Published in final edited form as:

J Am Stat Assoc. 2012 December ; 107(500): 1465–1479. doi:10.1080/01621459.2012.728154.

SURE Estimates for a Heteroscedastic Hierarchical Model

Xianchao Xie^{†,*}, S. C. Kou^{†,*}, and Lawrence D. Brown^{‡,*}

[†]Harvard University

[‡]University of Pennsylvania

Abstract

Hierarchical models are extensively studied and widely used in statistics and many other scientific areas. They provide an effective tool for combining information from similar resources and achieving partial pooling of inference. Since the seminal work by James and Stein (1961) and Stein (1962), shrinkage estimation has become one major focus for hierarchical models. For the homoscedastic normal model, it is well known that shrinkage estimators, especially the James-Stein estimator, have good risk properties. The heteroscedastic model, though more appropriate for practical applications, is less well studied, and it is unclear what types of shrinkage estimators are superior in terms of the risk. We propose in this paper a class of shrinkage estimators based on Stein's unbiased estimate of risk (SURE). We study asymptotic properties of various common estimators as the number of means to be estimated grows ($p \rightarrow \infty$). We establish the asymptotic optimality property for the SURE estimators. We then extend our construction to create a class of semi-parametric shrinkage estimators and establish corresponding asymptotic optimality results. We emphasize that though the form of our SURE estimators is partially obtained through a normal model at the sampling level, their optimality properties do not heavily depend on such distributional assumptions. We apply the methods to two real data sets and obtain encouraging results.

Key words and phrases

Heteroscedasticity; hierarchical model; shrinkage estimator; Stein's unbiased risk estimate (SURE); asymptotic optimality

1 Introduction

Hierarchical modeling has become an increasingly important statistical method in many scientific and engineering applications. It provides an effective tool to combine information and achieve partial pooling of inference. The application of hierarchical models usually involves simultaneous inference of some quantities of interest for different yet similar groups of populations. The earliest study of such problems in statistics is perhaps the simultaneous estimation of several normal means. Since the seminal work by James and Stein (1961), shrinkage estimation has been influential in the development of hierarchical

*S. C. Kou's research is supported in part by NIH/NIGMS grant R01GM090202 and NSF grant DMS-0449204. L. Brown's research is supported in part by NSF grant DMS-1007657. The authors thank Professor Philippe Rigollet at Princeton University for helpful discussion.

normal models. Stein (1962) described a hierarchical, empirical Bayes interpretation for this estimator. See also Lindley (1962). Efron and Morris (1973) further developed this empirical Bayes interpretation and proposed several competing parametric empirical Bayes estimators. A full Bayesian treatment of this problem can be found in Berger and Strawderman (1996). Recently, Brown and Greenshtein (2009) proposed a nonparametric empirical Bayes method.

There has been substantial research toward understanding the risk properties of shrinkage estimators for the homoscedastic hierarchical normal models (i.e., all the variances in the subpopulations are equal). Baranchik (1971) gave a general form of admissible minimax estimators. Strawderman (1971) studied a class of proper Bayes minimax estimators. Brown (1971) gave a sufficient condition for admissibility of generalized Bayes estimators. The use of loss other than the usual quadratic one is discussed in Brown (1975) and Berger (1976). The heteroscedastic case (i.e., the unequal variance case), on the other hand, is less well addressed, though it is more practical for real applications. Typical minimax estimators, like the one given in Hudson (1974) and Berger (1976), usually shrink the coordinates with lower variances more than those with higher ones, as opposed to the common intuition that more shrinkage should be applied to components with higher variance. The estimators considered in this paper do not exhibit this counter-intuitive behavior.

For real-world applications, parametric empirical Bayes estimators (Efron and Morris, 1975; Morris, 1983) are widely adopted. The application of parametric empirical Bayes models usually involves the specification of a second-level model and the estimation of the corresponding hyper-parameters. For example, for the normal case, the common practice is to choose the normal-normal hierarchical structure and estimate the hyper-parameters through maximum marginal likelihood (EBMLE) or method of moments (EBMM). There has also been substantial study on the application of hierarchical Bayes models and nonparametric empirical Bayes methods. Brown (2008) evaluated the performance of various shrinkage estimators using the data on batting average for Major League Baseball players over a single season. It was noted that the parametric empirical Bayes maximum likelihood and the hierarchical Bayes method tend to have a poor performance due to their heavy reliance on the parametric assumptions of the second-level model. Other methods like the empirical Bayes method of moments and nonparametric empirical Bayes method were shown to achieve a better performance. Motivated from such an empirical study, it is hence interesting to know whether it is possible to formally compare those different shrinkage estimators and identify the “optimal” shrinkage estimator.

For this purpose, we propose a class of shrinkage estimators that can be readily applied in the heteroscedastic hierarchical normal models. We name our shrinkage estimators SURE shrinkage estimators, since the method is inspired by Stein’s unbiased estimate of risk (SURE) (Stein, 1973; Stein 1981). We first focus on shrinkage estimators whose forms are derived from the classic normal-normal hierarchical model and show that our SURE shrinkage estimators possess asymptotic optimality properties within this (sub)class. The results are then generalized to a class of semi-parametric shrinkage estimators that only require the shrinkage factors to satisfy a certain monotone relationship. It is emphasized that this asymptotic optimality property neither depends on the specific distribution assumptions

nor requires that the sequence of group means be independent of the group variance, an assumption that is implicit in many of the classical empirical Bayes methods like EBMLE and EBMM. Therefore, there are scenarios where the SURE estimators *strictly dominate* the classical methods. Simulation studies are presented to compare the performance of the proposed estimators with several other shrinkage estimators. We apply our method to the baseball data analyzed in Brown (2008) and report encouraging results. We also use our method to analyze a housing data set and note some interesting phenomena when applying these methods.

The remainder of the paper is organized as follows: In Section 2, we introduce the basic setup and define the parametric SURE estimators along with a brief discussion of some other competing shrinkage estimators. The case of shrinking toward the origin and toward the grand mean is discussed in detail in Section 3 and Section 4. Section 5 considers general parametric SURE estimators, where in addition to the shrinkage factor, the shrinkage location is also determined by the data. Section 6 introduces a class of semi-parametric shrinkage estimators and discusses their optimality properties. We conduct a comprehensive simulation study in Section 7 and apply our method to analyzing two real datasets in Section 8. A brief summary is given in Section 9. The technical proofs are relegated to the appendix.

2 Basic setup

Consider the estimation problem

$$X_i|\theta_i \sim N(\theta_i, A_i), \quad i=1, 2, \dots, p, \quad (2.1)$$

where the X_i are independently distributed with known (potentially) distinct variances A_i . The classical conjugate hierarchical model puts a prior on θ_i

$$\theta_i \sim N(\mu, \lambda), \quad \text{independently for } i=1, 2, \dots, p,$$

where λ is an unknown hyper-parameter.

In this section and Section 3, we first assume the value of the prior mean as $\mu = 0$. The case of unknown prior mean will be the focus of later sections.

Application of Bayes formula (when $\mu = 0$) gives us

$$\theta_i|X_i \sim N\left(\frac{\lambda}{\lambda+A_i}X_i, \frac{\lambda A_i}{\lambda+A_i}\right), \quad X_i \sim N(0, \lambda+A_i),$$

which leads to the Bayes shrinkage estimator

$$\hat{\theta}_i^\lambda = \frac{\lambda}{\lambda + A_i} X_i.$$

The empirical Bayes method tries to estimate the unknown hyper-parameter λ using the marginal distribution of $\mathbf{X}$

$$f(\mathbf{X}|\lambda, \mathbf{A}) \propto \prod_i (\lambda + A_i)^{-1/2} \exp\{-X_i^2/(2(\lambda + A_i))\}. \quad (2.2)$$

The empirical Bayes (EB) MLE $\hat{\lambda}_{ML}$, which uniquely maximizes the above marginal MLE, can be obtained as the solution of

$$\sum_i \left[\frac{X_i^2}{(\lambda + A_i)^2} - \frac{1}{\lambda + A_i} \right] = 0, \quad (2.3)$$

whenever this equation has a solution. If (2.3) does not have a solution, i.e., it is negative when $\lambda = 0$, $\hat{\lambda}_{ML}$ is then zero. The corresponding empirical Bayes MLE for $\boldsymbol{\theta}$ is

$$\hat{\theta}_i^{ML} := \hat{\theta}_i^{\hat{\lambda}_{ML}} = \frac{\hat{\lambda}_{ML}}{\hat{\lambda}_{ML} + A_i} X_i.$$

Another estimate based on the marginal distribution (2.2) is the moment estimate

$$\hat{\lambda}_{MM} = \frac{1}{p} \sum_{i=1}^p (X_i^2 - A_i),$$

or its positive part

$$\hat{\lambda}_{MM}^+ = \left(\frac{1}{p} \sum_{i=1}^p (X_i^2 - A_i) \right)^+.$$

In the homoscedastic case, where $A_i = A$ for $i = 1, \dots, p$, we have $\hat{\lambda}_{ML} = \hat{\lambda}_{MM}^+$ and

$$\hat{\boldsymbol{\theta}}^{ML} = \hat{\boldsymbol{\theta}}^{MM+} = \left(1 - \frac{pA}{\sum_{i=1}^p X_i^2} \right)^+ \mathbf{X}.$$

Hence in this case these two estimators are closely related to the positive-part James-Stein estimator

$$\hat{\theta}_i^{JS+} = \left(1 - \frac{(p-2)A_i}{\sum_{i=1}^p X_i^2}\right)^+ X_i.$$

In this paper, instead of relying on the marginal distribution of $\mathbf{X}$ to estimate λ , we consider an alternative perspective. The motivation of our methods comes from Stein's unbiased risk estimate (SURE): under the sum of squared-error loss $l_p(\boldsymbol{\theta}, \hat{\boldsymbol{\theta}}) = \frac{1}{p} \sum_i (\hat{\theta}_i - \theta_i)^2$, if one uses the shrinkage estimator $\hat{\theta}_i^\lambda = \frac{\lambda}{\lambda + A_i} X_i$ to estimate $\boldsymbol{\theta}$ with a fixed λ , then an unbiased estimate for its risk

$$R_p(\boldsymbol{\theta}, \hat{\boldsymbol{\theta}}^\lambda) = E[l_p(\boldsymbol{\theta}, \hat{\boldsymbol{\theta}}^\lambda)] = \frac{1}{p} \sum_i \frac{A_i}{(A_i + \lambda)^2} (A_i \theta_i^2 + \lambda^2) \quad (2.4)$$

is

$$\text{SURE}(\lambda) = \frac{1}{p} \sum_i \left[\left(\frac{A_i}{A_i + \lambda} \right)^2 X_i^2 + \frac{A_i(\lambda - A_i)}{A_i + \lambda} \right]. \quad (2.5)$$

Note that equation (2.4) is just the usual bias-squared plus variance description of the risk; equation (2.5) can be derived from Stein's unbiased estimate of the risk or directly from (2.4) since $\theta_i^2 = E(X_i^2) - A_i$. This relationship suggests that we can estimate λ from the data as the minimizer of $\text{SURE}(\lambda)$:

$$\hat{\lambda}_{\text{SURE}} = \arg \min_{\lambda \geq 0} \text{SURE}(\lambda) = \arg \min_{\lambda \geq 0} \sum_i \left[\left(\frac{A_i}{A_i + \lambda} \right)^2 X_i^2 + \frac{A_i(\lambda - A_i)}{A_i + \lambda} \right]. \quad (2.6)$$

Setting $\text{SURE}'(\lambda) = 0$ yields an easily solved expression for $\hat{\lambda}_{\text{SURE}}$ as the solution to

$$\sum_i \left[\frac{A_i^2}{(A_i + \lambda)^3} X_i^2 - \frac{A_i^2}{(A_i + \lambda)^2} \right] = 0. \quad (2.7)$$

If (2.7) does not have a solution, $\hat{\lambda}_{\text{SURE}}$ is then zero. The corresponding SURE estimate for $\boldsymbol{\theta}$ is

$$\hat{\boldsymbol{\theta}}_i^{\text{SURE}} := \hat{\theta}_i^{\hat{\lambda}_{\text{SURE}}} = \frac{\hat{\lambda}_{\text{SURE}}}{\hat{\lambda}_{\text{SURE}} + A_i} X_i.$$

Again, it is worth pointing out that in the homoscedastic case, the three estimators $\hat{\lambda}_{ML}$, $\hat{\lambda}_{MM}^+$ and $\hat{\lambda}_{SURE}$ are identical and are closely related to the famous positive-part James-Stein estimator. But once the A_i are not all equal, $\hat{\lambda}_{ML}$, $\hat{\lambda}_{MM}^+$ and $\hat{\lambda}_{SURE}$ give distinct results.

The idea of minimizing the unbiased estimate of risk to obtain the estimate of tuning parameters has a considerable history in statistics. Li (1985, 1986, 1987) discussed the asymptotic properties of the SURE method and its connection to generalized cross validation in various scenarios. From a slightly different perspective, Johnstone (1987) discussed the admissibility properties of SURE and some alternative estimates of the risk. Kneip (1994) studied the property of SURE in a class of ordered linear smoothers. Donoho and Johnstone (1995) applied SURE to choose the threshold in their SureShrink method. Cavalier et al. (2002) established a nonasymptotic oracle inequality and used it to study the minimax adaptive results of SURE in some inverse problems. We emphasize that our results differ from the previous ones in that the model under consideration is heteroscedastic and our asymptotic results allow us to directly compare our SURE estimators with other shrinkage estimators. Numerical comparisons in Sections 7 and 8 indicate that our estimators have desirable risk properties relative to a number of other shrinkage estimators.

3 Risk properties of the SURE estimator

In this section, we consider the risk properties of the SURE estimator. We show that in the heteroscedastic case the SURE estimator $\hat{\theta}^{SURE}$ is optimal in an asymptotic sense, whereas it is not necessarily so for the other estimators, including the empirical Bayes ML and MM estimators.

Our first result concerns how well $SURE(\lambda)$ approximates $l_p(\theta, \hat{\theta})$.

Theorem 3.1—Assuming two conditions

$$(A) \quad \limsup_{p \rightarrow \infty} \frac{1}{p} \sum_{i=1}^p A_i^2 < \infty,$$

$$(B) \quad \limsup_{p \rightarrow \infty} \frac{1}{p} \sum_{i=1}^p A_i \theta_i^2 < \infty,$$

we have

$$\sup_{0 \leq \lambda \leq \infty} \left| SURE(\lambda) - l_p(\theta, \hat{\theta}^\lambda) \right| \rightarrow 0 \text{ in } L^2 \text{ and in probability, as } p \rightarrow \infty.$$

Conditions (A) and (B) are required mainly to facilitate a short proof of the above result. Though it is likely that conditions (A) and (B) can be further relaxed, they do not seem to be particularly restrictive and we thus do not seek the full generality here. Theorem 3.1 shows that the risk estimate $SURE(\lambda)$ is not only unbiased for $R_p(\theta, \hat{\theta})$, but, more importantly, is *uniformly* close to the actual loss $l_p(\theta, \hat{\theta})$. We thus expect that minimizing $SURE(\lambda)$ would lead to an estimate with competitive performance. To facilitate our discussion of the risk

properties of our SURE shrinkage estimator, we next introduce the oracle loss (OL) hyper-parameter:

$$\tilde{\lambda}^{OL} = \tilde{\lambda}^{OL}(\boldsymbol{\theta}; X_1, \dots, X_p) = \arg \min_{\lambda \geq 0} l_p(\boldsymbol{\theta}, \hat{\boldsymbol{\theta}}^\lambda) = \arg \min_{\lambda \geq 0} \frac{1}{p} \sum_{i=1}^p \left(\frac{\lambda}{\lambda + A_i} X_i - \theta_i \right)^2.$$

Correspondingly, we define the oracle loss “estimator” $\tilde{\boldsymbol{\theta}}^{OL}$ as

$$\tilde{\boldsymbol{\theta}}^{OL} = \frac{\tilde{\lambda}^{OL}}{\tilde{\lambda}^{OL} + \mathbf{A}} \mathbf{X}.$$

Of course, $\tilde{\boldsymbol{\theta}}^{OL}$ is not really an estimator since it depends on the unknown $\boldsymbol{\theta}$. (Hence we use the notation $\tilde{\boldsymbol{\theta}}^{OL}$ rather than $\hat{\boldsymbol{\theta}}^{OL}$.) Although not obtainable in practice, $\tilde{\boldsymbol{\theta}}^{OL}$ lays down the theoretical limit that one can ever hope to reach: no estimator within the class of estimators

of the form $\hat{\boldsymbol{\theta}}^\lambda = \frac{\lambda}{\lambda + \mathbf{A}} \mathbf{X}$ can have smaller achieved loss or risk. The performance of the SURE estimator, interestingly, comes close to the oracle one. The following theorem shows under very mild assumptions that our SURE estimator is asymptotically nearly as good as the oracle loss estimator.

Theorem 3.2—Assume conditions (A) and (B). Then

$$\lim_{p \rightarrow \infty} P \left(l_p(\boldsymbol{\theta}, \hat{\boldsymbol{\theta}}^{\text{SURE}}) \geq l_p(\boldsymbol{\theta}, \tilde{\boldsymbol{\theta}}^{OL}) + \varepsilon \right) = 0 \quad \text{for any fixed } \varepsilon > 0.$$

The results in the above theorem and all subsequence ones are for given A_i 's and θ_i 's; that is, the probabilities and expectations are evaluated given the sequence of (θ_i, A_i) . We require in Theorem 3.2 that ε is fixed. As one referee kindly pointed out, the result can be enhanced by letting ε approaches zero at some rate that depends on the sequence of A_i 's and θ_i 's. A direct consequence of the preceding theorem is that the SURE estimator has a loss that is asymptotically no larger than that of any other estimator in the general class.

Corollary 3.3—Assume conditions (A) and (B). Then for any estimator $\hat{\lambda}_p \rightarrow 0$ and the

corresponding $\hat{\boldsymbol{\theta}}^{\hat{\lambda}_p} = \frac{\hat{\lambda}_p}{\hat{\lambda}_p + \mathbf{A}} \mathbf{X}$, we always have

$$\lim_{p \rightarrow \infty} P \left(l_p(\boldsymbol{\theta}, \hat{\boldsymbol{\theta}}^{\text{SURE}}) \geq l_p(\boldsymbol{\theta}, \hat{\boldsymbol{\theta}}^{\hat{\lambda}_p}) + \varepsilon \right) = 0 \quad \text{for any fixed } \varepsilon > 0.$$

Theorem 3.2 shows that the loss of $\hat{\boldsymbol{\theta}}^{\text{SURE}}$ converges in probability to the optimum oracle value $l_p(\boldsymbol{\theta}, \tilde{\boldsymbol{\theta}}^{OL})$. We can actually show that under the same conditions $\hat{\boldsymbol{\theta}}^{\text{SURE}}$ is asymptotically as good as $\tilde{\boldsymbol{\theta}}^{OL}$ in terms of expected loss.

Theorem 3.4—Assume conditions (A) and (B). Then

$$\lim_{p \rightarrow \infty} \left[R_p(\theta, \hat{\theta}^{\text{SURE}}) - E \left(l_p(\theta, \tilde{\theta}^{OL}) \right) \right] = 0.$$

It follows from this theorem that $\hat{\theta}^{\text{SURE}}$ has an asymptotically oracle risk: its risk is asymptotically smaller than (at least no larger than) any other estimator in the general class.

Corollary 3.5—Assume conditions (A) and (B). Then for any estimator $\hat{\theta}^{\lambda_p}$ and the corresponding $\hat{\theta}^{\lambda_p} = \frac{\lambda_p}{\lambda_p + \mathbf{A}} \mathbf{X}$, we always have

$$\limsup_{p \rightarrow \infty} \left[R_p(\theta, \hat{\theta}^{\text{SURE}}) - R_p(\theta, \hat{\theta}^{\lambda_p}) \right] \leq 0.$$

Corollaries 3.3 and 3.5 suggest why $\hat{\theta}^{\text{SURE}}$ is generally better than either $\hat{\theta}^{\hat{M}L}$ or $\hat{\theta}^{\hat{M}M+}$ for heteroscedastic problems. Note that $\hat{\theta}^{\text{SURE}}$ is asymptotically as good as the oracle loss estimator. Any other asymptotically optimal estimator must have this same property. Theorem 3.1 indicates that this requires such an oracle estimator to asymptotically agree with $\hat{\theta}^{\text{SURE}}$. But in the heteroscedastic case $\hat{\theta}^{\hat{M}L}$ and $\hat{\theta}^{\hat{M}M+}$ satisfy different estimating equations from that of $\hat{\theta}^{\text{SURE}}$, as described in Section 2. Hence, for heteroscedastic problems neither $\hat{\theta}^{\hat{M}L}$ nor $\hat{\theta}^{\hat{M}M+}$ can generally be asymptotically optimal in the class of estimators of the form $\hat{\theta}^{\hat{\lambda}} = \frac{\hat{\lambda}}{\hat{\lambda} + \mathbf{A}} \mathbf{X}$. Sections 7 and 8 will illustrate this point through numerical examples.

4 Shrinkage toward the grand mean

The results in the previous section focus on the shrinkage estimators that shrink toward a pre-set value (taken to be zero above). In practice, it is often the case that, instead of a pre-set value, we want to shrink toward the grand mean $\bar{X}$. To utilize the previous result in this case, one might first center the data by subtracting the grand mean from each sample X_i , and then pretend that the resulting $X_i - \bar{X}$ are “independent” with “variance” A_i , and, following (2.6), one could minimize

$$\sum_i \left[\left(\frac{A_i}{A_i + \lambda} \right)^2 (X_i - \bar{X})^2 + \frac{A_i(\lambda - A_i)}{A_i + \lambda} \right]$$

to obtain the estimate $\hat{\lambda}'$. The estimate of θ_i then becomes

$$\hat{\theta}'_i = \frac{\hat{\lambda}'}{A_i + \hat{\lambda}'} X_i + \frac{A_i}{A_i + \hat{\lambda}'} \bar{X}, \quad (4.1)$$

which can be used in practice. However, our previous theoretical results are no longer directly applicable. In particular, the optimality property of the resulting estimator is no longer established, since neither $X_i - \bar{X}$ are independent nor the variances are exactly A_i .

Fortunately, similar ideas of using the unbiased risk estimate can still be applied. Consider the shrinkage estimator in the following form

$$\hat{\theta}_i^{\lambda, \bar{X}} = \frac{\lambda}{A_i + \lambda} X_i + \frac{A_i}{A_i + \lambda} \bar{X}.$$

Its risk is given by

$$\begin{aligned} R(\boldsymbol{\theta}, \hat{\boldsymbol{\theta}}^{\lambda, \bar{X}}) &= E[l_p(\boldsymbol{\theta}, \hat{\boldsymbol{\theta}}^{\lambda, \bar{X}})] \\ &= \frac{1}{p} \sum_{i=1}^p \frac{A_i^2}{(A_i + \lambda)^2} (\theta_i - \bar{\theta}_p)^2 + \frac{1}{p} \sum_{i=1}^p \frac{1}{(A_i + \lambda)^2} \left(\lambda^2 A_i + \frac{1}{p} A_i^2 (\bar{A}_p + 2\lambda) \right), \end{aligned}$$

where

$$\bar{A}_p = \frac{1}{p} \sum_{i=1}^p A_i, \quad \bar{\theta}_p = \frac{1}{p} \sum_{i=1}^p \theta_i.$$

An unbiased risk estimate is

$$\text{SURE}^G(\lambda) = \frac{1}{p} \sum_{i=1}^p \frac{A_i^2}{(A_i + \lambda)^2} (X_i - \bar{X})^2 + \frac{1}{p} \sum_{i=1}^p \frac{A_i}{A_i + \lambda} \left(\lambda - A_i + \frac{2}{p} A_i \right),$$

that is,

$$E[\text{SURE}^G(\lambda)] = R(\boldsymbol{\theta}, \hat{\boldsymbol{\theta}}^{\lambda, \bar{X}}).$$

Minimizing $\text{SURE}^G(\lambda)$ then leads to the grand-mean shrinkage estimator

$$\hat{\theta}_i^G = \frac{\hat{\lambda}_G}{A_i + \hat{\lambda}_G} X_i + \frac{A_i}{A_i + \hat{\lambda}_G} \bar{X}, \quad (4.2)$$

where

$$\hat{\lambda}_G = \arg \min_{\lambda \geq 0} \text{SURE}^G(\lambda).$$

Since this estimate is inspired by the Stein's risk identity, we still call it the SURE estimate. Parallel to the results in the previous section, the grand-mean SURE estimator also possesses asymptotic optimality properties. First, we have the following theorem, which tells us that $\text{SURE}^G(\lambda)$ is uniformly close to the achieved loss $l_p(\theta, \hat{\theta}^{\lambda, \bar{X}})$. Thus, one expects that minimizing $\text{SURE}^G(\lambda)$ would lead to a competitive estimate.

Theorem 4.1—Assume conditions (A), (B) and

$$(C) \quad \limsup_{p \rightarrow \infty} \frac{1}{p} \sum_{i=1}^p \theta_i < \infty.$$

Then

$$\sup_{0 \leq \lambda \leq \infty} \left| \text{SURE}^G(\lambda) - l_p(\theta, \hat{\theta}^{\lambda, \bar{X}}) \right| \rightarrow 0 \text{ in } L^1 \text{ and in probability, as } p \rightarrow \infty.$$

To establish the asymptotic optimality of our SURE estimator, similar to Section 3, we define the grand-mean oracle loss “estimator” $\tilde{\theta}^{GOL}$ as

$$\tilde{\theta}^{GOL} = \frac{\tilde{\lambda}^{GOL}}{\tilde{\lambda}^{GOL} + \mathbf{A}} \mathbf{X} + \frac{\mathbf{A}}{\tilde{\lambda}^{GOL} + \mathbf{A}} \bar{X},$$

where

$$\tilde{\lambda}^{GOL} = \arg \min_{\lambda \geq 0} l_p(\theta, \hat{\theta}^{\lambda, \bar{X}}) = \arg \min_{\lambda \geq 0} \frac{1}{p} \sum_{i=1}^p \left(\frac{\lambda}{\lambda + A_i} X_i + \frac{A_i}{\lambda + A_i} \bar{X} - \theta_i \right)^2.$$

No estimator within the class of estimators of the form $\hat{\theta}^{\hat{\lambda}, \bar{X}}$, $\hat{\lambda} = \hat{\lambda}(X_1, \dots, X_p)$, can have smaller achieved loss or risk than $\tilde{\theta}^{GOL}$. However, the performance of the SURE estimator comes close: under very mild assumptions our SURE estimator $\hat{\theta}^G$ is asymptotically nearly as good as the grand-mean oracle loss estimator, as shown in the next theorem.

Theorem 4.2—Assume conditions (A)–(C). Then

$$\lim_{p \rightarrow \infty} P \left(l_p(\theta, \hat{\theta}^G) \geq l_p(\theta, \tilde{\theta}^{GOL}) + \varepsilon \right) = 0 \quad \text{for any fixed } \varepsilon > 0.$$

Theorem 4.2 implies that the SURE estimator is asymptotically optimal:

Corollary 4.3—Assume conditions (A)–(C). Then for any estimator $\hat{\lambda}_p \rightarrow 0$ and the corresponding $\hat{\theta}^{\hat{\lambda}_p, \bar{X}} = \frac{\hat{\lambda}_p}{\hat{\lambda}_p + \mathbf{A}} \mathbf{X} + \frac{\mathbf{A}}{\hat{\lambda}_p + \mathbf{A}} \bar{X}$, we always have

$$\lim_{p \rightarrow \infty} P \left(l_p(\theta, \hat{\theta}^G) \geq l_p(\theta, \hat{\theta}^{\hat{\lambda}_p, \bar{X}}) + \varepsilon \right) = 0 \quad \text{for any fixed } \varepsilon > 0.$$

Theorem 4.2 and corollary 4.3 compare the estimators in term of the loss. Under the same mild assumptions we can show that the comparison can be extended to the expected loss.

Theorem 4.4—Assume conditions (A)–(C). Then

$$\lim_{p \rightarrow \infty} \left[R_p(\theta, \hat{\theta}^G) - E \left(l_p(\theta, \tilde{\theta}^{GOL}) \right) \right] = 0.$$

Corollary 4.5—Assume conditions (A)–(C). Then for any estimator $\hat{\lambda}_p \rightarrow 0$ and the corresponding $\hat{\theta}^{\hat{\lambda}_p, \bar{X}} = \frac{\hat{\lambda}_p}{\hat{\lambda}_p + \mathbf{A}} \mathbf{X} + \frac{\mathbf{A}}{\hat{\lambda}_p + \mathbf{A}} \bar{X}$, we always have

$$\limsup_{p \rightarrow \infty} \left[R_p(\theta, \hat{\theta}^G) - R_p(\theta, \hat{\theta}^{\hat{\lambda}_p, \bar{X}}) \right] \leq 0.$$

Therefore, in general the SURE estimator is asymptotically better than (or at least as good as) any estimator, including the empirical Bayes ones, for heteroscedastic problems.

5 Shrinkage toward a general data driven location

Instead of shrinking toward the origin or the grand mean, one might let the data determine where to shrink to. Specifically, we can consider the estimator in the form of

$$\hat{\theta}_i^{\lambda, \mu} = \frac{\lambda}{A_i + \lambda} X_i + \frac{A_i}{A_i + \lambda} \mu.$$

Its risk is

$$R(\theta, \hat{\theta}^{\lambda, \mu}) = \frac{1}{p} \sum_i \frac{A_i}{(A_i + \lambda)^2} (A_i (\theta_i - \mu)^2 + \lambda^2),$$

for which an unbiased estimate is

$$\text{SURE}^M(\lambda, \mu) = \frac{1}{p} \sum_i \frac{A_i}{(A_i + \lambda)^2} (A_i (X_i - \mu)^2 + \lambda^2 - A_i^2).$$

We can then estimate both μ and λ by minimizing $\text{SURE}^M(\lambda, \mu)$ to obtain

$$\hat{\theta}_i^M = \frac{\hat{\lambda}_M}{A_i + \hat{\lambda}_M} X_i + \frac{A_i}{A_i + \hat{\lambda}_M} \hat{\mu}_M, \quad (5.1)$$

where

$$(\hat{\lambda}_M, \hat{\mu}_M) = \arg \min_{\lambda \geq 0, \mu} \text{SURE}^M(\lambda, \mu).$$

As before, we expect the SURE estimator $\hat{\theta}^M$ to possess asymptotic optimality properties. The following theorem, parallel to Theorems 3.1 and 4.1, tells us that $\text{SURE}^M(\lambda, \mu)$ closely approximates $l_p(\theta, \hat{\theta}^{\lambda, \mu})$ in a uniform fashion. Thus, one expects that minimizing $\text{SURE}^M(\lambda, \mu)$ would again lead to a competitive estimate.

Theorem 5.1—Assume conditions (A), (B) and

$$(C') \quad \limsup_{p \rightarrow \infty} \frac{1}{p} \sum_{i=1}^p |\theta_i|^{2+\delta} < \infty \text{ for some } \delta > 0.$$

Then we have

$$\sup_{0 \leq \lambda \leq \infty, |\mu| \leq \max_i |X_i|} \left| \text{SURE}^M(\lambda, \mu) - l_p(\theta, \hat{\theta}^{\lambda, \mu}) \right| \rightarrow 0 \text{ in } L^1 \text{ and in probability, as } p \rightarrow \infty.$$

Note that condition (C') assumes that the $2+\delta$ th moment of θ is bounded; it is slightly stronger than condition (C). Note also that Theorem 5.1 restricts the shrinkage location μ to be within $[-\max_i |X_i|, \max_i |X_i|]$. This assumption is included for technical reasons to ease the proof in the Appendix. In practice it is harmless since no sensible shrinkage estimator would attempt to shrink toward a location that lies beyond the range of the data.

Next, parallel to the development of Sections 3 and 4, we define the general-mean oracle loss “estimator” $\tilde{\theta}^{MOL}$ as

$$\tilde{\theta}^{MOL} = \frac{\tilde{\lambda}^{MOL}}{\tilde{\lambda}^{MOL} + \mathbf{A}} \mathbf{X} + \frac{\mathbf{A}}{\tilde{\lambda}^{MOL} + \mathbf{A}} \tilde{\mu}^{MOL}$$

where

$$[\tilde{\lambda}^{MOL}, \tilde{\mu}^{MOL}] = \arg \min_{\lambda \geq 0, |\mu| \leq \max_i |X_i|} l_p(\theta, \hat{\theta}^{\lambda, \mu}) = \arg \min_{\lambda \geq 0, |\mu| \leq \max_i |X_i|} \left\| \frac{\lambda}{\lambda + \mathbf{A}} \mathbf{X} + \frac{\mathbf{A}}{\lambda + \mathbf{A}} \mu - \theta \right\|^2.$$

The next theorem and corollary show that the SURE estimator $\hat{\theta}^M$ is asymptotically nearly as good as the general-mean oracle loss estimator, and, consequently, it is asymptotically better than (or at least as good as) any other shrinkage estimator in terms of the achieved loss.

Theorem 5.2—Assume conditions (A), (B) and (C'). Then

$$\lim_{p \rightarrow \infty} P \left(l_p(\theta, \hat{\theta}^M) \geq l_p(\theta, \tilde{\theta}^{MOL}) + \varepsilon \right) = 0 \quad \text{for any fixed } \varepsilon > 0.$$

Corollary 5.3—Assume conditions (A), (B) and (C'). Then for any estimator

$$\hat{\theta}^{\hat{\lambda}_p, \hat{\mu}_p} = \frac{\hat{\lambda}_p}{\hat{\lambda}_p + \mathbf{A}} \mathbf{X} + \frac{\mathbf{A}}{\hat{\lambda}_p + \mathbf{A}} \hat{\mu}_p \text{ with } \hat{\lambda}_p \rightarrow 0 \text{ and } |\hat{\mu}_p| \leq \max_i |X_i|, \text{ we have}$$

$$\lim_{p \rightarrow \infty} P \left(l_p(\theta, \hat{\theta}^M) \geq l_p(\theta, \hat{\theta}^{\hat{\lambda}_p, \hat{\mu}_p}) + \varepsilon \right) = 0 \quad \text{for any fixed } \varepsilon > 0.$$

Under the same mild assumptions, the comparison of the estimators can be extended to the expected loss as well.

Theorem 5.4—Assume conditions (A), (B) and (C'). Then

$$\lim_{p \rightarrow \infty} \left[R(\theta, \hat{\theta}^M) - E \left(l_p(\theta, \tilde{\theta}^{MOL}) \right) \right] = 0.$$

Corollary 5.5—Assume conditions (A), (B) and (C'). Then for any estimator

$$\hat{\theta}^{\hat{\lambda}_p, \hat{\mu}_p} = \frac{\hat{\lambda}_p}{\hat{\lambda}_p + \mathbf{A}} \mathbf{X} + \frac{\mathbf{A}}{\hat{\lambda}_p + \mathbf{A}} \hat{\mu}_p \text{ with } \hat{\lambda}_p \rightarrow 0 \text{ and } |\hat{\mu}_p| \leq \max_i |X_i|, \text{ we have}$$

$$\limsup_{p \rightarrow \infty} \left[R(\theta, \hat{\theta}^M) - R(\theta, \hat{\theta}^{\hat{\lambda}_p, \hat{\mu}_p}) \right] \leq 0.$$

Theorem 5.2 and Theorem 5.4 tells us that the SURE estimator is asymptotically optimal: it has the smallest loss and risk among all shrinkage estimators of the form

$$\hat{\theta}^{\hat{\lambda}, \hat{\mu}} = \frac{\hat{\lambda}}{\hat{\lambda} + \mathbf{A}} \mathbf{X} + \frac{\mathbf{A}}{\hat{\lambda} + \mathbf{A}} \hat{\mu}. \text{ A special case is the comparison between the general-mean SURE estimator } \hat{\theta}^M \text{ and the grand-mean shrinkage estimator } \hat{\theta}^{\hat{G}}.$$

Corollary 5.6—Assume conditions (A), (B) and (C'). Then

$$\limsup_{p \rightarrow \infty} [R(\theta, \hat{\theta}^M) - R(\theta, \hat{\theta}^G)] \leq 0.$$

In other words, $\hat{\theta}^M$ asymptotically outperforms $\hat{\theta}^G$. This result provides a theoretical underpinning of the empirical result of Section 7, where we shall see that $\hat{\theta}^M$ encountered a smaller loss than $\hat{\theta}^G$.

Another possible variation of the SURE estimate is to consider the weighted loss function $l_w(\theta, \hat{\theta}) = \sum_i w_i (\hat{\theta}_i - \theta_i)^2$. An unbiased risk estimate for $\hat{\theta}^{\mu, \lambda}$ in this case is

$$\text{SURE}^W(\mu, \lambda) = \frac{1}{p} \sum_i \frac{w_i A_i}{(A_i + \lambda)^2} (A_i (X_i - \mu)^2 + \lambda^2 - A_i^2).$$

The theoretical properties (such as the optimality) of the resulting estimator would be an interesting question worth further investigation.

Note that when we take $w_i \propto 1/A_i$, the $\text{SURE}^W(\mu, \lambda)$ can be re-written as

$$\begin{aligned} \text{SURE}^W(\mu, \lambda) &= \frac{1}{p} \sum_{i=1}^p \left(\frac{(\hat{\theta}_i^{\mu, \lambda} - X_i)^2}{A_i} + 2 \cdot df_i - 1 \right) \\ &= \frac{1}{p} \sum_{i=1}^p \left[(\hat{\theta}_i^{\mu, \lambda} - X_i)^2 / A_i + \frac{2}{A_i} \text{Cov}(\hat{\theta}_i^{\mu, \lambda}, X_i) - 1 \right], \end{aligned}$$

where

$$\text{Cov}(\hat{\theta}_i^{\mu, \lambda}, X_i) = \frac{\lambda A_i}{A_i + \lambda}$$

is the unbiased estimate of the covariance penalty (Efron, 1986; Efron, 2004). The SURE criterion in this case coincides with Mallows's C_p (Mallow, 1973), or equivalently the AIC (Akaike, 1973), where the number of parameters is taken to be the generalized degree of freedom (Ye, 1998). The above results thus serve as a rigorous confirmation of the belief that AIC-type criteria usually lead to models that enjoy good risk properties.

Remark—In the discussion above we have assumed that at the sampling level, the model is normal: $X_i | \theta_i \sim N(\theta_i, A_i)$. It is noted here that such a distributional assumption is actually not necessary. With some minimum regularity conditions (such as the tail of the distribution does not decay too slowly) all the theorems and corollaries will remain valid. One assumption that we do make is that the variances are known or can be estimated independently.

6 Semi-parametric SURE shrinkage estimation

As we noted in the previous sections, the optimality properties of the SURE estimators do not depend on the hypothetical normal prior. However, the general form $\hat{\theta}_i = \frac{\lambda}{A_i + \lambda} X_i + \frac{A_i}{A_i + \lambda} \mu$ of the shrinkage estimators studied in the preceding section is indeed motivated from the normal prior. In this section, we consider a larger class of shrinkage estimators, generalize the SURE estimator in this larger setting and study its asymptotic optimality properties. This new class of shrinkage estimators enjoys a more flexible form. We shall see that the generalized SURE estimator is optimal among this larger class of shrinkage estimators. Because it is optimal within a larger class of estimators it automatically performs asymptotically at least as well as the SURE estimators in previous sections. There are circumstances in which it can strictly outperform those estimators, as explored in Sections 7 and 8.

To motivate this larger class of shrinkage estimators, let us consider the hierarchical setting of

$$\begin{aligned} \lambda &\sim \pi(\lambda) \\ \theta_i | \lambda; \mu &\stackrel{i.i.d.}{\sim} N(\mu, \lambda) \\ X_i | \theta_i; A_i &\stackrel{ind.}{\sim} N(\theta_i, A_i), \end{aligned}$$

where π is an unspecified hyper-prior on λ . The posterior mean of θ_i (assuming existence) is

$$E(\theta_i | \mathbf{X}) = E\left(\frac{\lambda}{A_i + \lambda} | \mathbf{X}\right) X_i + E\left(\frac{A_i}{A_i + \lambda} | \mathbf{X}\right) \mu. \quad (6.1)$$

We can interpret $E\left(\frac{\lambda}{A_i + \lambda} | \mathbf{X}\right)$, which is monotonically decreasing in A_i , as the shrinkage factor for the i -th component. This suggests us to consider general shrinkage estimators of the form

$$\hat{\theta}_i^{b_i, \mu} = (1 - b_i) \cdot X_i + b_i \cdot \mu,$$

where $b_i \in [0, 1]$, and in this general form we no longer require b_i to assume any parametric form: there is no hyperparameter λ . Clearly, without putting any constraint on the b_i 's, one expects that the resulting SURE shrinkage estimator may suffer from problems such as overfitting. One natural way to prevent this from happening is to require the following condition on the shrinkage factors

$$\text{Requirement (MON): } b_i \leq b_j \quad \text{for any } i \text{ and } j \text{ such that } A_i \leq A_j,$$

or equivalently b_i is non-decreasing in A_i . In other words, the larger the variance is, the stronger is the shrinkage. This requirement is quite intuitive, especially in light of expression (6.1). Note that this requirement is satisfied by all the previous parametric SURE estimators.

To derive our semi-parametric shrinkage estimator, we first observe that an unbiased risk estimate of $\hat{\theta}^{\mu}$ is

$$\text{SURE}^M(\mathbf{b}, \mu) = \frac{1}{p} \sum_{i=1}^p [b_i^2 (X_i - \mu)^2 + (1 - 2b_i) A_i].$$

Minimizing the SURE with respect to $(\mathbf{b}, \mu)$ then leads to our semi-parametric SURE shrinkage estimator

$$\hat{\theta}_i^{SM} = (1 - \hat{b}_i^{SM}) \cdot X_i + \hat{b}_i^{SM} \cdot \hat{\mu}^{SM}, \quad (6.2)$$

where

$$(\hat{\mathbf{b}}^{SM}, \hat{\mu}^{SM}) = \text{minimizer of } \text{SURE}^M(\mathbf{b}, \mu) \text{ subject to } b_i \in [0, 1] \text{ and Requirement (MON).}$$

Parallel to the parametric case, we can also consider the estimator that shrinks toward the grand mean, that is,

$$\hat{\theta}_i^{b_i, \bar{X}} = (1 - \bar{b}_i) \cdot X_i + \bar{b}_i \cdot \bar{X}.$$

An unbiased estimate of its risk is

$$\text{SURE}^G(\mathbf{b}) = \frac{1}{p} \sum_{i=1}^p \left[b_i^2 (X_i - \bar{X})^2 + \left(1 - 2\left(1 - \frac{1}{p}\right)b_i \right) A_i \right].$$

Minimizing the SURE^G with respect to $\mathbf{b}$ then leads to our semi-parametric SURE grand-mean shrinkage estimator

$$\hat{\theta}_i^{SG} = (1 - \hat{b}_i^{SG}) \cdot X_i + \hat{b}_i^{SG} \cdot \bar{X}, \quad (6.3)$$

where

$$\hat{\mathbf{b}}^{SG} = \text{minimizer of } \text{SURE}^G(\mathbf{b}) \text{ subject to } b_i \in [0, 1] \text{ and Requirement (MON).}$$

It is emphasized that even though we used (6.1) to motivate our methods, we do not actually impose any particular parametric form on our estimates of the shrinkage factor b_i other than the range and the monotonicity requirement. This is the reason we term our methods “semi-parametric”. The theoretical properties of the semi-parametric SURE shrinkage estimators are summarized as follows. To save space, we only discuss the asymptotic optimality of the general-mean SURE estimator $\hat{\theta}_i^{SM}$ below; the asymptotic property of $\hat{\theta}_i^{SG}$ can be similarly studied.

Theorem 6.1—Assuming conditions (A), (B) and (C'), we have

$$\sup \left| \text{SURE}^M(\mathbf{b}, \mu) - l_p(\boldsymbol{\theta}, \tilde{\boldsymbol{\theta}}^{\mathbf{b}, \mu}) \right| \rightarrow 0 \text{ in } L^1 \text{ and in probability, as } p \rightarrow \infty,$$

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

Theorem 6.2—Assume conditions (A), (B) and (C'). Then for any shrinkage estimator $\hat{\boldsymbol{\theta}}^{\hat{\mathbf{b}}_p, \hat{\mu}_p} = (1 - \hat{\mathbf{b}}_p) \cdot \mathbf{X} + \hat{\mathbf{b}}_p \cdot \hat{\mu}_p$, where $\hat{\mathbf{b}}_p \in [0, 1]$ satisfies Requirement (MON) and $|\hat{\mu}_p| \leq \max_i |X_i|$, we have

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

and

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

Theorem 6.2 shows that our semi-parametric SURE shrinkage estimator is optimal among the class of shrinkage estimators whose shrinkage factor is a non-decreasing function of the variance. In particular, the semi-parametric SURE shrinkage estimator is asymptotically superior than (at least no worse than) any hierarchical empirical Bayes estimator.

7 Simulation study

In this section, we conduct a number of simulations to study the performance of the SURE estimators. We consider $\hat{\boldsymbol{\theta}}^{\hat{\mathbf{b}}_p, \hat{\mu}_p}$ (see equations (4.2) and (5.1)) and the two semi-parametric shrinkage estimators $\hat{\boldsymbol{\theta}}^{\hat{\mathbf{b}}_p, \hat{\mu}_p}$ (equations (6.3) and (6.2)) and compare their performance with that of the empirical Bayes maximum likelihood estimator $\hat{\boldsymbol{\theta}}^{\hat{\mathbf{b}}_p, \hat{\mu}_p}$, the empirical Bayes method-of-moment estimator $\hat{\boldsymbol{\theta}}^{\hat{\mathbf{b}}_p, \hat{\mu}_p}$ and an extension of the James-Stein estimator $\hat{\theta}_i^{JS+}$. The empirical Bayes maximum likelihood estimator used here is given by

$$\hat{\theta}_i^{ML} := \hat{\theta}_i^{\hat{\lambda}_{ML}} = \frac{\hat{\lambda}_{ML}}{\hat{\lambda}_{ML} + A_i} X_i + \frac{A_i}{\hat{\lambda}_{ML} + A_i} \hat{\mu}_{ML}. \quad (7.1)$$

where $\hat{\lambda}_{ML}$ and $\hat{\mu}_{ML}$ are obtained by maximizing the marginal density

$$f(\mathbf{X}|\lambda, \mathbf{A}) \propto \prod_i (\lambda + A_i)^{-1/2} \exp\{-(X_i - \mu)^2 / (2(\lambda + A_i))\},$$

and the empirical Bayes method-of-moment estimator is given by

$$\hat{\theta}_i^{MM} := \hat{\theta}_i^{\hat{\lambda}_{MM}} = \frac{\hat{\lambda}_{MM}}{\hat{\lambda}_{MM} + A_i} X_i + \frac{A_i}{\hat{\lambda}_{MM} + A_i} \hat{\mu}_{MM}, \quad (7.2)$$

where $\hat{\lambda}_{MM}$ and $\hat{\mu}_{MM}$ are obtained as the root of the following equations

$$\mu = \frac{\sum_i X_i / (A_i + \lambda)}{\sum_i 1 / (A_i + \lambda)}, \quad \lambda = \frac{1}{p-1} \left(\sum_i (X_i - \mu)^2 - (p-1)/p \sum_i A_i \right)^+.$$

The extended James-Stein estimator is

$$\hat{\theta}_i^{JS+} := \hat{\mu}_{JS+} + \left(1 - \frac{p-3}{\sum_i (X_i - \hat{\mu}_{JS})^2 / A_i} \right)^+ (X_i - \hat{\mu}_{JS+}), \quad \hat{\mu}_{JS+} = \frac{\sum_i X_i / A_i}{\sum_i 1 / A_i}, \quad (7.3)$$

which has been discussed in Brown (2008).

For each simulation, we first draw (A_i, θ_i) ($i = 1, \dots, p$) independently from a distribution $\pi(A, \theta)$ and then draw X_i given (A_i, θ_i) . The shrinkage estimators are then found via the formulae described above. This process is repeated a large number of times ($N = 100,000$) to obtain an accurate estimate of the average risk for each estimator. The sample size p is chosen to vary from 20 to 500 at an interval of length 20.

In each example, we also calculate the oracle *risk* “estimator” $\tilde{\theta}^{OR}$, defined as

$$\tilde{\theta}^{OR} = \frac{\tilde{\lambda}^{OR}}{\tilde{\lambda}^{OR} + \mathbf{A}} \mathbf{X} + \frac{\mathbf{A}}{\tilde{\lambda}^{OR} + \mathbf{A}} \tilde{\mu}^{OR}$$

where

$$(\tilde{\lambda}^{OR}, \tilde{\mu}^{OR}) = \arg \min_{\lambda \geq 0, \mu} R_p(\theta, \hat{\theta}^{\lambda, \mu}) = \arg \min_{\lambda \geq 0, \mu} \sum_{i=1}^p \frac{1}{p} E \left[\left(\frac{\lambda}{\lambda + A_i} X_i + \frac{A_i}{\lambda + A_i} \mu - \theta_i \right)^2 \right].$$

Similar to the oracle loss estimators, the oracle risk estimator $\tilde{\theta}^R$ cannot be used without the knowledge of θ , but it does provide a sensible lower bound of the risk achievable by any shrinkage estimator with the given parametric form. An alternative oracle estimator, which we do not pursue here, is the hierarchical Bayes estimator where the correct hyperprior is used.

Note that since we have a large number ($N = 100,000$) of repetitions in our simulation, the averaged risk of the oracle risk estimator plotted against p will be essentially a flat line. For each shrinkage estimator considered here, the risk $R(\theta, \hat{\theta})$ will converge to a limit as $p \rightarrow \infty$. This limit can be calculated numerically. Table 1 below shows these limiting risks for each simulation example.

Example 1—We draw (A, θ, X) such that $A \sim \text{Unif}(0.1, 1)$ and $\theta \sim N(0, 1)$ independently, and $X \sim N(\theta, A)$. Note that we draw A from $\text{Unif}(0.1, 1)$ instead of from $\text{Unif}(0, 1)$ to make sure that the variances A_i are bounded away from 0. The oracle risk estimator $\tilde{\theta}^R$ is found to have $\lambda_0 = 1$ and $\mu_0 = 0$. The corresponding risk for $\tilde{\theta}^R$ is $R(\theta, \tilde{\theta}^R) = 1 - \ln(2/1.1)/0.9 \approx 0.3357$. The plot in Figure 1.1 shows the risks of the seven shrinkage estimators as the sample size p varies. Clearly the performance of all shrinkage estimators except the extended James-Stein estimator eventually approaches that of the oracle risk estimator. Table 1 confirms the picture. Note that when the sample size is relatively small, the four SURE estimators incur slightly larger risks compared with the two empirical Bayes estimators. This is due to the fact that the hierarchical distribution on A and θ is exactly the one assumed by the empirical Bayes estimators; in particular, the empirical Bayes maximum likelihood estimator relies on the parametric normal form of the prior, and the empirical Bayes method-of-moment estimator assumes independence between A and θ , both of which are satisfied here. The SURE estimators require neither of these conditions but still achieve rather competitive performance. When the sample size is moderately large, all six estimators well approach the limit given by the oracle risk estimator. The extended James-Stein estimator behaves far worse than the others.

Example 2—We draw (A, θ, X) such that $A \sim \text{Unif}(0.1, 1)$ and $\theta \sim \text{Unif}(0, 1)$ independently, and $X \sim N(\theta, A)$. In this example, θ no longer comes from a normal distribution, but θ and A are still independent. The oracle risk estimator is found to have $\lambda_0 \approx 0.0834$ and $\mu_0 = 0.5$. The corresponding risk for $\tilde{\theta}^R$ is $R(\theta, \tilde{\theta}^R) \approx 0.0697$. The plot in Figure 1.2 shows the risks of the seven shrinkage estimators as the sample size p varies. Again, as p gets large, the performance of all shrinkage estimators except the extended James-Stein estimator eventually approaches that of the oracle risk estimator, as confirmed by Table 1. This observation indicates that the parametric form of the prior on θ is not crucial as long as A and θ are independent.

Example 3— (A, θ, X) are drawn such that $A \sim \text{Unif}(0.1, 1)$, $\theta = A$ and $X \sim N(\theta, A)$. In this example, A and θ are no longer independent of each other. The oracle risk estimator is found to have $\lambda_0 \approx 0.0781$ and $\mu_0 \approx 0.5949$ numerically. The corresponding risk for $\theta^{\tilde{R}}$ is $R(\theta, \theta^{\tilde{R}}) \approx 0.0540$. The plot in Figure 1.3 shows the risks of the seven shrinkage estimators as functions of p , the sample size. As our theoretical result in Section 5 indicates, the performance of the SURE estimator $\theta^{\hat{M}}$ approaches that of the oracle risk estimator, which is seen in Figure 1.3. The limiting risks of the SURE grand-mean shrinkage estimator $\theta^{\hat{G}}$, the two empirical Bayes estimators and the extended James-Stein estimator, on the other hand, are strictly greater than the risk of the oracle estimator, as shown in Table 1. The main reason for the difference is that A and θ are no longer independent. It is quite interesting to note from Table 1 that the limiting risks of the two semi-parametric shrinkage estimators $\theta^{\hat{G}}$ and $\theta^{\hat{M}}$ are actually strictly smaller than the oracle risk (although due to the scale of the plot, it is not easy to spot). The reason for this “better-than-oracle” performance is that the semi-parametric estimators are not restricted to the specific parametric family that the oracle estimator assumes.

Example 4—In this example, (A, θ, X) are drawn as $A \sim \text{Inv-}\chi_{10}^2$, $\theta = A$ and $X \sim N(\theta, A)$. The inverse chi-square distribution is used here as it is the conjugate distribution for normal variance. The oracle risk estimator is found to have $\lambda_0 \approx 0.0032$ and $\mu_0 \approx 0.1266$ numerically. The corresponding risk for the oracle risk estimator is $R(\theta, \theta^{\tilde{R}}) \approx 0.0051$. The plot in Figure 1.4 shows the risks of the seven shrinkage estimators as functions of the sample size p . We see from this figure and Table 1 that the risks of the SURE estimators $\theta^{\hat{M}}$ and $\theta^{\hat{G}}$ approach that of the oracle risk estimator, whereas the limiting risks of the James-Stein and empirical Bayes estimators $\theta^{\hat{S}^+}$, $\theta^{\hat{M}L}$ and $\theta^{\hat{M}M}$ are strictly greater than the oracle risk. Note that the limiting risks of the two semi-parametric shrinkage estimators $\theta^{\hat{G}}$ and $\theta^{\hat{M}}$ are in fact strictly smaller than the oracle risk (although due to the scale of the plot, it is not easy to spot).

Example 5—The setting in this example is chosen in such a way that it reflects grouping in the data. We draw (A, θ, X) as $A \sim \frac{1}{2} \cdot 1_{\{A=0.1\}} + \frac{1}{2} \cdot 1_{\{A=0.5\}}$, (i.e., A is 0.1 or 0.5 with 50% probability each), $\theta|A = 0.1 \sim N(2, 0.1)$, $\theta|A = 0.5 \sim N(0, 0.5)$ and $X|\theta, A \sim N(\theta, A)$ so that there exist two groups in the data. The oracle risk estimator is found to have $\lambda_0 \approx 0.8347$ and $\mu_0 \approx 0.1506$. The corresponding risk for the oracle risk estimator is $R(\theta, \theta^{\tilde{R}}) \approx 0.1947$. Figure 1.5 plots the risks of the seven shrinkage estimators versus the sample size p . We see clearly that the risk of the SURE estimator $\theta^{\hat{M}}$ approaches that of the oracle risk estimator, whereas the risks of the other four parametric shrinkage estimators ($\theta^{\hat{G}}$, $\theta^{\hat{M}L}$, $\theta^{\hat{M}M}$ and $\theta^{\hat{S}^+}$) are notably greater than the oracle risk. The semi-parametric shrinkage estimator $\theta^{\hat{M}}$ is seen to achieve an even significant improvement over the oracle one, confirming the results of Section 6.

Example 6—In this example, we allow X to depart from the normal model, i.e., $X \sim N(\theta, A)$, to assess the sensitivity in performance of the estimators to the normality assumption. (A, θ, X) are drawn as $A \sim \text{Unif}(0.1, 1)$, $\theta = A$, and $X \sim \text{Unif}[\theta - \sqrt{3A}, \theta + \sqrt{3A}]$. The oracle risk estimator is found to have $\lambda_0 \approx 0.0781$ and $\mu_0 \approx 0.5949$ numerically. The corresponding

risk for the oracle risk estimator $\theta^{\tilde{R}}$ is $R(\theta, \theta^{\tilde{R}}) \approx 0.0540$. Figure 1.6 plots the risks of the seven shrinkage estimators versus the sample size p . We see that the performance of SURE estimator $\theta^{\hat{M}}$ approaches that of the oracle risk estimator, whereas the empirical Bayes estimators $\theta^{\hat{M}L}$ and $\theta^{\hat{M}M}$ and the extended James-Stein estimator $\theta^{\hat{S}+}$ do notably worse. Table 1 shows that the limiting risks of the two semi-parametric shrinkage estimators $\theta^{\hat{G}}$ and $\theta^{\hat{M}}$ are strictly smaller than the oracle risk (though the gap is not big enough to be easily seen on the plot).

8 Application to real data

8.1 Prediction of batting average

In this subsection, we apply the SURE estimators to the baseball data in Brown (2008) to assess their effectiveness. The data analyzed here are the batting records for all the Major League Baseball players in the season of 2005. Like in Brown (2008), we divide the data set into two half seasons and try to predict the batting average of each player for the second half using the data from the first half. We also carried out the necessary preprocessing steps proposed there. For example, we removed from the analysis the players whose number of at-bats is less than 11. For each player, let the number of at-bats be N and the successful number of batting be H , we have,

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

where $i = 1, 2$ is the season indicator and $j = 1, \dots, p$ is the player indicator. As in Brown (2008), the following variance-stabilizing transformation is used before applying the shrinkage estimators

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

resulting in

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

One error measurement, denoted as TSE, introduced in Brown (2008) are adopted here as the basis of comparison. TSE measures the sum of squared errors in terms of θ and X , the transformed values:

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

Table 2 below summarizes the result, where the shrinkage estimators are applied three times —to all the baseball players, the pitchers only, and the non-pitchers only. 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} . In the table EB-MM is the empirical Bayes method-of-moment estimator (7.2). EB-ML is the empirical Bayes maximum likelihood estimator (7.1). James-Stein corresponds to the extended James-Stein estimator (7.3). Since this particular data set has been widely studied, we also compare our methods with a number of more recently developed methods, including 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/without the covariate at bats effect) in Jiang and Zhang (2009, 2010). Results for those methods are from Brown (2008), Muralidharan (2010) and Jiang and Zhang (2009, 2010). The last group shows the results for our SURE estimators. SURE $\hat{\theta}^G$ is the SURE grand-mean shrinkage estimator (4.2). SURE $\hat{\theta}^M$ is the SURE estimator (5.1), where the shrinkage location μ is also determined from the data. The last two estimators are the semi-parametric SURE shrinkage estimators. SPSURE $\hat{\theta}^{SG}$ is the semi-parametric grand-mean shrinkage estimator (6.3); SPSURE $\hat{\theta}^{SM}$ is the semi-parametric SURE general-mean estimator (6.2).

The numerical results demonstrate that the SURE estimators have quite appealing performance. The total squared errors of the SURE estimators are significantly smaller than almost of their competitors', with the only exception being that the weighted general maximum likelihood methods achieves a better performance in the all players' case. The main reason, we believe, is that the baseball data contain features that may degrade the performance of classical empirical Bayes methods, as discussed in Brown (2008). For example, substantial evidence against the normal prior assumption is observed, and, furthermore, ignoring the correlation between the mean θ and the variance A is not justifiable here (a player with large p tends to play more games, resulting in large N). Both of these features can invalidate the use of empirical Bayes methods. On the other hand, our SURE shrinkage estimators, especially the semi-parametric ones, are shown to be robust and optimal in much more general circumstances, resulting in the superior numerical outcome.

Figure 2 plots the shrinkage factor for four of the estimators we have considered for the “all batters” data — the empirical Bayes method-of-moment estimator (EBMOM), the empirical Bayes maximum likelihood estimator (EBMLE), the parametric SURE estimator $\hat{\theta}^M$ and the semi-parametric SURE estimator $\hat{\theta}^{SM}$. In the parametric case, the shrinkage factor is $\lambda/(\lambda + A)$; in the semi-parametric case, it is $1 - b^{\hat{SM}}$ as in equation (6.2). Note that in each case the shrinkage factor ($\lambda/(\lambda + A)$ or $1 - b^{\hat{SM}}$) increases with N , the number of at-bats, as they should. This corresponds to a decrease in terms of $A = 1/4N$. Note that the shrinkage factor for the EB-ML estimator is much smaller than those for the other estimators, which corresponds to greater shrinkage to the central location, and this is intimately related to the relatively poor performance of this estimator for the current data set. Note also that $1 - b^{\hat{SM}}$ increases with N in a stepwise fashion. The fact that it is non-decreasing is a direct consequence of its definition. The stepwise property is an indirect consequence of its definition — monotone solutions to the minimization problem in (6.2) or (6.3) will always

be stepwise monotone. Finally, note that for large values of N (approximately $N = 130$), $1 - b^{\hat{SM}}$ has the largest value among the four shrinkage factors. Thus, for this data set, the SURE estimator $\hat{\theta}^{\hat{SM}}$ shrinks somewhat less than the EB-MM estimator or the parametric SURE estimator $\hat{\theta}^{\hat{M}}$ when N is large, but shrinks comparably to these estimators for smaller N . It is also true that the estimates of central tendency differ for these estimators, but the differences are small to moderate. The corresponding values of $\hat{\mu}$ are $\hat{\mu} = 0.528, 0.538, 0.456, 0.529$, respectively.

8.2 Estimation of housing price

In this subsection, we apply the SURE estimators to a housing data set. The goal is to estimate the average housing price in each town of Scheffield, England from a small fraction of the data, as would be the case of a survey sampling. The data was produced by the Land Registry of the United Kingdom. It contains the information about all houses sold in Scheffield, England from 2000 to mid 2008. The sale price, the sale time, the postcode that identifies the location of the house and other relevant statistics about the sales are available for each house that has been sold during this time period. Nagaraja et al. (2011) discusses various analysis of similar, larger data sets from the US. We here confine our interest mainly to the estimation of average housing prices for each town in Scheffield, which has a distinct postcode. As conventional, the logarithm of the housing prices are used throughout the study to better approximate the normality assumption. Our analysis starts by removing the inflation from year to year by subtracting from each sample the overall year effect. (A more sophisticated method might build a two-way model with the year effect treated as a fixed effect and the area effect as a random effect.) We then randomly draw a small fraction of the data. This small fraction serves as a survey sample, from which we want to estimate the average housing price of each town of the entire data set. One particularly interesting feature of the data set is that the number of towns is small (around 20), while the number of house sales in most towns are large (above 500). To have a clear picture, we let the survey sample size range from 10% to 20% of the entire data. We exclude the towns with less than 20 house sales so that we would at least have three data points in the sample for each town.

To compare the performance of different shrinkage estimators, we run the simulation $N = 10,000$ times and report the average results in Table 3. The variances A_i are estimated from the sample variance of each town. As in the previous example, we use the naive estimator, the sample mean of each town, as the benchmark. Each number in the table refers to the ratio of the squared error of a particular estimator to that of the naive estimator. TSE stands for the total squared estimation error on the logarithm scale, while TSEP corresponds to that on the original scale. Note that unlike the baseball data, the parameter of interest θ_j (the average housing prices) can be directly obtained here, we can therefore evaluate the actual TSE and TSEP instead of estimating them through adjusting the prediction errors. There have also been discussion on alternatives other than squared error loss in the study of housing price, see Varian (1974) for one such example.

There are several interesting observations. First, the improvement of shrinkage estimators over the naive estimator as a group is not as impressive as in the baseball data case, though significant error reduction is still achieved. This is due to the fact that the number of groups

here is significantly smaller (20 here compared to around 500 in the baseball data). Second, as the sample size increases, the relative performance of shrinkage estimators decreases. This is because the variance of each sample mean decreases, resulting in smaller shrinkage. Third, overall speaking, the SURE shrinkage estimators achieve better performance compared with the other shrinkage estimators. The good performance of the semi-parametric SURE estimator $\hat{\theta}^G$ is particularly noteworthy. Fourth, when the number of groups p is small (around 20 here), it is not necessarily always beneficial to simultaneously estimate μ , the shrinkage location, and the shrinkage factors, since the asymptotic result is yet to take effect. Shrinking the estimates toward a predetermined location such as the grand mean could give better result.

9 Summary

Inspired by Stein's unbiased risk estimate (SURE), we propose in this paper a class of shrinkage estimators for the heteroscedastic hierarchical model, which is arguably more realistic in practical applications. We show that each SURE shrinkage estimator is asymptotically optimal in its own class. This includes the parametric SURE estimators, whose forms are derived from the classical parametric hierarchical model, as well as semi-parametric SURE estimators, which only assume that the individual shrinkage factor is monotone in the variance. We note that the asymptotic optimality of the SURE shrinkage estimators do not depend on the specific distribution assumptions, such as the normal assumption. We test the SURE estimators in comprehensive simulation studies and two real data sets, observing encouraging results: the SURE estimators offer numerically superior performance compared to the classical empirical Bayes and James-Stein estimators. The semi-parametric SURE estimators appear to be particularly competitive. We recommend the use of the semi-parametric SURE estimator $\hat{\theta}^M$ (where the shrinkage location is simultaneously estimated), when the number of groups are large. For data with small number of groups, we recommend the semi-parametric SURE estimator $\hat{\theta}^G$, which shrinks toward the grand mean.

There are several relevant research questions not fully addressed in this paper. For example, the sparse normal means problem (Johnstone and Silverman; 2004) has become increasingly important in statistics. It therefore would be of interest to study the performance of the proposed methods under this setting. It could also be of interest to study the extent to which the proposed estimators are minimax by using the techniques discussed in Maruyama and Strawderman(2005). It would also be of interest to study whether these estimators are ensemble minimax in the sense of Efron and Morris (1973) and Brown, Nie and Xie (2012). The peculiar features in the baseball data suggest that models that explicitly consider the dependence between θ_i 's and A_i 's might be more appropriate. For example, we can consider a hierarchical Bayes model where θ_i explicitly depends on A_i . It is interesting to see how the performance of such estimators is compared with the SURE estimators proposed in this paper.

References

- Akaike, H. Information theory as an extension of the maximum likelihood principle. In: Petrov, BN.; Csaki, F., editors. Second International Symposium on Information Theory. Akademiai Kiado; Budapest: 1973. p. 267-281.
- Baranchik AJ. A family of estimators of the mean of a multivariate normal distribution. *Annals of Mathematical Statistics*. 1970; 41:642–645.
- Berger J. Admissible minimax estimation of a multivariate normal mean with arbitrary quadratic loss. *Annals of Statistics*. 1976; 4:223–226.
- Berger J, Strawderman WE. Choice of hierarchical priors: Admissibility in estimation of normal means. *Annals of Statistics*. 1996; 24:931–951.
- Brown LD. Estimation with incomplete specified loss functions (the case with several location parameters). *Journal of the American Statistical Association*. 1975; 70:417–427.
- Brown LD. In-season prediction of batting averages: a field test of empirical Bayes and Bayes methodologies. *Annals of Applied Statistics*. 2008; 2:113–152.
- Brown LD, Greenshtein E. Nonparametric empirical Bayes and compound decision approaches to estimation of a high-dimensional vector of means. *Annals of Statistics*. 2009; 37(4):1685–1704.
- Brown LD, Nie H, Xie X. Ensemble minimax estimation for multivariate normal means. 2012 Submitted.
- Cavalier L, Golubev GK, Picard D, Tsybakov AB. Oracle inequalities for inverse problems. *Annals of Statistics*. 2002; 30(3):843–874.
- Donoho DL, Johnstone IM. Adapting to Unknown Smoothness via Wavelet Shrinkage. *Journal of the American Statistical Association*. 1995; 90:1200–1224.
- Efron B. How biased is the apparent error rate of a prediction rule? *Journal of the American Statistical Association*. 1986; 81:461–470.
- Efron B. The estimation of prediction error: Covariance penalties and cross- validation. *Journal of American Statistical Association*. 2004; 99:619–642.
- Efron B, Morris C. Stein's estimation rule and its competitors – an empirical Bayes approach. *Journal of the American Statistical Association*. 1973; 68:117–130.
- Efron B, Morris C. Data analysis using Stein's estimator and its generalizations. *Journal of American Statistical Association*. 1975; 70:311–319.
- Embrenchts, P.; Klüppelberg, C.; Mikosch, T. *Modelling Extreme Events*. Springer; New York: 1997.
- Hudson, HM. Technical Report No 58. Department of Statistics, Stanford University; 1974. Empirical Bayes estimation.
- Johnstone, IL. On the admissibility of some unbiased estimates of loss. In: Gupta, GS.; Berger, J., editors. *Statistical decision theory and related topics. Vol. IV*. Springer Verlag; New York, NY: 1987. p. 361-380.
- James, W.; Stein, CM. Estimation with quadratic loss. *Proceedings of 4th Berkeley Symposium on Probability and Statistics, I*; 1961. p. 367-379.
- Jiang W, Zhang CH. General maximum likelihood empirical Bayes estimation of normal means. *Annals of Statistics*. 2009; 37(4):1647–1684.
- Jiang, W.; Zhang, C-H. Empirical Bayes in-season prediction of baseball batting averages. In: Berger, J.; Cai, T.; Iain, MJ., editors. *Borrowing Strength: Theory Powering Applications - A Festschrift for Lawrence D. Brown*. Beachwood; Ohio: 2010. p. 263-273.
- Johnstone, IM. On the admissibility of some unbiased estimates of loss. In: Gupta, GS.; Berger, J., editors. *Statistical decision theory and related topics. Vol. IV*. Springer Verlag; New York, NY: 1987. p. 361-380.
- Johnstone IM, Silverman BM. Needles and straw in haystacks: Empirical Bayes estimates of possibly sparse sequences. *Annals of Statistics*. 2004; 32(4):1594–1649.
- Kneip A. Ordered linear smoothers. *Annals of Statistics*. 1994; 22(2):835–866.
- Li K-C. From Stein's unbiased risk estimates to the method of generalized cross-validation. *Annals of Statistics*. 1985; 13(4):1352–1377.

- Li K-C. Asymptotic optimality of C_L and generalized cross-validation in ridge regression with application to spline Smoothing. *Annals of Statistics*. 1986; 14(3):1101–1112.
- Li KC. Asymptotic optimality for CP, CL, cross-validation, and generalized cross-validation: Discrete index set. *Annals of Statistics*. 1987; 15(3):958–975.
- Lindley DV. Discussion of a paper by C. Stein. *Journal of the Royal Statistical Society, Series B*. 1962; 24:285–287.
- Mallows CL. Some Comments on CP. *Technometrics*. 1973; 15 (4):661–675.
- Maruyama Y, Strawderman W. A new class of generalized Bayes minmax ridge regression estimators. *Annals of Statistics*. 2005; 33:1753–1770.
- Morris C. Natural exponential families with quadratic variance functions. *Annals of Statistics*. 1982; 10(1):65–80.
- Morris C. Parametric empirical Bayes inference: theory and applications. *Journal of the American Statistical Association*. 1983; 78:47–55.
- Muralidharan O. An empirical Bayes mixture method for effect size and false discovery rate estimation. *Annals of Applied Statistics*. 2010; 4(1):422–438.
- Nagaraja CH, Brown LD, Zhao LH. An autoregressive approach to house price modeling. *Annals of Applied Statistics*. 2011; 5(1):124–149.
- Stein CM. Confidence sets for the mean of a multivariate normal distribution (with discussion). *Journal of the Royal Statistical Society, Series B*. 1962; 24:265–296.
- Stein, CM. Estimation of the mean of a multivariate distribution. *Proceedings Prague Symposium on Asymptotic Statistics*; 1973. p. 345–381.
- Stein CM. Estimation of the Mean of a Multivariate Normal Distribution. *The Annals of Statistics*. 1981; 9(6):1135–1151.
- Strawderman W. Proper Bayes estimators of the multivariate normal mean. *Annals of Mathematical Statistics*. 1971; 42:385–388.
- Varian, H. A Bayesian Approach to Real Estate Assessment. In: Fienberg, S.; Zellner, A., editors. *Studies in Bayesian Econometrics and Statistics*. 1975. p. 195–208.
- Ye J. On Measuring and Correcting the Effects of Data Mining and Model Selection. *Journal of the American Statistical Association*. 1998; 93:120–131.

Appendix: Proofs

Proof of Theorem 3.1

We only need to show the L^2 convergence. Since

$$\text{SURE}(\lambda) - l_p(\boldsymbol{\theta}, \hat{\boldsymbol{\theta}}^\lambda) = \frac{1}{p} \sum_i \left(X_i^2 - A_i - \theta_i^2 - \frac{2\lambda}{A_i + \lambda} (X_i^2 - X_i \theta_i - A_i) \right),$$

we know

$$\sup_{0 \leq \lambda \leq \infty} \left| \text{SURE}(\lambda) - l_p(\boldsymbol{\theta}, \hat{\boldsymbol{\theta}}^\lambda) \right| \leq \sup_{0 \leq \lambda \leq \infty} \left| \frac{1}{p} \sum_i \frac{2\lambda}{A_i + \lambda} (X_i^2 - X_i \theta_i - A_i) \right| + \left| \frac{1}{p} \sum_i (X_i^2 - A_i - \theta_i^2) \right|.$$

We consider the two terms separately. For the first term, without loss of generality, let us assume $A_1 \leq A_2 \leq \dots \leq A_p$. Then we know

$$\sup_{0 \leq \lambda \leq \infty} \left| \frac{1}{p} \sum_i \frac{2\lambda}{A_i + \lambda} (X_i^2 - X_i \theta_i - A_i) \right| \leq \sup_{1 \leq c_1 \leq \dots \leq c_p \leq 0} \left| \frac{2}{p} \sum_{i=1}^p c_i (X_i^2 - X_i \theta_i - A_i) \right|.$$

As in Lemma 2.1 in Li (1986), observe that

$$\sup_{1 \leq c_1 \leq \dots \leq c_p \leq 0} \left| \frac{2}{p} \sum_{i=1}^p c_i (X_i^2 - X_i \theta_i - A_i) \right| = \max_{1 \leq j \leq p} \left| \frac{2}{p} \sum_{i=1}^j (X_i^2 - X_i \theta_i - A_i) \right|.$$

Let $M_j = \sum_{i=1}^j (X_i^2 - X_i \theta_i - A_i)$. Then $\{M_j; j = 1, 2, \dots\}$ forms a martingale. The L^p maximum inequality implies

$$E(\max_{1 \leq j \leq p} M_j^2) \leq 4E(M_p^2) = 4 \sum_{i=1}^p (2A_i^2 + A_i \theta_i^2).$$

Regularity conditions (A) and (B) thus guarantee that $E(\max_j (\frac{2}{p} M_j)^2) \rightarrow 0$, which yields

$$\sup_{0 \leq \lambda \leq \infty} \left| \frac{1}{p} \sum_i \frac{2\lambda}{A_i + \lambda} (X_i^2 - X_i \theta_i - A_i) \right| \rightarrow 0 \text{ in } L^2, \text{ as } p \rightarrow \infty.$$

For the second term $\frac{1}{p} \sum_i (X_i^2 - A_i - \theta_i^2)$, a direct calculation gives

$$E \left[\left(\frac{1}{p} \sum_i (X_i^2 - A_i - \theta_i^2) \right)^2 \right] = \frac{2}{p^2} \sum_{i=1}^p (A_i^2 + 2A_i \theta_i^2) \rightarrow 0,$$

by conditions (A) and (B). This completes the proof.

Proof of Theorem 3.2

Since $\text{SURE}(\hat{\lambda}_{\text{SURE}}) = \text{SURE}(\hat{\lambda}^{OL})$ by definition, and we know from the preceding theorem that $\sup_{\lambda} |\text{SURE}(\lambda) - l_p(\theta, \hat{\theta})| \rightarrow 0$ in probability, it follows that for any $\varepsilon > 0$

$$P \left(l_p(\theta, \hat{\theta}^{\text{SURE}}) \geq l_p(\theta, \tilde{\theta}^{OL}) + \varepsilon \right) \leq P \left(l_p(\theta, \hat{\theta}^{\text{SURE}}) - \text{SURE}(\hat{\lambda}_{\text{SURE}}) \geq l_p(\theta, \tilde{\theta}^{OL}) - \text{SURE}(\tilde{\lambda}^{OL}) + \varepsilon \right) \rightarrow 0,$$

which completes the proof.

Proof of Corollary 3.3

This is a direct consequence of the definition of $\tilde{\theta}^{OL}$ and Theorem 3.2.

Proof of Theorem 3.4

Since

$$\begin{aligned} l_p(\theta, \hat{\theta}^{SURE}) - l_p(\theta, \tilde{\theta}^{OL}) &= \left(l_p(\theta, \hat{\theta}^{SURE}) - \text{SURE}(\hat{\lambda}_{SURE}) \right) + \left(\text{SURE}(\hat{\lambda}_{SURE}) - \text{SURE}(\tilde{\lambda}^{OL}) \right) + \left(\text{SURE}(\tilde{\lambda}^{OL}) - l_p(\theta, \tilde{\theta}^{OL}) \right) \\ &\leq 2 \sup_{0 \leq \lambda \leq \infty} \left| \text{SURE}(\lambda) - l_p(\theta, \hat{\theta}^\lambda) \right|, \end{aligned}$$

we know from Theorem 3.1 that

$$l_p(\theta, \hat{\theta}^{SURE}) - l_p(\theta, \tilde{\theta}^{OL}) \rightarrow 0 \text{ in } L^2 \text{ and in } L^1.$$

Therefore,

$$\lim_{p \rightarrow \infty} \left[R_p(\theta, \hat{\theta}^{SURE}) - R_p(\theta, \tilde{\theta}^{OL}) \right] = 0.$$

Proof of Corollary 3.5

This is a direct consequence of the definition of $\tilde{\theta}^{OL}$ and Theorem 3.4.

Proof of Theorem 4.1

Since

$$\text{SURE}^G(\lambda) - l_p(\theta, \hat{\theta}^{\lambda, \bar{X}}) = \frac{1}{p} \sum_i \left((X_i^2 - A_i - \theta_i^2) - \frac{2\lambda}{A_i + \lambda} (X_i^2 - X_i \theta_i - A_i) - \frac{2A_i}{A_i + \lambda} (\bar{X}(X_i - \theta_i) - \frac{A_i}{p}) \right),$$

we have

$$\begin{aligned} &\sup_{0 \leq \lambda \leq \infty} \left| \text{SURE}^G(\lambda) - l_p(\theta, \hat{\theta}^{\lambda, \bar{X}}) \right| \\ &\leq \left| \frac{1}{p} \sum_i (X_i^2 - A_i - \theta_i^2) \right| + \sup_{0 \leq \lambda \leq \infty} \left| \frac{1}{p} \sum_i \frac{2\lambda}{A_i + \lambda} (X_i^2 - X_i \theta_i - A_i) \right| \\ &\quad + \sup_{0 \leq \lambda \leq \infty} \left| \frac{1}{p} \sum_i \frac{2A_i}{A_i + \lambda} \left(\bar{X}(X_i - \theta_i) - \frac{A_i}{p} \right) \right|. \end{aligned}$$

The convergence of the first two terms in L^2 has already been established in the proof of Theorem 3.1. We only need to show that the last term converges to 0 in L^1 . But

$$\sup_{0 \leq \lambda \leq \infty} \left| \frac{1}{p} \sum_i \frac{2A_i(\bar{X}(X_i - \theta_i) - A_i/p)}{A_i + \lambda} \right| + \sup_{0 \leq \lambda \leq \infty} \left| \frac{1}{p} \sum_i \frac{2A_i(X_i - \theta_i)}{A_i + \lambda} \right| \cdot |\bar{X}| + \frac{2}{p^2} \sum_i A_i.$$

Following the technique in the proof of Theorems 3.1, it can be shown that

$$\sup_{0 \leq \lambda \leq \infty} \left| \frac{1}{p} \sum_{i=1}^p \frac{2A_i(X_i - \theta_i)}{A_i + \lambda} \right| \rightarrow 0 \text{ in } L^2.$$

We also know that $E\bar{X}^2 = \frac{1}{p^2} \sum_i A_i + (\frac{1}{p} \sum_i \theta_i)^2$, which is bounded by Conditions (A) and (C). Therefore, we have

$$\sup_{0 \leq \lambda \leq \infty} \left| \frac{1}{p} \sum_i \frac{2A_i(X_i - \theta_i)}{A_i + \lambda} \right| \cdot |\bar{X}| \rightarrow 0 \text{ in } L^1,$$

by Cauchy-Schwartz inequality, and this completes the proof.

Proof of Theorem 4.2

With Theorem 4.1 established, the proof is almost identical to that of Theorem 3.2.

Proof of Corollary 4.3

This is a direct consequence of the definition of $\tilde{\theta}^{GOL}$ and Theorem 4.2.

Proof of Theorem 4.4

Since

$$\begin{aligned} l_p(\theta, \hat{\theta}^G) - l_p(\theta, \tilde{\theta}^{GOL}) &= (l_p(\theta, \hat{\theta}^G) - \text{SURE}^G(\hat{\lambda}_G)) + (\text{SURE}^G(\hat{\lambda}_G) - \text{SURE}^G(\tilde{\lambda}^{GOL})) + (\text{SURE}^G(\tilde{\lambda}^{GOL}) - l_p(\theta, \tilde{\theta}^{GOL})) \\ &\leq 2 \sup_{0 \leq \lambda \leq \infty} \left| \text{SURE}^G(\lambda) - l_p(\theta, \hat{\theta}^{\lambda, \bar{X}}) \right|, \end{aligned}$$

we know from Theorem 4.1 that

$$l_p(\theta, \hat{\theta}^{\text{SURE}}) - l_p(\theta, \tilde{\theta}^{OL}) \rightarrow 0 \text{ in } L^1.$$

Therefore,

$$\lim_{p \rightarrow \infty} [R_p(\boldsymbol{\theta}, \hat{\boldsymbol{\theta}}^{\text{SURE}}) - R_p(\boldsymbol{\theta}, \tilde{\boldsymbol{\theta}}^{OL})] = 0.$$

Proof of Corollary 4.5

This is a direct consequence of the definition of $\tilde{\boldsymbol{\theta}}^{OL}$ and Theorem 4.4.

To prove Theorem 5.1, we need the following lemma.

Lemma A.1

Assume conditions (A), (B) and (C'). Then we have

$$E(\max_{1 \leq i \leq p} X_i^2) = O(p^{2/(2+\delta^*)}),$$

where $\delta^* = \min(1, \delta)$.

Proof

We can write $X_i = \sqrt{A_i}Z_i + \theta_i$, where Z_i are i.i.d. standard normal random variables. It follows from $X_i^2 = A_i Z_i^2 + \theta_i^2 + 2\sqrt{A_i}\theta_i Z_i$ that

$$\max_{1 \leq i \leq p} X_i^2 \leq \max_i A_i \cdot \max_i Z_i^2 + \max_i \theta_i^2 + 2 \max_i \sqrt{A_i} |\theta_i| \cdot \max_i |Z_i|. \quad (\text{A.1})$$

Condition (A) implies that $\max_i A_i^2 \leq \sum_i A_i^2 = O(p)$. Thus, $\max_i A_i = O(p^{1/2})$. Similarly,

Condition (B) implies $\max_i \sqrt{A_i} |\theta_i| = O(p^{1/2})$. Condition (C') implies that $\sum_i |\theta_i|^{2+\alpha} = O(p)$ for all $0 < \alpha \leq \delta$; in particular, $\sum_i |\theta_i|^{2+\delta^*} = O(p)$. Since $\max_i |\theta_i|^{2+\delta^*} \leq \sum_i |\theta_i|^{2+\delta^*} = O(p)$, we know that $\max_i \theta_i^2 = O(p^{2/(2+\delta^*)})$. It is well known (see, for example, Embrechts, et al., 1997, Chapter 3) that

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

Taking them into (A.1), we obtain

$$E(\max_{1 \leq i \leq p} X_i^2) = O(p^{1/2} \log p) + O(p^{2/(2+\delta^*)}) + O(p^{1/2} \sqrt{\log p}) = O(p^{2/(2+\delta^*)}).$$

Proof of Theorem 5.1

Since

$$\text{SURE}^M(\lambda, \mu) - l_p(\boldsymbol{\theta}, \hat{\boldsymbol{\theta}}^{\lambda, \mu}) = \text{SURE}(\lambda) - l_p(\boldsymbol{\theta}, \hat{\boldsymbol{\theta}}^\lambda) - \frac{2\mu}{p} \sum_i \frac{A_i}{A_i + \lambda} (X_i - \theta_i),$$

it follows that

$$\begin{aligned} & \sup_{0 \leq \lambda \leq \infty, |\mu| \leq \max_i |X_i|} \left| \text{SURE}^M(\lambda, \mu) - l_p(\boldsymbol{\theta}, \hat{\boldsymbol{\theta}}^{\lambda, \mu}) \right| \\ & \leq \sup_{0 \leq \lambda \leq \infty} \left| \text{SURE}(\lambda) - l_p(\boldsymbol{\theta}, \hat{\boldsymbol{\theta}}^\lambda) \right| + \frac{2}{p} \max_{1 \leq i \leq p} |X_i| \cdot \sup_{0 \leq \lambda \leq \infty} \left| \sum_i \frac{A_i(X_i - \theta_i)}{A_i + \lambda} \right|. \end{aligned} \quad (\text{A.2})$$

We know from Theorem 3.1 that

$$\sup_{0 \leq \lambda \leq \infty} \left| \text{SURE}(\lambda) - l_p(\boldsymbol{\theta}, \hat{\boldsymbol{\theta}}^\lambda) \right| \rightarrow 0 \text{ in } L^2.$$

It remains to show that the second term in (A.2) converges to zero in L^1 .

Following the same steps as in the proof of Theorem 3.1, we can show that

$$\sup_{0 \leq \lambda \leq \infty} \left| \sum_i \frac{A_i(X_i - \theta_i)}{A_i + \lambda} \right| \leq \max_{1 \leq j \leq p} \left| \sum_{i=j}^p (X_i - \theta_i) \right|.$$

Therefore, by the L^p maximum inequality on martingales, we have

$$\begin{aligned} E \left\{ \sup_{0 \leq \lambda \leq \infty} \left[\sum_i \frac{A_i(X_i - \theta_i)}{A_i + \lambda} \right]^2 \right\} & \leq E \left\{ \max_{1 \leq j \leq p} \left[\sum_{i=j}^p (X_i - \theta_i) \right]^2 \right\} \leq 4E \left(\sum_{i=1}^p (X_i - \theta_i) \right)^2 \\ & = 4 \sum_i A_i = O(p). \end{aligned}$$

Combining this with Lemma A.1, we obtain by Cauchy-Schwartz inequality

$$\begin{aligned} \frac{1}{p} E \left(\max_{1 \leq i \leq p} |X_i| \cdot \sup_{0 \leq \lambda \leq \infty} \left| \sum_i \frac{A_i(X_i - \theta_i)}{A_i + \lambda} \right| \right) & \leq \frac{1}{p} \left(E \left(\max_{1 \leq i \leq p} X_i^2 \right) \cdot E \left\{ \sup_{0 \leq \lambda \leq \infty} \left[\sum_i \frac{A_i(X_i - \theta_i)}{A_i + \lambda} \right]^2 \right\} \right)^{1/2} \\ & = O(p^{-\frac{\delta^*}{2(2+\delta^*)}}) = o(1), \end{aligned}$$

which completes the proof.

Proof of Theorem 5.2

With Theorem 5.1 established, the proof of

$$\lim_{p \rightarrow \infty} P \left(l_p(\boldsymbol{\theta}, \hat{\boldsymbol{\theta}}^M) \geq l_p(\boldsymbol{\theta}, \tilde{\boldsymbol{\theta}}^{MOL}) + \varepsilon \right) = 0 \quad \text{for any } \varepsilon > 0$$

is essentially identical to the proof of Theorem 3.2.

Proof of Corollary 5.3

This is a direct consequence of the definition of $\tilde{\boldsymbol{\theta}}^{MOL}$ and Theorem 5.2.

Proof of Theorem 5.4

Since

$$l_p(\boldsymbol{\theta}, \hat{\boldsymbol{\theta}}^M) - l_p(\boldsymbol{\theta}, \tilde{\boldsymbol{\theta}}^{MOL}) \leq 2 \sup_{0 \leq \lambda \leq \infty, |\mu| \leq \max_i |X_i|} \left| \text{SURE}^M(\lambda, \mu) - l_p(\boldsymbol{\theta}, \tilde{\boldsymbol{\theta}}^{\lambda, M}) \right|,$$

the result follows from Theorem 5.1.

Proof of Corollary 5.5

This is a direct consequence of the definition of $\tilde{\boldsymbol{\theta}}^{MOL}$ and Theorem 5.4.

Proof of Corollary 5.6

This is a special case of Corollary 5.5.

Proof of Theorem 6.1

First, we have

$$\left| \text{SURE}^M(\mathbf{b}, \mu) - l_p(\boldsymbol{\theta}, \hat{\boldsymbol{\theta}}^{b, \mu}) \right| \leq \frac{1}{p} \left| \sum_{i=1}^p (X_i^2 - \theta_i^2 - A_i) \right| + \frac{1}{p} \left| \sum_{i=1}^p 2(1 - b_i) (X_i^2 - X_i \theta_i - A_i) \right| + \frac{1}{p} \left| \sum_{i=1}^p 2b_i (X_i - \theta_i) \cdot \mu \right|.$$

Note that the order of b_i is determined by that of A_i , which is not random. The rest of the proof follows essentially the same steps as in that of Theorem 5.1 upon using Lemma A.1.

Proof of Theorem 6.2

With Theorem 6.1 established, the proof of

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

is the same as the proof of Theorem 3.2. Likewise, to show

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

we use the inequality

$$l_p(\boldsymbol{\theta}, \hat{\boldsymbol{\theta}}^{SM}) - l_p(\boldsymbol{\theta}, \hat{\boldsymbol{\theta}}^{\tilde{\mathbf{b}}_p, \hat{\mu}_p}) \leq 2 \sup \left| \text{SURE}^M(\mathbf{b}, \mu) - l_p(\boldsymbol{\theta}, \hat{\boldsymbol{\theta}}^{\mathbf{b}, \mu}) \right|,$$

and Theorem 6.1.

Figure 1.

Comparing the risks of different shrinkage estimators. (1)–(6) correspond to the six simulation examples.

Figure 2.

Plot of the shrinkage factors $\lambda/(\lambda + A)$ or $1 - b$ for all-batters. Four estimators are compared: EB-MM, EB-MLE, the parametric SURE $\hat{\theta}^M$ and the semi-parametric SURE estimator $\hat{\theta}^{SM}$.

Table 1

The limiting risk $\lim_{p \rightarrow \infty} R(\theta, \hat{\theta})$ of different shrinkage estimators. The six columns (1)–(6) correspond to the six simulation examples.

	(1)	(2)	(3)	(4)	(5)	(6)
EBMLE	0.3357	0.0697	0.0775	0.0057	0.2470	0.0775
EBMOM	0.3357	0.0697	0.0755	0.0058	0.2434	0.0755
J-S	0.3632	0.0737	0.0797	0.0056	0.2594	0.0797
Oracle	0.3357	0.0697	0.0540	0.0051	0.1947	0.0540
SURE $\hat{\theta}^{\hat{\tau}}$	0.3357	0.0697	0.0553	0.0051	0.2337	0.0553
SURE $\hat{\theta}^{\hat{g}}$	0.3357	0.0697	0.0540	0.0051	0.1947	0.0540
SURE $\hat{\theta}^{\hat{G}}$	0.3357	0.0697	0.0523	0.0050	0.2335	0.0523
SURE $\hat{\theta}^{\hat{M}}$	0.3357	0.0697	0.0491	0.0050	0.1739	0.0491

Table 2

Prediction errors of batting averages.

	ALL	Pitchers	NonPitchers
naive	1	1	1
Grand mean	0.852	0.127	0.378
Parametric EB-MM	0.593	0.129	0.387
Parametric EB-ML	0.902	0.117	0.398
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
SURE $\hat{\theta}^{\hat{\sigma}}$	0.505	0.123	0.278
SURE $\hat{\theta}^{\hat{M}}$	0.422	0.123	0.282
SPSURE $\hat{\theta}^{\hat{\sigma}}$	0.409	0.081	0.261
SPSURE $\hat{\theta}^{\hat{M}}$	0.419	0.077	0.278

Table 3

Estimation errors for housing prices under different sample sizes.

	sampling 10%		sampling 15%		sampling 20%	
	TSE	TSEP	TSE	TSEP	TSE	TSEP
naive	1.000	1.000	1.000	1.000	1.000	1.000
EB(MM)	0.735	0.748	0.779	0.788	0.816	0.824
EB(ML)	0.734	0.747	0.776	0.785	0.813	0.821
J-S	0.991	0.992	0.994	0.994	0.996	0.996
SURE θ^F	0.746	0.793	0.772	0.820	0.819	0.872
SURE θ^H	0.997	1.074	1.050	1.131	1.100	1.184
SPSURE θ^G	0.556	0.574	0.518	0.534	0.522	0.536
SPSURE θ^M	0.879	0.881	0.863	0.863	0.865	0.865