

OPTIMAL SHRINKAGE ESTIMATION OF MEAN PARAMETERS IN FAMILY OF DISTRIBUTIONS WITH QUADRATIC VARIANCE

BY XIANCHAO XIE, S. C. KOU¹ AND LAWRENCE BROWN²

*Two Sigma Investments LLC, Harvard University and
University of Pennsylvania*

This paper discusses the simultaneous inference of mean parameters in a family of distributions with quadratic variance function. We first introduce a class of semiparametric/parametric shrinkage estimators and establish their asymptotic optimality properties. Two specific cases, the location-scale family and the natural exponential family with quadratic variance function, are then studied in detail. We conduct a comprehensive simulation study to compare the performance of the proposed methods with existing shrinkage estimators. We also apply the method to real data and obtain encouraging results.

1. Introduction. The simultaneous inference of several mean parameters has interested statisticians since the 1950s and has been widely applied in many scientific and engineering problems ever since. [Stein \(1956\)](#) and [James and Stein \(1961\)](#) discussed the homoscedastic (equal variance) normal model and proved that shrinkage estimators can have uniformly smaller risk compared to the ordinary maximum likelihood estimate. This seminal work inspired a broad interest in the study of shrinkage estimators in hierarchical normal models. [Efron and Morris \(1972, 1973\)](#) studied the James–Stein estimators in an empirical Bayes framework and proposed several competing shrinkage estimators. [Berger and Strawderman \(1996\)](#) discussed this problem from a hierarchical Bayesian perspective. For applications of shrinkage techniques in practice, see [Efron and Morris \(1975\)](#), [Rubin \(1981\)](#), [Morris \(1983\)](#), [Green and Strawderman \(1985\)](#), [Jones \(1991\)](#) and [Brown \(2008\)](#).

There has also been substantial discussion on simultaneous inference for non-Gaussian cases. [Brown \(1966\)](#) studied the admissibility of invariance estimators in general location families. [Johnstone \(1984\)](#) discussed inference in the context of Poisson estimation. [Clevenson and Zidek \(1975\)](#) showed that Stein’s effect also exists in the Poisson case, while a more general treatment of discrete exponential families is given in [Hwang \(1982\)](#). The application of non-Gaussian

Received February 2015; revised August 2015.

¹Supported in part by grants from NIH and NSF.

²Supported in part by grants from NSF.

MSC2010 subject classifications. 60K35.

Key words and phrases. Hierarchical model, shrinkage estimator, unbiased estimate of risk, asymptotic optimality, quadratic variance function, NEF-QVF, location-scale family.

hierarchical models has also flourished. Mosteller and Wallace (1964) used a hierarchical Bayesian model based on the negative-binomial distribution to study the authorship of the Federalist Papers. Berry and Christensen (1979) analyzed the binomial data using a mixture of dirichlet processes. Gelman and Hill (2007) gave a comprehensive discussion of data analysis using hierarchical models. Recently, nonparametric empirical Bayes methods have been proposed by Brown and Greenshtein (2009), Jiang and Zhang (2009, 2010) and Koenker and Mizera (2014).

In this article, we focus on the simultaneous inference of the mean parameters for families of distributions with quadratic variance function. These distributions include many common ones, such as the normal, Poisson, binomial, negative-binomial and gamma distributions; they also include location-scale families, such as the t , logistic, uniform, Laplace, Pareto and extreme value distributions. We ask the question: among all the estimators that estimate the mean parameters by shrinking the within-group sample mean toward a central location, is there an optimal one, subject to the intuitive constraint that more shrinkage is applied to observations with larger variances (or smaller sample sizes)? We propose a class of semiparametric/parametric shrinkage estimators and show that there is indeed an asymptotically optimal shrinkage estimator; this estimator is explicitly obtained by minimizing an unbiased estimate of the risk. We note that similar types of estimators are found in Xie, Kou and Brown (2012) in the context of the heteroscedastic (unequal variance) hierarchical normal model. The treatment in this article, however, is far more general, as it covers a much wider range of distributions. We illustrate the performance of our shrinkage estimators by a comprehensive simulation study on both exponential and location-scale families. We apply our shrinkage estimators on the baseball data obtained by Brown (2008), observing quite encouraging results.

The remainder of the article is organized as follows: In Section 2, we introduce the general class of semiparametric URE shrinkage estimators, identify the asymptotically optimal one and discuss its properties. We then study the special case of location-scale families in Section 3 and the case of natural exponential families with quadratic variance (NEF-QVF) in Section 4. A systematic simulation study is conducted in Section 5, along with the application of the proposed methods to the baseball data set in Section 6. We give a brief discussion in Section 7 and the technical proofs are placed in the Appendix.

2. Semiparametric estimation of mean parameters of distributions with quadratic variance function. Consider simultaneously estimating the mean parameters of p independent observations Y_i , $i = 1, \dots, p$. Assume that the observation Y_i comes from a distribution with quadratic variance function, that is, $E(Y_i) = \theta_i \in \Theta$ and $\text{Var}(Y_i) = V(\theta_i)/\tau_i$ such that

$$V(\theta_i) = v_0 + v_1\theta_i + v_2\theta_i^2$$

with v_k ($k = 0, 1, 2$) being known constants. The set Θ of allowable parameters is a subset of $\{\theta : V(\theta) \geq 0\}$. τ_i is assumed to be known here and can be interpreted as the within-group sample size (i.e., when Y_i is the sample average of the i th group) or as the (square root) inverse-scale of Y_i . It is worth emphasizing that distributions with quadratic variance function include many common ones, such as the normal, Poisson, binomial, negative-binomial and gamma distributions as well as location-scale families. We introduce the general theory on distributions with quadratic variance function in this section and will specifically treat the cases of location-scale family and exponential family in the next two sections.

In a simultaneous inference problem, hierarchical models are often used to achieve partial pooling of information among different groups. For example, in the famous normal-normal hierarchical model $Y_i \overset{\text{ind.}}{\sim} N(\theta_i, 1/\tau_i)$, one often puts a conjugate prior distribution $\theta_i \overset{\text{i.i.d.}}{\sim} N(\mu, \lambda)$ and uses the posterior mean

$$(2.1) \quad E(\theta_i | \mathbf{Y}; \mu, \lambda) = \frac{\tau_i}{\tau_i + 1/\lambda} \cdot Y_i + \frac{1/\lambda}{\tau_i + 1/\lambda} \cdot \mu$$

to estimate θ_i . Similarly, if Y_i represents the within-group average of Poisson observations $\tau_i Y_i \overset{\text{ind.}}{\sim} \text{Poisson}(\tau_i \theta_i)$, then with a conjugate gamma prior distribution $\theta_i \overset{\text{i.i.d.}}{\sim} \Gamma(\alpha, \lambda)$, the posterior mean

$$(2.2) \quad E(\theta_i | \mathbf{Y}; \alpha, \lambda) = \frac{\tau_i}{\tau_i + 1/\lambda} \cdot Y_i + \frac{1/\lambda}{\tau_i + 1/\lambda} \cdot \alpha\lambda$$

is often used to estimate θ_i . The hyper-parameters, (μ, λ) or (α, λ) above, are usually first estimated from the marginal distribution of Y_i and then plugged into the above formulae to form an empirical Bayes estimate.

One potential drawback of the formal parametric empirical Bayes approaches lies in its explicit parametric assumption on the prior distribution. It can lead to undesirable results if the explicit parametric assumption is violated in real applications—we will see a real-data example in Section 6. Aiming to provide more flexible and, at the same time, efficient simultaneous estimation procedures, we consider in this section a class of semiparametric shrinkage estimators.

To motivate these estimators, let us go back to the normal and Poisson examples (2.1) and (2.2). It is seen that the Bayes estimate of each mean parameter θ_i is the weighted average of Y_i and the prior mean μ (or $\alpha\lambda$). In other words, θ_i is estimated by shrinking Y_i toward a central location (μ or $\alpha\lambda$). It is also noteworthy that the amount of shrinkage is governed by τ_i , the sample size: the larger the sample size, the less is the shrinkage toward the central location. This feature makes intuitive sense. We will see in Section 4.2 that in fact these observations hold not only for normal and Poisson distributions, but also for general natural exponential families.

With these observations in mind, we consider in this section shrinkage estimators of the following form:

$$(2.3) \quad \hat{\theta}_i^{\mathbf{b},\mu} = (1 - b_i) \cdot Y_i + b_i \cdot \mu$$

with $b_i \in [0, 1]$ satisfying

$$(2.4) \quad \text{Requirement (MON)} : b_i \leq b_j \text{ for any } i \text{ and } j \text{ such that } \tau_i \geq \tau_j.$$

Requirement (MON) asks the estimator to shrink the group mean with a larger sample size (or smaller variance) less toward the central location μ . Other than this intuitive requirement, we do not put on any restriction on b_i . Therefore, this class of estimators is semiparametric in nature.

The question we want to investigate is, for such a general estimator $\hat{\theta}^{\mathbf{b},\mu}$, whether there exists an optimal choice of \mathbf{b} and μ . Note that the two parametric estimates (2.1) and (2.2) are special cases of the general class with $b_i = \frac{1/\lambda}{\tau_i + 1/\lambda}$. We will see shortly that such an optimal choice indeed exists asymptotically (i.e., as $p \rightarrow \infty$) and this asymptotically optimal choice is specified by an unbiased risk estimate (URE).

For a general estimator $\hat{\theta}^{\mathbf{b},\mu}$ with fixed \mathbf{b} and μ , under the sum of squared-error loss

$$l_p(\theta, \hat{\theta}^{\mathbf{b},\mu}) = \frac{1}{p} \sum_{i=1}^p (\hat{\theta}_i^{\mathbf{b},\mu} - \theta_i)^2,$$

an unbiased estimate of its risk $R_p(\theta, \hat{\theta}^{\mathbf{b},\mu}) = E[l_p(\theta, \hat{\theta}^{\mathbf{b},\mu})]$ is given by

$$\text{URE}(\mathbf{b}, \mu) = \frac{1}{p} \sum_{i=1}^p \left[b_i^2 \cdot (Y_i - \mu)^2 + (1 - 2b_i) \cdot \frac{V(Y_i)}{\tau_i + \nu_2} \right],$$

because

$$\begin{aligned} E[\text{URE}(\mathbf{b}, \mu)] &= \frac{1}{p} \sum_{i=1}^p \{ b_i^2 [\text{Var}(Y_i) + (\theta_i - \mu)^2] + (1 - 2b_i) \text{Var}(Y_i) \} \\ &= \frac{1}{p} \sum_{i=1}^p [(1 - b_i)^2 \text{Var}(Y_i) + b_i^2 (\theta_i - \mu)^2] = R_p(\theta, \hat{\theta}^{\mathbf{b},\mu}). \end{aligned}$$

Note that the idea and results below can be easily extended to the case of weighted quadratic loss, with the only difference being that the regularity conditions will then involve the corresponding weight sequence.

Ideally the “best” choice of \mathbf{b} and μ is the one that minimizes $R_p(\theta, \hat{\theta}^{\mathbf{b},\mu})$, which is, however, unobtainable as the risk depends on the unknown θ . To bypass this impracticability, we minimize URE, the unbiased estimate, with respect to (\mathbf{b}, μ) instead. This gives our semiparametric URE shrinkage estimator:

$$(2.5) \quad \hat{\theta}_i^{\text{SM}} = (1 - \hat{b}_i) \cdot Y_i + \hat{b}_i \cdot \hat{\mu}^{\text{SM}},$$

where

$$\begin{aligned}
 (\hat{\mathbf{b}}^{\text{SM}}, \hat{\mu}^{\text{SM}}) &= \text{minimizer of } \text{URE}(\mathbf{b}, \mu) \\
 &\text{subject to } b_i \in [0, 1], \mu \in \left[-\max_i |Y_i|, \max_i |Y_i| \right] \cap \Theta \\
 &\text{and Requirement (MON)}.
 \end{aligned}$$

We require $|\mu| \leq \max_i |Y_i|$, since no sensible estimators shrink the observations to a location completely outside the range of the data. Intuitively, the URE shrinkage estimator would behave well if $\text{URE}(\mathbf{b}, \mu)$ is close to the risk $R_p(\boldsymbol{\theta}, \hat{\boldsymbol{\theta}}^{\mathbf{b}, \mu})$.

To investigate the properties of the semiparametric URE shrinkage estimator, we now introduce the following regularity conditions:

- (A) $\limsup_{p \rightarrow \infty} \frac{1}{p} \sum_{i=1}^p \text{Var}(Y_i) < \infty$;
- (B) $\limsup_{p \rightarrow \infty} \frac{1}{p} \sum_{i=1}^p \text{Var}(Y_i) \cdot \theta_i^2 < \infty$;
- (C) $\limsup_{p \rightarrow \infty} \frac{1}{p} \sum_{i=1}^p \text{Var}(Y_i^2) < \infty$;
- (D) $\sup_i \left(\frac{\tau_i}{\tau_i + \nu_2}\right)^2 < \infty, \sup_i \left(\frac{\nu_1}{\tau_i + \nu_2}\right)^2 < \infty$;
- (E) $\limsup_{p \rightarrow \infty} \frac{1}{p^{1-\varepsilon}} E(\max_{1 \leq i \leq p} Y_i^2) < \infty$ for some $\varepsilon > 0$.

The theorem below shows that $\text{URE}(\mathbf{b}, \mu)$ not only unbiasedly estimates the risk, but also serves as a good approximation of the actual loss $l_p(\boldsymbol{\theta}, \hat{\boldsymbol{\theta}}^{\mathbf{b}, \mu})$, which is a much stronger property. In fact, $\text{URE}(\mathbf{b}, \mu)$ is asymptotically uniformly close to the actual loss. Therefore, we expect that minimizing $\text{URE}(\mathbf{b}, \mu)$ would lead to an estimate with competitive risk properties.

THEOREM 2.1. *Assuming regularity conditions (A)–(E), we have*

$$\sup |\text{URE}(\mathbf{b}, \mu) - l_p(\boldsymbol{\theta}, \hat{\boldsymbol{\theta}}^{\mathbf{b}, \mu})| \rightarrow 0 \quad \text{in } L^1 \text{ and in probability, as } p \rightarrow \infty,$$

where the supremum is taken over $b_i \in [0, 1], |\mu| \leq \max_i |Y_i|$ and Requirement (MON).

The following result compares the asymptotic behavior of our URE shrinkage estimator (2.5) with other shrinkage estimators from the general class. It establishes the asymptotic optimality of our URE shrinkage estimator.

THEOREM 2.2. *Assume regularity conditions (A)–(E). Then for any shrinkage estimator $\hat{\boldsymbol{\theta}}^{\hat{\mathbf{b}}, \hat{\mu}} = (1 - \hat{\mathbf{b}}) \cdot \mathbf{Y} + \hat{\mathbf{b}} \cdot \hat{\mu}$, where $\hat{\mathbf{b}} \in [0, 1]$ satisfies Requirement (MON), and $|\hat{\mu}| \leq \max_i |Y_i|$, we always have*

$$\lim_{p \rightarrow \infty} P(l_p(\boldsymbol{\theta}, \hat{\boldsymbol{\theta}}^{\text{SM}}) \geq l_p(\boldsymbol{\theta}, \hat{\boldsymbol{\theta}}^{\hat{\mathbf{b}}, \hat{\mu}}) + \varepsilon) = 0 \quad \text{for any } \varepsilon > 0$$

and

$$\limsup_{p \rightarrow \infty} [R(\boldsymbol{\theta}, \hat{\boldsymbol{\theta}}^{\text{SM}}) - R(\boldsymbol{\theta}, \hat{\boldsymbol{\theta}}^{\hat{\mathbf{b}}, \hat{\mu}})] \leq 0.$$

As a special case of Theorem 2.2, the semiparametric URE shrinkage estimator asymptotically dominates the parametric empirical Bayes estimators, like (2.1) or (2.2). It is worth noting that the asymptotic optimality of our semiparametric URE shrinkage estimators does not assume any prior distribution on the mean parameters θ_i , nor does it assume any parametric form on the distribution of \mathbf{Y} (other than the quadratic variance function). Therefore, the results enjoy a large extent of robustness. In fact, the individual Y_i 's do not even have to be from the same distribution family as long as the regularity conditions (A)–(E) are met. (However, whether shrinkage estimation in that case is a good idea or not becomes debatable.)

2.1. *Shrinking toward the grand mean.* In the previous development, the central shrinkage location μ is determined by minimizing URE. The joint minimization of \mathbf{b} and μ offers asymptotic optimality in the class of estimators. For small or moderate p (the number of Y_i 's), however, it is not necessarily true that the semiparametric URE shrinkage estimator will always be the optimal one. In this setting, it might be beneficial to set μ by a predetermined rule and only optimize \mathbf{b} , as it might reduce the variability of the resulting estimate. In this subsection, we consider shrinking toward the grand mean:

$$\hat{\mu} = \bar{Y} = \frac{1}{p} \sum_{i=1}^p Y_i.$$

The particular reason why the grand average is chosen instead of the weighted average $\bar{Y}_w = (\sum_{i=1}^p \tau_i Y_i) / (\sum_{i=1}^p \tau_i)$ is that the latter might be biased when θ_i and τ_i are dependent. In the case where such dependence is not a concern, the idea and results obtained below can be similarly derived.

The corresponding estimator becomes

$$(2.6) \quad \hat{\theta}_i^{\mathbf{b}, \bar{Y}} = (1 - b_i) Y_i + b_i \bar{Y},$$

where $b_i \in [0, 1]$ satisfies Requirement (MON). To find the asymptotically optimal choice of \mathbf{b} , we start from an unbiased estimate of the risk of $\hat{\theta}^{\mathbf{b}, \bar{Y}}$.

It is straightforward to verify that for fixed \mathbf{b} an unbiased estimate of the risk of $\hat{\theta}^{\mathbf{b}, \bar{Y}}$ is

$$\text{URE}^G(\mathbf{b}) = \frac{1}{p} \sum_{i=1}^p \left[b_i^2 (Y_i - \bar{Y})^2 + \left(1 - 2 \left(1 - \frac{1}{p} \right) b_i \right) \frac{V(Y_i)}{\tau_i + \nu_2} \right].$$

Note that we use the superscript “G”, which stands for “grand mean”, to distinguish it from the previous $\text{URE}(\mathbf{b}, \mu)$. Like what we did previously, minimizing the URE^G with respect to \mathbf{b} then leads to our semiparametric URE grand-mean shrinkage estimator

$$(2.7) \quad \hat{\theta}_i^{\text{SG}} = (1 - \hat{b}_i^{\text{SG}}) \cdot Y_i + \hat{b}_i^{\text{SG}} \cdot \bar{Y},$$

where

$$\hat{\mathbf{b}}^{\text{SG}} = \text{minimizer of } \text{URE}^G(\mathbf{b})$$

subject to $b_i \in [0, 1]$ and Requirement (MON).

Again, we expect that the URE estimator $\hat{\theta}^{\text{SG}}$ would be competitive if URE^G is close to the risk function or the loss function. The next theorem confirms the uniform closeness.

THEOREM 2.3. *Under regularity conditions (A)–(E), we have*

$$\sup |\text{URE}^G(\mathbf{b}) - l_p(\theta, \hat{\theta}^{\mathbf{b}, \bar{Y}})| \rightarrow 0 \quad \text{in } L^1 \text{ and in probability, as } p \rightarrow \infty,$$

where the supremum is taken over $b_i \in [0, 1]$ and Requirement (MON).

Consequently, $\hat{\theta}^{\text{SG}}$ is asymptotically optimal among all shrinkage estimators $\hat{\theta}^{\mathbf{b}, \bar{Y}}$ that shrink toward the grand mean \bar{Y} , as shown in the next theorem.

THEOREM 2.4. *Assume regularity conditions (A)–(E). Then for any shrinkage estimator $\hat{\theta}^{\mathbf{b}, \bar{Y}} = (1 - \hat{\mathbf{b}}) \cdot \mathbf{Y} + \hat{\mathbf{b}} \cdot \bar{Y}$, where $\hat{\mathbf{b}} \in [0, 1]$ satisfies Requirement (MON), we have*

$$\lim_{p \rightarrow \infty} P(l_p(\theta, \hat{\theta}^{\text{SG}}) \geq l_p(\theta, \hat{\theta}^{\mathbf{b}, \bar{Y}}) + \varepsilon) = 0 \quad \text{for any } \varepsilon > 0$$

and

$$\limsup_{p \rightarrow \infty} [R(\theta, \hat{\theta}^{\text{SG}}) - R(\theta, \hat{\theta}^{\mathbf{b}, \bar{Y}})] \leq 0.$$

3. Simultaneous estimation of mean parameters in location-scale families.

In this section, we focus on location-scale families, which are a special case of distributions with quadratic variance functions. We show how the regularity conditions can be simplified in this case.

For a location-scale family, we can write

$$(3.1) \quad Y_i = \theta_i + Z_i / \sqrt{\tau_i},$$

where the standard variates Z_i are i.i.d. with mean zero and variance ν_0 . The constants ν_1 and ν_2 in the quadratic variance function $V(\theta_i) = \nu_0 + \nu_1\theta_i + \nu_2\theta_i^2$ are zero for the location-scale family. $1/\sqrt{\tau_i}$ is the scale of Y_i .

The next lemma simplifies the regularity conditions for a location-scale family.

LEMMA 3.1. *For $Y_i, i = 1, \dots, p$, independently from a location-scale family (3.1), the following four conditions imply the regularity conditions (A)–(E) in Section 2:*

- (i) $\limsup_{p \rightarrow \infty} \frac{1}{p} \sum_{i=1}^p 1/\tau_i^2 < \infty$;
- (ii) $\limsup_{p \rightarrow \infty} \frac{1}{p} \sum_{i=1}^p \theta_i^2/\tau_i < \infty$;
- (iii) $\limsup_{p \rightarrow \infty} \frac{1}{p} \sum_{i=1}^p |\theta_i|^{2+\varepsilon} < \infty$ for some $\varepsilon > 0$;
- (iv) the standard variate Z satisfies $P(|Z| > t) \leq Dt^{-\alpha}$ for constants $D > 0$, $\alpha > 4$.

Note that (i)–(iv) in Lemma 3.1 covers the common location-scale families, including the t (degree of freedom > 4), normal, uniform, logistic, Laplace, Pareto ($\alpha > 4$) and extreme value distributions.

Lemma 3.1, together with Theorems 2.1 and 2.2, immediately yields the following corollaries: the semiparametric URE shrinkage estimator is asymptotically optimal for location-scale families.

COROLLARY 3.2. For Y_i , $i = 1, \dots, p$, independently from a location-scale family (3.1), under conditions (i)–(iv) in Lemma 3.1, we have

$$\sup |URE(\mathbf{b}, \mu) - l_p(\boldsymbol{\theta}, \hat{\boldsymbol{\theta}}^{\mathbf{b}, \mu})| \rightarrow 0 \quad \text{in } L^1 \text{ and in probability, as } p \rightarrow \infty,$$

where the supremum is taken over $b_i \in [0, 1]$, $|\mu| \leq \max_i |Y_i|$ and Requirement (MON).

COROLLARY 3.3. Let Y_i , $i = 1, \dots, p$, be independent from a location-scale family (3.1). Assume conditions (i)–(iv) in Lemma 3.1. Then for any shrinkage estimator $\hat{\boldsymbol{\theta}}^{\hat{\mathbf{b}}, \hat{\mu}} = (1 - \hat{\mathbf{b}}) \cdot \mathbf{Y} + \hat{\mathbf{b}} \cdot \hat{\mu}$, where $\hat{\mathbf{b}} \in [0, 1]$ satisfies Requirement (MON), and $|\hat{\mu}| \leq \max_i |Y_i|$, we always have

$$\lim_{p \rightarrow \infty} P(l_p(\boldsymbol{\theta}, \hat{\boldsymbol{\theta}}^{\text{SM}}) \geq l_p(\boldsymbol{\theta}, \hat{\boldsymbol{\theta}}^{\hat{\mathbf{b}}, \hat{\mu}}) + \varepsilon) = 0 \quad \text{for any } \varepsilon > 0$$

and

$$\limsup_{p \rightarrow \infty} [R(\boldsymbol{\theta}, \hat{\boldsymbol{\theta}}^{\text{SM}}) - R(\boldsymbol{\theta}, \hat{\boldsymbol{\theta}}^{\hat{\mathbf{b}}, \hat{\mu}})] \leq 0.$$

In the case of shrinking toward the grand mean \bar{Y} , the corresponding semiparametric URE grand-mean shrinkage estimator is also asymptotically optimal.

COROLLARY 3.4. For Y_i , $i = 1, \dots, p$, independently from a location-scale family (3.1), under conditions (i)–(iv) in Lemma 3.1, we have

$$\sup |URE^G(\mathbf{b}) - l_p(\boldsymbol{\theta}, \hat{\boldsymbol{\theta}}^{\mathbf{b}, \bar{Y}})| \rightarrow 0 \quad \text{in } L^1 \text{ and in probability, as } p \rightarrow \infty,$$

where the supremum is taken over $b_i \in [0, 1]$ and Requirement (MON).

COROLLARY 3.5. Let Y_i , $i = 1, \dots, p$, be independent from a location-scale family (3.1). Assume conditions (i)–(iv) in Lemma 3.1. Then for any shrinkage

estimator $\hat{\theta}^{\hat{\mathbf{b}}, \bar{Y}} = (1 - \hat{\mathbf{b}}) \cdot \mathbf{Y} + \hat{\mathbf{b}} \cdot \bar{Y}$, where $\hat{\mathbf{b}} \in [0, 1]$ satisfies Requirement (MON), we have

$$\lim_{p \rightarrow \infty} P(l_p(\theta, \hat{\theta}^{\text{SG}}) \geq l_p(\theta, \hat{\theta}^{\hat{\mathbf{b}}, \bar{Y}}) + \varepsilon) = 0 \quad \text{for any } \varepsilon > 0$$

and

$$\limsup_{p \rightarrow \infty} [R(\theta, \hat{\theta}^{\text{SG}}) - R(\theta, \hat{\theta}^{\hat{\mathbf{b}}, \bar{Y}})] \leq 0.$$

4. Simultaneous estimation of mean parameters in natural exponential family with quadratic variance function.

4.1. *Semiparametric URE shrinkage estimators.* In this section, we focus on natural exponential families with quadratic variance functions (NEF-QVF), as they incorporate the most common distributions that one encounters in practice. We show how the regularity conditions (A)–(E) can be significantly simplified and offer concrete examples.

It is well known that there are in total six distinct distributions that belong to NEF-QVF [Morris (1982)]: the normal, binomial, Poisson, negative-binomial, Gamma and generalized hyperbolic secant (GHS) distributions. We represent in general an NEF-QVF as

$$Y_i \sim \text{NEF-QVF}[\theta_i, V(\theta_i)/\tau_i],$$

where $\theta_i = E(Y_i) \in \Theta$ is the mean parameter and τ_i is the convolution parameter (or within-group sample size). For example, in the binomial case $Y_i \sim \text{Bin}(n_i, p_i)/n_i$, $\theta_i = p_i$, $V(\theta_i) = \theta_i(1 - \theta_i)$ and $\tau_i = n_i$.

The next result provides easy-to-check conditions that considerably simplify those in Section 2. As the case of heteroscedastic normal data has been studied in Xie, Kou and Brown (2012), we concentrate on the other five NEF-QVF distributions.

LEMMA 4.1. *For the five non-Gaussian NEF-QVF distributions, Table 1 lists the respective conditions, under which regularity conditions (A)–(E) in Section 2 are satisfied. For example, for the binomial distribution, the condition is $\tau_i = n_i \geq 2$ for all i .*

Lemma 4.1 and the general theory in Section 2 yield the following optimality results for our semiparametric URE shrinkage estimator in the case of NEF-QVF.

COROLLARY 4.2. *Let $Y_i \stackrel{ind.}{\sim} \text{NEF-QVF}[\theta_i, V(\theta_i)/\tau_i]$, $i = 1, \dots, p$, be non-Gaussian. Under the respective conditions listed in Table 1, we have*

$$\sup | \text{URE}(\mathbf{b}, \mu) - l_p(\theta, \hat{\theta}^{\mathbf{b}, \mu}) | \rightarrow 0 \quad \text{in } L^1 \text{ and in probability, as } p \rightarrow \infty,$$

where the supremum is taken over $b_i \in [0, 1]$, $|\mu| \leq \max_i |Y_i|$ and Requirement (MON).

TABLE 1
 The conditions for the five nonnormal NEF-QVF distributions to guarantee the regularity conditions (A)–(E) in Section 2

Distribution	Data	(ν_0, ν_1, ν_2)	Note	Conditions
Binomial	$Y_i \sim \text{Bin}(n_i, p_i)/n_i$	(0, 1, -1)	$\tau_i = n_i$	$n_i \geq 2$ for all $i = 1, \dots, p$
Poisson	$Y_i \sim \text{Poi}(\tau_i \theta_i)/\tau_i$	(0, 1, 0)		(i) $\inf_i \tau_i > 0, \inf_i (\tau_i \theta_i) > 0$ (ii) $\sum_i \theta_i^3 = O(p)$
Neg-binomial	$Y_i \sim N \text{Bin}(n_i, p_i)/n_i$	(0, 1, 1)	$\tau_i = n_i$ $\theta_i = \frac{p_i}{1-p_i}$	(i) $\inf_i (n_i p_i) > 0$ (ii) $\sum_i (\frac{p_i}{1-p_i})^4 = O(p)$
Gamma	$Y_i \sim \Gamma(\tau_i \alpha, \lambda_i)/\tau_i$	(0, 0, 1/ α)	$\theta_i = \alpha \lambda_i$	(i) $\inf_i \tau_i > 0$ (ii) $\sum_i \lambda_i^4 = O(p)$
GHS	$Y_i \sim \text{GHS}(\tau_i \alpha, \lambda_i)/\tau_i$	($\alpha, 0, 1/\alpha$)	$\theta_i = \alpha \lambda_i$	(i) $\inf_i \tau_i > 0$ (ii) $\sum_i \lambda_i^4 = O(p)$

COROLLARY 4.3. Let $Y_i \stackrel{ind.}{\sim} \text{NEF-QVF}[\theta_i, V(\theta_i)/\tau_i], i = 1, \dots, p$, be non-Gaussian. Assume the respective conditions listed in Table 1. Then for any shrinkage estimator $\hat{\theta}^{\hat{\mathbf{b}}, \hat{\mu}} = (1 - \hat{\mathbf{b}}) \cdot \mathbf{Y} + \hat{\mathbf{b}} \cdot \hat{\mu}$, where $\hat{\mathbf{b}} \in [0, 1]$ satisfies Requirement (MON), and $|\hat{\mu}| \leq \max_i |Y_i|$, we always have

$$\lim_{p \rightarrow \infty} P(l_p(\theta, \hat{\theta}^{\text{SM}}) \geq l_p(\theta, \hat{\theta}^{\hat{\mathbf{b}}, \hat{\mu}}) + \varepsilon) = 0 \quad \text{for any } \varepsilon > 0$$

and

$$\limsup_{p \rightarrow \infty} [R(\theta, \hat{\theta}^{\text{SM}}) - R(\theta, \hat{\theta}^{\hat{\mathbf{b}}, \hat{\mu}})] \leq 0.$$

For shrinking toward the grand mean \bar{Y} , the corresponding semiparametric URE grand mean shrinkage estimator is also asymptotically optimal (within the smaller class).

COROLLARY 4.4. Let $Y_i \stackrel{ind.}{\sim} \text{NEF-QVF}[\theta_i, V(\theta_i)/\tau_i], i = 1, \dots, p$, be non-Gaussian. Under the respective conditions listed in Table 1, we have

$$\sup | \text{URE}^G(\mathbf{b}) - l_p(\theta, \hat{\theta}^{\mathbf{b}, \bar{Y}}) | \rightarrow 0 \quad \text{in } L^1 \text{ and in probability, as } p \rightarrow \infty,$$

where the supremum is taken over $b_i \in [0, 1]$ and Requirement (MON).

COROLLARY 4.5. Let $Y_i \stackrel{ind.}{\sim} \text{NEF-QVF}[\theta_i, V(\theta_i)/\tau_i], i = 1, \dots, p$, be non-Gaussian. Assume the respective conditions listed in Table 1. Then for any shrinkage estimator $\hat{\theta}^{\hat{\mathbf{b}}, \bar{Y}} = (1 - \hat{\mathbf{b}}) \cdot \mathbf{Y} + \hat{\mathbf{b}} \cdot \bar{Y}$, where $\hat{\mathbf{b}} \in [0, 1]$ satisfies Requirement (MON), we have

$$\lim_{p \rightarrow \infty} P(l_p(\theta, \hat{\theta}^{\text{SG}}) \geq l_p(\theta, \hat{\theta}^{\hat{\mathbf{b}}, \bar{Y}}) + \varepsilon) = 0 \quad \text{for any } \varepsilon > 0$$

TABLE 2
Conjugate priors for the five well-known NEF-QVF distributions

Distribution	Data	Conjugate prior	γ and μ
Normal	$Y_i \sim N(\theta_i, 1/\tau_i)$	$\theta_i \overset{\text{i.i.d.}}{\sim} N(\mu, \lambda)$	$\gamma = 1/\lambda$
Binomial	$Y_i \sim \text{Bin}(\tau_i, \theta_i)/\tau_i$	$\theta_i \overset{\text{i.i.d.}}{\sim} \text{Beta}(\alpha, \beta)$	$\gamma = \alpha + \beta, \mu = \frac{\alpha}{\alpha + \beta}$
Poisson	$Y_i \sim \text{Poi}(\tau_i \theta_i)/\tau_i$	$\theta_i \overset{\text{i.i.d.}}{\sim} \Gamma(\alpha, \lambda)$	$\gamma = 1/\lambda, \mu = \alpha\lambda$
Neg-binomial	$Y_i \sim N \text{Bin}(\tau_i, p_i)/\tau_i$	$p_i \overset{\text{i.i.d.}}{\sim} \text{Beta}(\alpha, \beta), \theta_i = \frac{p_i}{1-p_i}$	$\gamma = \beta - 1, \mu = \frac{\alpha}{\beta - 1}$
Gamma	$Y_i \sim \Gamma(\tau_i \alpha, \lambda_i)/\tau_i$	$\lambda_i \overset{\text{i.i.d.}}{\sim} \text{inv-}\Gamma(\alpha_0, \beta_0), \theta_i = \alpha\lambda_i$	$\gamma = \frac{\alpha_0 - 1}{\alpha}, \mu = \frac{\alpha\beta_0}{\alpha_0 - 1}$

and

$$\limsup_{p \rightarrow \infty} [R(\theta, \hat{\theta}^{\text{SG}}) - R(\theta, \hat{\theta}^{\text{b}, \bar{Y}})] \leq 0.$$

4.2. *Parametric URE shrinkage estimators and conjugate priors.* For Y_i from an exponential family, hierarchical models based on the conjugate prior distributions are often used; the hyper-parameters in the prior distribution are often estimated from the marginal distribution of Y_i in an empirical Bayes way. Two questions arise naturally in this scenario. First, are there other choices to estimate the hyper-parameters? Second, is there an optimal choice? We will show in this subsection that our URE shrinkage idea applies to the parametric conjugate prior case; the resulting *parametric* URE shrinkage estimators are asymptotically optimal, and thus asymptotically dominate the traditional empirical Bayes estimators.

Let $Y_i \sim \text{NEF-QVF}[\theta_i, V(\theta_i)/\tau_i], i = 1, \dots, p$, be independent. If θ_i are i.i.d. from the conjugate prior, the Bayesian estimate of θ_i is then

$$(4.1) \quad \hat{\theta}_i^{\gamma, \mu} = \frac{\tau_i}{\tau_i + \gamma} \cdot Y_i + \frac{\gamma}{\tau_i + \gamma} \cdot \mu,$$

where γ and μ are functions of the hyper-parameters in the prior distribution. Table 2 details the conjugate priors for the five well-known NEF-QVF distributions—the normal, binomial, Poisson, negative-binomial and gamma distributions—and the corresponding expressions of γ and μ in terms of the hyper-parameters. For example, in the binomial case, $Y_i \overset{\text{ind.}}{\sim} \text{Bin}(\tau_i, \theta_i)/\tau_i$ and the conjugate prior is $\theta_i \overset{\text{i.i.d.}}{\sim} \text{Beta}(\alpha, \beta); \gamma = \alpha + \beta, \mu = \alpha/(\alpha + \beta)$. Although it can be shown that for the sixth NEF-QVF distribution—the GHS distribution—taking a conjugate prior also gives (4.1), the conjugate prior distribution does not have a clean expression and is rarely encountered in practice. We thus omit it from Table 2.

We now apply our URE idea to formula (4.1) to estimate (γ, μ) , in contrast to the conventional empirical Bayes method that determines the hyper-parameters

through the marginal distribution. For fixed γ and μ , an unbiased estimate for the risk of $\hat{\theta}^{\gamma, \mu}$ is given by

$$\text{URE}^P(\gamma, \mu) = \frac{1}{p} \sum_{i=1}^p \left[\frac{\gamma^2}{(\tau_i + \gamma)^2} \cdot (Y_i - \mu)^2 + \frac{\tau_i - \gamma}{\tau_i + \gamma} \cdot \frac{V(Y_i)}{\tau_i + \nu_2} \right],$$

where we use the superscript “ P ” to stand for “parametric”. Minimizing $\text{URE}^P(\gamma, \mu)$ leads to our parametric URE shrinkage estimator

$$(4.2) \quad \hat{\theta}_i^{\text{PM}} = \frac{\tau_i}{\tau_i + \hat{\gamma}^{\text{PM}}} \cdot Y_i + \frac{\hat{\gamma}^{\text{PM}}}{\tau_i + \hat{\gamma}^{\text{PM}}} \cdot \hat{\mu}^{\text{PM}},$$

where

$$(\hat{\gamma}^{\text{PM}}, \hat{\mu}^{\text{PM}}) = \arg \min \text{URE}^P(\gamma, \mu) \quad \text{over } \left\{ 0 \leq \gamma \leq \infty, |\mu| \leq \max_i |Y_i|, \mu \in \Theta \right\}.$$

Parallel to Theorems 2.1 and 2.2, the next two results show that the parametric URE shrinkage estimator gives the asymptotically optimal choice of (γ, μ) , if one wants to use estimators of the form (4.1).

THEOREM 4.6. *Let $Y_i \stackrel{\text{ind.}}{\sim} \text{NEF-QVF}[\theta_i, V(\theta_i)/\tau_i]$, $i = 1, \dots, p$, be non-Gaussian. Under the respective conditions listed in Table 1, we have*

$$\sup | \text{URE}^P(\gamma, \mu) - l_p(\theta, \hat{\theta}^{\gamma, \mu}) | \rightarrow 0 \quad \text{in } L^1 \text{ and in probability, as } p \rightarrow \infty,$$

where the supremum is taken over $\{0 \leq \gamma \leq \infty, |\mu| \leq \max_i |Y_i|, \mu \in \Theta\}$.

THEOREM 4.7. *Let $Y_i \stackrel{\text{ind.}}{\sim} \text{NEF-QVF}[\theta_i, V(\theta_i)/\tau_i]$, $i = 1, \dots, p$, be non-Gaussian. Assume the respective conditions listed in Table 1. Then for any estimator $\hat{\theta}^{\hat{\gamma}, \hat{\mu}} = \frac{\tau}{\tau + \hat{\gamma}} \mathbf{Y} + \frac{\hat{\gamma}}{\tau + \hat{\gamma}} \hat{\mu}$, where $\hat{\gamma} \geq 0$ and $|\hat{\mu}| \leq \max_i |Y_i|$, we always have*

$$\lim_{p \rightarrow \infty} P(l_p(\theta, \hat{\theta}^{\text{PM}}) \geq l_p(\theta, \hat{\theta}^{\hat{\gamma}, \hat{\mu}}) + \varepsilon) = 0 \quad \text{for any } \varepsilon > 0$$

and

$$\limsup_{p \rightarrow \infty} [R_p(\theta, \hat{\theta}^{\text{PM}}) - R_p(\theta, \hat{\theta}^{\hat{\gamma}, \hat{\mu}})] \leq 0.$$

In the case of shrinking toward the grand mean \bar{Y} (when p , the number of Y_i 's, is small or moderate), we have the following parametric results parallel to the semiparametric ones.

First, for the grand-mean shrinkage estimator

$$(4.3) \quad \hat{\theta}^{\gamma, \bar{Y}} = \frac{\tau_i}{\tau_i + \gamma} \cdot Y_i + \frac{\gamma}{\tau_i + \gamma} \cdot \bar{Y},$$

with a fixed γ , an unbiased estimate of its risk is

$$\text{URE}^{\text{PG}}(\gamma) = \frac{1}{p} \sum_{i=1}^p \left[\frac{\gamma^2}{(\tau_i + \gamma)^2} (Y_i - \bar{Y})^2 + \left(1 - 2 \left(1 - \frac{1}{p} \right) \frac{\gamma}{\tau_i + \gamma} \right) \frac{V(Y_i)}{\tau_i + \nu_2} \right].$$

Minimizing it yields our parametric URE grand-mean shrinkage estimator

$$(4.4) \quad \hat{\theta}_i^{\text{PG}} = \frac{\tau_i}{\tau_i + \hat{\gamma}^{\text{PG}}} \cdot Y_i + \frac{\hat{\gamma}^{\text{PG}}}{\tau_i + \hat{\gamma}^{\text{PG}}} \cdot \bar{Y},$$

where

$$\hat{\gamma}^{\text{PG}} = \arg \min_{0 \leq \gamma \leq \infty} \text{URE}^{\text{PG}}(\gamma).$$

Similar to Theorems 4.6 and 4.7, the next two results show that in the case of shrinking toward the grand mean under the formula (4.3), the parametric URE grand-mean shrinkage estimator is asymptotically optimal.

THEOREM 4.8. *Let $Y_i \stackrel{\text{ind.}}{\sim} \text{NEF-QVF}[\theta_i, V(\theta_i)/\tau_i]$, $i = 1, \dots, p$, be non-Gaussian. Under the respective conditions listed in Table 1, we have*

$$\sup_{0 \leq \gamma < \infty} |\text{URE}^{\text{PG}}(\gamma) - l_p(\boldsymbol{\theta}, \hat{\boldsymbol{\theta}}^{\gamma, \bar{Y}})| \rightarrow 0 \quad \text{in } L^1 \text{ and in probability, as } p \rightarrow \infty.$$

THEOREM 4.9. *Let $Y_i \stackrel{\text{ind.}}{\sim} \text{NEF-QVF}[\theta_i, V(\theta_i)/\tau_i]$, $i = 1, \dots, p$, be non-Gaussian. Assume the respective conditions listed in Table 1. Then for any estimator $\hat{\boldsymbol{\theta}}^{\hat{\gamma}, \bar{Y}} = \frac{\boldsymbol{\tau}}{\boldsymbol{\tau} + \hat{\gamma}} \mathbf{Y} + \frac{\hat{\gamma}}{\boldsymbol{\tau} + \hat{\gamma}} \bar{Y}$, where $\hat{\gamma} \geq 0$, we have*

$$\lim_{p \rightarrow \infty} P(l_p(\boldsymbol{\theta}, \hat{\boldsymbol{\theta}}^{\text{PG}}) \geq l_p(\boldsymbol{\theta}, \hat{\boldsymbol{\theta}}^{\hat{\gamma}, \bar{Y}}) + \varepsilon) = 0 \quad \text{for any } \varepsilon > 0$$

and

$$\liminf_{p \rightarrow \infty} [R_p(\boldsymbol{\theta}, \hat{\boldsymbol{\theta}}^{\text{PG}}) - R_p(\boldsymbol{\theta}, \hat{\boldsymbol{\theta}}^{\hat{\gamma}, \bar{Y}})] \leq 0.$$

5. Simulation study. In this section, we conduct a number of simulations to investigate the performance of the URE estimators and compare their performance to that of other existing shrinkage estimators. For each simulation, we first draw (θ_i, τ_i) , $i = 1, \dots, p$, independently from a distribution and then generate Y_i given (θ_i, τ_i) . This process is repeated a large number of times ($N = 100,000$) to obtain an accurate estimate of the risk for each estimator. The sample size p is chosen to vary from 20 to 500 at an interval of length 20. For notational convenience, in this section, we write $A_i = 1/\tau_i$ so that A_i is (essentially) the variance.

5.1. *Location-scale family.* For the location-scale families, we consider three non-Gaussian cases: the Laplace [where the standard variate Z has density $f(z) = \frac{1}{2} \exp(-|z|)$], logistic [where the standard variate Z has density $f(z) = e^{-z}/(1 + e^{-z})^2$] and Student- t distributions with 7 degrees of freedom. To evaluate the performance of the semiparametric URE estimator $\hat{\theta}^{\text{SM}}$, we compare it to the naive estimator

$$\hat{\theta}_i^{\text{Naive}} = Y_i$$

and the extended James–Stein estimator

$$(5.1) \quad \hat{\theta}_i^{\text{JS}^+} = \hat{\mu}^{\text{JS}^+} + \left(1 - \frac{(p-3)}{\sum_{i=1}^p (Y_i - \hat{\mu}^{\text{JS}^+})^2 / A_i}\right)^+ \cdot Y_i,$$

where $A_i = 1/\tau_i$ and $\hat{\mu}^{\text{JS}^+} = \sum_{i=1}^p (X_i/A_i) / \sum_{i=1}^p 1/A_i$.

For each of the three distributions we study four different setups to generate $(\theta_i, A_i = 1/\tau_i)$ for $i = 1, \dots, p$. We then generate Y_i via (3.1) except for scenario (4) below.

Scenario (1). (θ_i, A_i) are drawn from $A_i \sim \text{Unif}(0.1, 1)$ and $\theta_i \sim N(0, 1)$ independently. In this scenario, the location and scale are independent of each other. Panels (a) in Figures 1–3 plot the performance of the three estimators. The risk function of the naive estimator $\hat{\theta}_i^{\text{Naive}}$, being a constant for all p , is way above the other two. The risk of the semiparametric URE estimator is significantly smaller than that of both the extended James–Stein estimator and the naive estimator, particularly so when the sample size $p > 40$.

Scenario (2). (θ_i, A_i) are drawn from $A_i \sim \text{Unif}(0.1, 1)$ and $\theta_i = A_i$. This scenario tests the case that the location and scale have a strong correlation. Panels (b) in Figures 1–3 show the performance of the three estimators. The risk of the semiparametric URE estimator is significantly smaller than that of both the extended James–Stein estimator and the native estimator. The naive estimator $\hat{\theta}_i^{\text{Naive}}$, with a constant risk, performs the worst. This example indicates that the semiparametric URE estimator behaves robustly well even when there is a strong correlation between the location and the scale. This is because the semiparametric URE estimator does not make any assumption on the relationship between θ_i and τ_i .

Scenario (3). (θ_i, A_i) are drawn such that $A_i \sim \frac{1}{2} \cdot 1_{\{A_i=0.1\}} + \frac{1}{2} \cdot 1_{\{A_i=0.5\}}$ —that is, A_i is 0.1 or 0.5 with 50% probability each—and that conditioning on A_i being 0.1 or 0.5, $\theta_i | A_i = 0.1 \sim N(2, 0.1)$; $\theta_i | A_i = 0.5 \sim N(0, 0.5)$. This scenario tests the case that there are two underlying groups in the data. Panels (c) in Figures 1–3 compare the performance of the semiparametric URE estimator to that of the native estimator and the extended James–Stein estimator. The semiparametric URE estimator is seen to significantly outperform the other two estimators.

Scenario (4). (θ_i, A_i) are drawn from $A_i \sim \text{Unif}(0.1, 1)$ and $\theta_i = A_i$. Given θ_i and A_i , Y_i are generated from $Y_i \sim \text{Unif}(\theta_i - \sqrt{3A_i}\sigma, \theta + \sqrt{3A_i}\sigma)$, where σ is the standard deviation of the standard variate Z , that is, $\sigma = \sqrt{2}$, $\pi/\sqrt{3}$ and $\sqrt{7/5}$

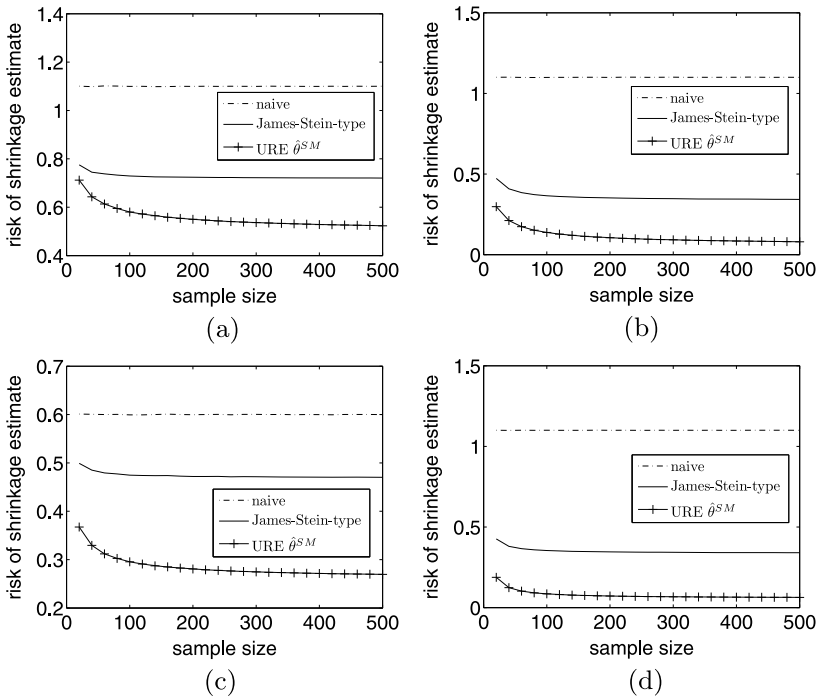


FIG. 1. Comparison of the risks of different shrinkage estimators for the Laplace case. (a) $A \sim \text{Unif}(0.1, 1)$, $\theta \sim N(0, 1)$ independently; $Y = \theta + \sqrt{A} \cdot Z$. (b) $A \sim \text{Unif}(0.1, 1)$, $\theta = A$; $Y = \theta + \sqrt{A} \cdot Z$. (c) $A \sim \frac{1}{2} \cdot 1_{\{A=0.1\}} + \frac{1}{2} \cdot 1_{\{A=0.5\}}$, $\theta|A = 0.1 \sim N(2, 0.1)$, $\theta|A = 0.5 \sim N(0, 0.5)$; $Y = \theta + \sqrt{A} \cdot Z$. (d) $A \sim \text{Unif}(0.1, 1)$, $\theta = A$; $Y \sim \text{Unif}[\theta - \sqrt{6A}, \theta + \sqrt{6A}]$.

for the Laplace, logistics and t -distribution ($df = 7$), respectively. This scenario tests the case of model mis-specification and, hence, the robustness of the estimators. Note that Y_i is drawn from a uniform distribution, not from the Laplace, logistic or t distribution. Panels (d) in Figures 1–3 show the performance of the three estimators. It is seen that the naive estimator behaves the worst and that the semiparametric URE estimator clearly outperforms the other two. This example indicates the robust performance of the semiparametric URE estimator even when the model is incorrectly specified. This is because URE estimator essentially only involves the first two moments of Y_i ; it does not rely on the specific density function of the distribution.

5.2. Exponential family. We consider exponential family in this subsection, conducting simulation evaluations on the beta-binomial and Poisson-gamma models.

5.2.1. Beta-binomial hierarchical model. For binomial observations $Y_i \stackrel{\text{ind.}}{\sim} \text{Bin}(\tau_i, \theta_i)/\tau_i$, classical hierarchical inference typically assumes the conjugate

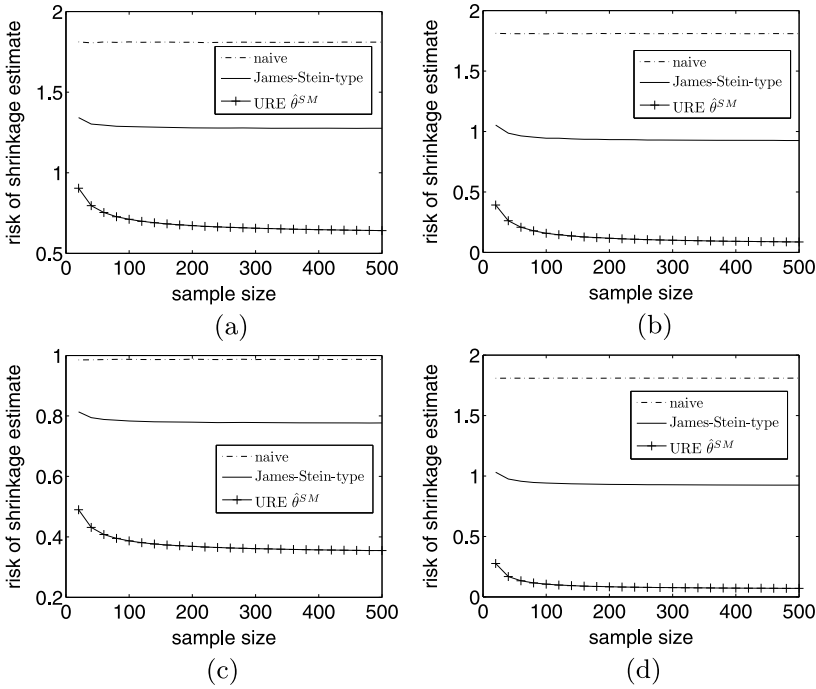


FIG. 2. Comparison of the risks of different shrinkage estimators for the logistic case. (a) $A \sim \text{Unif}(0.1, 1)$, $\theta \sim N(0, 1)$ independently; $Y = \theta + \sqrt{A} \cdot Z$. (b) $A \sim \text{Unif}(0.1, 1)$, $\theta = A$; $Y = \theta + \sqrt{A} \cdot Z$. (c) $A \sim \frac{1}{2} \cdot I_{\{A=0.1\}} + \frac{1}{2} \cdot I_{\{A=0.5\}}$, $\theta|A=0.1 \sim N(2, 0.1)$, $\theta|A=0.5 \sim N(0, 0.5)$; $Y = \theta + \sqrt{A} \cdot Z$. (d) $A \sim \text{Unif}(0.1, 1)$, $\theta = A$; $Y \sim \text{Unif}[\theta - \pi\sqrt{A}, \theta + \pi\sqrt{A}]$.

prior $\theta_i \stackrel{\text{i.i.d.}}{\sim} \text{Beta}(\alpha, \beta)$. The marginal distribution of Y_i is used by classical empirical Bayes methods to estimate the hyper-parameters. Plugging the estimate of these hyper-parameters into the posterior mean of θ_i given Y_i yields the empirical Bayes estimate of θ_i . In this subsection, we consider both the semiparametric URE estimator $\hat{\theta}^{\text{SM}}$ [equation (2.5)] and the parametric URE estimator $\hat{\theta}^{\text{PM}}$ [equation (4.2)], and compare them with the parametric empirical Bayes maximum likelihood estimator $\hat{\theta}^{\text{ML}}$ and the parametric empirical Bayes method-of-moment estimator $\hat{\theta}^{\text{MM}}$. The empirical Bayes maximum likelihood estimator $\hat{\theta}^{\text{ML}}$ is given by

$$\hat{\theta}_i^{\text{ML}} = \frac{\tau_i}{\tau_i + \hat{\gamma}^{\text{ML}}} \cdot Y_i + \frac{\hat{\gamma}^{\text{ML}}}{\tau_i + \hat{\gamma}^{\text{ML}}} \cdot \hat{\mu}^{\text{ML}},$$

where $(\hat{\gamma}^{\text{ML}}, \hat{\mu}^{\text{ML}})$ maximizes the marginal likelihood of Y_i :

$$(\hat{\gamma}^{\text{ML}}, \hat{\mu}^{\text{ML}}) = \arg \max_{\gamma \geq 0, \mu} \prod_i \frac{\Gamma(\gamma\mu + \tau_i y_i) \Gamma(\gamma(1 - \mu) + (1 - y_i)\tau_i) \Gamma(\gamma)}{\Gamma(\gamma + \tau_i) \Gamma(\gamma\mu) \Gamma(\gamma(1 - \mu))},$$

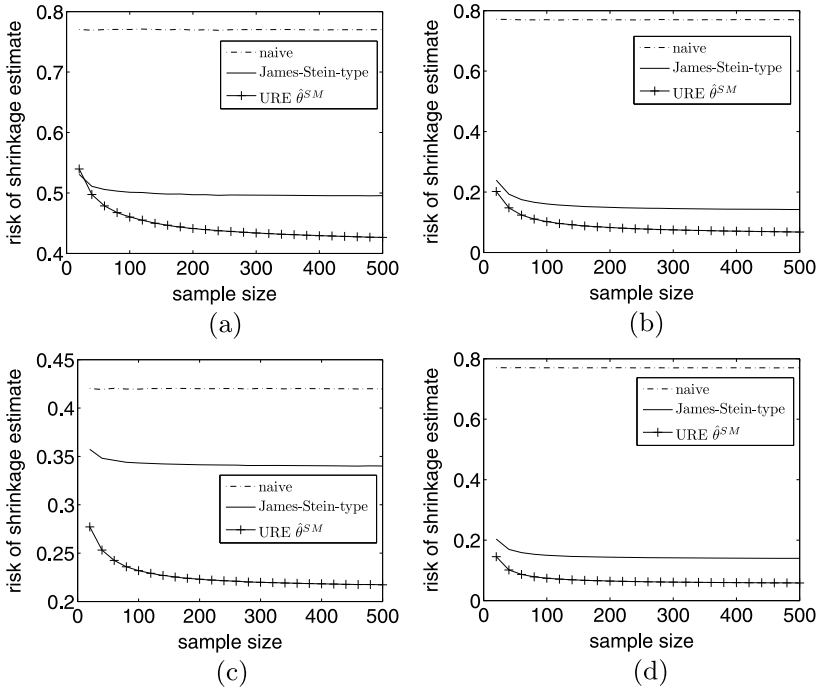


FIG. 3. Comparison of the risks of different shrinkage estimators for the t -distribution ($df = 7$). (a) $A \sim \text{Unif}(0.1, 1)$, $\theta \sim N(0, 1)$ independently; $Y = \theta + \sqrt{A} \cdot Z$. (b) $A \sim \text{Unif}(0.1, 1)$, $\theta = A$; $Y = \theta + \sqrt{A} \cdot Z$. (c) $A \sim \frac{1}{2} \cdot 1_{\{A=0.1\}} + \frac{1}{2} \cdot 1_{\{A=0.5\}}$, $\theta|A = 0.1 \sim N(2, 0.1)$, $\theta|A = 0.5 \sim N(0, 0.5)$; $Y = \theta + \sqrt{A} \cdot Z$. (d) $A \sim \text{Unif}(0.1, 1)$, $\theta = A$; $Y \sim \text{Unif}[\theta - \sqrt{21A/5}, \theta + \sqrt{21A/5}]$.

where $\mu = \alpha/(\alpha + \beta)$ and $\gamma = \alpha + \beta$ as in Table 2. Likewise, the empirical Bayes method-of-moment estimator $\hat{\theta}^{\text{MM}}$ is given by

$$\hat{\theta}_i^{\text{MM}} = \frac{\tau_i}{\tau_i + \hat{\gamma}^{\text{MM}}} \cdot Y_i + \frac{\hat{\gamma}^{\text{MM}}}{\tau_i + \hat{\gamma}^{\text{MM}}} \cdot \hat{\mu}^{\text{MM}},$$

where

$$\hat{\mu}^{\text{MM}} = \bar{Y} = \frac{1}{p} \sum_{i=1}^p Y_i,$$

$$\hat{\gamma}^{\text{MM}} = \frac{\bar{Y}(1 - \bar{Y}) \cdot \sum_{i=1}^p (1 - 1/\tau_i)}{[\sum_{i=1}^p (Y_i^2 - \bar{Y}/\tau_i - \bar{Y}^2(1 - 1/\tau_i))]^{+}}.$$

There are in total four different simulation setups in which we study the four different estimators. In addition, in each case, we also calculate the oracle *risk* “estimator” $\tilde{\theta}^{\text{OR}}$, defined as

$$(5.2) \quad \tilde{\theta}_i^{\text{OR}} = \frac{\tau_i}{\tau_i + \tilde{\gamma}^{\text{OR}}} \cdot Y_i + \frac{\tilde{\gamma}^{\text{OR}}}{\tau_i + \tilde{\gamma}^{\text{OR}}} \cdot \tilde{\mu}^{\text{OR}},$$

where

$$\begin{aligned} (\tilde{\gamma}^{\text{OR}}, \tilde{\mu}^{\text{OR}}) &= \arg \min_{\gamma \geq 0, \mu} R_p(\boldsymbol{\theta}, \hat{\boldsymbol{\theta}}^{\gamma, \mu}) \\ &= \arg \min_{\gamma \geq 0, \mu} \sum_{i=1}^p \frac{1}{p} E \left[\left(\frac{\tau_i}{\tau_i + \gamma} Y_i + \frac{\gamma}{\tau_i + \gamma} \mu - \theta_i \right)^2 \right]. \end{aligned}$$

Clearly, the oracle risk estimator $\tilde{\boldsymbol{\theta}}^{\text{OR}}$ cannot be used in practice, since it depends on the unknown $\boldsymbol{\theta}$, but it does provide a sensible lower bound of the risk achievable by any shrinkage estimator with the given parametric form.

EXAMPLE 1. We generate $\tau_i \sim \text{Poisson}(3) + 2$ and $\theta_i \sim \text{Beta}(1, 1)$ independently, and draw $Y_i \sim \text{Bin}(\tau_i, \theta_i)/\tau_i$. The oracle estimator $\tilde{\boldsymbol{\theta}}^{\text{OR}}$ is found to have $\tilde{\gamma}^{\text{OR}} = 2$ and $\tilde{\mu}^{\text{OR}} = 0.5$. The corresponding risk for the oracle estimator is numerically found to be $R_p(\boldsymbol{\theta}, \tilde{\boldsymbol{\theta}}^{\text{OR}}) \approx 0.0253$. The plot in Figure 4(a) shows the risks of the five shrinkage estimators as the sample size p varies. It is seen that the performance of all four shrinkage estimators approaches that of the parametric oracle estimator, the “best estimator” one can hope to get under the parametric form. Note that the two empirical Bayes estimators converges to the oracle estimator faster than the two URE shrinkage estimators. This is because the hierarchical distribution on τ_i and θ_i are exactly the one assumed by the empirical Bayes estimators. In contrast, the URE estimators do not make any assumption on the hierarchical distribution but still achieve rather competitive performance. When the sample size is moderately large, all four estimators well approach the limit given by the parametric oracle estimator.

EXAMPLE 2. We generate $\tau_i \sim \text{Poisson}(3) + 2$ and $\theta_i \sim \frac{1}{2} \text{Beta}(1, 3) + \frac{1}{2} \text{Beta}(3, 1)$ independently, and draw $Y_i \sim \text{Bin}(\tau_i, \theta_i)/\tau_i$. In this example, θ_i no longer comes from a beta distribution, but θ_i and τ_i are still independent. The oracle estimator is found to have $\tilde{\gamma}^{\text{OR}} \approx 1.5$ and $\tilde{\mu}^{\text{OR}} = 0.5$. The corresponding risk for the oracle estimator $\hat{\boldsymbol{\theta}}^{\tau_0, \theta}$ is $R_p(\boldsymbol{\theta}, \tilde{\boldsymbol{\theta}}^{\text{OR}}) \approx 0.0248$. The plot in Figure 4(b) shows the risks of the five shrinkage estimators as the sample size p varies. Again, as p gets large, the performance of all shrinkage estimators eventually approaches that of the oracle estimator. This observation indicates that the parametric form of the prior on θ_i is not crucial as long as τ_i and θ_i are independent.

EXAMPLE 3. We generate $\tau_i \sim \text{Poisson}(3) + 2$ and let $\theta_i = 1/\tau_i$, and then we draw $Y_i \sim \text{Bin}(\tau_i, \theta_i)/\tau_i$. In this case, there is a (negative) correlation between θ_i and τ_i . The parametric oracle estimator is found to have $\tilde{\gamma}^{\text{OR}} \approx 23.0898$ and $\tilde{\mu}^{\text{OR}} \approx 0.2377$ numerically; the corresponding risk is $R_p(\boldsymbol{\mu}, \tilde{\boldsymbol{\theta}}^{\text{OR}}) \approx 0.0069$. The plot in Figure 4(c) shows the risks of the five shrinkage estimators as functions of the sample size p . Unlike the previous examples, the two empirical Bayes estimators no longer converge to the parametric oracle estimator, that is, the limit of their

risk (as $p \rightarrow \infty$) is *strictly* above the risk of the parametric oracle estimator. On the other hand, the risk of the parametric URE estimator $\hat{\theta}^{\text{PM}}$ still converges to the risk of the parametric oracle estimator. It is interesting to note that the limiting risk of the semiparametric URE estimators $\hat{\theta}^{\text{SM}}$ is actually *strictly* smaller than the risk of the parametric oracle estimator (although the difference between the two is not easy to spot due to the scale of the plot).

EXAMPLE 4. We generate (τ_i, θ_i, Y_i) as follows. First, we draw I_i from Bernoulli(1/2), and then generate $\tau_i \sim I_i \cdot \text{Poisson}(10) + (1 - I_i) \cdot \text{Poisson}(1) + 2$ and $\theta_i \sim I_i \cdot \text{Beta}(1, 3) + (1 - I_i) \cdot \text{Beta}(3, 1)$. Given (τ_i, θ_i) , we draw $Y_i \sim \text{Bin}(\tau_i, \theta_i)/\tau_i$. In this example, there exist two groups in the data (indexed by I_i). It thus serves to test the different estimators in the presence of grouping. The parametric oracle estimator is found to have $\tilde{\gamma}^{\text{OR}} \approx 0.3108$ and $\tilde{\mu}^{\text{OR}} \approx 2.0426$; the corresponding risk is $R_p(\mu, \tilde{\theta}^{\text{OR}}) \approx 0.0201$. Figure 4(d) plots the risks of the five shrinkage estimators versus the sample size p . The two empirical Bayes estimators clearly encounter much greater risk than the URE estimators, and the limiting risks of the two empirical Bayes estimators are significantly larger than the risk of the parametric oracle estimator. The risk of the parametric URE estimator $\hat{\theta}^{\text{PM}}$ converges to that of the parametric oracle estimator. It is quite noteworthy that the semiparametric URE estimator $\hat{\theta}^{\text{SM}}$ achieves a significant improvement over the parametric oracle one.

5.2.2. *Poisson–Gamma hierarchical model.* For Poisson observations $Y_i \stackrel{\text{ind.}}{\sim} \text{Poisson}(\tau_i\theta_i)/\tau_i$, the conjugate prior is $\theta_i \stackrel{\text{i.i.d.}}{\sim} \Gamma(\alpha, \lambda)$. Like in the previous subsection, we compare five estimators: the empirical Bayes maximum likelihood estimator, the empirical Bayes method-of-moment estimator, the parametric and semiparametric URE estimators and the parametric oracle “estimator” (5.2). The empirical Bayes maximum likelihood estimator $\hat{\theta}^{\text{ML}}$ is given by

$$\hat{\theta}_i^{\text{ML}} = \frac{\tau_i}{\tau_i + \hat{\gamma}^{\text{ML}}} \cdot Y_i + \frac{\hat{\gamma}^{\text{ML}}}{\tau_i + \hat{\gamma}^{\text{ML}}} \cdot \hat{\mu}^{\text{ML}},$$

where $(\hat{\gamma}^{\text{ML}}, \hat{\mu}^{\text{ML}})$ maximizes the marginal likelihood of Y_i :

$$(\hat{\gamma}^{\text{ML}}, \hat{\mu}^{\text{ML}}) = \arg \max_{\gamma \geq 0, \mu} \prod_i \frac{\gamma^{\gamma\mu} \Gamma(\gamma\mu + \tau_i y_i)}{(\tau_i + \gamma)^{\tau_i y_i + \gamma\mu} \Gamma(\gamma\mu)},$$

where $\mu = \alpha\lambda$ and $\gamma = 1/\lambda$ as in Table 2. The empirical Bayes method-of-moment estimator $\hat{\theta}^{\text{MM}}$ is given by

$$\hat{\theta}_i^{\text{MM}} = \frac{\tau_i}{\tau_i + \hat{\gamma}^{\text{MM}}} \cdot Y_i + \frac{\hat{\gamma}^{\text{MM}}}{\tau_i + \hat{\gamma}^{\text{MM}}} \cdot \hat{\mu}^{\text{MM}},$$

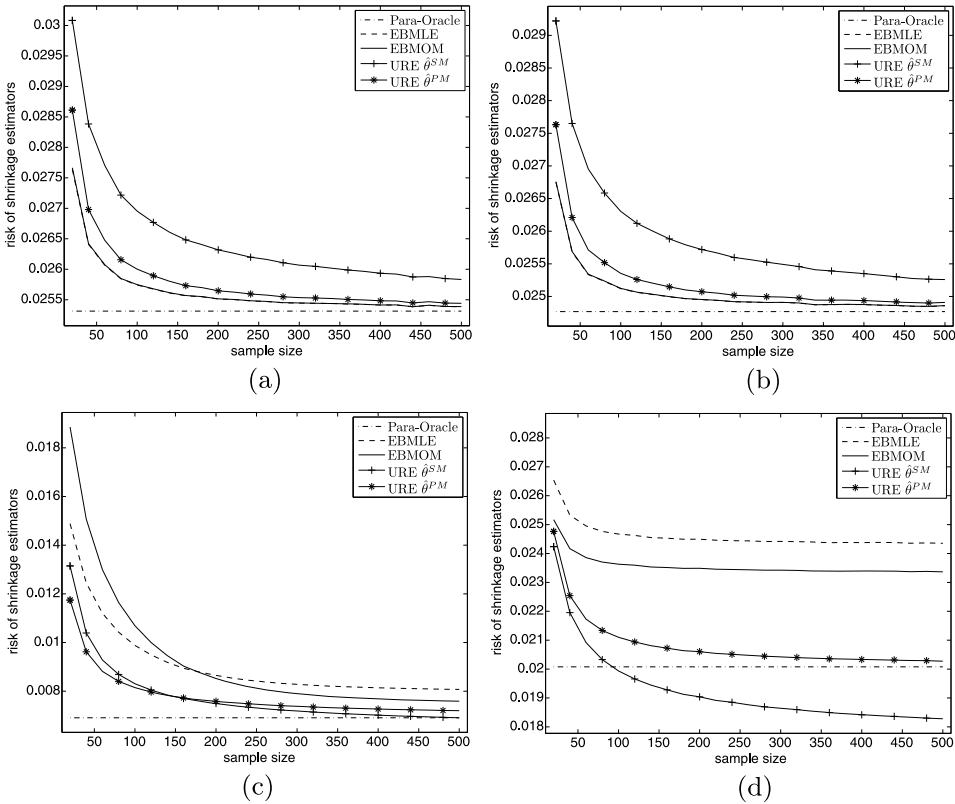


FIG. 4. Comparison of the risks of shrinkage estimators in the Beta-Binomial hierarchical models. (a) $\tau \sim \text{Poisson}(3) + 2$, $\theta \sim \text{Beta}(1, 1)$ independently; $Y \sim \text{Bin}(\tau, \theta)/\tau$. (b) $\tau \sim \text{Poisson}(3) + 2$, $\theta \sim \frac{1}{2} \cdot \text{Beta}(1, 3) + \frac{1}{2} \cdot \text{Beta}(3, 1)$ independently; $Y \sim \text{Bin}(\tau, \theta)/\tau$. (c) $\tau \sim \text{Poisson}(3) + 2$, $\theta = 1/\tau$, $Y \sim \text{Bin}(\tau, \theta)/\tau$. (d) $I \sim \text{Bern}(1/2)$, $\tau \sim I \cdot \text{Poisson}(10) + (1 - I) \cdot \text{Poisson}(1) + 2$, $\theta \sim I \cdot \text{Beta}(1, 3) + (1 - I) \cdot \text{Beta}(3, 1)$; $Y \sim \text{Bin}(\tau, \theta)/\tau$.

where

$$\hat{\mu}^{\text{MM}} = \bar{Y} = \frac{1}{p} \sum_{i=1}^p Y_i,$$

$$\hat{\gamma}^{\text{MM}} = \frac{p \cdot \bar{Y}}{[\sum_{i=1}^p (Y_i^2 - \bar{Y}/\tau_i - \bar{Y}^2)]^+}.$$

We consider four different simulation settings.

EXAMPLE 5. We generate $\tau_i \sim \text{Poisson}(3) + 2$ and $\theta_i \sim \Gamma(1, 1)$ independently, and draw $Y_i \sim \text{Poisson}(\tau_i \theta_i)/\tau_i$. The plot in Figure 5(a) shows the risks of the five shrinkage estimators as the sample size p varies. Clearly, the performance of all shrinkage estimators approaches that of the parametric oracle estimator. As

in the beta-binomial case, the two empirical Bayes estimators converge to the oracle estimator faster than the two URE shrinkage estimators. Again, this is because the hierarchical distribution on τ_i and θ_i are exactly the one assumed by the empirical Bayes estimators. The URE estimators, without making any assumption on the hierarchical distribution, still achieve rather competitive performance.

EXAMPLE 6. We generate $\tau_i \sim \text{Poisson}(3) + 2$ and $\theta_i \sim \text{Unif}(0.1, 1)$ independently, and draw $Y_i \sim \text{Poisson}(\tau_i \theta_i) / \tau_i$. In this setting, θ_i no longer comes from a gamma distribution, but θ_i and τ_i are still independent. The plot in Figure 5(b) shows the risks of the five shrinkage estimators as the sample size p varies. As p gets large, the performance of all shrinkage estimators eventually approaches that of the oracle estimator. Like in the beta-binomial case, the picture indicates that the parametric form of the prior on θ_i is not crucial as long as τ_i and θ_i are independent.

EXAMPLE 7. We generate $\tau_i \sim \text{Poisson}(3) + 2$ and let $\theta_i = 1/\tau_i$, and then we draw $Y_i \sim \text{Poisson}(\tau_i \theta_i) / \tau_i$. In this setting, there is a (negative) correlation between θ_i and τ_i . The plot in Figure 5(c) shows the risks of the five shrinkage estimators as functions of the sample size p . Unlike the previous two examples, the two empirical Bayes estimators no longer converge to the parametric oracle estimator—the limit of their risk is *strictly* above the risk of the parametric oracle estimator. The risk of the parametric URE estimator $\hat{\theta}^{\text{PM}}$, on the other hand, still converges to the risk of the parametric oracle estimator. The limiting risk of the semiparametric URE estimators $\hat{\theta}^{\text{SM}}$ is actually *strictly* smaller than the risk of the parametric oracle estimator (although it is not easy to spot it on the plot).

EXAMPLE 8. We generate (τ_i, θ_i) by first drawing $I_i \sim \text{Bernoulli}(1/2)$ and then $\tau_i \sim I_i \cdot \text{Poisson}(10) + (1 - I_i) \cdot \text{Poisson}(1) + 2$ and $\theta_i \sim I_i \cdot \Gamma(1, 1) + (1 - I_i) \cdot \Gamma(5, 1)$. With (τ_i, θ_i) obtained, we draw $Y_i \sim \text{Poisson}(\tau_i \theta_i) / \tau_i$. This setting tests the case that there is grouping in the data. Figure 5(d) plots the risks of the five shrinkage estimators versus the sample size p . It is seen that the two empirical Bayes estimators have the largest risk, and that the parametric URE estimator $\hat{\theta}^{\text{PM}}$ achieves the risk of the parametric oracle estimator in the limit. The semiparametric URE estimator $\hat{\theta}^{\text{SM}}$ notably outperforms the parametric oracle estimator, when $p > 100$.

6. Application to the prediction of batting average. In this section, we apply the URE shrinkage estimators to a baseball data set, collected and discussed in Brown (2008). This data set consists of the batting records for all the Major League Baseball players in the season of 2005. Following Brown (2008), the data are divided into two half seasons; the goal is to use the data from the first half

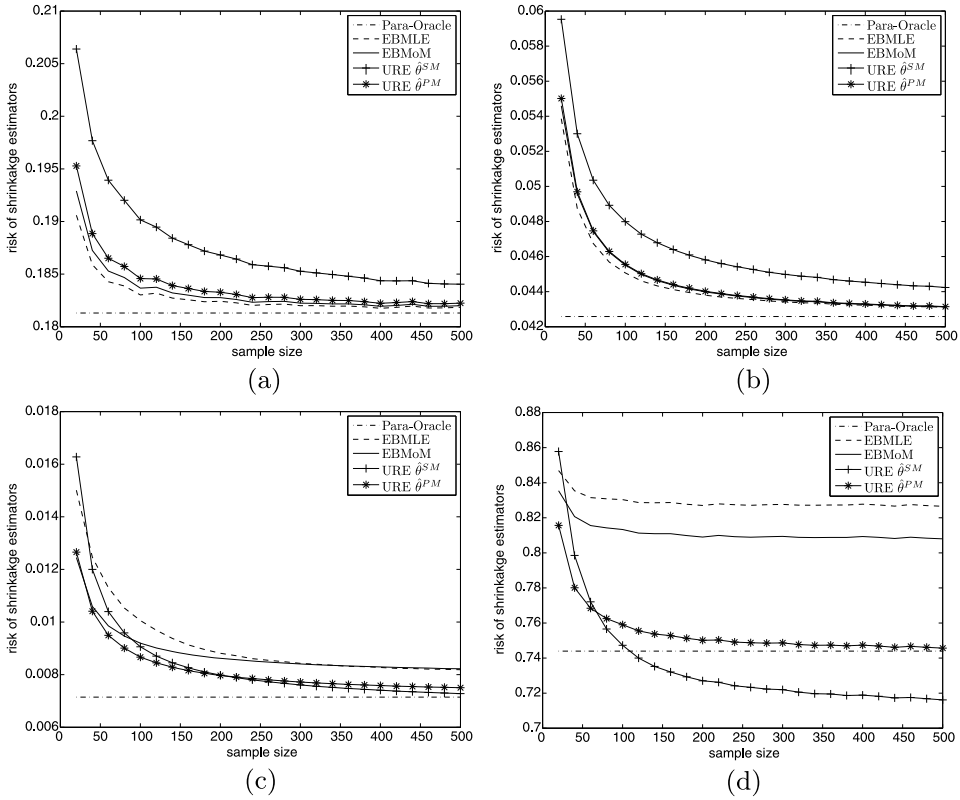


FIG. 5. Comparison of the risks of shrinkage estimators in Poisson–Gamma hierarchical model. (a) $\tau \sim \text{Poisson}(3) + 2$, $\theta \sim \Gamma(1, 1)$ independently; $Y \sim \text{Poisson}(\tau\theta)/\tau$. (b) $\tau \sim \text{Poisson}(3) + 2$, $\theta \sim \text{Unif}(0.1, 1)$ independently; $Y \sim \text{Poisson}(\tau\theta)/\tau$. (c) $\tau \sim \text{Poisson}(3) + 2$, $\theta = 1/\tau$, $Y \sim \text{Poisson}(\tau\theta)/\tau$. (d) $I \sim \text{Bern}(1/2)$, $\tau \sim I \cdot \text{Poisson}(10) + (1 - I) \cdot \text{Poisson}(1) + 2$, $\theta \sim I \cdot \text{Gamma}(1, 1) + (1 - I) \cdot \text{Gamma}(5, 1)$; $Y \sim \text{Poisson}(\tau\theta)/\tau$.

season to predict the players' batting average in the second half season. The prediction can then be compared against the actual record of the second half season. The performance of different estimators can thus be directly evaluated.

For each player, let the number of at-bats be N and the successful number of batting be H ; we then have

$$H_{ij} \sim \text{Binomial}(N_{ij}, p_j),$$

where $i = 1, 2$ is the season indicator, $j = 1, 2, \dots$, is the player indicator, and p_j corresponds to the player's hitting ability. Let Y_{ij} be the observed proportion:

$$Y_{ij} = H_{ij}/N_{ij}.$$

For this binomial setup, we apply our method to obtain the semiparametric URE estimators \hat{p}^{SM} and \hat{p}^{SG} , defined in (2.5) and (2.7), respectively, and the parametric URE estimators \hat{p}^{PM} and \hat{p}^{PG} , defined in (4.2) and (4.4), respectively.

To compare the prediction accuracy of different estimators, we note that most shrinkage estimators in the literature assume normality of the underlying data. Therefore, for sensible evaluation of different methods, we can apply the following variance-stabilizing transformation as discussed in [Brown \(2008\)](#):

$$X_{ij} = \arcsin \sqrt{\frac{H_{ij} + 1/4}{N_{ij} + 1/2}},$$

which gives

$$X_{ij} \sim N\left(\theta_j, \frac{1}{4N_{ij}}\right), \quad \theta_j = \arcsin(\sqrt{p_j}).$$

To evaluate an estimator $\hat{\theta}$ based on the transformed X_{ij} , we measure the total sum of squared prediction errors (TSE) as

$$\text{TSE}(\hat{\theta}) = \sum_j (X_{2j} - \hat{\theta}_j)^2 - \sum_j \frac{1}{4N_{2j}}.$$

To conform to the above transformation (as used by most shrinkage estimators), we calculate $\hat{\theta}_j = \arcsin(\sqrt{\hat{p}_j})$, where \hat{p}_j is a URE estimator of the binomial probability, so that the TSE of the URE estimators can be calculated and compared with other (normality based) shrinkage estimators.

Table 3 below summarizes the numerical results of our URE estimators with a collection of competing shrinkage estimators. The values reported are the ratios of the error of a given estimator to that of the benchmark naive estimator, which simply uses the first half season X_{1j} to predict the second half X_{2j} . All shrinkage estimators are applied three times—to all the baseball players, the pitchers only, and the nonpitchers only. The first group of shrinkage estimators in Table 3 are the classical ones based on normal theory: two empirical Bayes methods (applied to X_{1j}), the grand mean and the extended James–Stein estimator (5.1). The second group includes a number of more recently developed methods: the nonparametric shrinkage methods in [Brown and Greenshtein \(2009\)](#), the binomial mixture model in [Muralidharan \(2010\)](#) and the weighted least squares and general maximum likelihood estimators (with or without the covariate of at-bats effect) in [Jiang and Zhang \(2009, 2010\)](#). The numerical results for these methods are from [Brown \(2008\)](#), [Muralidharan \(2010\)](#) and [Jiang and Zhang \(2009, 2010\)](#). The last group corresponds to the results from our binomial URE methods: the first two are the parametric methods and the last two are the semiparametric ones.

It is seen that our URE shrinkage estimators, especially the semiparametric ones, achieve very competitive prediction result among all the estimators. We think the primary reason is that the baseball data contain unique features that violate the underlying assumptions of the classical empirical Bayes methods. Both the normal prior assumption and the implicit assumption of the uncorrelatedness between the

TABLE 3
Prediction errors of batting averages for the baseball data

	ALL	Pitchers	Non-Pitchers
Naive	1	1	1
Grand mean \bar{X}_1	0.852	0.127	0.378
Parametric EB-MM	0.593	0.129	0.387
Parametric EB-ML	0.902	0.117	0.398
Extended James–Stein	0.525	0.164	0.359
Nonparametric EB	0.508	0.212	0.372
Binomial mixture	0.588	0.156	0.314
Weighted least square (Null)	1.074	0.127	0.468
Weighted generalized MLE (Null)	0.306	0.173	0.326
Weighted least square (AB)	0.537	0.087	0.290
Weighted generalized MLE (AB)	0.301	0.141	0.261
Parametric URE $\hat{\theta}^{PG}$	0.515	0.105	0.278
Parametric URE $\hat{\theta}^{PM}$	0.421	0.105	0.276
Semiparametric URE $\hat{\theta}^{SG}$	0.414	0.045	0.259
Semiparametric URE $\hat{\theta}^{SM}$	0.422	0.041	0.273

binomial probability p and the sample size τ are not justified here. To illustrate the last point, we present a scatter plot of \log_{10} (number of at bats) versus the observed batting average y for the nonpitcher group in Figure 6.

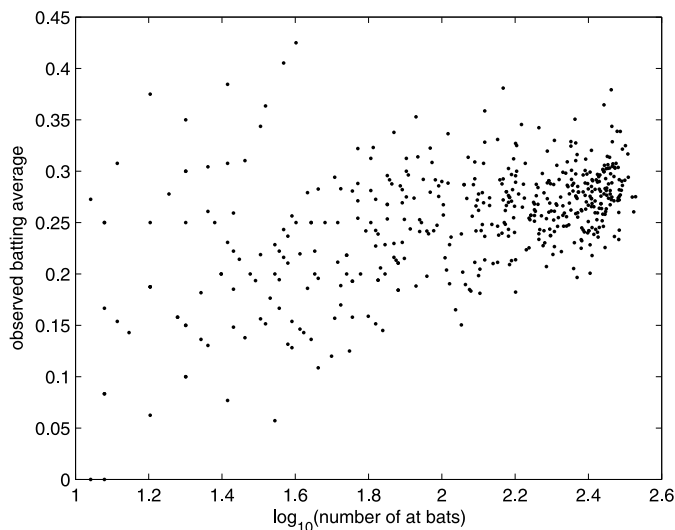


FIG. 6. *Scatter plot of \log_{10} (number of at bats) versus observed batting average*

7. Summary. In this paper, we develop a general theory of URE shrinkage estimation in family of distributions with quadratic variance function. We first discuss a class of semiparametric URE estimator and establish their optimality property. Two specific cases are then carefully studied: the location-scale family and the natural exponential families with quadratic variance function. In the latter case, we also study a class of parametric URE estimators, whose forms are derived from the classical conjugate hierarchical model. We show that each URE shrinkage estimator is asymptotically optimal in its own class and their asymptotic optimality do not depend on the specific distribution assumptions, and more importantly, do not depend on the implicit assumption that the group mean θ and the sample size τ are uncorrelated, which underlies many classical shrinkage estimators. The URE estimators are evaluated in comprehensive simulation studies and one real data set. It is found that the URE estimators offer numerically superior performance compared to the classical empirical Bayes and many other competing shrinkage estimators. The semiparametric URE estimators appear to be particularly competitive.

It is worth emphasizing that the optimality properties of the URE estimators is not in contradiction with well established results on the nonexistence of Stein's paradox in simultaneous inference problems with finite sample space [Gutmann (1982)], since the results we obtained here are asymptotic ones when p approaches infinity. A question that naturally arises here is then how large p needs to be for the URE estimators to become superior compared with their competitors. Even though we did not develop a formal finite-sample theory for such comparison, our comprehensive simulation indicates that p usually does not need to be large— p can be as small as 100—for the URE estimators to achieve competitive performance.

The theory here extends the one on the normal hierarchical models in Xie, Kou and Brown (2012). There are three critical features that make the generalization of previous results possible here: (i) the use of quadratic risk; (ii) the linear form of the shrinkage estimator and (iii) the quadratic variance function of the distribution. The three features together guarantee the existence of an unbiased risk estimate. For the hierarchical models where an unbiased risk estimate does not exist, similar idea can still be applied to some estimate of risk, for example, the bootstrap estimate [Efron (2004)]. However, a theory is in demand to justify the performance of the resulting shrinkage estimators.

It would also be an important area of research to study confidence intervals for the URE estimators obtained here. Understanding whether the optimality of the URE estimators implies any optimality property of the estimators of the hyperparameters under certain conditions is an interesting related question. However, such topics are out of scope for the current paper and we will need to address them in future research.

APPENDIX: PROOFS

PROOF OF THEOREM 2.1. Throughout this proof, when a supremum is taken, it is over $b_i \in [0, 1]$, $|\mu| \leq \max_i |Y_i|$ and Requirement (MON), unless explicitly

stated otherwise. Since

$$\begin{aligned} \text{URE}(\mathbf{b}, \mu) - l(\boldsymbol{\theta}, \hat{\boldsymbol{\theta}}^{\mathbf{b}, \mu}) \\ = \frac{1}{p} \sum_{i=1}^p (1 - 2b_i) \left(\frac{V(Y_i)}{\tau_i + \nu_2} - (Y_i - \theta_i)^2 \right) \\ + \frac{2}{p} \sum_{i=1}^p b_i (Y_i - \theta_i)(\theta_i - \mu), \end{aligned}$$

it follows that

$$\begin{aligned} \sup |\text{URE}(\mathbf{b}, \mu) - l(\boldsymbol{\theta}, \hat{\boldsymbol{\theta}}^{\mathbf{b}, \mu})| &\leq \frac{1}{p} \left| \sum_{i=1}^p \left(\frac{V(Y_i)}{\tau_i + \nu_2} - (Y_i - \theta_i)^2 \right) \right| \\ \text{(A.1)} \quad &+ \frac{2}{p} \sup \left| \sum_{i=1}^p b_i \left(\frac{V(Y_i)}{\tau_i + \nu_2} - (Y_i - \theta_i)^2 \right) \right| \\ &+ \frac{2}{p} \sup \left| \sum_{i=1}^p b_i (Y_i - \theta_i)(\theta_i - \mu) \right|. \end{aligned}$$

For the first term on the right-hand side, we note that

$$\begin{aligned} \frac{V(Y_i)}{\tau_i + \nu_2} - (Y_i - \theta_i)^2 \\ = -\frac{\tau_i}{\tau_i + \nu_2} (Y_i^2 - EY_i^2) + \left(2\theta_i - \frac{\nu_1}{\tau_i + \nu_2} \right) (Y_i - \theta_i). \end{aligned}$$

Thus,

$$\begin{aligned} E \left(\left(\frac{V(Y_i)}{\tau_i + \nu_2} - (Y_i - \theta_i)^2 \right)^2 \right) \\ \leq 2 \left(\left(\frac{\tau_i}{\tau_i + \nu_2} \right)^2 \text{Var}(Y_i^2) + \left(2\theta_i - \frac{\nu_1}{\tau_i + \nu_2} \right)^2 \text{Var}(Y_i) \right) \\ \leq 2 \left(\frac{\tau_i}{\tau_i + \nu_2} \right)^2 \text{Var}(Y_i^2) + 16\theta_i^2 \text{Var}(Y_i) + 4 \left(\frac{\nu_1}{\tau_i + \nu_2} \right)^2 \text{Var}(Y_i). \end{aligned}$$

It follows that by conditions (A)–(D)

$$\text{(A.2)} \quad \frac{1}{p} \sum_{i=1}^p \left(\frac{V(Y_i)}{\tau_i + \nu_2} - (Y_i - \theta_i)^2 \right) \rightarrow 0 \quad \text{in } L^2 \text{ as } p \rightarrow \infty.$$

For the term $\frac{2}{p} \sup |\sum_i b_i (\frac{V(Y_i)}{\tau_i + \nu_2} - (Y_i - \theta_i)^2)|$ in (A.1), without loss of generality, let us assume $\tau_1 \leq \dots \leq \tau_p$; we then know from Requirement (MON) that $b_1 \geq$

... $\geq b_2$. As in Lemma 2.1 in Li (1986), we observe that

$$\begin{aligned} & \sup \frac{2}{p} \left| \sum_{i=1}^p b_i \left(\frac{V(Y_i)}{\tau_i + \nu_2} - (Y_i - \theta_i)^2 \right) \right| \\ &= \sup_{1 \geq b_1 \geq \dots \geq b_p \geq 0} \frac{2}{p} \left| \sum_{i=1}^p b_i \left(\frac{V(Y_i)}{\tau_i + \nu_2} - (Y_i - \theta_i)^2 \right) \right| \\ &= \max_{1 \leq j \leq p} \frac{2}{p} \left| \sum_{i=1}^j \left(\frac{V(Y_i)}{\tau_i + \nu_2} - (Y_i - \theta_i)^2 \right) \right|. \end{aligned}$$

Let $M_j = \sum_{i=1}^j \left(\frac{V(Y_i)}{\tau_i + \nu_2} - (Y_i - \theta_i)^2 \right)$. Then $\{M_j; j = 1, 2, \dots\}$ forms a martingale. The L^p maximum inequality implies

$$E \left(\max_{1 \leq j \leq p} M_j^2 \right) \leq 4E(M_p^2) = 4 \sum_{i=1}^p E \left(\frac{V(Y_i)}{\tau_i + \nu_2} - (Y_i - \theta_i)^2 \right)^2,$$

which implies by (A.2) that

$$(A.3) \quad \sup \frac{2}{p} \left| \sum_{i=1}^p b_i \left(\frac{V(Y_i)}{\tau_i + \nu_2} - (Y_i - \theta_i)^2 \right) \right| \rightarrow 0 \quad \text{in } L^2 \text{ as } p \rightarrow \infty.$$

For the last term $\frac{2}{p} \sup \left| \sum_i b_i (Y_i - \theta_i)(\theta_i - \mu) \right|$ in (A.1), we note that

$$\frac{1}{p} \sum_{i=1}^p b_i (Y_i - \theta_i)(\theta_i - \mu) = \frac{1}{p} \sum_{i=1}^p b_i \theta_i (Y_i - \theta_i) - \frac{\mu}{p} \sum_{i=1}^p b_i (Y_i - \theta_i).$$

Using the same argument as in the proof of (A.3), we can show that

$$\begin{aligned} & \sup \frac{1}{p} \left| \sum_i b_i \theta_i (Y_i - \theta_i) \right| \rightarrow 0 \quad \text{in } L^2, \\ & E \left(\sup \left| \sum_i b_i (Y_i - \theta_i) \right|^2 \right) = O(p). \end{aligned}$$

Applying condition (E) and the Cauchy–Schwarz inequality, we obtain

$$\begin{aligned} & \frac{1}{p} E \left(\sup \left| \mu \sum_{i=1}^p b_i (Y_i - \theta_i) \right| \right) \\ &= \frac{1}{p} E \left(\max_{1 \leq i \leq p} |Y_i| \cdot \sup \left| \sum_i b_i (Y_i - \theta_i) \right| \right) \\ &\leq \frac{1}{p} \left(E \left(\max_{1 \leq i \leq p} Y_i^2 \right) \cdot E \left\{ \sup \left| \sum_i b_i (Y_i - \theta_i) \right|^2 \right\} \right)^{1/2} \\ &= O(p^{(1-\varepsilon+1)/2-1}) = O(p^{-\varepsilon/2}). \end{aligned}$$

Therefore,

$$\frac{1}{p} \sup \left| \mu \sum_{i=1}^p b_i (Y_i - \theta_i) \right| \rightarrow 0 \text{ in } L^1.$$

This completes the proof, since each term on the right-hand side of (A.1) converges to zero in L^1 . \square

PROOF OF THEOREM 2.2. Throughout this proof, when a supremum is taken, it is over $b_i \in [0, 1]$, $|\mu| \leq \max_i |Y_i|$ and Requirement (MON). Note that

$$\text{URE}(\hat{\mathbf{b}}^{\text{SM}}, \hat{\mu}^{\text{SM}}) \leq \text{URE}(\hat{\mathbf{b}}, \hat{\mu})$$

and we know from Theorem 2.1 that

$$\sup |\text{URE}(\mathbf{b}, \mu) - l_p(\boldsymbol{\theta}, \hat{\boldsymbol{\theta}}^{\mathbf{b}, \mu})| \rightarrow 0 \quad \text{in probability.}$$

It follows that for any $\varepsilon > 0$

$$\begin{aligned} P(l_p(\boldsymbol{\theta}, \hat{\boldsymbol{\theta}}^{\text{SM}}) \geq l_p(\boldsymbol{\theta}, \hat{\boldsymbol{\theta}}^{\hat{\mathbf{b}}, \hat{\mu}}) + \varepsilon) \\ \leq P(l_p(\boldsymbol{\theta}, \hat{\boldsymbol{\theta}}^{\text{SM}}) - \text{URE}(\hat{\mathbf{b}}^{\text{SM}}, \hat{\mu}^{\text{SM}}) \geq l_p(\boldsymbol{\theta}, \hat{\boldsymbol{\theta}}^{\hat{\mathbf{b}}, \hat{\mu}}) - \text{URE}(\hat{\mathbf{b}}, \hat{\mu}) + \varepsilon) \\ \leq P\left(|l_p(\boldsymbol{\theta}, \hat{\boldsymbol{\theta}}^{\text{SM}}) - \text{URE}(\hat{\mathbf{b}}^{\text{SM}}, \hat{\mu}^{\text{SM}})| \geq \frac{\varepsilon}{2}\right) \\ + P\left(|l_p(\boldsymbol{\theta}, \hat{\boldsymbol{\theta}}^{\hat{\mathbf{b}}, \hat{\mu}}) - \text{URE}(\hat{\mathbf{b}}, \hat{\mu})| \geq \frac{\varepsilon}{2}\right) \rightarrow 0. \end{aligned}$$

Next, to show that

$$\limsup_{p \rightarrow \infty} [R_p(\boldsymbol{\theta}, \hat{\boldsymbol{\theta}}^{\text{SM}}) - R_p(\boldsymbol{\theta}, \hat{\boldsymbol{\theta}}^{\hat{\mathbf{b}}, \hat{\mu}})] \leq 0,$$

we note that

$$\begin{aligned} l_p(\boldsymbol{\theta}, \hat{\boldsymbol{\theta}}^{\text{SM}}) - l_p(\boldsymbol{\theta}, \hat{\boldsymbol{\theta}}^{\hat{\mathbf{b}}, \hat{\mu}}) \\ = (l_p(\boldsymbol{\theta}, \hat{\boldsymbol{\theta}}^{\text{SM}}) - \text{URE}(\hat{\mathbf{b}}^{\text{SM}}, \hat{\mu}^{\text{SM}})) + (\text{URE}(\hat{\mathbf{b}}^{\text{SM}}, \hat{\mu}^{\text{SM}}) - \text{URE}(\hat{\mathbf{b}}, \hat{\mu})) \\ + (\text{URE}(\hat{\mathbf{b}}, \hat{\mu}) - l_p(\boldsymbol{\theta}, \hat{\boldsymbol{\theta}}^{\hat{\mathbf{b}}, \hat{\mu}})) \\ \leq 2 \sup |\text{URE}(\mathbf{b}, \mu) - l_p(\boldsymbol{\theta}, \hat{\boldsymbol{\theta}}^{\mathbf{b}, \mu})|. \end{aligned}$$

Theorem 2.1 then implies that

$$\limsup_{p \rightarrow \infty} [R(\boldsymbol{\theta}, \hat{\boldsymbol{\theta}}^{\text{SM}}) - R(\boldsymbol{\theta}, \hat{\boldsymbol{\theta}}^{\hat{\mathbf{b}}, \hat{\mu}})] \leq 0. \quad \square$$

PROOF OF THEOREM 2.3. Throughout this proof, when a supremum is taken, it is over $b_i \in [0, 1]$ and Requirement (MON). Since

$$\begin{aligned} & \text{URE}^G(\mathbf{b}) - l_p(\boldsymbol{\theta}, \hat{\boldsymbol{\theta}}^{\mathbf{b}, \bar{Y}}) \\ &= \frac{1}{p} \sum_{i=1}^p \left(1 - 2 \left(1 - \frac{1}{p} \right) b_i \right) \left(\frac{V(Y_i)}{\tau_i + \nu_2} - (Y_i - \theta_i)^2 \right) \\ & \quad + \frac{2}{p} \sum_{i=1}^p b_i \left(\theta_i (Y_i - \theta_i) + \frac{1}{p} (Y_i - \theta_i)^2 - (Y_i - \theta_i) \bar{Y} \right), \end{aligned}$$

it follows that

$$\begin{aligned} & \sup |\text{URE}^G(\mathbf{b}) - l_p(\boldsymbol{\theta}, \hat{\boldsymbol{\theta}}^{\mathbf{b}, \bar{Y}})| \\ & \leq \frac{1}{p} \left| \sum_i \left(\frac{V(Y_i)}{\tau_i + \nu_2} - (Y_i - \theta_i)^2 \right) \right| \\ \text{(A.4)} \quad & + \frac{2}{p} \left(1 - \frac{1}{p} \right) \sup \left| \sum_i b_i \left(\frac{V(Y_i)}{\tau_i + \nu_2} - (Y_i - \theta_i)^2 \right) \right| \\ & + \frac{2}{p} \sup \left| \sum_i b_i \theta_i (Y_i - \theta_i) \right| \\ & + \frac{2}{p^2} \sum_i (Y_i - \theta_i)^2 + \frac{2}{p} |\bar{Y}| \cdot \sup \left| \sum_i b_i (Y_i - \theta_i) \right|. \end{aligned}$$

We have already shown in the proof of Theorem 2.1 that the first three terms on the right-hand side converge to zero in L^2 . It only remains to manage the last two terms:

$$E \left(\frac{2}{p^2} \sum_i (Y_i - \theta_i)^2 \right) = \frac{2}{p^2} \sum_i \text{Var}(Y_i) \rightarrow 0$$

by regularity condition (A),

$$\frac{1}{p} E \left(|\bar{Y}| \cdot \sup \left| \sum_i b_i (Y_i - \theta_i) \right| \right) \leq \frac{1}{p} E \left(\max_{1 \leq i \leq p} |Y_i| \cdot \sup \left| \sum_i b_i (Y_i - \theta_i) \right| \right) \rightarrow 0,$$

as was shown in the proof of Theorem 2.1. Therefore, the last two terms of (A.4) converge to zero in L^1 , and this completes the proof. \square

PROOF OF THEOREM 2.4. With Theorem 2.3 established, the proof is almost identical to that of Theorem 2.2. \square

PROOF OF LEMMA 3.1. It is straightforward to check that (i)–(iii) imply conditions (A)–(D) in Section 2, so we only need to verify condition (E). Since

$Y_i^2 = Z_i^2/\tau_i + \theta_i^2 + 2\theta_i Z_i/\sqrt{\tau_i}$, we know that

$$(A.5) \quad \max_{1 \leq i \leq p} Y_i^2 \leq \max_i \frac{1}{\tau_i} \cdot \max_i Z_i^2 + \max_i \theta_i^2 + 2 \max_i |\theta_i/\sqrt{\tau_i}| \cdot \max_i |Z_i|.$$

(i) and (ii) imply $\max_i 1/\tau_i = O(p^{1/2})$ and $\max_i |\theta_i/\sqrt{\tau_i}| = O(p^{1/2})$. (iii) gives $\max_i \theta_i^2 = O(p^{2/(2+\varepsilon)})$. We next bound $E(\max_{1 \leq i \leq p} |Z_i|^k)$ for $k = 1, 2$. (iv) implies that for $k = 1, 2$,

$$\begin{aligned} E\left(\max_i |Z_i|^k\right) &= \int_0^\infty kt^{k-1} P\left(\max_i |Z_i| > t\right) dt \\ &\leq \int_0^\infty kt^{k-1} (1 - (1 - Dt^{-\alpha})^p) dt \\ &= \int_0^{p^{1/\alpha}} kt^{k-1} (1 - (1 - Dt^{-\alpha})^p) dt + \int_{p^{1/\alpha}}^\infty kt^{k-1} (1 - (1 - Dt^{-\alpha})^p) dt \\ &= O(p^{k/\alpha}) + p^{k/\alpha} \int_1^\infty kz^{k-1} \left(1 - \left(1 - \frac{1}{p} Dz^{-\alpha}\right)^p\right) dz, \end{aligned}$$

where a change of variable $z = t/p^{1/\alpha}$ is applied. We know by the monotone convergence theorem that for $k \geq 1$, as $p \rightarrow \infty$,

$$\int_1^\infty z^{k-1} \left(1 - \left(1 - \frac{1}{p} Dz^{-\alpha}\right)^p\right) dz \rightarrow \int_1^\infty z^{k-1} (1 - \exp(-Dz^{-\alpha})) dz < \infty.$$

It then follows that

$$E\left(\max_{1 \leq i \leq p} |Z_i|^k\right) = O(p^{k/\alpha}), \quad \text{for } k = 1, 2.$$

Taking it back to (A.5) gives

$$E\left(\max_{1 \leq i \leq p} Y_i^2\right) = O(p^{1/2+2/\alpha}) + O(p^{2/(2+\varepsilon)}) + O(p^{1/2+1/\alpha}),$$

which verifies condition (E). \square

To prove Lemma 4.1, we need the following lemma.

LEMMA A.1. *Let Y_i be independent from one of the six NEF-QVFs. Then condition (B) in Section 2 and*

$$(F) \quad \limsup_{p \rightarrow \infty} \frac{1}{p} \sum_{i=1}^p |\theta_i|^{2+\varepsilon} < \infty \text{ for some } \varepsilon > 0;$$

$$(G) \quad \limsup_{p \rightarrow \infty} \frac{1}{p} \sum_{i=1}^p \text{Var}^2(Y_i) < \infty;$$

$$(H) \quad \sup_i \text{skew}(Y_i) = \sup_i \frac{1}{\sqrt{\tau_i}} \frac{v_1 + 2v_2\theta_i}{(v_0 + v_1\theta_i + v_2\theta_i^2)^{1/2}} < \infty;$$

imply condition (E).

PROOF OF LEMMA A.1. Let us denote $\sigma_i^2 = \text{Var}(Y_i)$. We can write $Y_i = \sigma_i Z_i + \theta_i$, where Z_i are independent with mean zero and variance one. It follows from $Y_i^2 = \sigma_i^2 Z_i^2 + \theta_i^2 + 2\sigma_i \theta_i Z_i$ that

$$(A.6) \quad \max_{1 \leq i \leq p} Y_i^2 \leq \max_i \sigma_i^2 \cdot \max_i Z_i^2 + \max_i \theta_i^2 + 2 \max_i \sigma_i |\theta_i| \cdot \max_i |Z_i|.$$

Condition (B) implies $\max_i \sigma_i^2 \theta_i^2 \leq \sum_i \sigma_i^2 \theta_i^2 = O(p)$. Thus, $\max_i \sigma_i |\theta_i| = O(p^{1/2})$. Similarly, condition (G) implies that $\max_i \sigma_i^2 = O(p^{1/2})$. Condition (F) implies that $\max_i |\theta_i|^{2+\varepsilon} \leq \sum_i |\theta_i|^{2+\varepsilon} = O(p)$, which gives $\max_i \theta_i^2 = O(p^{2/(2+\varepsilon)})$. If we can show that

$$(A.7) \quad E\left(\max_{1 \leq i \leq p} |Z_i|\right) = O(\log p), \quad E\left(\max_{1 \leq i \leq p} Z_i^2\right) = O(\log^2 p),$$

then we establish (E), since

$$E\left(\max_{1 \leq i \leq p} Y_i^2\right) = O(p^{1/2} \log^2 p + p^{2/(2+\varepsilon)} + p^{1/2} \log p) = O(p^{2/(2+\varepsilon^*)}),$$

where $\varepsilon^* = \min(\varepsilon, 1)$. To prove (A.7), we begin from

$$(A.8) \quad E\left(\max_i |Z_i|^k\right) = k \int_0^\infty t^{k-1} P\left(\max_i |Z_i| > t\right) dt \quad \text{for all } k > 0.$$

The large deviation results for NEF-QVF in Morris [(1982), Section 9] and condition (H) (i.e., Y_i have bounded skewness) imply that for all $t > 1$, the tail probabilities $P(|Z_i| > t)$ are uniformly bounded exponentially: there exists a constant $c_0 > 0$ such that

$$P(|Z_i| > t) \leq e^{-c_0 t} \quad \text{for all } i.$$

Taking it into (A.8), we have

$$(A.9) \quad \begin{aligned} E\left(\max_i |Z_i|^k\right) &\leq \int_0^\infty k t^{k-1} (1 - (1 - e^{-c_0 t})^p) dt \\ &= \int_0^{\log p/c_0} k t^{k-1} (1 - (1 - e^{-c_0 t})^p) dt \\ &\quad + \int_{\log p/c_0}^\infty k t^{k-1} (1 - (1 - e^{-c_0 t})^p) dt \\ &= O(\log^k p) + \int_0^\infty k \left(z + \frac{1}{c_0} \log p\right)^{k-1} \left(1 - \left(1 - \frac{1}{p} e^{-c_0 z}\right)^p\right) dz, \end{aligned}$$

where in the last line a change of variable $z = t - \log p/c_0$ is applied. We know by the monotone convergence theorem that for $k \geq 1$, as $p \rightarrow \infty$,

$$\int_0^\infty z^{k-1} \left(1 - \left(1 - \frac{1}{p} e^{-c_0 z}\right)^p\right) dz \rightarrow \int_0^\infty z^{k-1} (1 - \exp(-e^{-c_0 z})) dz < \infty.$$

It then follows from (A.9) that

$$E\left(\max_{1 \leq i \leq p} |Z_i|^k\right) = O(\log^k p), \quad \text{for } k = 1, 2,$$

which completes our proof. \square

PROOF OF LEMMA 4.1. We go over the five distributions one by one.

(1) Binomial. Since $\theta_i = p_i$, $\text{Var}(Y_i) = p_i(1 - p_i)/n_i$, and $\text{Var}(Y_i^2) \leq EY_i^4 \leq 1$, it is straightforward to verify that $n_i \geq 2$ for all i guarantees conditions (A)–(E) in Section 2.

(2) Poisson. $\text{Var}(Y_i) = \theta_i/\tau_i$, and $\text{Var}(Y_i^2) = (4\tau_i^2\theta_i^3 + 6\tau_i\theta_i^2 + \theta_i)/\tau_i^3$. It is straightforward to verify that $\inf_i \tau_i > 0$, $\inf_i \tau_i\theta_i > 0$ and $\sum_i \theta_i^3 = O(p)$ imply conditions (A)–(D) in Section 2 and conditions (F)–(H) in Lemma A.1.

(3) Negative-binomial. $\theta_i = \frac{p_i}{1-p_i}$, $\text{Var}(Y_i) = \frac{1}{n_i} \frac{p_i}{(1-p_i)^2} = \frac{1}{n_i}(\theta_i + \theta_i^2)$, so $v_0 = 0$, $v_1 = v_2 = 1$. $\text{Var}(Y_i^2) = \frac{1}{n_i^3(1-p_i)^4}(p_i + 4p_i^2 + 6n_i p_i^2 + p_i^3 + 4n_i p_i^3 + 4n_i^2 p_i^3)$. From these, we know that $\sum_{i=1}^p \text{Var}^2(Y_i) = \sum_i \frac{1}{(n_i p_i)^2} \left(\frac{p_i}{1-p_i}\right)^4 = O(\sum_i \left(\frac{p_i}{1-p_i}\right)^4) = O(p)$, which verifies conditions (A) and (G). $\sum_{i=1}^p \text{Var}(Y_i)\theta_i^2 = \sum_i \frac{1}{n_i p_i} \left(\frac{p_i}{1-p_i}\right)^4 = O(\sum_i \left(\frac{p_i}{1-p_i}\right)^4) = O(p)$, which verifies condition (B). For condition (C), since $n_i \geq 1$ and $0 \leq p_i \leq 1$, we only need to verify that $\sum_i \frac{1}{n_i^3(1-p_i)^4}(p_i + 6n_i p_i^2) = O(p)$. This is true, since $\sum_i \frac{1}{n_i^3(1-p_i)^4}(p_i + 6n_i p_i^2) = \sum_i \left(\frac{1}{(n_i p_i)^3} \left(\frac{p_i}{1-p_i}\right)^4 + \frac{6}{(n_i p_i)^2} \left(\frac{p_i}{1-p_i}\right)^4\right) = O(\sum_i \left(\frac{p_i}{1-p_i}\right)^4) = O(p)$. Condition (D) is automatically satisfied. For condition (F), consider $\sum_{i=1}^p \theta_i^4$. It is $\sum_i \left(\frac{p_i}{1-p_i}\right)^4 = O(p)$. For condition (H), note that $\text{skew}(Y_i) = \frac{1}{\sqrt{n_i}} \frac{p_i+1}{\sqrt{p_i}} \leq 1 + 1/\sqrt{n_i p_i}$. Thus, $\sup_i \text{skew}(Y_i) < \infty$ by (i).

(4) Gamma. $\theta_i = \alpha\lambda_i$, $\text{Var}(Y_i) = \alpha\lambda_i^2/\tau_i$, so $v_0 = v_1 = 0$, $v_2 = 1/\alpha$. $\text{skew}(Y_i) = 2/\sqrt{\tau_i\alpha}$. $\text{Var}(Y_i^2) = \frac{\alpha}{\tau_i} \lambda_i^4 \left(\alpha + \frac{1}{\tau_i}\right) \left(4\alpha + \frac{6}{\tau_i}\right)$. It is straightforward to verify that (i) and (ii) imply conditions (A)–(D) in Section 2 and conditions (F)–(H) in Lemma A.1.

(5) GHS. $\theta_i = \alpha\lambda_i$, $\text{Var}(Y_i) = \alpha(1 + \lambda_i^2)/\tau_i$, so $v_0 = \alpha$, $v_1 = 0$, $v_2 = 1/\alpha$. $\text{skew}(Y_i) = (2/\sqrt{\tau_i\alpha})\lambda_i/(1 + \lambda_i^2)^{1/2} \leq 2/\sqrt{\tau_i\alpha}$. $\text{Var}(Y_i^2) = \frac{2\alpha}{\tau_i}(1 + \lambda_i^2)\left(\alpha + \frac{1}{\tau_i}\right)\left(\frac{1}{\tau_i} + \frac{3}{\tau_i}\lambda_i^2 + 2\alpha\lambda_i^2\right)$. It is then straightforward to verify that (i) and (ii) imply conditions (A)–(D) in Section 2 and conditions (F)–(H) in Lemma A.1. \square

PROOF OF THEOREM 4.6. We note that the set over which the supremum is taken is a subset of that of Theorem 2.1. The desired result thus automatically holds. \square

PROOF OF THEOREM 4.7. With Theorem 4.6 established, the proof is almost identical to that of Theorem 2.2. \square

PROOF OF THEOREM 4.8. We note that the set over which the supremum is taken is a subset of that of Theorem 2.3. The desired result thus holds. \square

PROOF OF THEOREM 4.9. The proof essentially follows the same steps in that of Theorem 2.2. \square

Acknowledgments. The views expressed herein are the authors alone and are not necessarily the views of Two Sigma Investments, LLC or any of its affiliates.

REFERENCES

- BERGER, J. O. and STRAWDERMAN, W. E. (1996). Choice of hierarchical priors: Admissibility in estimation of normal means. *Ann. Statist.* **24** 931–951. [MR1401831](#)
- BERRY, D. A. and CHRISTENSEN, R. (1979). Empirical Bayes estimation of a binomial parameter via mixtures of Dirichlet processes. *Ann. Statist.* **7** 558–568. [MR0527491](#)
- BROWN, L. D. (1966). On the admissibility of invariant estimators of one or more location parameters. *Ann. Math. Statist.* **37** 1087–1136. [MR0216647](#)
- BROWN, L. D. (2008). In-season prediction of batting averages: A field test of empirical Bayes and Bayes methodologies. *Ann. Appl. Stat.* **2** 113–152. [MR2415597](#)
- BROWN, L. D. and GREENSHTEIN, E. (2009). Nonparametric empirical Bayes and compound decision approaches to estimation of a high-dimensional vector of normal means. *Ann. Statist.* **37** 1685–1704. [MR2533468](#)
- CLEVENSON, M. L. and ZIDEK, J. V. (1975). Simultaneous estimation of the means of independent Poisson laws. *J. Amer. Statist. Assoc.* **70** 698–705. [MR0394962](#)
- EFRON, B. and MORRIS, C. N. (1975). Data analysis using Stein's estimator and its generalizations. *J. Amer. Statist. Assoc.* **70** 311–319.
- EFRON, B. (2004). The estimation of prediction error: Covariance penalties and cross-validation. *J. Amer. Statist. Assoc.* **99** 619–642. [MR2090899](#)
- EFRON, B. and MORRIS, C. (1972). Empirical Bayes on vector observations: An extension of Stein's method. *Biometrika* **59** 335–347. [MR0334386](#)
- EFRON, B. and MORRIS, C. (1973). Stein's estimation rule and its competitors—An empirical Bayes approach. *J. Amer. Statist. Assoc.* **68** 117–130. [MR0388597](#)
- GELMAN, A. and HILL, J. (2007). *Data Analysis Using Regression and Multilevel/Hierarchical Models*. Cambridge Univ. Press, Cambridge.
- GREEN, J. and STRAWDERMAN, W. E. (1985). The use of Bayes/empirical Bayes estimation in individual tree volume development. *Forest Science* **31** 975–990.
- GUTMANN, S. (1982). Stein's paradox is impossible in problems with finite sample space. *Ann. Statist.* **10** 1017–1020. [MR0663454](#)
- HWANG, J. T. (1982). Improving upon standard estimators in discrete exponential families with applications to Poisson and negative binomial cases. *Ann. Statist.* **10** 857–867. [MR0663437](#)
- JAMES, W. and STEIN, C. (1961). Estimation with quadratic loss. In *Proc. 4th Berkeley Sympos. Math. Statist. and Prob., Vol. 1* 361–379. Univ. California Press, Berkeley, CA. [MR0133191](#)
- JIANG, W. and ZHANG, C.-H. (2009). General maximum likelihood empirical Bayes estimation of normal means. *Ann. Statist.* **37** 1647–1684. [MR2533467](#)
- JIANG, W. and ZHANG, C.-H. (2010). Empirical Bayes in-season prediction of baseball batting averages. In *Borrowing Strength: Theory Powering Applications—A Festschrift for Lawrence D. Brown. Inst. Math. Stat. Collect.* **6** 263–273. IMS, Beachwood, OH. [MR2798524](#)
- JOHNSTONE, I. (1984). Admissibility, difference equations and recurrence in estimating a Poisson mean. *Ann. Statist.* **12** 1173–1198. [MR0760682](#)

- JONES, K. (1991). Specifying and estimating multi-level models for geographical research. *Transactions of the Institute of British Geographers* **16** 148–159.
- KOENKER, R. and MIZERA, I. (2014). Convex optimization, shape constraints, compound decisions, and empirical Bayes rules. *J. Amer. Statist. Assoc.* **109** 674–685. [MR3223742](#)
- LI, K.-C. (1986). Asymptotic optimality of C_L and generalized cross-validation in ridge regression with application to spline smoothing. *Ann. Statist.* **14** 1101–1112. [MR0856808](#)
- MORRIS, C. N. (1982). Natural exponential families with quadratic variance functions. *Ann. Statist.* **10** 65–80. [MR0642719](#)
- MORRIS, C. N. (1983). Parametric empirical Bayes inference: Theory and applications. *J. Amer. Statist. Assoc.* **78** 47–65. [MR0696849](#)
- MOSTELLER, F. and WALLACE, D. L. (1964). *Inference and Disputed Authorship: The Federalist*. Addison-Wesley, Reading, MA. [MR0175668](#)
- MURALIDHARAN, O. (2010). An empirical Bayes mixture method for effect size and false discovery rate estimation. *Ann. Appl. Statist.* **4** 422–438. [MR2758178](#)
- RUBIN, D. (1981). Using empirical Bayes techniques in the law school validity studies. *J. Amer. Statist. Assoc.* **75** 801–827.
- STEIN, C. (1956). Inadmissibility of the usual estimator for the mean of a multivariate normal distribution. In *Proceedings of the Third Berkeley Symposium on Mathematical Statistics and Probability, 1954–1955, Vol. I* 197–206. Univ. California Press, Berkeley, CA. [MR0084922](#)
- XIE, X., KOU, S. C. and BROWN, L. D. (2012). SURE estimates for a heteroscedastic hierarchical model. *J. Amer. Statist. Assoc.* **107** 1465–1479. [MR3036408](#)

X. XIE
TWO SIGMA INVESTMENTS LLC
100 AVENUE OF THE AMERICAS, FLOOR 16
NEW YORK, NEW YORK 10013
USA
E-MAIL: xie1981@gmail.com

S. C. KOU
DEPARTMENT OF STATISTICS
HARVARD UNIVERSITY
CAMBRIDGE, MASSACHUSETTS 02138
USA
E-MAIL: kou@stat.harvard.edu

L. BROWN
THE WHARTON SCHOOL
UNIVERSITY OF PENNSYLVANIA
PHILADELPHIA, PENNSYLVANIA 19104
USA
E-MAIL: lbrown@wharton.upenn.edu