and $w_1'(y, s^2; m)$ is a non-decreasing function of s^2 . Then under Model (1) and for $m_i > 3$, we have,*

$$r_0(\boldsymbol{\delta}_{(0)}^\pi, \mathcal{G}) \leq r_0(\boldsymbol{\delta}^{\text{TF}}, \mathcal{G}) < r_0(\mathbf{Y}, \mathcal{G}).$$

However, when a large number of units are investigated simultaneously, traditional sample variance estimators may suffer from selection bias (Jing et al. 2016a, Kwon and Zhao 2022), and their direct use may lead to severe deterioration in the MSE for estimating the means. We present Proposition EC.2 which shows that under the squared error loss function the risk of the oracle NEST estimator is uniformly smaller than that of the Pseudo-Tweedie's estimator $\boldsymbol{\delta}^{\text{TF}}$ (Definition EC.1).

We impose the following regularity conditions for comparing the estimation risks of $\boldsymbol{\delta}_{(1)}^\pi$ and $\boldsymbol{\delta}^{\text{TF}}$ in Proposition EC.2.

ASSUMPTION EC.1. *The shrinkage factor $\gamma(y, s^2, m)$ is non-increasing in s^2 .*

ASSUMPTION EC.2. *Let $\omega(y, s^2; m) := w_1(y, s^2; m) \frac{\partial}{\partial y} w_2(y, s^2; m)$. Then, $\omega(y, s^2; m) \leq 0$.*

Assumptions EC.1 and EC.2 are regularity conditions on the behavior of the shrinkage factor γ and the score functions w_1, w_2 . For instance, Assumption EC.1 enforces monotonicity on shrinkage factor $\gamma(y, s^2, m)$, which is satisfied, for example, when (μ, τ) have a conjugate prior under Model (1). Similarly, Assumption EC.2 holds under conjugate priors and is also true when the prior on τ

is discrete with just one mass point. In Remarks EC.1 and EC.2 we discuss these assumptions in more details. Proposition EC.1 can be extended as follows.

PROPOSITION EC.2. *Suppose $f_m(y, s^2)$ is a log-concave density and $w_1'(y, s^2; m)$ is a non-decreasing function of s^2 . If Assumptions EC.1 – EC.2 hold then, under Model (1) and for $m_i > 5$, we have,*

$$r_0(\boldsymbol{\delta}_{(0)}^\pi, \mathcal{G}) \leq r_0(\boldsymbol{\delta}_{(1)}^\pi, \mathcal{G}) \leq r_0(\boldsymbol{\delta}^{\text{TF}}, \mathcal{G}) < r_0(\mathbf{Y}, \mathcal{G}).$$

Note that the oracle NEST estimator $\delta_{i,(1)}^\pi$ is, in general, different from $\delta_{i,(0)}^\pi$: $\delta_{i,(1)}^\pi$ is the Bayes estimator under the weighted squared error loss, while $\delta_{i,(0)}^\pi$ is the Bayes estimator under the usual squared error loss. This follows since from Equation (3), and dropping subscript i ,

$$\delta_{(1)}^\pi = \frac{\mathbb{E}(\tau\mu|y, s^2)}{\mathbb{E}(\tau|y, s^2)} = \frac{\mathbb{E}[\tau\mathbb{E}(\mu|y, s^2, \tau)|y, s^2]}{\mathbb{E}(\tau|y, s^2)} \neq \mathbb{E}(\mu|y, s^2) = \delta_{(0)}^\pi.$$

However, $r_0(\boldsymbol{\delta}_{(0)}^\pi, \mathcal{G}) = r_0(\boldsymbol{\delta}_{(1)}^\pi, \mathcal{G})$ when μ and τ are conditionally independent given y and s^2 , in which case $\delta_{(1)}^\pi = \delta_{(0)}^\pi$. In particular, Corollary EC.1 shows that when $\mathbb{E}(\mu_i|y_i, s_i^2, \tau_i)$ is independent of τ_i then the data-driven NEST estimator is asymptotically close to $\boldsymbol{\delta}_{(0)}^\pi$.

COROLLARY EC.1. *Consider hierarchical Model (1) and suppose $\mathbb{E}(\mu_i|y_i, s_i^2, \tau_i)$ is independent of τ_i . Then, under the conditions of Theorem 3, we have $\frac{c_n}{n} \|\boldsymbol{\delta}_n^{\text{ds}}(\lambda) - \boldsymbol{\delta}_{(0)}^\pi\|_2^2 = o_p(1)$ as $n \rightarrow \infty$. Furthermore, under the same conditions, $c_n^{1/2} |l_n^{(0)}\{\boldsymbol{\delta}_n^{\text{ds}}(\lambda), \boldsymbol{\mu}; \boldsymbol{\tau}\} - l_n^{(0)}\{\boldsymbol{\delta}_{(0)}^\pi, \boldsymbol{\mu}; \boldsymbol{\tau}\}| = o_p(1)$ as $n \rightarrow \infty$.*

Corollary EC.1 is a straightforward consequence of Theorem 3 and the fact that when $\mathbb{E}(\mu_i|y_i, s_i^2, \tau_i)$ is independent of τ_i , $\delta_{i,(1)}^\pi = \delta_{i,(0)}^\pi$. A popular setting where $\mathbb{E}(\mu_i|y_i, s_i^2, \tau_i)$ is indeed independent of τ_i is the scenario where $G_\mu(\cdot|\tau)$ and $H_\tau(\cdot)$ belong to the family of conjugate priors under Model (1). Under this setting the posterior expectation of μ is a linear combination of the prior expectation of μ and the maximum likelihood estimate y (Diaconis and Ylvisaker 1979), and the weights in this linear combination are proportional to m and the prior sample size. For instance, if $G_\mu(\cdot|\tau)$ is $N(\mu_0, 1/\tau)$ and $H_\tau(\cdot)$ is $\Gamma(\text{shape} = \alpha, \text{rate} = \beta)$ in Model (1) then standard calculations give $\delta_{(0)}^\pi = (\mu_0 + my)/(m+1)$ where μ_0 is the prior expectation of μ and $m_0 = 1$ is the prior sample size. Moreover, in this setting we also have

$$\log f_m(y_i, s_i^2) = c_0 + \frac{m-3}{2} \log s_i^2 - (\alpha + m/2) \log \left\{ \beta + 0.5(m-1)s_i^2 + 0.5 \frac{m}{m+1} (y_i - \mu_0)^2 \right\},$$

where c_0 is a constant independent of (y_i, s_i^2) . From the above display,

$$w_1(y_i, s_i^2) = -\frac{(\alpha + m/2)\{m/(m+1)\}(y_i - \mu_0)}{\beta + 0.5(m-1)s_i^2 + 0.5\{m/(m+1)\}(y_i - \mu_0)^2}$$

$$w_2(y_i, s_i^2) = \frac{(m-3)}{2s_i^2} - \frac{0.5(\alpha + m/2)(m-1)}{\beta + 0.5(m-1)s_i^2 + 0.5\{m/(m+1)\}(y_i - \mu_0)^2}.$$

Substituting these expressions for $w_1(y_i, s_i^2), w_2(y_i, s_i^2)$ in Equation (6) gives $\delta_{i,(1)}^\pi = \delta_{i,(0)}^\pi$. Another example where $\delta_{(1)}^\pi = \delta_{(0)}^\pi$ is when $G_\mu(\cdot|\tau)$ and $H_\tau(\cdot)$ are independent discrete distributions with just one mass point, respectively. These two examples are also the settings, among others, under which the first inequality in the left hand side of Proposition EC.1 is strict, that is, $r_0(\delta_{(0)}^\pi, \mathcal{G}) < r_0(\delta^{\text{TF}}, \mathcal{G})$.

In Section EC.3.1 we present several numerical experiments to compare the data-driven NEST estimator against other shrinkage estimators for compound estimation of the Normal means from Model (1) under the squared error loss. Our empirical results corroborate Proposition EC.2 and suggest that the efficiency gain of the data-driven NEST estimator over competing linear shrinkage methods (Jing et al. 2016b, Weinstein et al. 2018) and Pseudo-Tweedie's formula (Definition EC.1) is substantial across many settings. We end this section with the following remarks:

REMARK EC.1. Here we provide three examples of \mathcal{G} that satisfy Assumptions EC.1 and EC.2.

1. $G_\mu(\cdot|\tau)$ and $H_\tau(\cdot)$ belong to the family of conjugate priors under Model (1).
2. $H_\tau(\cdot)$ is a distribution with just one mass point, say τ_0 , and $G_\mu(\cdot|\tau_0)$ has a log-concave density.
3. $H_\tau(\cdot)$ has a log-concave density and $G_\mu(\cdot|\tau)$ is a discrete distribution with just one mass point,

say μ_0 .

In Figure EC.1 we consider specific choices of \mathcal{G} for each of these three examples. In particular, the left panel of Figure EC.1 provides a contour plot of \mathcal{G} for example 1 where $G_\mu(\cdot|\tau)$ is $N(0, 1/\tau)$ and $H_\tau(\cdot)$ is $\Gamma(\text{shape} = 10, \text{rate} = 5)$. The center panel presents the contour plot of the marginal density f_m for example 2 where $G_\mu(\cdot|\tau)$ is $N(0, 1/\tau)$ and $H_\tau(\cdot)$ has all its mass concentrated at $\tau_0 = 1$. Finally, the right panel presents the contour plot of f_m for example 3 where $\mu_0 = 0$ and $H_\tau(\cdot)$ is $\Gamma(\text{shape} = 10, \text{rate} = 5)$.




Figure EC.1 Remark EC.1 contour plots with $m = 5$. Left to Right: contour plot of \mathcal{G} for example 1 where $G_\mu(\cdot|\tau)$ is $N(0, 1/\tau)$ and $H_\tau(\cdot)$ is $\Gamma(\text{shape} = 10, \text{rate} = 5)$. Contour plot of the marginal density f_m for example 2 where $G_\mu(\cdot|\tau)$ is $N(0, 1/\tau)$ and $H_\tau(\cdot)$ has all its mass concentrated at $\tau_0 = 1$. Contour plot of f_m for example 3 where $\mu_0 = 0$ and $H_\tau(\cdot)$ is $\Gamma(\text{shape} = 10, \text{rate} = 5)$.

REMARK EC.2. Here we provide two examples of \mathcal{G} that induce marginal densities f_m which are no longer log-concave and, hence, do not satisfy Assumption EC.1. In these examples, log-concavity of f_m is lost because $w_1'(y, s^2; m)$ may be positive on its domain.

1. $H_\tau(\cdot)$ is a distribution with just one mass point, say τ_0 , and $G_\mu(\cdot|\tau_0)$ has mass points at $2\tau_0$ and $-2\tau_0$ with equal probability.
2. $H_\tau(\cdot)$ has mass points at 1 and 2 with equal probability, and $\mu = 4/\tau$.

Figure EC.2 presents the contour plots of f_m for these two examples.

EC.2. Proofs

EC.2.1. Proof of Proposition EC.1

Recall from Definition EC.1 that

$$\delta_i^{\text{TF}} := \delta^{\text{TF}}(y_i, s_i^2, m_i) = y_i + \frac{s_i^2}{m_i} w_1(y_i, s_i^2; m_i),$$

where $w_1(y, s^2; m) := \frac{\partial}{\partial y} \log f_m(y, s^2)$. Since the n study units are independent, we will focus on unit i .

We note that the inequality in the left hand side of Proposition EC.1 follows from the fact that $\delta_{i,(0)}^\pi$ is the Bayes estimator of μ_i that uniquely minimizes the expected squared error loss under




Figure EC.2 Remark EC.2 contour plots. Left: contour plot of the marginal density f_m for example 1 where $\tau_0 = 1$ and $m = 5$. Right: contour plot of f_m for example 2 with $m = 5$.

Model (1). To prove the inequality in the right hand side of Proposition EC.1, we proceed as follows.

Denote $w_{1,i} := w_1(y_i, s_i^2; m_i)$. We have,

$$\begin{aligned} \mathbb{E}(\mu_i - \delta_i^{\text{TF}})^2 &= \frac{1}{m_i} \mathbb{E}(1/\tau_i) + \frac{1}{m_i^2} \mathbb{E}(S_i^2 w_{1,i})^2 + \frac{2}{m_i} \mathbb{E}(Y_i - \mu_i) S_i^2 w_{1,i} \\ &= \mathbb{E}(\mu_i - Y_i)^2 + \frac{1}{m_i^2} \mathbb{E}(S_i^2 w_{1,i})^2 + \frac{2}{m_i^2} \mathbb{E} \frac{S_i^2}{\tau_i} w'_{1,i}, \end{aligned} \quad (\text{EC.1})$$

where $w'_{1,i} := w'_1(y_i, s_i^2; m_i)$ and $w'_1(y, s^2; m) = \frac{\partial}{\partial y} w_1(y, s^2; m)$. The equality in Equation (EC.1)

follows from integration by parts and the fact that $Y_i | \mu_i, \tau_i \sim N(\mu_i, 1/(m_i \tau_i))$. Consider the term

$T_1 := \frac{1}{m_i^2} \mathbb{E}(S_i^2 w_{1,i})^2 + \frac{2}{m_i^2} \mathbb{E} \left(\frac{S_i^2}{\tau_i} w'_{1,i} \right)$ and note that

$$\begin{aligned} T_1 &= \frac{1}{m_i^2} \mathbb{E} \left(S_i^2 w_{1,i} \right)^2 + \frac{1}{m_i^2} \mathbb{E} \left[(S_i^2)^2 w'_{1,i} \right] + \frac{2}{m_i^2} \mathbb{E} \left(\frac{S_i^2}{\tau_i} w'_{1,i} \right) - \frac{1}{m_i^2} \mathbb{E} \left[(S_i^2)^2 w'_{1,i} \right] \\ &= \frac{1}{m_i^2} \mathbb{E} \left[\frac{2S_i^2}{\tau_i} - (S_i^2)^2 \right] w'_{1,i}. \end{aligned} \quad (\text{EC.2})$$

The equality in Equation (EC.2) follows because, dropping subscript i ,

$$\mathbb{E} \left[(S^2)^2 (w_1^2 + w'_1) \right] = \mathbb{E} \left[(S^2)^2 \frac{f''_{m,(1)}(Y, S^2)}{f_m(Y, S^2)} \right] = 0,$$

where $f''_{m,(1)}(y, s^2)$ is the second order partial derivative of $f_m(y, s^2)$ with respect to y . Now, we can re-write Equation (EC.2) as

$$T_1 = \frac{1}{m_i^2} \mathbb{E}_{\mu_i, \tau_i} \mathbb{E}_{Y_i, S_i^2 | \mu_i, \tau_i} \left\{ \left[\frac{2S_i^2}{\tau_i} - (S_i^2)^2 \right] w'_{1,i} \right\} = \frac{1}{m_i^2} \mathbb{E}_{\mu_i, \tau_i} (T_2) + \frac{1}{m_i^2} \mathbb{E}_{\mu_i, \tau_i} (T_3), \quad (\text{EC.3})$$

where $\mathbb{E}_{\mu, \tau}$ is the expectation with respect to the joint distribution of (μ, τ) , $\mathbb{E}_{Y, S^2 | \mu, \tau}$ is the expectation with respect to the joint distribution of (Y, S^2) conditional on (μ, τ) and

$$\begin{aligned} T_2 &= \mathbb{E}_{Y_i, S_i^2 | \mu_i, \tau_i} \left\{ \left[\frac{2S_i^2}{\tau_i} - (S_i^2)^2 \right] w'_{1,i} \mid S_i^2 < \frac{2}{\tau_i} \right\} \mathbb{P} \left(S_i^2 < \frac{2}{\tau_i} \mid \tau_i \right), \\ T_3 &= \mathbb{E}_{Y_i, S_i^2 | \mu_i, \tau_i} \left\{ \left[\frac{2S_i^2}{\tau_i} - (S_i^2)^2 \right] w'_{1,i} \mid S_i^2 > \frac{2}{\tau_i} \right\} \mathbb{P} \left(S_i^2 > \frac{2}{\tau_i} \mid \tau_i \right). \end{aligned}$$

Denote $p := \mathbb{P}(S_i^2 > 2/\tau_i \mid \tau_i)$ and let c_i be the partial derivative of $w_1(y, s^2; m)$ with respect to y and evaluated at $(y_i, 2/\tau_i, m_i)$. Now, using equations (EC.2) and (EC.3) in Equation (EC.1), we get

$$\mathbb{E}(\mu_i - \delta_i^{\text{TF}})^2 = \mathbb{E}(\mu_i - Y_i)^2 + \frac{1}{m_i^2} \mathbb{E}_{\mu_i, \tau_i} (T_2 + T_3). \quad (\text{EC.4})$$

We will show that $T_2 + T_3 < 0$ which will be enough to prove the statement of Proposition EC.1 using Equation (EC.4).

We first state a few results that are straightforward consequences of Model (1) and the assumptions of Proposition EC.1. We have,

1. Under Model 1, $(m_i - 1)S_i^2 \tau_i \sim \chi_{m_i - 1}^2$.
2. Additionally,

$$\mathbb{E}_{Y_i, S_i^2 | \mu_i, \tau_i} \left[\frac{2S_i^2}{\tau_i} - (S_i^2)^2 \right] = \frac{m_i - 3}{(m_i - 1)\tau_i^2} > 0,$$

since $m_i > 3$ in the statement of Proposition EC.1.

3. Since $f_m(y, s^2)$ is a log-concave density, $w'_{1,i} \leq 0$ and so $T_2 \leq 0$ while $T_3 \geq 0$.

Assume, without loss of generality, $w'_{1,i} < 0$. Since $\mathbb{E}_{Y_i, S_i^2 | \mu_i, \tau_i} [2S_i^2/\tau_i - (S_i^2)^2] > 0$, we have

$$(1 - p) \mathbb{E}_{Y_i, S_i^2 | \mu_i, \tau_i} \left[\frac{2S_i^2}{\tau_i} - (S_i^2)^2 \mid S_i^2 < \frac{2}{\tau_i} \right] > -p \mathbb{E}_{Y_i, S_i^2 | \mu_i, \tau_i} \left[\frac{2S_i^2}{\tau_i} - (S_i^2)^2 \mid S_i^2 > \frac{2}{\tau_i} \right] > 0. \quad (\text{EC.5})$$

Furthermore, as $w'_{1,i} < 0$ and $w'_{1,i}$ is a non-decreasing function of s_i^2 ,

$$T_2 \leq (1-p)\mathbb{E}_{Y_i, S_i^2 | \mu_i, \tau_i} \left\{ \left[\frac{2S_i^2}{\tau_i} - (S_i^2)^2 \right] c_i \middle| S_i^2 < \frac{2}{\tau_i} \right\}, \quad (\text{EC.6})$$

Note that $c_i < 0$, and as defined earlier, it is the partial derivative of $w_1(y, s^2; m)$ with respect to y and evaluated at $(y_i, 2/\tau_i, m_i)$. Therefore, using equations (EC.5) and (EC.6),

$$T_2 \leq (1-p)\mathbb{E}_{Y_i, S_i^2 | \mu_i, \tau_i} \left\{ \left[\frac{2S_i^2}{\tau_i} - (S_i^2)^2 \right] c_i \middle| S_i^2 < \frac{2}{\tau_i} \right\} < -p\mathbb{E}_{Y_i, S_i^2 | \mu_i, \tau_i} \left\{ \left[\frac{2S_i^2}{\tau_i} - (S_i^2)^2 \right] c_i \middle| S_i^2 > \frac{2}{\tau_i} \right\} \quad (\text{EC.7})$$

Now, we consider the term T_3 . Recall that

$$T_3 = p\mathbb{E}_{Y_i, S_i^2 | \mu_i, \tau_i} \left\{ \left[\frac{2S_i^2}{\tau_i} - (S_i^2)^2 \right] w'_{1,i} \middle| S_i^2 > \frac{2}{\tau_i} \right\} > 0.$$

Since $w'_{1,i} < 0$ and $w_{1,i}'$ is a non-decreasing function of s_i^2 ,

$$0 < T_3 \leq -p\mathbb{E}_{Y_i, S_i^2 | \mu_i, \tau_i} \left\{ \left[\frac{2S_i^2}{\tau_i} - (S_i^2)^2 \right] |c_i| \middle| S_i^2 > \frac{2}{\tau_i} \right\}. \quad (\text{EC.8})$$

So, using equations (EC.7) and (EC.8)

$$T_2 + T_3 < -p\mathbb{E}_{Y_i, S_i^2 | \mu_i, \tau_i} \left\{ \left[\frac{2S_i^2}{\tau_i} - (S_i^2)^2 \right] (c_i + |c_i|) \middle| S_i^2 > \frac{2}{\tau_i} \right\} = 0,$$

Hence the desired result follows from the display above and Equation (EC.4). ■

EC.2.2. Proof of Proposition EC.2

We will first collect a few notations that will be used throughout the proof. Denote $w_{1,i} := w_1(y_i, s_i^2; m_i)$, $w_{2,i} := w_2(y_i, s_i^2; m_i)$ and $\gamma_i := \gamma(y_i, s_i^2, m_i)$. Let $w'_1(y, s^2; m) = \frac{\partial}{\partial y} w_1(y, s^2; m)$ and denote $w'_{1,i} := w'_1(y_i, s_i^2; m_i)$. Similarly, $\nu(y, s^2; m) = \frac{1}{f_m(y, s^2)} \frac{\partial^2}{\partial y^2} f_m(y, s^2)$ and denote $\nu_i := \nu(y_i, s_i^2; m_i)$. Finally, let $\gamma'(y, s^2, m) = \frac{\partial}{\partial y} \gamma(y, s^2, m)$ and denote $\gamma'_i := \gamma'(y_i, s_i^2, m_i)$.

The proof of Proposition EC.2 will use the following two lemmata.

LEMMA EC.1. *Suppose $f_m(y, s^2)$ is a log concave density and Assumption EC.1 holds. Then under Model (1) and $m_i > 5$, we have,*

$$\mathbb{E} \left\{ \left[\frac{2S_i^2}{\tau_i} (1 - \gamma_i) - (S_i^2)^2 (1 - \gamma_i^2) \right] w'_{1,i} \right\} \geq 0.$$

LEMMA EC.2. Under Assumption EC.2 and Model (1), we have,

$$\frac{1}{m_i^2} \mathbb{E} \left[S_i^2 \left(S_i^2 \gamma_i^2 \nu_i + 2 \frac{w_{1,i} \gamma_i'}{\tau_i} \right) \right] \leq 0.$$

Lemmata EC.1 and EC.2 are proved in Sections EC.2.2.1 and EC.2.2.2 respectively. We now prove Proposition EC.2. Recall from Definition 1 that the oracle NEST estimator for μ_i is

$$\delta_{i,(1)}^\pi = y_i + \frac{s_i^2}{m_i} \gamma(y_i, s_i^2, m_i) w_1(y_i, s_i^2, m_i),$$

where

$$\gamma(y_i, s_i^2, m_i) = \frac{m_i - 1}{m_i - 3 - 2s_i^2 w_2(y_i, s_i^2, m_i)}.$$

Since the n study units are independent, we will focus on unit i .

Under the squared error loss, $\delta_{i,(0)}^\pi$ is the Bayes estimator of μ_i in Model (1). This establishes the first inequality on the left hand side of Proposition EC.2. Together with Proposition EC.1, we only need to show $r_0(\delta_{(1)}^\pi, \mathcal{G}) \leq r_0(\delta^{\text{TF}}, \mathcal{G})$. First note that

$$\begin{aligned} r_0(\delta^{\text{TF}}, \mathcal{G}) - r_0(\delta_{(1)}^\pi, \mathcal{G}) &= \frac{1}{m_i^2} \mathbb{E} \left[(S_i^2)^2 (1 - \gamma_i^2) w_{1,i}^2 \right] + \frac{2}{m_i} \mathbb{E} \left[(Y_i - \mu_i) S_i^2 (1 - \gamma_i) w_{1,i} \right] \\ &= \frac{1}{m_i^2} \mathbb{E} \left[(S_i^2)^2 (1 - \gamma_i^2) w_{1,i}^2 \right] + \frac{2}{m_i^2} \mathbb{E} \left\{ \frac{S_i^2}{\tau_i} \left[(1 - \gamma_i) w_{1,i}' - w_{1,i} \gamma_i' \right] \right\}. \end{aligned} \quad (\text{EC.9})$$

The equality in equation (EC.9) follows from integration by parts and the fact that $Y_i \sim N\left(\mu_i, \frac{1}{m_i \tau_i}\right)$. We can re-write Equation (EC.9) as,

$$\begin{aligned} &\frac{1}{m_i^2} \mathbb{E} \left[(S_i^2)^2 (1 - \gamma_i^2) w_{1,i}^2 \right] + \frac{1}{m_i^2} \mathbb{E} \left[(S_i^2)^2 (1 - \gamma_i^2) w_{1,i}' \right] - \frac{2}{m_i^2} \mathbb{E} \left(\frac{S_i^2}{\tau_i} w_{1,i} \gamma_i' \right) \\ &\quad + \frac{1}{m_i^2} \mathbb{E} \left\{ \left[\frac{2S_i^2}{\tau_i} (1 - \gamma_i) - (S_i^2)^2 (1 - \gamma_i^2) \right] w_{1,i}' \right\}. \end{aligned} \quad (\text{EC.10})$$

From Lemma EC.1, the last term in Equation (EC.10) is non-negative. Let us consider the first three terms in Equation (EC.10) and denote them by,

$$T := \frac{1}{m_i^2} \mathbb{E} \left[(S_i^2)^2 (1 - \gamma_i^2) w_{1,i}^2 \right] + \frac{1}{m_i^2} \mathbb{E} \left[(S_i^2)^2 (1 - \gamma_i^2) w_{1,i}' \right] - \frac{2}{m_i^2} \mathbb{E} \left(\frac{S_i^2}{\tau_i} w_{1,i} \gamma_i' \right).$$

As shown in the proof of Proposition EC.1, $\mathbb{E}[(S^2)^2(w_1^2 + w_1')] = 0$. The above display involving the term T can be written as

$$T = -\frac{1}{m_i^2} \mathbb{E} \left[(S_i^2)^2 \gamma_i^2 \nu_i \right] - \frac{2}{m_i^2} \mathbb{E} \left(\frac{S_i^2}{\tau_i} w_{1,i} \gamma_i' \right). \quad (\text{EC.11})$$

From Lemma EC.2, the term T in Equation (EC.11) is non-negative. This establishes the inequality on the right hand side of $r_0(\boldsymbol{\delta}_{(1)}^\pi, \mathcal{G})$ in Proposition EC.2 and completes the proof. ■

EC.2.2.1. Proof of Lemma EC.1 - Denote

$$Z_i := \frac{2S_i^2}{\tau_i}(1 - \gamma_i) - (S_i^2)^2(1 - \gamma_i^2) = \frac{Q_i}{\tau_i},$$

where $Q_i = 2S_i^2(1 - \gamma_i) - (S_i^2)^2(1 - \gamma_i^2)\tau_i$. Since $f_m(y, s^2)$ is a log-concave density, $w'_{1,i} \leq 0$. We will show that $Q_i \leq 0$ which will be sufficient to prove $Z_i w'_{1,i} \geq 0$ and hence the statement of Lemma EC.1.

Dropping subscript i ,

$$\mathbb{E}(Q) = \mathbb{E}_{Y, S^2} \left[2S^2(1 - \gamma) - (S^2)^2(1 - \gamma^2)\hat{\tau}^\pi \right] = \mathbb{E}_{Y, S^2} [R(Y, S^2)],$$

where $\hat{\tau}^\pi = \mathbb{E}(\tau|y, s^2, m) = (s^2\gamma)^{-1}$ from Equation (8) and \mathbb{E}_{Y, S^2} is expectation with respect to the joint marginal distribution of (Y, S^2) . Suppose, if possible, $Q_i > 0$ for all $(\tau_i, y_i, s_i^2) \in \mathbb{R}^+ \times \mathbb{R} \times \mathbb{R}^+$. We will show that $\mathbb{E}(Q) \leq 0$ which will present a contradiction to $Q > 0$ for all $(\tau, y, s^2) \in \mathbb{R}^+ \times \mathbb{R} \times \mathbb{R}^+$. Fix a $y \in \mathbb{R}$ and consider the following cases.

Case 1 – Suppose $0 < \gamma(y, s^2, m) \leq 1/(2c)$ where $c \geq 1$ is a constant. Then, we have $s^2\hat{\tau}^\pi \geq 2c$ and consequently $s^2(1 + \gamma)\hat{\tau}^\pi > 2$. So $R(y, s^2) < 0$. Now, from Assumption EC.1, γ is a continuous and non-increasing function of s^2 . Therefore, there exists $c_1(y) \in \mathbb{R}^+$, depending on y , such that $s^2 > c_1(y)$ whenever $0 < \gamma(y, s^2, m) \leq 1/(2c)$. Thus, $R(y, s^2) < 0$ for $s^2 > c_1(y)$.

Case 2 – Next, suppose $1/2 < \gamma(y, s^2, m) \leq (m-1)/(2m-6)$. Then $(2m-6)/(m-1) \leq s^2\hat{\tau}^\pi < 2$ and $s^2(1 + \gamma)\hat{\tau}^\pi \geq (3m-7)/(m-1)$. Since $m > 5$, $(3m-7)/(m-1) > 2$ and so $s^2(1 + \gamma)\hat{\tau}^\pi > 2$. Thus $R(y, s^2) < 0$ and using Assumption EC.1, $R(y, s^2) < 0$ for $c_2(y) < s^2 \leq c_1(y)$, where $c_2(y)$ is such that $s^2 > c_2(y)$ whenever $\gamma(y, s^2, m) \leq (m-1)/(2m-6)$.

The remaining four cases proceed in a similar manner as follows:

Case 3 – Suppose $(m-1)/(2m-6) < \gamma(y, s^2, m) \leq (2m-6)/(3m-11)$. Then $(3m-11)/(2m-6) \leq s^2\hat{\tau}^\pi$. So $s^2(1 + \gamma)\hat{\tau}^\pi \geq (5m-17)/(2m-6)$. Since $(5m-17)/(2m-6) > 2$ if $m > 5$, we have

$s^2(1 + \gamma)\hat{\tau}^\pi > 2$ and so $R(y, s^2) < 0$ for $c_3(y) < s^2 \leq c_2(y)$. Similarly, we can show that $R(y, s^2) \leq 0$ for $c_4(y) < s^2 \leq c_3(y)$ where $s^2 \in (c_4(y), c_3(y)]$ whenever $(2m - 6)/(3m - 11) < \gamma(y, s^2, m) \leq 1$.

Case 4 – Now suppose, $1 < \gamma(y, s^2, m) \leq (2m - 5)/(2m - 6)$. Then $(2m - 6)/(2m - 5) \leq s^2\hat{\tau}^\pi < 1$ and $s^2(1 + \gamma)\hat{\tau}^\pi < 2$. Note that here $\gamma > 1$ as opposed to $\gamma \leq 1$ in the earlier cases. Therefore, $s^2(1 + \gamma)\hat{\tau}^\pi < 2$ implies $R(y, s^2) < 0$ for $c_5(y) < s^2 \leq c_4(y)$.

Case 5 – Similarly, if γ is in the intervals $((2m - 5)/(2m - 6), (m - 2)/(m - 3)]$, $((m - 2)/(m - 3), (m - 1)/(m - 3)]$ and $((m - 1)/(m - 3), (m + 1)/(m - 3)]$ then we have $\gamma \geq 1$ and $s^2(1 + \gamma)\hat{\tau}^\pi < 2$. Thus, on each of the corresponding intervals for s^2 , $R(y, s^2) < 0$.

Case 6 – Denote $r_1 = 1$ and $r_t = 2r_{t-1} + 3$ for $t = 2, 3, \dots$. Suppose $(m + r_{t-1})/(m - 3) < \gamma(y, s^2, m) \leq (m + r_t)/(m - 3)$. Then for each of these intervals indexed by $t \geq 2$, we have $\gamma \geq 1$, $s^2(1 + \gamma)\hat{\tau}^\pi < 2$ and so $R(y, s^2) < 0$ on the corresponding intervals for s^2 .

So from these six cases, $R(y, s^2) \leq 0$ for all $s^2 > 0$ and consequently $\mathbb{E}_{Y, S^2}[R(Y, S^2)] \leq 0$. Therefore, $\mathbb{E}(Q) \leq 0$ which contradicts that $Q > 0$ for all $(\tau, y, s^2) \in \mathbb{R}^+ \times \mathbb{R} \times \mathbb{R}^+$. Now suppose that for some $\Omega \subset \mathbb{R}^+ \times \mathbb{R} \times \mathbb{R}^+$, $Q > 0$ whenever $(\tau, y, s^2) \in \Omega$. However, Assumption EC.1 and the aforementioned six cases imply that $\mathbb{E}(Q|\Omega) \leq 0$. Thus, $Q \leq 0$ for all $(\tau, y, s^2) \in \mathbb{R}^+ \times \mathbb{R} \times \mathbb{R}^+$. So, we have $Zw'_1 \geq 0$ and this completes the proof of Lemma EC.1. ■

EC.2.2.2. Proof of Lemma EC.2 - Dropping subscript i , denote,

$$T := -\frac{1}{m^2}\mathbb{E}\left[(S^2\gamma(Y, S^2, m))^2\nu(Y, S^2, m)\right] - \frac{2}{m^2}\mathbb{E}\left[\frac{S^2}{\tau}w_1(Y, S^2, m)\gamma'(Y, S^2, m)\right].$$

We will show that $T \geq 0$.

Let $\gamma := \gamma(y, s^2, m)$, $\nu := \nu(y, s^2, m)$, $w_1 := w_1(y, s^2, m)$ and $\gamma' := \gamma'(y, s^2, m)$. First note that using standard integration by parts, we have

$$-\frac{1}{m^2}\mathbb{E}\left[(S^2\gamma)^2\nu\right] = \frac{2}{m^2}\mathbb{E}\left[(S^2)^2\gamma w_1\gamma'\right].$$

So, we can write

$$T = \frac{2}{m^2}\mathbb{E}\left\{S^2w_1\gamma'\left[S^2\gamma - \frac{1}{\tau}\right]\right\}.$$

Furthermore, from Definition 1,

$$\gamma' = \frac{2}{m-1} \gamma^2 s^2 w_2',$$

where $w_2' := \frac{\partial}{\partial y} w_2(y, s^2, m)$. So, we have

$$T = \frac{4}{m^2(m-1)} \mathbb{E} \left\{ (S^2 \gamma)^2 w_1 w_2' \left[S^2 \gamma - \frac{1}{\tau} \right] \right\}.$$

Now conditional on (Y, S^2) , note that from Jensen's inequality,

$$s^2 \gamma = \frac{1}{\mathbb{E}(\tau | y, s^2)} \leq \mathbb{E} \left(\frac{1}{\tau} \middle| y, s^2 \right).$$

Furthermore, since $w_1 w_2' \leq 0$ from Assumption EC.2, we have

$$T = \frac{4}{m^2(m-1)} \mathbb{E}_{Y, S^2} \left\{ (S^2 \gamma)^2 w_1 w_2' \left[S^2 \gamma - \mathbb{E} \left(\frac{1}{\tau} \middle| Y, S^2 \right) \right] \right\} \geq 0.$$

This completes the proof of Lemma EC.2. ■

EC.2.3. Proof of Theorem 1

Consider the sample criterion $\hat{\mathbb{M}}_{\lambda, n}(\tilde{\mathcal{W}})$ (Equation (16)) and population criterion $\mathbb{M}_{\lambda}(\tilde{\mathcal{W}})$ (Equation (15)). We have

$$\begin{aligned} \left| \hat{\mathbb{M}}_{\lambda, n}(\tilde{\mathcal{W}}) - \mathbb{M}_{\lambda}(\tilde{\mathcal{W}}) \right| &\leq \left| \frac{1}{n^2} \sum_{i=1}^n \sum_{j=1}^n \kappa_{\lambda}[\tilde{\mathbf{w}}(\mathbf{X}^i), \tilde{\mathbf{w}}(\mathbf{X}^j)](\mathbf{X}^i, \mathbf{X}^j) \mathbb{I}(i \neq j) - \mathbb{M}_{\lambda}(\tilde{\mathcal{W}}) \right| + \\ &\quad \left| \frac{1}{n^2} \sum_{i=1}^n \kappa_{\lambda}[\tilde{\mathbf{w}}(\mathbf{X}^i), \tilde{\mathbf{w}}(\mathbf{X}^i)](\mathbf{X}^i, \mathbf{X}^i) \right| \tag{EC.12} \\ &:= I_1 + I_2. \end{aligned}$$

Define $\bar{\mathbb{M}}_{\lambda, n}(\tilde{\mathcal{W}}) = [n(n-1)]^{-1} \sum_{i=1}^n \sum_{j=1}^n \kappa_{\lambda}[\tilde{\mathbf{w}}(\mathbf{X}^i), \tilde{\mathbf{w}}(\mathbf{X}^j)](\mathbf{X}^i, \mathbf{X}^j) \mathbb{I}(i \neq j)$. Then

$$I_1 \leq |\bar{\mathbb{M}}_{\lambda, n}(\tilde{\mathcal{W}}) - \mathbb{M}_{\lambda}(\tilde{\mathcal{W}})| + n^{-1} |\bar{\mathbb{M}}_{\lambda, n}(\tilde{\mathcal{W}})|.$$

By Assumption 1, $n^{-1} |\bar{\mathbb{M}}_{\lambda, n}(\tilde{\mathcal{W}})|$ is $O_p(n^{-1})$. Now, note that $\bar{\mathbb{M}}_{\lambda, n}(\tilde{\mathcal{W}})$ is an unbiased estimator of $\mathbb{M}_{\lambda}(\tilde{\mathcal{W}})$ and is a U-statistic with a symmetric kernel function $\kappa_{\lambda}[\tilde{\mathbf{w}}(\mathbf{X}^i), \tilde{\mathbf{w}}(\mathbf{X}^j)](\mathbf{X}^i, \mathbf{X}^j)$. From Assumption 1, $\kappa_{\lambda}[\tilde{\mathbf{w}}(\mathbf{X}^i), \tilde{\mathbf{w}}(\mathbf{X}^j)](\mathbf{X}^i, \mathbf{X}^j)$ has finite second moments. Moreover, from theorem 4.1 of Liu et al. (2016), $\bar{\mathbb{M}}_{\lambda, n}(\tilde{\mathcal{W}})$ is a non-degenerate U-statistic whenever $f \neq \tilde{f}$. Thus, from the CLT for U-statistics (Serfling (2009) section 5.5), $|\bar{\mathbb{M}}_{\lambda, n}(\tilde{\mathcal{W}}) - \mathbb{M}_{\lambda}(\tilde{\mathcal{W}})|$ is $O_p(n^{-1/2})$. Moreover, from Assumption 1, I_2 is $O_p(n^{-1})$. Theorem 1 is thus proved by combining these results. ■

EC.2.4. Proof of Theorem 2

We begin by introducing some notations and a lemma that will be useful for proving Theorem 2.

Recall from Equation (11) that $w_{k,i}$ denotes the $(i,k)^{th}$ element of the true score matrix \mathcal{W}_0 where $k = 1, 2$. Similarly, from Definition 2, $\hat{w}_{\lambda,n}^{(k)}(i)$ denotes the $(i,k)^{th}$ element of the estimated score matrix $\hat{\mathcal{W}}_n(\lambda)$. Let

$$\Delta_{\lambda,n}^{(p)}(i,k) = \mathbb{E} \left| \hat{w}_{\lambda,n}^{(k)}(i) - w_{k,i} \right|^p \text{ and } \bar{\Delta}_{\lambda,n}(i,k) = \mathbb{E} \left\{ \left| \hat{w}_{\lambda,n}^{(k)}(i) - w_{k,i} \right|^2 f(\mathbf{X}_i) \right\},$$

where the expectation is taken with respect to the joint distribution of $\mathbf{X}_1, \dots, \mathbf{X}_n$ and \mathbf{X}_i 's are i.i.d with density f . Denote $\mathbf{w}_0^{(k)}$ and $\hat{\mathbf{w}}_{\lambda,n}^{(k)}$ as, respectively, the k^{th} columns of \mathcal{W}_0 and $\hat{\mathcal{W}}_n(\lambda)$. Let

$$\hat{M}_{\lambda,n}^{(k)}(\mathbf{w}_0^{(k)}) = [\mathbf{w}_0^{(k)}]^T \mathbf{K}_\lambda \mathbf{w}_0^{(k)} + 2[\mathbf{w}_0^{(k)}]^T \nabla \mathbf{K}_\lambda^{(k)} + \frac{1}{n^2} \sum_{i=1}^n \sum_{j=1}^n \{[\nabla_{\mathbf{X}_i} \nabla_{\mathbf{X}_j} \mathcal{K}_\lambda(\mathbf{X}_i, \mathbf{X}_j)] \mathbf{e}_k\},$$

where $\nabla \mathbf{K}_\lambda^{(k)}$ is the k^{th} column of $\nabla \mathbf{K}_\lambda$ which was introduced in Section 3.1 and \mathbf{e}_k is the canonical basis vector with 1 at coordinate k . So from Equation (16), $\hat{\mathbb{M}}_{\lambda,n}(\mathcal{W}_0) = \hat{M}_{\lambda,n}^{(1)}(\mathbf{w}_0^{(1)}) + \hat{M}_{\lambda,n}^{(2)}(\mathbf{w}_0^{(2)})$.

Similarly, denote

$$M_\lambda^{(k)}(\hat{\mathbf{w}}_{\lambda,n}^{(k)}) = \mathbb{E}[D_\lambda(\mathbf{X}_i, \mathbf{X}_n)],$$

where $D_\lambda(\mathbf{X}_i, \mathbf{X}_n) = \{w_{k,n} - \hat{w}_{\lambda,n}^{(k)}(n)\} \mathcal{K}_\lambda(\mathbf{X}_i, \mathbf{X}_n) \{w_{k,i} - \hat{w}_{\lambda,n}^{(k)}(i)\}$. Notice that from the definition of the KSD in Equation (15), $\mathbb{M}_\lambda\{\hat{\mathcal{W}}_n(\lambda)\} = M_\lambda^{(1)}(\hat{\mathbf{w}}_{\lambda,n}^{(1)}) + M_\lambda^{(2)}(\hat{\mathbf{w}}_{\lambda,n}^{(2)})$. Finally, for any two sequences a_n and b_n , we will use $a_n \lesssim b_n$ to denote $a_n/b_n = O(1)$ as $n \rightarrow \infty$.

LEMMA EC.3. *Under Assumption 2, the following statements hold:*

(a) *We have,*

$$\Delta_{\lambda,n}^{(1)}(i,k) \lesssim \sqrt{\{\log n\}^2 \bar{\Delta}_{\lambda,n}(i,k)} + \frac{1}{n} \text{ as } n \rightarrow \infty.$$

(b) *For any $\lambda > 0$ and some constant $c > 0$,*

$$\bar{\Delta}_{\lambda,n}(i,k) \lesssim \lambda^{-2} M_\lambda^{(k)}(\hat{\mathbf{w}}_{\lambda,n+1}^{(k)}) + \lambda^2 (\log n)^4 + \lambda (\log n) \Delta_{\lambda,n}^{(2)}(i,k) + \frac{\lambda^{-2}}{n^c \log n}.$$

(c) *If $\mathcal{K}_\lambda(\cdot, \cdot)$ is the RBF kernel with bandwidth parameter $\lambda \in \Lambda$ and Λ is a compact subset of \mathbb{R}^+ bounded away from zero, then*

$$M_\lambda^{(k)}(\hat{\mathbf{w}}_{\lambda,n}^{(k)}) \leq \frac{\mathbb{E}\{w_{k,i}\}^2 - \mathbb{E}\{\hat{w}_{\lambda,n}^{(k)}(i)\}^2}{n-1}.$$

The proof of Lemma EC.3 is available in Section EC.2.5. We will now begin the proof of Theorem 2.

The main idea of the proof of Theorem 2 is to construct an appropriate bound for $\Delta_{\lambda,n}^{(2)}(i, k)$. As we will see, this bound depends on $\Delta_{\lambda,n}^{(1)}(i, k)$, $\bar{\Delta}_{\lambda,n}(i, k)$ and $M_{\lambda}^{(k)}(\hat{\mathbf{w}}_{\lambda,n+1}^{(k)})$. Once this bound is established, Markov's inequality will provide the desired convergence in probability in the statement of Theorem 2. Our approach for analyzing $\Delta_{\lambda,n}^{(2)}(i, k)$ relies on the technique developed in Luo et al. (2022) for proving their Theorem 1. We first note that since the density f is a convolution with a Gaussian density, there exists some constant $c > 0$ such that for all large $\|\mathbf{X}_i\|_2$, $|w_{k,i}|/\|\mathbf{X}_i\|_2 \leq c$. Furthermore, we will consider a constrained version of our optimization problem such that the solution satisfies $|\hat{w}_{\lambda,n}^{(k)}(i)| = O(\|\mathbf{X}_i\|_2)$. So, $|\hat{w}_{\lambda,n}^{(k)}(i) - w_{k,i}|$ is bounded above by $O(\|\mathbf{X}_i\|_2)$. Thus,

$$\mathbb{E} \left[\left\{ \hat{w}_{\lambda,n}^{(k)}(i) - w_{k,i} \right\}^2 \middle| \|\mathbf{X}_i\|_2 \leq 2c_0 \log n \right] \lesssim \Delta_{\lambda,n}^{(1)}(i, k) \log n, \quad (\text{EC.13})$$

where $c_0 > 0$ is a constant. Similarly,

$$\mathbb{E} \left[\left\{ \hat{w}_{\lambda,n}^{(k)}(i) - w_{k,i} \right\}^2 \middle| \|\mathbf{X}_i\|_2 > 2c_0 \log n \right] \lesssim \mathbb{E} \left[\|\mathbf{X}_i\|_2^2 \middle| \|\mathbf{X}_i\|_2 > 2c_0 \log n \right]. \quad (\text{EC.14})$$

Since \mathbf{X}_i is sub-exponential, the right hand side of Equation (EC.14) is bounded by $O(n^{-1})$.

Therefore, from equations (EC.13) and (EC.14),

$$\Delta_{\lambda,n}^{(2)}(i, k) \lesssim \Delta_{\lambda,n}^{(1)}(i, k) \log n + \frac{1}{n}. \quad (\text{EC.15})$$

But, from statements (a) and (b) of Lemma EC.3

$$\Delta_{\lambda,n}^{(1)}(i, k) \lesssim \sqrt{(\log n)^2 \left[\lambda^{-2} M_{\lambda}^{(k)}(\hat{\mathbf{w}}_{\lambda,n+1}^{(k)}) + \lambda^2 (\log n)^4 + \lambda \log n \Delta_{\lambda,n}^{(2)}(i, k) + \frac{\lambda^{-2}}{n^{c \log n}} \right]} + \frac{1}{n} \quad (\text{EC.16})$$

So, using Equation (EC.15) and $\lambda \asymp n^{-1/3}$ in Equation (EC.16) we get

$$\Delta_{\lambda,n}^{(1)}(i, k) \lesssim \sqrt{(\log n)^2 \left[n^{2/3} M_{\lambda}^{(k)}(\hat{\mathbf{w}}_{\lambda,n+1}^{(k)}) + n^{-2/3} (\log n)^4 + n^{-1/3} (\log n)^2 \Delta_{\lambda,n}^{(1)}(i, k) \right]}. \quad (\text{EC.17})$$

Now, from statement (c) of Lemma EC.3, we have

$$M_{\lambda}^{(k)}(\hat{\mathbf{w}}_{\lambda,n+1}^{(k)}) \leq \frac{\mathbb{E}\{w_{k,i}\}^2 - \mathbb{E}\{\hat{w}_{\lambda,n}^{(k)}(i)\}^2}{n}. \quad (\text{EC.18})$$

Note that in Equation (EC.18), the numerator $\mathbb{E}\{w_{k,i}\}^2 - \mathbb{E}\{\hat{w}_{\lambda,n}^{(k)}(i)\}^2 \lesssim \Delta_{\lambda,n+1}^{(1)}(i,k) \log n$ using the arguments developed for equations (EC.13) and (EC.14). Therefore,

$$M_\lambda^{(k)}(\hat{\mathbf{w}}_{\lambda,n+1}^{(k)}) \lesssim \frac{\log n}{n} \Delta_{\lambda,n+1}^{(1)}(i,k).$$

Substituting the above display in Equation (EC.17) we get

$$\Delta_{\lambda,n}^{(1)}(i,k) \lesssim \sqrt{(\log n)^2 \left[n^{-1/3} \log n \Delta_{\lambda,n+1}^{(1)}(i,k) + n^{-2/3} (\log n)^4 + n^{-1/3} (\log n)^2 \Delta_{\lambda,n}^{(1)}(i,k) \right]}.$$

Since $\Delta_{\lambda,n}^{(1)}(i,k)$ is bounded, the above display implies that that $\Delta_{\lambda,n}^{(1)}(i,k) \rightarrow 0$ as $n \rightarrow \infty$.

We will now establish the rate of convergence of $\Delta_{\lambda,n}^{(1)}(i,k)$. Let $A_n = \max\{\Delta_{\lambda,n}^{(1)}(i,k), n^{-1/3} (\log n)^4\}$. Using equations (EC.17) and (EC.18), we can derive the following inequality for all large n :

$$A_n \leq C (\log n)^{3/2} n^{-1/6} \sqrt{A_{n+1}}, \quad (\text{EC.19})$$

where the constant C is independent of n . Recursively applying Equation (EC.19) $s > 0$ times, we get

$$A_n \leq \left[C (\log n)^{3/2} n^{-1/6} \right]^{1+\dots+2^{-s}} A_{n+s+1}^{\frac{1}{2^s+1}}.$$

But for large n , $A_n < 1$ and this implies $A_n \leq [C (\log n)^{3/2} n^{-1/6}]^{1+\dots+2^{-s}}$. Therefore, letting $s \rightarrow \infty$, we have $A_n \leq C (\log n)^3 n^{-1/3}$ and so $\Delta_{\lambda,n}^{(1)}(i,k) \lesssim (\log n)^3 n^{-1/3}$. Moreover, using Equation (EC.15), $\Delta_{\lambda,n}^{(2)}(i,k) \lesssim (\log n)^4 n^{-1/3}$. The statement of Theorem 2 then follows using Markov's inequality and by noting that $\mathbb{E}\|\mathcal{W}_n(\lambda) - \mathcal{W}_0\|_F^2 = \sum_{i=1}^n \{\Delta_{\lambda,n}^{(2)}(i,1) + \Delta_{\lambda,n}^{(2)}(i,2)\}$. ■

EC.2.5. Proof of Lemma EC.3

Proof of statement (a) – From Cauchy-Schwartz inequality we get,

$$\mathbb{E} \left[|\hat{w}_{\lambda,n}^{(k)}(i) - w_{k,i}| \mathbb{I}\{\|\mathbf{X}_i\|_2 \leq 2c_0 \log n\} \right] \leq \left[c_1 (\log n)^2 \bar{\Delta}_{\lambda,n}(i,k) \right]^{0.5},$$

where c_1 is a constant depending only on c_0 . Also, using arguments similar to that of Equation (EC.14), we have

$$\mathbb{E} \left[|\hat{w}_{\lambda,n}^{(k)}(i) - w_{k,i}| \mathbb{I}\{\|\mathbf{X}_i\|_2 > 2c_0 \log n\} \right] = O(n^{-1}).$$

Thus, the proof of statement (a) then follows by combining the above two displays. \blacksquare

Proof of statement (b) – Suppose we have an i.i.d sample $(\mathbf{X}_1, \dots, \mathbf{X}_{n+1})$ of size $n + 1$ from f .

Now, the definition of $D_\lambda(\mathbf{X}_i, \mathbf{X}_{n+1})$ implies

$$D_\lambda(\mathbf{X}_i, \mathbf{X}_{n+1}) = \mathcal{K}_\lambda(\mathbf{X}_i, \mathbf{X}_{n+1}) \{w_{k,n+1} - \hat{w}_{\lambda,n+1}^{(k)}(n+1)\} \{w_{k,i} - \hat{w}_{\lambda,n+1}^{(k)}(i)\},$$

where $\hat{w}_{\lambda,n+1}^{(k)}(i)$ is the $(i, k)^{th}$ element of the estimated score matrix $\hat{\mathcal{W}}_{n+1}(\lambda)$ that is constructed based on the sample of size $n + 1$. Denote $\|\mathbf{X}_i - \mathbf{X}_n\| = \sqrt{(\mathbf{X}_i - \mathbf{X}_n)^T (\mathbf{X}_i - \mathbf{X}_n)}$ and fix $\epsilon = \lambda \log n$.

To prove statement (b), we will analyze the following four terms:

$$I_1 := \mathbb{E} \left[D_\lambda(\mathbf{X}_i, \mathbf{X}_{n+1}) \mathbb{I} \{ \|\mathbf{X}_i - \mathbf{X}_{n+1}\| < \epsilon \} \right] \quad (\text{EC.20})$$

$$I_2 := \mathbb{E} \left[\mathcal{K}_\lambda(\mathbf{X}_i, \mathbf{X}_{n+1}) \{w_{k,i} - \hat{w}_{\lambda,n+1}^{(k)}(i)\}^2 \mathbb{I} \{ \|\mathbf{X}_i - \mathbf{X}_{n+1}\| < \epsilon \} \right] \quad (\text{EC.21})$$

$$I_3 := \mathbb{E} \left[\mathcal{K}_\lambda(\mathbf{X}_i, \mathbf{X}_{n+1}) \{w_{k,i} - \hat{w}_{\lambda,n}^{(k)}(i)\}^2 \mathbb{I} \{ \|\mathbf{X}_i - \mathbf{X}_{n+1}\| < \epsilon \} \right] \quad (\text{EC.22})$$

$$I_4 := \mathbb{E} \int f(\mathbf{X}_i) \mathcal{K}_\lambda(\mathbf{X}_i, \mathbf{X}_{n+1}) \{w_{k,i} - \hat{w}_{\lambda,n}^{(k)}(i)\}^2 \mathbb{I} \{ \|\mathbf{X}_i - \mathbf{X}_{n+1}\| < \epsilon \} d\mathbf{X}_{n+1}. \quad (\text{EC.23})$$

For our analyses, we will assume that the eigen values of Ω are finite and bounded away from zero.

We begin by analyzing the term I_1 in Equation (EC.20). Now,

$$\mathbb{E}[D_\lambda(\mathbf{X}_i, \mathbf{X}_{n+1})] - I_1 = \mathbb{E}[D_\lambda(\mathbf{X}_i, \mathbf{X}_{n+1}) \mathbb{I} \{ \|\mathbf{X}_i - \mathbf{X}_{n+1}\| \geq \epsilon \}].$$

Since \mathcal{K}_λ in Equation (10) is the RBF kernel and $\epsilon = \lambda \log n$, we have $\mathcal{K}_\lambda(\mathbf{X}_i, \mathbf{X}_{n+1}) \mathbb{I} \{ \|\mathbf{X}_i - \mathbf{X}_{n+1}\| \geq \epsilon \} \lesssim n^{-c \log n}$ for some constant $c > 0$. Thus $\mathbb{E}[D_\lambda(\mathbf{X}_i, \mathbf{X}_{n+1}) \mathbb{I} \{ \|\mathbf{X}_i - \mathbf{X}_{n+1}\| \geq \epsilon \}] \lesssim \Delta_{\lambda,n+1}^{(2)}(i, k) n^{-c \log n}$. Since $\Delta_{\lambda,n+1}^{(2)}(i, k)$ is bounded, we have established,

$$I_1 \lesssim \mathbb{E}[D_\lambda(\mathbf{X}_i, \mathbf{X}_{n+1})] + O(n^{-c \log n}). \quad (\text{EC.24})$$

Consider term I_2 in Equation (EC.21). We will show that

$$I_2 \lesssim I_1 + O(\epsilon^2). \quad (\text{EC.25})$$

From Assumption 2, the score function $w_{k,i}$ is L_f - Lipschitz continuous. Moreover, for small ϵ and when $\|\mathbf{X}_i - \mathbf{X}_{n+1}\| < \epsilon$, $\hat{w}_{\lambda,n+1}^{(k)}(n+1)$ is $L_{n,\epsilon}$ - Lipschitz continuous since

$$\left| \hat{w}_{\lambda,n+1}^{(k)}(n+1) - \hat{w}_{\lambda,n+1}^{(k)}(i) \right| \leq L_{n,\epsilon} \epsilon,$$

where $\mathbb{E}L_{n,\epsilon}^2 < \infty$. Then $I_2 - I_1$ is bounded above by

$$\mathbb{E} \left[\epsilon (L_f + L_{n,\epsilon}) \mathcal{K}_\lambda(\mathbf{X}_i, \mathbf{X}_{n+1}) \left| w_{k,i} - \hat{w}_{\lambda,n+1}^{(k)}(i) \right| \mathbb{I}\{\|\mathbf{X}_i - \mathbf{X}_{n+1}\| < \epsilon\} \right].$$

Using Cauchy-Schwartz inequality, the square of the above display has the following upper bound:

$$\epsilon^2 \mathbb{E} \left[(L_f + L_{n,\epsilon})^2 \mathbb{I}\{\|\mathbf{X}_i - \mathbf{X}_{n+1}\| < \epsilon\} \right] \mathbb{E} \left[\mathcal{K}_\lambda^2(\mathbf{X}_i, \mathbf{X}_{n+1}) \left| w_{k,i} - \hat{w}_{\lambda,n+1}^{(k)}(i) \right|^2 \mathbb{I}\{\|\mathbf{X}_i - \mathbf{X}_{n+1}\| < \epsilon\} \right].$$

But, $\mathbb{E}[(L_f + L_{n,\epsilon})^2 \mathbb{I}\{\|\mathbf{X}_i - \mathbf{X}_{n+1}\| < \epsilon\}]$ is bounded above by

$$C_f \frac{\pi}{\Gamma(2)} \epsilon^2 \mathbb{E} \left[(L_f + L_{n,\epsilon})^2 \right],$$

where $\Gamma(\cdot)$ is the Gamma function. Finally, noticing that $\mathcal{K}_\lambda^2(\mathbf{X}_i, \mathbf{X}_{n+1}) \leq \mathcal{K}_\lambda(\mathbf{X}_i, \mathbf{X}_{n+1})$, we have

$$I_2 \lesssim I_1 + \epsilon \sqrt{\epsilon^2 \Delta_{\lambda,n+1}^{(2)}(i, k)},$$

which establishes Equation (EC.25).

Similarly the difference between I_3 (Equation (EC.22)) and I_2 can be bounded by $4\epsilon^2 \mathbb{E}[L_{n,\epsilon}^2 \mathcal{K}_\lambda(\mathbf{X}_i, \mathbf{X}_{n+1}) \mathbb{I}\{\|\mathbf{X}_i - \mathbf{X}_{n+1}\| < \epsilon\}]$, which implies that

$$I_3 \lesssim I_2 + O(\epsilon^4). \tag{EC.26}$$

Next, we will show that for the term I_4 in Equation (EC.23), the following holds:

$$\lambda^2 \bar{\Delta}_{\lambda,n}(i, k) \lesssim I_4 \lesssim I_3 + \lambda^3 \Delta_{\lambda,n}^{(2)}(i, k) \log n. \tag{EC.27}$$

This inequality along with equations (EC.26), (EC.25) and (EC.24) will establish statement (b).

First, note that

$$I_4 = \mathbb{E} \left[\left\{ w_{k,i} - \hat{w}_{\lambda,n}^{(k)}(i) \right\}^2 f(\mathbf{X}_i) \right] \int \mathcal{K}_\lambda(\mathbf{X}_i, \mathbf{X}_{n+1}) \mathbb{I}\{\|\mathbf{X}_i - \mathbf{X}_{n+1}\| < \epsilon\} d\mathbf{X}_{n+1},$$

since $\hat{w}_{\lambda,n}^{(k)}(i)$ is independent of \mathbf{X}_{n+1} for $i < n + 1$. Now, with $\epsilon = \lambda \log n$, we have

$$\int \mathcal{K}_\lambda(\mathbf{X}_i, \mathbf{X}_{n+1}) \mathbb{I}\{\|\mathbf{X}_i - \mathbf{X}_{n+1}\| < \epsilon\} d\mathbf{X}_{n+1} \geq c\lambda^2, \tag{EC.28}$$

for some constant $c > 0$ when n is large. So, Equation (EC.28) implies $\lambda^2 \bar{\Delta}_{\lambda,n}(i, k) \lesssim I_4$ and establishes the lower bound in Equation (EC.27). To establish the upper bound we will assume that the density f is L_f - Lipschitz continuous (Assumption 2). We have,

$$\begin{aligned} f(\mathbf{X}_i) \int \mathcal{K}_\lambda(\mathbf{X}_i, \mathbf{X}_{n+1}) \mathbb{I}\{\|\mathbf{X}_i - \mathbf{X}_{n+1}\| < \epsilon\} d\mathbf{X}_{n+1} &\lesssim \\ \int \mathcal{K}_\lambda(\mathbf{X}_i, \mathbf{X}_{n+1}) \mathbb{I}\{\|\mathbf{X}_i - \mathbf{X}_{n+1}\| < \epsilon\} f(\mathbf{X}_{n+1}) d\mathbf{X}_{n+1} &+ L_f \epsilon \int \mathcal{K}_\lambda(\mathbf{X}_i, \mathbf{X}_{n+1}) \mathbb{I}\{\|\mathbf{X}_i - \mathbf{X}_{n+1}\| < \epsilon\} d\mathbf{X}_{n+1}. \end{aligned}$$

Then, pre-multiplying the above display with $\{w_{k,i} - \hat{w}_{\lambda,n}^{(k)}(i)\}^2$ and taking expectations, we get

$$I_4 - I_3 \lesssim L_f \epsilon \Delta_{\lambda,n}^{(2)}(i, k) \int \mathcal{K}_\lambda(\mathbf{X}_i, \mathbf{X}_{n+1}) \mathbb{I}\{\|\mathbf{X}_i - \mathbf{X}_{n+1}\| < \epsilon\} d\mathbf{X}_{n+1} \leq L_f \epsilon \Delta_{\lambda,n}^{(2)}(i, k) C \lambda^2,$$

for some constant C and n large. This implies, $I_4 \lesssim I_3 + \lambda^3 \Delta_{\lambda,n}^{(2)}(i, k)$ and establishes the upper bound in Equation (EC.27). \blacksquare

Proof of statement (c) – We first note that $\hat{M}_{\lambda,n}^{(k)}(\hat{\mathbf{w}}_{\lambda,n}^{(k)})$ and $\hat{M}_{\lambda,n}^{(k)}(\mathbf{w}_0^{(k)})$ are, respectively V-statistics of $M_\lambda^{(k)}(\hat{\mathbf{w}}_{\lambda,n}^{(k)})$ and $M_\lambda^{(k)}(\mathbf{w}_0^{(k)})$. Moreover, we have

$$\hat{M}_{\lambda,n}^{(k)}(\hat{\mathbf{w}}_{\lambda,n}^{(k)}) \leq \hat{M}_{\lambda,n}^{(k)}(\mathbf{w}_0^{(k)}).$$

So, taking expectations on both sides of the above display we get

$$\frac{n^2 - n}{n^2} M_\lambda^{(k)}(\hat{\mathbf{w}}_{\lambda,n}^{(k)}) + \frac{1}{n} \mathbb{E}\{\hat{w}_{\lambda,n}^{(k)}(i)\}^2 \leq \frac{n^2 - n}{n^2} M_\lambda^{(k)}(\mathbf{w}_0^{(k)}) + \frac{1}{n} \mathbb{E}\{w_{k,i}\}^2.$$

But $M_\lambda^{(k)}(\mathbf{w}_0^{(k)}) = 0$ and so from the above display, we have

$$M_\lambda^{(k)}(\hat{\mathbf{w}}_{\lambda,n}^{(k)}) \leq \frac{\mathbb{E}\{w_{k,i}\}^2 - \mathbb{E}\{\hat{w}_{\lambda,n}^{(k)}(i)\}^2}{n - 1},$$

which completes the proof of statement (c). \blacksquare

EC.2.6. Proof of Theorem 3

We first state two lemmata that are needed for proving Theorem 3. Denote c_0, c_1, \dots some generic positive constants which may vary in different statements.

LEMMA EC.4. *If Assumption 3 holds, then with probability tending to 1, $C_1/\log n \leq \tau \leq C_2 \log n$ and $|\mu| \leq C_3 \log n$ for some positive constants C_1, C_2 and C_3 .*

LEMMA EC.5. Consider Model (1). Suppose Assumption 3 holds. Then with probability tending to 1, $\mathbb{E}(\tau_i | y_i, s_i^2) \geq c_0 / \log n$.

Lemmata EC.4 and EC.5 are proved in Sections EC.2.6.1 and EC.2.6.2, respectively. We now prove Theorem 3.

To establish the first part of Theorem 3, note that

$$\begin{aligned} \frac{1}{n} \|\boldsymbol{\delta}_n^{\text{ds}}(\lambda) - \boldsymbol{\delta}_{(1)}^\pi\|_2^2 &= \frac{1}{nm^2} \sum_{i=1}^n \left| \frac{w_{1,i}}{\hat{\tau}_i^\pi} - \frac{\hat{w}_{\lambda,n}^{(1)}(i)}{\tau_{i,n}^{\text{ds}}(\lambda)} \right|^2 \\ &\leq \frac{2}{nm^2} \sum_{i=1}^n \frac{1}{[\tau_{i,n}^{\text{ds}}(\lambda)]^2} \left| w_{1,i} - \hat{w}_{\lambda,n}^{(1)}(i) \right|^2 + \frac{2}{nm^2} \sum_{i=1}^n \left| w_{1,i} \right|^2 \left| \frac{1}{\tau_{i,n}^{\text{ds}}(\lambda)} - \frac{1}{\hat{\tau}_i^\pi} \right|^2 \\ &:= T_1 + T_2. \end{aligned}$$

Consider the first term T_1 . From the discussion in Section 3.2, there is a positive constant c_0 such that $\tau_{i,n}^{\text{ds}}(\lambda) > c_0 > 0$ for all $i = 1, \dots, n$. It follows that for some constant $c_1 > 0$ depending on the fixed m ,

$$T_1 \leq \frac{c_1}{n} \left\| \boldsymbol{w}_0^{(1)} - \hat{\boldsymbol{w}}_{\lambda,n}^{(1)} \right\|_2^2 \leq \frac{c_1}{n} \left\| \mathcal{W}_0 - \hat{\mathcal{W}}_n(\lambda) \right\|_F^2, \quad (\text{EC.29})$$

where $\hat{\boldsymbol{w}}_{\lambda,n}^{(1)}$ and $\boldsymbol{w}_0^{(1)}$ are, respectively, the first column of $\hat{\mathcal{W}}_n(\lambda)$ and \mathcal{W}_0 . From Theorem 2 the last term on the right hand side of the inequality in equation (EC.29) is $O_p\{n^{-1/3}(\log n)^4\}$.

Next consider the second term T_2 . We have

$$T_2 \leq \frac{c_2}{n} \sum_{i=1}^n \left| \frac{w_{1,i}}{\hat{\tau}_i^\pi} \right|^2 \left| w_{2,i} - \hat{w}_{\lambda,n}^{(2)}(i) \right|^2. \quad (\text{EC.30})$$

We will use Lemma EC.5 to bound the terms $|w_{1,i}/\hat{\tau}_i^\pi|$ in equation (EC.30). First note that since \boldsymbol{X}_i is sub-exponential, $|w_{1,i}| \leq c_1 \log n$ with high probability from the discussion in the proof of Theorem 2. Next, Model (1), Assumption 3 and Lemma EC.4 imply that with probability tending to 1, $|Y_i| \leq c_2 \log n$ and $(m-1)n^{-1} \leq (m-1)S_i^2 \tau_i \leq (m-1) + 2\sqrt{(m-1)\log n} + 2\log n$ [cf. Lemma 1 of Laurent and Massart (2000)]. So conditional on these events, we have, from Lemma EC.5, $|w_{1,i}/\hat{\tau}_i^\pi| \leq c_3(\log n)^2$. Thus,

$$T_2 \leq \frac{c_4(\log n)^4}{n} \sum_{i=1}^n \left| w_{2,i} - \hat{w}_{\lambda,n}^{(2)}(i) \right|^2,$$

which is $O_p\{n^{-1/3}(\log n)^8\}$ from Theorem 2. Thus $n^{-1}\|\boldsymbol{\delta}_n^{\text{ds}}(\lambda) - \boldsymbol{\delta}_{(1)}^\pi\|_2^2$ is $O_p\{n^{-1/3}(\log n)^8\}$.

Now we will prove the second part of Theorem 3. Observe that $|l_n^{(1)}(\boldsymbol{\mu}, \boldsymbol{\delta}_{(1)}^\pi; \boldsymbol{\tau}) - l_n^{(1)}\{\boldsymbol{\mu}, \boldsymbol{\delta}_n^{\text{ds}}(\lambda); \boldsymbol{\tau}\}|$ equals

$$\left| \sqrt{l_n^{(1)}(\boldsymbol{\mu}, \boldsymbol{\delta}_{(1)}^\pi; \boldsymbol{\tau})} - \sqrt{l_n^{(1)}\{\boldsymbol{\mu}, \boldsymbol{\delta}_n^{\text{ds}}(\lambda); \boldsymbol{\tau}\}} \right| \left| \sqrt{l_n^{(1)}(\boldsymbol{\mu}, \boldsymbol{\delta}_{(1)}^\pi; \boldsymbol{\tau})} + \sqrt{l_n^{(1)}\{\boldsymbol{\mu}, \boldsymbol{\delta}_n^{\text{ds}}(\lambda); \boldsymbol{\tau}\}} \right|$$

and Triangle inequality implies

$$\begin{aligned} \left| \sqrt{l_n^{(1)}(\boldsymbol{\mu}, \boldsymbol{\delta}_{(1)}^\pi; \boldsymbol{\tau})} - \sqrt{l_n^{(1)}\{\boldsymbol{\mu}, \boldsymbol{\delta}_n^{\text{ds}}(\lambda); \boldsymbol{\tau}\}} \right| &\leq \left| \sqrt{\frac{1}{n} \sum_{i=1}^n \tau_i \{\delta_{i,(1)}^\pi - \delta_{i,n}^{\text{ds}}(\lambda)\}^2} \right| \\ &\leq c_0 \sqrt{\frac{\log n}{n}} \|\boldsymbol{\delta}_n^{\text{ds}}(\lambda) - \boldsymbol{\delta}_{(1)}^\pi\|_2, \end{aligned} \quad (\text{EC.31})$$

where the last inequality in Equation (EC.31) follows from Lemma EC.4. Thus, from the first part of Theorem 3 and Lemma EC.4, the quantity on the right hand side of the inequality in equation (EC.31) is $O_p\{(\log n)^{9/2}n^{-1/6}\}$. Thus, it follows from equation (EC.31) that

$$\sqrt{l_n^{(1)}\{\boldsymbol{\mu}, \boldsymbol{\delta}_n^{\text{ds}}(\lambda); \boldsymbol{\tau}\}} \leq \sqrt{l_n^{(1)}(\boldsymbol{\mu}, \boldsymbol{\delta}_{(1)}^\pi; \boldsymbol{\tau})} + O_p\{(\log n)^{9/2}n^{-1/6}\}, \text{ and}$$

$$|l_n^{(1)}(\boldsymbol{\mu}, \boldsymbol{\delta}_{(1)}^\pi; \boldsymbol{\tau}) - l_n^{(1)}\{\boldsymbol{\mu}, \boldsymbol{\delta}_n^{\text{ds}}(\lambda); \boldsymbol{\tau}\}| \leq 4\sqrt{l_n^{(1)}(\boldsymbol{\mu}, \boldsymbol{\delta}_{(1)}^\pi; \boldsymbol{\tau})} \left| \sqrt{l_n^{(1)}(\boldsymbol{\mu}, \boldsymbol{\delta}_{(1)}^\pi; \boldsymbol{\tau})} - \sqrt{l_n^{(1)}\{\boldsymbol{\mu}, \boldsymbol{\delta}_n^{\text{ds}}(\lambda); \boldsymbol{\tau}\}} \right| \{1 + o_p(1)\}.$$

Now $\boldsymbol{\delta}_{(1)}^\pi$ is the Bayes estimator of $\boldsymbol{\mu}$ under the weighted squared error loss and so its risk $\mathbb{E}l_n^{(1)}(\boldsymbol{\mu}, \boldsymbol{\delta}_{(1)}^\pi; \boldsymbol{\tau}) < \infty$. This implies $l_n^{(1)}(\boldsymbol{\mu}, \boldsymbol{\delta}_{(1)}^\pi; \boldsymbol{\tau})$ is $O_p(1)$. Thus, from equation (EC.31), the first part of Theorem 3 and the display above, we have the desired result. \blacksquare

EC.2.6.1. Proof of Lemma EC.4 The proof of Lemma EC.4 follows directly from Assumption 3 and Markov's inequality. For example, fix a $\nu > 0$ and note that, for $r = \epsilon_2^{-\nu} > 1$,

$$\mathbb{P}\left(\tau \leq \frac{\epsilon_2^{1+\nu}}{\log n}\right) \leq \frac{\mathbb{E}_H\left\{\exp(\epsilon_2/\tau)\right\}}{n^r}.$$

\blacksquare

EC.2.6.2. Proof of Lemma EC.5 Recall that $f(y_i, s_i^2) = \int_{\mathbb{R}^+} \int_{\mathbb{R}} f_1(y_i|\mu, \tau) f_2(s_i^2|\tau) g(\mu|\tau) h(\tau) d\mu d\tau$, where $g(\cdot|\tau)$ and $h(\cdot)$ are, respectively, the density functions associated with the distribution functions $G_\mu(\cdot|\tau)$ and $H_\tau(\cdot)$ in Model (1), f_1 is the density of a Gaussian random variable with mean μ and variance $1/(m\tau)$ and f_2 is the density of S^2 where $(m-1)S^2\tau \sim \mathcal{X}_{m-1}^2$.

We will first analyze the behavior of $f(y_i, s_i^2)$. Gaussian concentration implies that with high probability $\{(Y_i - \mu_i)^2 m\tau_i\} \leq 2\log n$ and so, conditional on this event,

$$f_1(y_i|\mu, \tau) \geq c_0 \frac{\sqrt{\tau}}{n}. \quad (\text{EC.32})$$

Moreover, using the Chi-square concentration in Lemma 1 of Laurent and Massart (2000), $(m-1)S_i^2\tau_i \leq (m-1) + 2\sqrt{(m-1)\log n} + 2\log n$ and $S_i^2\tau_i \geq n^{-1}$ with high probability. It follows that

$$f_2(s_i^2|\tau) \geq c_1 \tau \frac{a_n}{n^{(m-1)/2}}, \quad (\text{EC.33})$$

conditional on this event, where $a_n = \exp\{-\sqrt{(m-1)\log n}\}$. Using equations (EC.32) and (EC.33), we have

$$f(y_i, s_i^2) \geq c_2 \frac{a_n}{n^{(m+1)/2}} \int_{\mathbb{R}^+} \tau^{3/2} h(\tau) d\tau. \quad (\text{EC.34})$$

Now, use Assumption 3 and Lemma EC.4 on the quantity $\int_{\mathbb{R}^+} \tau^{3/2} h(\tau) d\tau$ in equation (EC.34) to conclude that

$$\int_{\mathbb{R}^+} \tau^{3/2} h(\tau) d\tau \geq \frac{c_3}{(\log n)^{3/2}} \mathbb{P}\left(\tau \geq C_2/\log n\right). \quad (\text{EC.35})$$

So, with equations (EC.35), (EC.34) and Lemma EC.4, we have with high probability,

$$f(y_i, s_i^2) \geq c_4 \frac{a_n}{n^{(m+1)/2} (\log n)^{3/2}}. \quad (\text{EC.36})$$

Now we proceed to prove the statement of Lemma EC.5. Fix $\nu > 0$ such that $\epsilon^{-\nu} > 2m$ in Lemma EC.4 and let $C_1 = \epsilon^{1+\nu}$. First note that from Assumption 3, $\mathbb{P}(\tau \leq C_1/\log n) \leq c_0/n^{2m}$. Now, Markov's inequality implies,

$$\mathbb{E}(\tau|y_i, s_i^2) \geq \frac{C_1}{\log n} \left\{ 1 - \mathbb{P}\left[\tau \leq C_1(\log n)^{-1} \mid y_i, s_i^2\right] \right\}.$$

Moreover,

$$\mathbb{P}\left(\tau \leq \frac{C_1}{\log n} \mid y_i, s_i^2\right) = \{f(y_i, s_i^2)\}^{-1} \int_0^{C_2/\log n} h(\tau) \left\{ \int_{\mathbb{R}} g(\mu|\tau) f_1(y_i|\mu, \tau) f_2(s_i^2|\tau) d\mu \right\} d\tau.$$

Now $f_1(y_i|\mu, \tau) \leq c_0\sqrt{\tau}$ and $f_2(s_i^2|\tau) \leq c_1\tau$ where $c_0, c_1 > 0$ are constants. So for some positive constant c_2 ,

$$\begin{aligned} \mathbb{P}(\tau \leq C_2/\log n \mid y_i, s_i^2) &\leq \frac{c_2}{f(y_i, s_i^2)} \int_0^{C_2/\log n} \tau^{3/2} h(\tau) \left\{ \int_{\mathbb{R}} g(\mu|\tau) d\mu \right\} d\tau \\ &\leq \frac{c_3}{f(y_i, s_i^2)(\log n)^{3/2}} \int_0^{C_2/\log n} h(\tau) d\tau = \frac{c_3}{f(y_i, s_i^2)(\log n)^{3/2}} \mathbb{P}(\tau \leq C_2/\log n). \end{aligned}$$

Thus, from the above display, Assumption 3 and Lemma EC.4,

$$\mathbb{E}(\tau \mid y_i, s_i^2) \geq \frac{C_2}{\log n} \left\{ 1 - c_3 \frac{n^{-2m}}{f(y_i, s_i^2)(\log n)^{3/2}} \right\}.$$

Finally, equation (EC.36) and the above display prove the statement of Lemma EC.5. ■

EC.2.7. Proof of Equations (7) and (8)

The proof follows by first recalling that under the hierarchical model of Equation (1),

$$f_{m_i}(y_i, s_i^2 \mid \mu_i, \tau_i) \propto \exp\left\{-\frac{\tau_i}{2} [m_i y_i^2 + (m_i - 1) s_i^2] + m_i \tau_i \mu_i y_i - \frac{m_i}{2} \tau_i \mu_i^2 + \frac{m_i - 3}{2} \log s_i^2\right\}.$$

Therefore from Bayes theorem,

$$f_{m_i}(\mu_i, \tau_i \mid y_i, s_i^2) = \frac{f_{m_i}(y_i, s_i^2 \mid \mu_i, \tau_i)}{f_{m_i}(y_i, s_i^2)} g(\mu_i|\tau_i) h(\tau_i) \propto \exp\left\{\boldsymbol{\eta}_i^T \mathbf{T}(\mu_i, \tau_i) - A(\boldsymbol{\eta}_i)\right\} g(\mu_i|\tau_i) h(\tau_i).$$

Here $\boldsymbol{\eta}_i = (m_i y_i, -m_i y_i^2 - (m_i - 1) s_i^2) := (\eta_{1i}, \eta_{2i})$, $\mathbf{T}(\mu_i, \tau_i) = (\tau_i \mu_i, \tau_i/2)$ and

$$\begin{aligned} A(\boldsymbol{\eta}_i) &= -0.5(m_i - 3) \log \gamma(\eta_{1i}, \eta_{2i}) + \log f_{m_i} \{m_i^{-1} \eta_{1i}, \gamma(\eta_{1i}, \eta_{2i})\}, \\ \gamma(\eta_{1i}, \eta_{2i}) &= \frac{-\eta_{2i} - m_i^{-1} \eta_{1i}^2}{m_i - 1}, \end{aligned}$$

with $f_m(y, s^2) = \int \int f_m(y, s^2 \mid \mu, \tau) g_\mu(\mu|\tau) h_\tau(\tau) d\mu d\tau$ being the marginal density function of (Y, S^2) .

So the posterior distribution of (μ_i, τ_i) belongs to a 2-parameter exponential family and using the properties of exponential family distributions we have, dropping subscript i ,

$$\hat{\tau}^\pi := \hat{\tau}^\pi(y, s^2, m) = \mathbb{E}(\tau \mid y, s^2, m) = 2 \frac{\partial A(\boldsymbol{\eta})}{\partial \eta_2} = \frac{m - 3}{(m - 1) s^2} - \frac{2}{m - 1} w_2(y, s^2; m).$$

Furthermore, with $\zeta = \tau\mu$,

$$\begin{aligned}\hat{\zeta}^\pi &:= \hat{\zeta}^\pi(y, s^2, m) = \mathbb{E}(\zeta|y, s^2, m) = \frac{\partial A(\boldsymbol{\eta})}{\partial \eta_1} = \frac{(m-3)y}{(m-1)s^2} + m^{-1}w_1(y, s^2; m) - \frac{2}{m-1}yw_2(y, s^2; m) \\ &= y\mathbb{E}(\tau|y, s^2, m) + m^{-1}w_1(y, s^2; m).\end{aligned}$$

■

EC.3. Additional Numerical Experiments

EC.3.1. Compound estimation of Normal means under squared error loss

We focus on the hierarchical Model of Equation (1) and compare six approaches for estimating $\boldsymbol{\mu}$ under the squared error loss when the variances $\sigma_i = 1/\tau_i$ are assumed to be unknown. These approaches can be categorized into three types: the first consists of the NEST method (NEST Orc. λ), which estimates λ by minimizing the true loss, the proposed data-driven NEST method and Tweedie's formula (TF) that uses sample variances. For both NEST and TF, λ is chosen using the modified cross-validation approach described in Section 3.2 with $\vartheta_n(\lambda; \mathcal{U}, \mathcal{V})$ defined as

$$\vartheta_n(\lambda; \mathcal{U}, \mathcal{V}) = \frac{1}{n} \sum_{i=1}^n \{\bar{V}_i - \delta_i^{\text{ds}}(\bar{U}_i; \mathcal{U}, \lambda)\}^2,$$

to reflect the squared error loss. The second are linear shrinkage methods: the group linear estimator (Grp Linear) of [Weinstein et al. \(2018\)](#) and the semi-parametric linear shrinkage rule (Jing.SM) from [Jing et al. \(2016b\)](#). Finally, the third type is the g-modelling approach of [Gu and Koenker \(2017a,b\)](#). For Grp. Linear we use code provided by [Weinstein et al. \(2018\)](#) while for Jing.SM we write our own routine in R. We continue to rely on the function WGLVmix in the R package REBayes ([Koenker and Gu 2017](#)) for NPMLE. We note that amongst the six methods considered here, NEST and NEST Orc. λ are designed to estimate the means under the weighted squared error loss while the remaining four approaches target the squared error loss. Furthermore, with the exception of Grp. Linear, all other methods considered here estimate $\boldsymbol{\mu}$ when the variances are unknown. Grp. Linear, on the other hand, assumes full knowledge of the unknown variances for shrinkage estimation of the means and here we use sample variances for its implementation.

The aforementioned six approaches are evaluated on five different simulation settings, with the goal of assessing the relative performance of the competing estimators as the heterogeneity in the variances σ_i^2 is varied while keeping the sample sizes m_i fixed at m . The five simulation settings can be categorized into three types: a setting where mean and variances are independent; three settings where mean and variance are correlated; and a setting that represents departure from the Normal data-generating model. For each setting we set $n = 1,000$ and compute the average squared error risk for each competing estimator of $\boldsymbol{\mu}$ across 50 Monte Carlo repetitions. Figures EC.3 to EC.7 plot the relative risk which is the ratio of the average squared error risk for any competing estimator to that of oracle Bayes estimator $\boldsymbol{\delta}_{(0)}^\pi$ of $\boldsymbol{\mu}$ (Equation (5)) so that a ratio bigger than 1 represents a poorer risk performance of the competing estimator relative to the Bayes oracle.

The first setting, Figure EC.3, corresponds to the independent case. Here, for each $i = 1, \dots, n$, $\mu_i \stackrel{i.i.d}{\sim} 0.7 N(0, 0.1) + 0.15 N(1, 3) + 0.15 N(-1, 3)$ and $\sigma_i^2 \stackrel{i.i.d}{\sim} U(0.5, u)$ where we let u vary across five levels, $\{1, 2, 3, 4, 5\}$. The three plots in Figure EC.3 show the relative risks as u varies for $m = 10, 15$ and 20 (left to right). We see that for $m = 10$, the competing methods split into two levels of performance. The group with the lowest relative risks consists of NPMLE, TF and NEST while the two linear shrinkage methods exhibit substantially higher relative risks. Moreover, we also see that as heterogeneity increases with increasing u , the gap between the two groups' relative risks increases, indicating that NPMLE, TF and the proposed NEST method are particularly useful for compound estimation of normal means when the variances are unknown and heterogeneous, and the sample size for estimating those variances are themselves small. As m increases, the performance of the two linear shrinkage methods and TF improve which is expected as there are now more replicates per unit of study to construct a relatively reliable estimate of the unknown variances. However, the performance of NEST improves too and particularly at $m = 20$ (Figure EC.3 right), NPMLE exhibits a slightly higher relative risk than NEST and TF.

The second setting, Figure EC.4, corresponds to the correlated case. The precisions $\tau_i = 1/\sigma_i^2$ are generated independently from a gamma mixture, with an even chance of drawing




Figure EC.3 Comparison of relative risks when (μ_i, σ_i^2) are independent. Here $\mu_i \stackrel{i.i.d}{\sim} 0.7 N(0, .1) + 0.15 N(1, 3) + 0.15 N(-1, 3)$ and $\sigma_i^2 \stackrel{i.i.d}{\sim} U(0.5, u)$. Plots show $m = 10, 15, 20$ left to right.

$\Gamma(20, \text{rate} = 20)$ or $\Gamma(20, \text{rate} = u)$ and given τ_i , the means μ_i are independently $0.5 N(0.5/\tau_i, 0.5^2) + 0.5 N(-0.5/\tau_i, 0.5^2)$. In this setting, the magnitude of the variances increase with u and the means grow with the variances. We note from Figure EC.4 that Grp. Linear and Jing.SM exhibit improved performance particularly for small values of u . As u increases, TF, NPMLE and NEST perform well although their relative risk profiles are substantially away from 1 as u increases. For TF and NEST this behavior is expected given the statements of Propositions EC.1 and EC.2. The improved performance of Jing.SM in this setting is potentially related to the observation that when u is small, the rate mixture of Gamma distributions on τ_i can be well approximated by a single Gamma distribution and that coincides with the parametric prior that Jing et al. (2016b) use on the precision to derive their empirical Bayes estimator for the means.

In the third setting, Figure EC.5, (μ_i, τ_i) continue to be correlated and have a conjugate prior distribution under Model (1). The precisions τ_i are drawn from $\Gamma(20, \text{rate} = u)$ and conditional on τ_i , μ_i are independently $N(0, 0.5/\tau_i)$. Under this data generating scheme, the posterior mean of μ_i is $my_i/(m+2)$ which is independent of u . This is the reason that the relative risks of the competing estimators in Figure EC.5 do not vary with the heterogeneity in the variances. Compared to the first two settings, we see that the linear shrinkage estimators have a relatively better performance and Jing.SM dominates all other shrinkage estimators. This is expected because in this setting the posterior mean of μ_i is indeed a linear function of the sample mean y_i . For $m = 10$ and 15, we




Figure EC.4 Comparison of relative risks for correlated (μ_i, τ_i) . Here $\tau_i \stackrel{i.i.d.}{\sim} 0.5\Gamma(20, \text{rate} = 20) + 0.5\Gamma(20, \text{rate} = u)$ and $\mu_i | \tau_i \stackrel{ind.}{\sim} 0.5 N(0.5/\tau_i, 0.5^2) + 0.5 N(-0.5/\tau_i, 0.5^2)$. Plots show $m = 10, 15, 20$ left to right.

notice that the relative risk of NEST is marginally better than the competing estimators while at $m = 20$ Grp Linear and NEST have similar risk performance.




Figure EC.5 Comparison of relative risks when (μ_i, τ_i) have conjugate priors. Here $\mu_i | \tau_i \stackrel{ind.}{\sim} N(0, 0.5/\tau_i)$ and $\tau_i \stackrel{i.i.d.}{\sim} \Gamma(20, \text{rate} = u)$. Plots show $m = 10, 15, 20$ left to right.

Figure EC.6 presents the fourth setting where the precisions τ_i are drawn from the gamma mixture of Setting 2 and $\mu_i | \tau_i \stackrel{ind.}{\sim} 0.7 N(0, 0.01) + 0.15 N(0.5/\tau_i, 1) + 0.15 N(-0.5/\tau_i, 1)$. We see a similar pattern to that in Figure EC.3 at $m = 10$. For $m = 15$ and 20, we notice that the relative risks of the linear shrinkage methods are now higher than their levels at $m = 10$. This is not surprising for in this setting, while the risk performance of all methods have improved with larger sample sizes, NPMLE, TF and NEST exhibit a bigger improvement in risk than those of Grp Linear and Jing.SM.




Figure EC.6 Comparison of relative risks when μ is sparse. Here $\tau_i \stackrel{i.i.d}{\sim} 0.5 \Gamma(20, \text{rate} = 20) + 0.5 \Gamma(20, \text{rate} = u)$ and $\mu_i | \tau_i \stackrel{ind.}{\sim} 0.7 N(0, 0.01) + 0.15 N(0.5/\tau_i, 1) + 0.15 N(-0.5/\tau_i, 1)$. Plots show $m = 10, 15, 20$ left to right.


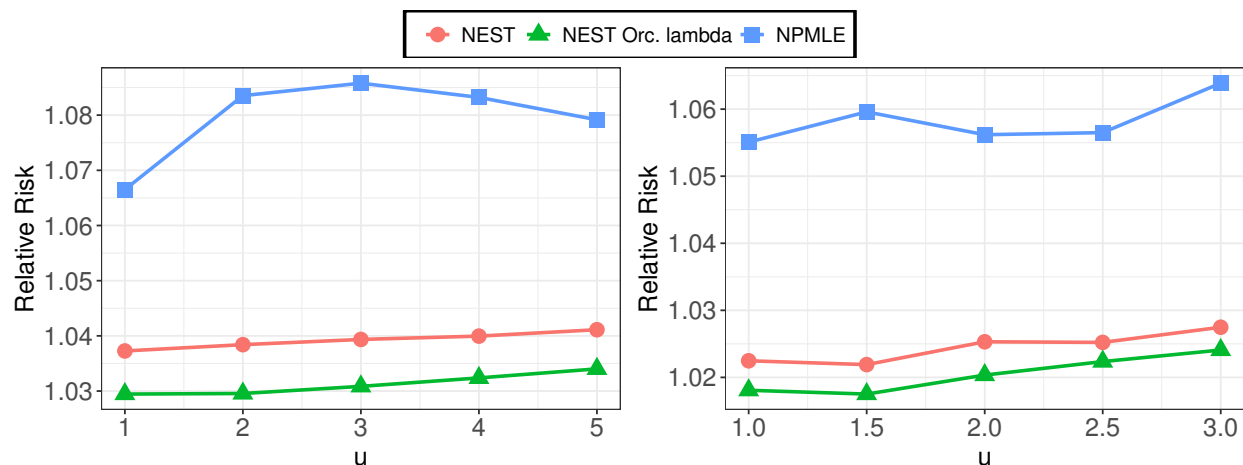


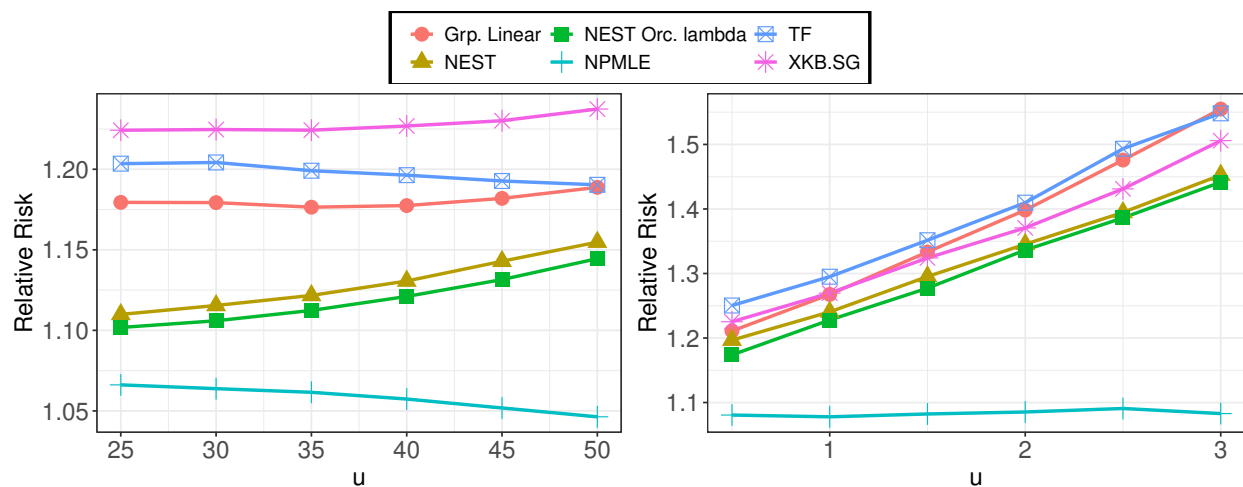
Figure EC.7 Comparison of relative risk for non-normal data. Here $Y_{ij} | \mu_i, \sigma_i \stackrel{i.i.d}{\sim} U(\mu_i - \sqrt{3\sigma_i^2}, \mu_i + \sqrt{3\sigma_i^2})$, σ_i^2 are sampled independently from $N(u, 1)$ truncated below at 0.1 and $\mu_i | \sigma_i^2 \stackrel{ind.}{\sim} 0.8 N(0.25\sigma_i^2, 0.25) + 0.2 N(\sigma_i^2, 1)$. Plots show $m = 10, 15, 20$ left to right.

The fifth setting, Figure EC.7, corresponds to the setting where the data $Y_{ij} | (\mu_i, \sigma_i^2)$ are not normally distributed. Here the proposed NEST method demonstrates robustness to departures from the Normal model particularly in comparison to TF, Grp. Linear and Jing.SM.

Overall, the results of the preceding five simulation settings corroborate the statement of Proposition EC.2 and reveal that when the variances are unknown, the NEST estimation framework enjoys a relatively better risk performance for estimating the means under the squared error loss than the linear shrinkage methods and Tweedie's formula that rely on sample variances.



(a) Simulation settings 4 (left) and 7 (right) from Section 5 where estimation is conducted under the weighted squared error loss.



(b) Simulation settings 2 (left) and 5 (right) from Section EC.3.1 where estimation is conducted under the squared error loss.

Figure EC.8 Compound estimation of means under small and unequal sample sizes \mathbf{m} . Here \mathbf{m} is a fixed vector of size n with elements sampled randomly from $(4, 5, 6)$ with replacement.

EC.3.2. Numerical Experiments with unequal sample sizes m_i

In this section, we present the risk performance of the competing approaches of sections 5 and EC.3.1 when the sample sizes m_i are small and differ across the $n = 1000$ units of study. We use the following four simulation settings: Settings 4 and 7 from Section 5 and Settings 2 and 5

from Section EC.3.1. However, we change how $\mathbf{m} = (m_1, \dots, m_n)$ are generated in these settings. Figures EC.8a and EC.8b present the relative risks of the competing estimators across these four scenarios. Here \mathbf{m} is generated as a fixed vector of size n with elements sampled randomly from $\{4, 5, 6\}$ with replacement. The case where $m \in \{4, 5\}$ represents a particularly challenging scenario for NEST which is based on the following observation: in Equation (8) $(m_i - 3)\{(m_i - 1)S_i^2\}^{-1}$ is an unbiased estimator of τ_i which follows from the fact that $\{(m_i - 1)S_i^2\tau_i\}^{-1}$ has an inverse Chi-square distribution with $m_i - 1$ degrees of freedom. However the variance of this distribution does not exist unless $m_i > 5$.

Figure EC.8a represents Settings 4 (left) and 7 (right) from Section 5. We note that NEST continues to dominate NPMLE under the weighted squared error loss even when m_i are small. However, at such small sample sizes the risk of the NEST estimator is relatively larger than that of $\delta_{(1)}^\pi$.

Figure EC.8b exhibits Settings 2 (left) and 5 (right) from Section EC.3.1 where the estimation is conducted under the squared error loss. Here, we use the semi-parametric monotonically constrained SURE estimator that shrinks towards the grand mean, XKB.SG, from Xie et al. (2012) in place of Jing.SM from Section EC.3.1 as the latter was originally designed for the case $m_i = m$. While the extension of Jing.SM to unequal m_i is straightforward, we do not pursue that direction in this article and, instead, use XKB.SG in its place.

In Figure EC.8b we note that NPMLE dominates NEST and NEST exhibits a relatively better risk performance than TF, Grp. Linear and XKB.SG. However, in Setting 5 (right panel of Figure EC.8b), where the data $Y_{ij} | (\mu_i, \sigma_i^2)$ are not normally distributed, the performance of NEST is substantially poorer than NPMLE when u is large. The two main reasons for this behavior are related to (1) Proposition EC.2 where the oracle NEST estimator $\delta_{(1)}^\pi$ is not, in general, the optimal estimator of the means under the squared error loss, and (2) when the sample size m_i is less than 6 the variance of $\{(m_i - 1)S_i^2\tau_i\}^{-1}$ under Model (1) does not exist.

EC.3.3. Compound Estimation of Ratios

In this section we demonstrate the use of the NEST estimation framework for compound estimation of n ratios $\theta_i = \sqrt{m_i}\mu_i/\sigma_i$ which represent a popular financial metric for assessing mutual fund performance (see Section EC.4.2 for a related real data application involving compound estimation of mutual fund Sharpe ratios.). We evaluate the performance of the same six methods under the




Figure EC.9 Comparison of relative risk for estimating θ . Here $\sigma_i^2 \stackrel{i.i.d.}{\sim} U(0.1, u)$ and $\mu_i|\sigma_i \stackrel{i.i.d.}{\sim} 0.5 N(0.5\sigma_i^2, 0.5^2) + 0.5 N(-0.5\sigma_i^2, 0.5^2)$. Plots show $m = 10, 15, 20$ left to right.

squared error loss with n fixed at 1000 and $m_i = m$ for $i = 1, \dots, n$. We consider two simulation settings and plot the risk performance of the competing estimators of $\theta = (\theta_1, \dots, \theta_n)$ relative to the optimal Bayes estimator that estimates θ using the vector of posterior means $\mathbb{E}(\theta_i|y_i, s_i^2)$ for $i = 1, \dots, n$. Note that amongst the six methods considered here, NPMLE is the only method that is designed to estimate these posterior means. The other methods estimate θ_i by separately estimating μ_i and σ_i , and then take their ratio to construct an estimate. Grp. Linear and TF, in particular, rely on the sample standard deviation for estimating σ_i while NEST and Jing.SM employ their respective empirical Bayes estimators of σ_i^2 .

Setting 1 is presented in Figure EC.9. The data Y_{ij} are generated independently from $N(\mu_i, \sigma_i^2)$, the variances σ_i^2 are simulated uniformly between 0.1 and u , and the means are independently drawn from a mixture model with half chance $N(-\sigma_i^2/2, 0.5^2)$ and the other half $N(\sigma_i^2/2, 0.5^2)$. We continue to see that NEST has a lower relative risk than Grp. Linear and TF, both of which

use sample variances. NEST also dominates Jing.SM for almost all values of u while NPMLE dominates NEST for small values of u . As u increases the heterogeneity in the data grows and we see that the relative risks of Grp. Linear and Tweedie’s Formula across all m first decrease and then increase. The shift in the behavior of these estimators is related to the observation that as u increases, the centers of the mixture model that generates μ_i , are on average, further away from one another. This makes estimating the numerator of the ratio easier for all methods up until a point. As heterogeneity increases further, the risks of these methods that use the sample standard deviation in the denominator of θ_i are relatively worse than the risk of NEST and NPMLE.




Figure EC.10 Comparison of relative risk for estimating θ . Here $\tau_i \stackrel{i.i.d.}{\sim} 0.5\Gamma(20, \text{rate} = 20) + 0.5\Gamma(20, \text{rate} = u)$ and $\mu_i | \tau_i \stackrel{ind.}{\sim} 0.5 N(0.5/\tau_i, 0.5^2) + 0.5 N(-0.5/\tau_i, 0.5^2)$. Plots show $m = 10, 15, 20$ left to right.

Setting 2 is presented in Figure EC.10 where the means are generated according to Setting 1 but the precisions are independently drawn from a mixture model with half chance $\Gamma(\text{shape} = 20, \text{rate} = 20)$ and the other half $\Gamma(\text{shape} = 20, \text{rate} = u)$. We note that while NEST dominates TF and Grp. Linear, NPMLE and Jing.SM dominate NEST when the heterogeneity is relatively smaller and when $m = 10$. The improved performance of Jing.SM in this setting is potentially related to the observation that when u is small, the rate mixture of Gamma distributions on τ_i can be well approximated by a single Gamma distribution and that coincides with the parametric prior that Jing et al. (2016b) use on the precision to derive their empirical Bayes estimator for the variances. In Figures EC.9 and EC.10 we note that for $m = 10$, the relative risk of NEST is substantially

higher than NEST Orc. λ . This is not unexpected because NEST Orc. λ estimates λ by minimizing the true squared error loss involving θ while the data-driven NEST estimator relies on the modified cross-validation approach described in Section 3.2 with $\vartheta_n(\lambda; \mathcal{U}, \mathcal{V})$ defined as

$$\vartheta_n(\lambda; \mathcal{U}, \mathcal{V}) = \frac{1}{n} \sum_{i=1}^n \{ \bar{V}_i - \delta_i^{\text{ds}}(\bar{U}_i; \mathcal{U}, \lambda) \}^2,$$

to choose λ .

EC.4. Real Data Analyses

EC.4.1. Baseball Data

We analyze the monthly data on the number of “at bats” and “hits” for all U.S Major League baseball players over the regular seasons from 2002 until 2011. In this analysis we focus on both pitchers and non-pitchers using an approach similar to that of [Gu and Koenker \(2017a\)](#). The data are available from the R package `REBayes` and have been aggregated into half seasons to produce an unbalanced panel. It includes observations on 932 players who have at least ten at bats in any half season and appear in no fewer than five half-seasons (note that there are a total of 20 half-seasons that a player can appear in).

Following [Brown \(2008\)](#), let the transformed batting average Y_{ij} for player $i (= 1, \dots, n)$ at time $j (= 1, \dots, m_i)$ be denoted by $Y_{ij} = \arcsin \left(\sqrt{\frac{H_{ij} + 0.25}{N_{ij} + 0.5}} \right)$ where H_{ij} denotes the number of “hits” and N_{ij} denotes the number of “at bats” at time j for player i . We assume that $Y_{ij} \sim N(\mu_i, v_{ij}^2/\tau_i)$ where $\mu_i = \arcsin(\sqrt{p_i})$, p_i being player i ’s batting success probability, and $v_{ij}^2 = 1/(4N_{ij})$. Here $1/\tau_i$ are player specific scale parameters as described in [Gu and Koenker \(2017a\)](#). Under this setup, the sufficient statistics are

$$\hat{\mu}_i = \left(\sum_{j=1}^{m_i} 1/v_{ij}^2 \right)^{-1} \sum_{j=1}^{m_i} Y_{ij}/v_{ij}^2 \sim N(\mu_i, v_i^2/\tau_i) \text{ with } v_i^2 = \left(4 \sum_{j=1}^{m_i} N_{ij} \right)^{-1},$$

$$S_i^2 = \frac{1}{m_i - 1} \sum_{j=1}^{m_i} (Y_{ij} - \hat{\mu}_i)^2/v_{ij}^2 \text{ with } (m_i - 1)S_i^2\tau_i \sim \mathcal{X}_{m_i-1}^2.$$

In this analysis, the goal is to use the 2002-2011 data to predict the batting averages of the players in 2012. Players are divided into three categories: all, non-pitchers, and pitchers. We consider

Table EC.1 Performance of the competing estimators relative to the performance of the naive estimator \mathbf{Y}_{2011} .

Here $R\text{-}TSE(\boldsymbol{\delta}) = TSE(\boldsymbol{\delta})/TSE(\mathbf{Y}_{2011})$ with similar definitions for R-NSE and R-TSEp. The smallest two relative errors are bolded

		NPMLE-Indep	NPMLE-Dep	$\hat{\boldsymbol{\mu}}$	\bar{Y}^{2011}	NEST	NEST orc λ
All							
n for estimation: 932	R-TSE	0.463	0.491	0.352	1.958	0.350	0.350
n for prediction: 370	R-NSE	0.668	0.678	0.676	1.495	0.670	0.670
	R-TSEp	0.582	0.610	0.501	1.783	0.498	0.498
Nonpitchers							
n for estimation: 792	R-TSE	0.535	0.559	0.503	0.551	0.488	0.484
n for prediction: 325	R-NSE	0.656	0.665	0.679	0.973	0.667	0.659
	R-TSEp	0.651	0.677	0.619	0.682	0.604	0.601
Pitchers							
n for estimation: 140	R-TSE	0.659	0.662	0.629	0.804	0.628	0.620
n for prediction: 45	R-NSE	0.659	0.663	0.649	0.769	0.649	0.638
	R-TSEp	0.124	0.124	0.133	0.337	0.133	0.122

the following seven estimators of μ_i : two non-parametric maximum likelihood based estimators, denoted NPMLE-Indep and NPMLE-Dep which assume, respectively, independent and dependent priors on $(\mu_i, 1/\tau_i)$, the sufficient statistics $\hat{\boldsymbol{\mu}} = (\hat{\mu}_1, \dots, \hat{\mu}_n)$ of $\boldsymbol{\mu}$, the grand mean across all players in the 2011 season $\bar{Y}^{2011} = n^{-1} \sum_{i=1}^n Y_{i,2011}$, the proposed NEST estimator and its counterpart with an oracle choice for the tuning parameter λ (NEST orc λ), and the naive estimator that uses 2011 batting averages $\mathbf{Y}_{2011} = (Y_{1,2011}, \dots, Y_{n,2011})$. To assess how well these methods predict 2012 batting averages $\mathbf{Y}_{2012} = (Y_{1,2012}, \dots, Y_{n,2012})$, we consider three criteria for evaluating any estimate δ_i of μ_i : total squared error from [Brown \(2008\)](#) and defined as $TSE(\boldsymbol{\delta}) = \sum_{i=1}^n \left\{ (Y_{i,2012} - \delta_i)^2 - (4N_{i,2012})^{-1} \right\}$, normalized squared error from [Gu and Koenker \(2017a\)](#) and defined as $NSE(\boldsymbol{\delta}) = \sum_{i=1}^n \left\{ 4N_{i,2012} (Y_{i,2012} - \delta_i)^2 \right\}$, and total squared error on a probability scale from [Jiang et al. \(2010\)](#) which is defined as $TSEp(\hat{\boldsymbol{p}}) = \sum_{i=1}^n \left\{ (p_{i,2012} - \hat{p}_i)^2 - p_{i,2012}(1 - p_{i,2012})(4N_{i,2012})^{-1} \right\}$. Here $\hat{p}_i = \sin^2(\delta_i)$ and $p_{i,2012} = \sin^2(Y_{i,2012})$.

In Table EC.1, we report the performance of the competing estimators relative to the performance of the naive estimator \mathbf{Y}_{2011} wherein $R-TSE(\boldsymbol{\delta}) = TSE(\boldsymbol{\delta})/TSE(\mathbf{Y}_{2011})$ with similar definitions for $R-NSE$ and $R-TSE_p$. Thus, a smaller value of $R-TSE$, $R-NSE$ or $R-TSE_p$ indicates a relatively better prediction error. Across “All” and “Nonpitchers”, NEST exhibits the best relative risk for two of the three performance metrics. It is interesting to note that the sufficient statistics $\hat{\boldsymbol{\mu}}$ are quite competitive in this example while NPMLE with independent priors dominate the one with dependent priors across “All”, “Nonpitchers” and “Pitchers”. The compound estimation problem for “Pitchers” is an example of a setting where n is relatively small and NEST demonstrates a better risk performance than NPMLE for total squared error and normalized squared error losses.

EC.4.2. Mutual Fund Sharpe Ratios

In this section we analyze a dataset on $n_1 = 5,000$ monthly mutual fund returns spanning 12 months from January 2014 to December 2014. This data are sourced from the Wharton research data services (Wharton School 1993). The goal in this analysis is to use Sharpe ratios constructed using the data on the first $m_1 = 6$ months, January 2014 - June 2014, to predict the corresponding Sharpe ratios for the next 6 months. Formally, let Y_{ij} denote the excess return of fund $i (= 1, \dots, n_1)$ in month $j (= 1, \dots, m_1)$ over the return on the 3 month treasury yield. Denote $\bar{Y}_i = m_1^{-1} \sum_{j=1}^{m_1} Y_{ij}$, $S_i^2 = (m_1 - 1)^{-1} \sum_{j=1}^{m_1} (Y_{ij} - \bar{Y}_i)^2$ and $\delta_i^{\text{naive}} = \bar{Y}_i / \sqrt{S_i^2}$ to be, respectively, the sample mean, the sample variance and the observed Sharpe ratio of the monthly excess returns. Of the 5,000 funds available during these first 6 months, there are $n_2 = 4,958$ funds that appear in the next 6 months, July 2014 - December 2014, and have at least 3 months of returns available during this period. For our prediction, we consider these n_2 funds to assess the performance of various estimators for predicting $\theta_i = \mu_i / \sigma_i$ where μ_i and σ_i are the sample mean and sample standard deviation of the excess returns of the n_2 funds during the next 6 months.

We consider the following estimators of $\boldsymbol{\theta} = (\theta_1, \dots, \theta_n)$: NEST, NEST orc. λ , Tweedie’s formula (TF), Grp Linear, Jing.SM and NPMLE from Section EC.3.1. Additionally, we consider the SURE estimators XKB.SG and XKB.G from Xie et al. (2012). Note that for predicting θ_i , TF, Grp

Table EC.2 Performance of the competing estimators relative to the performance of the naive estimator δ^{naive} .

Here $R\text{-}TSE(\delta) = TSE(\delta)/TSE(\delta^{\text{naive}})$ with similar definitions for R-WSE and R-WAE. The smallest two relative errors are bolded

	n_1	n_2	Grp Lin.	XKB.G	XKB.SG	NPMLE	Jing.SM	TF	NEST	NEST orc. λ
R - TSE	5000	4958	0.896	0.931	0.922	0.842	0.720	0.997	0.688	0.686
R - WSE	5000	4958	0.893	0.930	0.920	0.837	0.716	0.997	0.686	0.684
R - WAE	5000	4958	1.067	0.974	0.983	0.479	0.171	0.999	0.782	0.780

Linear, XKB.SG and XKB.G rely on the sample variances S_i^2 . To evaluate the performance of these estimators for predicting θ , we consider the following three criteria with $m_{i,2} \in [3, 6]$: Total Squared Error : $TSE(\delta) = \sum_{i=1}^{n_2} (\theta_i - \delta_i)^2$; weighted Squared Error : $WSE(\delta) = \sum_{i=1}^{n_2} m_{i,2} (\theta_i - \delta_i)^2$; and weighted Absolute Error : $WAE(\delta) = \sum_{i=1}^{n_2} m_{i,2} |1 - \delta_i/\theta_i|$. In Table EC.2, we present the performance of the competing estimators relative to the performance of the naive estimator $\delta^{\text{naive}} = (\delta_i^{\text{naive}} : 1 \leq i \leq n)$ so that a smaller value of R-TSE, R-WSE or R-WAE indicates a relatively better prediction error.

Along the performance measures of Total Squared Error and Weighted Squared Error, NEST has the smallest relative risk among all competing estimators considered in this example. With respect to the Weighted Absolute Error, Jing.SM has a substantially smaller relative risk than NEST while Grp Linear appears to be doing relatively worse and exhibits an R-WAE bigger than 1. When compared against the three linear shrinkage methods considered here, NEST and NPMLE demonstrate an overall value in joint shrinkage estimation of the means μ_i and the variances σ_i^2 for predicting θ .

EC.5. Extensions

This section considers the extension of our methodology to several well known members in the two-parameter exponential family. We will focus on examples where the nuisance parameter is known. Our proposed estimation framework is motivated by the double shrinkage idea, but the approach nonetheless handles the case with known nuisance parameters. We discuss four examples, in each

of which we derive the Bayes estimator of the natural parameter under the squared error loss. The Bayes estimator in these examples relies on the unknown score function (of the marginal density of the sufficient statistic), which can be consistently estimated using the ideas in Section 3.

EXAMPLE EC.1 (LOCATION MIXTURE OF GAUSSIANS). Consider the following hierarchical model

$$Y_i \mid \mu_i, \tau_i \stackrel{i.i.d.}{\sim} N(\mu_i, 1/\tau_i), \quad \mu_i \stackrel{i.i.d.}{\sim} G_\mu(\cdot), \quad \text{for } i = 1, \dots, n, \quad (\text{EC.37})$$

where τ_i are known and $G_\mu(\cdot)$ is an unspecified prior. Equation (EC.37) represents the heteroskedastic normal means problem with known variances $1/\tau_i$ [see for example Weinstein et al. (2018)]. In this setting, the sufficient statistic for μ_i is Y_i and the Bayes estimator of μ_i under the squared error loss is given by

$$\mu_i^\pi := \mathbb{E}(\mu_i \mid y_i, \tau_i) = y_i + \frac{1}{\tau_i} \frac{\partial}{\partial y_i} \log f(y_i \mid \tau_i),$$

where $f(\cdot \mid \tau_i)$ is the pdf of the distribution of Y_i given τ_i marginalizing out μ_i . From Section 3 and with $m_i = 1$, $\mathbf{x}_i = (y_i, \tau_i)$, the NEST estimate of μ_i is given by $\delta_{i,n}^{\text{nest}}(\lambda) = y_i + \frac{1}{\tau_i} \hat{w}_{\lambda,n}^{(1)}(i)$.

EXAMPLE EC.2 (SCALE MIXTURE OF GAMMA DISTRIBUTIONS). Consider the following model

$$Y_{ij} \mid \alpha_i, 1/\beta_i \stackrel{i.i.d.}{\sim} \Gamma(\alpha_i, 1/\beta_i), \quad 1/\beta_i \stackrel{i.i.d.}{\sim} G(\cdot),$$

where the shape parameters α_i are known and $G(\cdot)$ is an unspecified prior distribution on scale parameters $1/\beta_i$. Here $T_i = \sum_{j=1}^m Y_{ij}$ is a sufficient statistic and $T_i \mid \alpha_i, \beta_i \stackrel{i.i.d.}{\sim} \Gamma(m\alpha_i, 1/\beta_i)$. The posterior distribution of $1/\beta_i$ belongs to a one-parameter exponential family with density

$$f(1/\beta_i \mid T_i, \alpha_i) \propto \exp \left\{ -T_i \beta_i + (m\alpha_i - 1) \log T_i - \log f(T_i \mid \alpha_i) \right\}, \quad (\text{EC.38})$$

where $f(\cdot \mid \alpha_i)$ is the pdf of the distribution of T_i given α_i (marginalizing out $1/\beta_i$). From Equation (EC.38), the Bayes estimator of β_i under the squared error loss is given by

$$\beta_i^\pi := \mathbb{E}(\beta_i \mid T_i, \alpha_i) = \frac{m\alpha_i - 1}{T_i} - \frac{\partial}{\partial T_i} \log f(T_i \mid \alpha_i).$$

With $\mathbf{x}_i = (T_i, \alpha_i)$, the NEST estimate of β_i is given by $\delta_{i,n}^{\text{nest}}(\lambda) = \frac{m\alpha_i - 1}{T_i} - \hat{w}_{\lambda,n}^{(1)}(i)$.

EXAMPLE EC.3 (SHAPE MIXTURE OF GAMMA DISTRIBUTIONS). We consider the following model:

$$Y_{ij} \mid \alpha_i, 1/\beta_i \stackrel{i.i.d.}{\sim} \Gamma(\alpha_i, 1/\beta_i), \quad \alpha_i \stackrel{i.i.d.}{\sim} G(\cdot),$$

where the scale parameters $1/\beta_i$ are known and $G(\cdot)$ is an unspecified prior distribution on the shape parameters α_i . Let $Y_i = \sum_{j=1}^m Y_{ij}$. Then $Y_i \mid \alpha_i, \beta_i \stackrel{ind.}{\sim} \Gamma(m\alpha_i, 1/\beta_i)$ and $T_i = \log Y_i$ is a sufficient statistic. Moreover, the posterior distribution of α_i belongs to a one-parameter exponential family with density

$$f(\alpha_i \mid T_i, 1/\beta_i) \propto \exp \left\{ (m\alpha_i)T_i - \beta_i \exp(T_i) - \log f(T_i \mid 1/\beta_i) \right\}, \quad (\text{EC.39})$$

where $f(\cdot \mid 1/\beta_i)$ is the density of the distribution of T_i given $1/\beta_i$ marginalizing out α_i . From Equation (EC.39), the Bayes estimator of α_i under the squared error loss is given by

$$\alpha_i^\pi := \mathbb{E}(\alpha_i \mid T_i, 1/\beta_i) = \frac{\beta_i \exp(T_i)}{m} + \frac{1}{m} \frac{\partial}{\partial T_i} \log f(T_i \mid 1/\beta_i).$$

With $\mathbf{x}_i = (T_i, 1/\beta_i)$, the NEST estimate of α_i is $\delta_{i,n}^{\text{nest}}(\lambda) = \frac{\beta_i \exp(T_i)}{m} + \frac{1}{m} \hat{w}_{\lambda,n}^{(1)}(i)$.

EXAMPLE EC.4 (SCALE MIXTURE OF WEIBULLS). We consider the following model:

$$Y_{ij} \mid k_i, \beta_i \stackrel{i.i.d.}{\sim} \text{Weibull}(k_i, \beta_i), \quad \beta_i \stackrel{i.i.d.}{\sim} G(\cdot). \quad (\text{EC.40})$$

We have $f(y \mid k, \beta) = \beta k y^{k-1} \exp(-\beta y^k)$. In Equation (EC.40) the shape parameters k_i are known, $G(\cdot)$ is an unspecified prior distribution on the scale parameters β_i , $T_i = \sum_{j=1}^m \{Y_{ij}\}^{k_i}$ is a sufficient statistic, and $T_i \mid k_i, 1/\beta_i \stackrel{ind.}{\sim} \Gamma(m, 1/\beta_i)$. From Example 2, the Bayes estimator of β_i is

$$\beta_i^\pi := \mathbb{E}(\beta_i \mid T_i, k_i) = \frac{m-1}{T_i} + \frac{\partial}{\partial T_i} \log f(T_i \mid k_i).$$

With $\mathbf{x}_i = (T_i, k_i)$, the NEST estimate of β_i is given by $\delta_{i,n}^{\text{nest}}(\lambda) = \frac{m-1}{T_i} - \hat{w}_{\lambda,n}^{(1)}(i)$.

The preceding examples present a setting with known nuisance parameter. When both parameters are unknown, extensions of our estimation framework to an arbitrary member of the two-parameter exponential family is difficult. The main reason is that in the Gaussian case the sufficient statistics




Figure EC.11 Running time comparison of NEST and NPMLE. Left: r fixed at 30. Right: n fixed at 5000.

are independent and their marginal distributions are known. However, for other distributions such as the Gamma and Beta, the joint distribution of the two sufficient statistics is generally unknown. This impedes a full generalization of our approach. We anticipate that an iterative scheme that conducts shrinkage estimation on the primary and nuisance coordinates in turn may be developed by combining the ideas in Examples EC.3 and EC.4 above. We do not pursue those extensions in this article.

EC.6. Computational complexity

Here we discuss the computational complexities of NEST and NPMLE, and provide a comparison of their running time.

In contrast to the linear shrinkage estimators, such as Group Linear Weinstein et al. (2018), SURE estimators of Xie et al. (2012) and Jing et al. (2016b), NEST and NPMLE are similar in the sense that both these approaches rely on solving a convex optimization problem to estimate the means. An implementation of the convex optimization problem for estimating the joint prior distribution of the means and the variances via the NPMLE is available in the R package REBayes (Koenker and Gu 2017). This package relies on interior point methods for solving the convex problem. See Koenker and Mizera (2014) for more details. The main effort in solving this problem depends on computing a Hessian and the computational complexity of that is $O(n^2r^2)$ where r

is the number of grid points at which the prior masses for each of the two prior distributions will be estimated. Usually $n \gg r$. Recently, [Kim et al. \(2020\)](#) propose a fast sequential quadratic programming based algorithm for estimating a one dimensional prior distribution using the NPMLE however their algorithm is not available for the case where the prior is two dimensional, such as the scenario when both mean and variances are unknown. For NEST, the convex quadratic optimization problem described in Equation (12) depends on the $n \times n$ matrix \mathbf{K}_λ and uses the interior-point optimizer in MOSEK. The worst case computational complexity for evaluating the hessian of the underlying objective function is $O(n^3)$, which is substantially larger than that of NPMLE. [Figure EC.11](#) provides a comparison of the running time of NEST and NPMLE for increasing n and r . In the left panel of [Figure EC.11](#) we fix $r = 30$, the default choice implemented in the function `WGLVmix` in `REBayes`, and vary n . In the right panel, n is fixed at 5000 with r varying. We see that for problems with large n and r fixed, NPMLE is substantially faster than NEST and NEST, in its current implementation, may not be as scalable as NPMLE is to large n problems. Our future research efforts will be geared towards developing faster first-order methods to solve Equation (12) rather than relying on the interior-point solver in MOSEK.