

EMPIRICAL BAYES ESTIMATES FOR A TWO-WAY CROSS-CLASSIFIED MODEL

BY LAWRENCE D. BROWN¹, GOURAB MUKHERJEE²
AND ASAF WEINSTEIN¹

*University of Pennsylvania, University of Southern California and
Stanford University*

We develop an empirical Bayes procedure for estimating the cell means in an unbalanced, two-way additive model with fixed effects. We employ a hierarchical model, which reflects exchangeability of the effects within treatment and within block but not necessarily between them, as suggested before by Lindley and Smith [*J. R. Stat. Soc., B* **34** (1972) 1–41]. The hyperparameters of this hierarchical model, instead of considered fixed, are to be substituted with data-dependent values in such a way that the point risk of the empirical Bayes estimator is small. Our method chooses the hyperparameters by minimizing an unbiased risk estimate and is shown to be asymptotically optimal for the estimation problem defined above, under suitable conditions. The usual empirical Best Linear Unbiased Predictor (BLUP) is shown to be substantially different from the proposed method in the unbalanced case and, therefore, performs suboptimally. Our estimator is implemented through a computationally tractable algorithm that is scalable to work under large designs. The case of missing cell observations is treated as well.

1. Introduction. Multilevel cross-classified models are pervasive in statistics, with applications ranging from detecting sources of variability in medical research [Goldstein, Browne and Rasbash (2002)] to understanding micro–macro linkages in social studies [Mason, Wong and Entwisle (1983), Zaccarin and Rivellini (2002)]. These models offer a natural and flexible approach to specify meaningful latent structures and, importantly, a systematic way to use all information for simultaneously analyzing the effects of more than one factor [Rasbash and Goldstein (1994)]. Hierarchical cross-classified models have classically been used to decompose the total variability of the response into individual sources and for prediction in random-effects models. Nevertheless, ever since the appearance of the James–Stein estimator [James and Stein (1961)] and its Bayesian interpretation [Lindley (1962), Stein (1962)], the usefulness of such models in estimation problems involving multiple *nonrandom* effects has been well recognized.

Received May 2016; revised February 2017.

¹Supported in part by grant NSF DMS-15-12084.

²Supported in part by the Zumberge individual award from the University of Southern California’s James H. Zumberge faculty research and innovation fund.

MSC2010 subject classifications. Primary 62C12; secondary 62C25, 62F10, 62J07.

Key words and phrases. Shrinkage estimation, empirical Bayes, two-way ANOVA, oracle optimality, Stein’s unbiased risk estimate (SURE), empirical BLUP.

Hierarchical models have been used to facilitate shrinkage estimators in linear regression models since the early 1970s [Efron and Morris (1972)]. In both theoretical and more applied work, various authors have employed hierarchical models to produce estimators that shrink *toward* a subspace [e.g., Sclove (1968), Oman (1982), Jiang, Nguyen and Rao (2011), Tan (2016)] or *within* a subspace [e.g., Lindley and Smith (1972), Rolph (1976), Kou and Yang (2017)]; see Section 2 of the last reference for a discussion on the difference between the two types of resulting estimators. Cross-classified additive models are in a sense the most immediate extension of Stein's canonical example. Specifically, unlike in a general linear model, the symmetries of within-batch effects can be regarded as a priori information, which suggest the use of exchangeable priors, such as those proposed by Lindley and Smith (1972) and Efron and Morris (1973). In the case of balanced design, the properties of resulting shrinkage estimators are by now well understood and have a close relationship to the James–Stein estimator. Indeed, when all cell counts are equal, multiple one-way, homoscedastic estimation problems emerge; for these the James–Stein estimator has optimality properties under many criteria. But in the unbalanced case, the problems of estimating the effects corresponding to different batches are intertwined due to lack of orthogonality in the design matrix. Hence, the situation in the case of unbalanced design is substantially different.

This paper deals with empirical Bayes (EB) estimation of the cell means in the two-way *fixed* effects additive model with *unbalanced* design. We consider a family of Bayes estimators resulting from a normal hierarchical model, which reflects within-batch exchangeability and is indexed by a set of hyperparameters governing the prior. Any corresponding estimator that substitutes *data-dependent* values for the hyperparameters is referred to as an empirical Bayes estimator. We propose an empirical Bayes procedure that is asymptotically optimal for the estimation of the cell means under squared loss. In our asymptotic analysis, the number of row and column levels tends to infinity. Importantly, so-called empirical BLUP (Best Linear Unbiased Predictors) estimators, using the usual maximum-likelihood approach in estimating the hyperparameters, are shown to perform suboptimally in the unbalanced case. Instead of using the maximum-likelihood criterion, we choose the values for the hyperparameters by minimizing an unbiased estimate of the risk (URE), which leads to estimates that are different in an essential way. The proposed approach is appealing in the fixed effects case, because it uses a criterion directly related to the risk instead of using the likelihood under the postulated hierarchical model.

Using the URE criterion to calibrate tuning parameters has been proposed in many previous works and in a broad range of parametric and nonparametric estimation problems [Li (1986), Ghosh, Nickerson and Sen (1987), Donoho et al. (1995), Johnstone and Silverman (2004), Candès, Sing-Long and Trzasko (2013), to name a few]. Recently, Xie, Kou and Brown (2012) employed URE minimization to construct alternative empirical Bayes estimators to the usual ones in the Gaussian mean problem with known *heteroscedastic* variances and showed that it produces asymptotically uniformly better estimates.

Our work can be viewed as a generalization of [Xie, Kou and Brown \(2012\)](#) from the one-way unbalanced layout to the two-way unbalanced layout. The two-way unbalanced problem presents various new challenges. The basis for the difference, of course, lies in the fact that the two-way case imposes structure on the mean vector, which is nontrivial to handle due to missingness or imbalance in the design. Some of the implications are that the analysis of the performance of EB methods is substantially more involved than in the one-way scenario. In addition, the computation of the URE estimator, which is trivial in the one-way scenario, becomes a cause of concern, especially with a growing number of factor levels. We offer an implementation of the corresponding URE estimate that in the all-cells-filled case has comparable computational performance to that of the standard empirical BLUP in the popular R package `lme4` of [Bates \(2010\)](#). Our theoretical analysis of the two-way case differs in fundamental aspects from the optimality proof techniques usually used in the one-way Normal mean estimation problem. To tackle the difficulties encountered in the two-way problem, where computations involving matrices are generally unavoidable, we developed a flexible approach for proving asymptotic optimality based on efficient pointwise risk estimation; this essentially reduces our task to controlling the moments of Gaussian quadratic forms.

We would also like to point out that the current work is different from the recent extensions of [Kou and Yang \(2017\)](#) of the URE approach to the general Gaussian linear model. While the setup considered in that paper formally includes our setup as a special case, their results have limited implications for additive cross-classified models. For example, the covariance matrix used in their second level of the hierarchy is not general enough to accommodate the within-batch exchangeable structure we employ and is instead governed by a single hyperparameter. Moreover, their asymptotic results require keeping the dimension of the linear subspace fixed, whereas the number of factor levels is increasing in our setup.

Organization of the paper. In Section 2, we describe our estimation setup: we begin with the all-cells-filled situation, and then present a more general model which allows missing observations. In Section 3, we show that our proposed estimation methodology is asymptotically optimal and is capable of recovering the directions and magnitude for optimal shrinkage. In Section 4, we report the results from extensive numerical experiments, and also demonstrate the applicability of our proposed method on a real-world problem concerning the estimation of the average nitrate levels in water sources based on location and time of day.

2. Model setup and estimation methods.

2.1. *Basic model and Bayes estimators.* Consider the following basic two-way cross-classified additive model with fixed effects:

$$(1) \quad \begin{aligned} y_{ij} &= \eta_{ij} + \varepsilon_{ij}, & \eta_{ij} &= \mu + \alpha_i + \beta_j, \\ \varepsilon_{ij} &\sim N(0, \sigma^2 K_{ij}^{-1}), & 1 \leq i \leq r, 1 \leq j \leq c, \end{aligned}$$

K_{ij} is the count in the (i, j) th cell; $\sigma^2 > 0$ is assumed to be known; and ε_{ij} is an independent Gaussian noise term, which can be considered as the average of K_{ij} homoscedastic error terms. Throughout the paper, we write vectors in bold and matrices in capital letters. The parameters $\mu, \alpha = (\alpha_1, \dots, \alpha_r)^\top, \beta = (\beta_1, \dots, \beta_c)^\top$ are not identifiable without imposing further side conditions, but the vector of cell means $\eta = (\eta_{11}, \eta_{12}, \dots, \eta_{rc})^\top$ always is. Our goal is to estimate η under sum-of-squares loss:

$$(2) \quad L_{r,c}(\eta, \hat{\eta}) = \frac{1}{rc} \|\hat{\eta} - \eta\|^2 = \frac{1}{rc} \sum_{i=1}^r \sum_{j=1}^c (\hat{\eta}_{ij} - \eta_{ij})^2.$$

In model (1), α_i and β_j will be referred to as the i th row and the j th column effects, respectively. In the *all-cells-filled* model, $K_{ij} \geq 1$ for $1 \leq i \leq r$ and $1 \leq j \leq c$. The more general model, which allows empty cells, is presented in Section 2.3. We would like to emphasize the focus in this section on the loss (2) rather than the weighted quadratic loss $L_{r,c}^{\text{wgt}}(\eta, \hat{\eta}) = (rc)^{-1} \sum_{i=1}^r \sum_{j=1}^c K_{ij} (\hat{\eta}_{ij} - \eta_{ij})^2$, which is sometimes called the *prediction* loss, and under which asymptotically optimal estimation has been investigated before [Dicker (2013)]. Nevertheless, in later sections results are presented for a general quadratic loss, which includes the weighted loss as a special case.

The usual estimator of η is the weighted least squares (WLS) estimator, also the maximum-likelihood estimator under (1). The WLS estimator is unbiased and minimax but can be substantially improved on by shrinkage estimators, particularly when $r, c \rightarrow \infty$ [Draper and Van Nostrand (1979)]. As the starting point for the shrinkage estimators proposed in this paper, we consider a family of Bayes estimators with respect to the conjugate prior

$$\alpha_1, \dots, \alpha_r \stackrel{\text{i.i.d.}}{\sim} N(0, \sigma_A^2), \quad \beta_1, \dots, \beta_c \stackrel{\text{i.i.d.}}{\sim} N(0, \sigma_B^2),$$

where σ_A^2, σ_B^2 are hyperparameters. This prior extends the conjugate normal prior in the one-way case and was proposed by Lindley and Smith (1972) to reflect exchangeability within rows and columns separately.³ Employing this prior is standard in Bayesian Analysis-of-Variance [e.g., Gelman et al. (2004), Chapter 15.6; Gelman (2005)]. In vector form, the two-level hierarchical model can be written as

$$(3) \quad \begin{aligned} \text{Level 1:} & \quad \mathbf{y}|\eta \sim N_{rc}(\eta, \sigma^2 M), & \quad \eta = \mathbf{1}\mu + Z\theta, & \quad \theta^\top = (\alpha^\top, \beta^\top), \\ \text{Level 2:} & \quad \theta \sim N_{r+c}(0, \sigma^2 \Lambda \Lambda^\top), \end{aligned}$$

³As in Lindley and Smith (1972) and Gelman (2005), we refer to within-batch effects in our model as exchangeable, although they are really i.i.d.

where $M = \text{diag}(K_{11}^{-1}, K_{12}^{-1}, \dots, K_{rc}^{-1})$ is an $rc \times rc$ matrix and $Z = [Z_A Z_B]$ with $Z_A = I_r \otimes 1_c$ and $Z_B = 1_r \otimes I_c$. The $(r + c) \times (r + c)$ matrix

$$\Lambda = \begin{bmatrix} \sqrt{\lambda_A} I_r & 0 \\ 0 & \sqrt{\lambda_B} I_c \end{bmatrix}$$

is written in terms of the *relative* variance components $\lambda_A = \sigma_A^2/\sigma^2$ and $\lambda_B = \sigma_B^2/\sigma^2$. Henceforth, for notational simplicity, the dependence of Λ on the model hyperparameters will be kept implicit. The marginal covariance of \mathbf{y} in (5) is $\sigma^2 \Sigma$ where

$$(4) \quad \Sigma = Z \Lambda \Lambda^T Z^T + M = \lambda_A Z_A Z_A^T + \lambda_B Z_B Z_B^T + M.$$

The following is a standard result and is proved in Section S.2 of the supplementary materials [Brown, Mukherjee and Weinstein (2018)].

LEMMA 2.1. *For any fixed $\mu \in \mathbb{R}$ and nonnegative λ_A, λ_B , the Bayes estimate of $\boldsymbol{\eta}$ in (3) is given by*

$$(5) \quad E[\boldsymbol{\eta}|\mathbf{y}] = \mathbf{y} - M \Sigma^{-1}(\mathbf{y} - \mathbf{1}\mu),$$

where Σ depends on the hyperparameters λ_A, λ_B through Λ .

Instead of fixing the values of $\mu, \lambda_A, \lambda_B$ in advance, we may now return to model (1) and consider the parametric family of estimators obtained by fixing the values of the hyperparameters in (5),

$$(6) \quad \begin{aligned} \mathcal{S}(a, b) = \{ & \hat{\boldsymbol{\eta}}(\mu, \lambda_A, \lambda_B) = \mathbf{y} - M \Sigma^{-1}(\mathbf{y} - \mathbf{1}\mu) : \mu \in [a, b], \\ & \lambda_A \geq 0, \lambda_B \geq 0 \} \end{aligned}$$

for $a, b \in \mathbb{R}$. Note that above, μ is restricted to lie in $[a, b]$. In practice, we will consider only estimators in $\mathcal{S}[\tau] := \mathcal{S}(\hat{a}_\tau, \hat{b}_\tau)$ where $\hat{a}_\tau(\mathbf{y}) = \text{quantile}\{y_{ij} : 1 \leq i \leq r, 1 \leq j \leq c; \tau/2\}$ and $\hat{b}_\tau(\mathbf{y}) = \text{quantile}\{y_{ij} : 1 \leq i \leq r, 1 \leq j \leq c; 1 - \tau/2\}$, the $\tau/2$ and $(1 - \tau/2)$ quantiles of the observations. The constraint on the location hyperparameter μ is imposed for technical reasons but is moderate enough to be well justified. Indeed, an estimator that shrinks toward a point that lies near the periphery or outside the range of the data is at the risk of being nonrobust; it also seems as an undesirable choice for a Bayes estimator corresponding to (3), which models $\boldsymbol{\alpha}$ and $\boldsymbol{\beta}$ as having zero means. In practice, τ may be taken to be 1% or 5%.

An empirical Bayes estimator is any estimator that plugs data-dependent values $\hat{\mu}, \hat{\lambda}_A, \hat{\lambda}_B$ into (5), with the restriction that $\hat{\mu}$ is in the allowable range. In the next section, we propose a specific criterion for estimating the hyperparameters.

2.2. *Empirical Bayes estimators.* The usual empirical Bayes estimators are derived relying on hierarchical model (3). The fixed effect μ and the relative variance components λ_A and λ_B are treated as unknown fixed parameters to be estimated based on the marginal distribution of \mathbf{y} and substituted into (5). For any set of estimates substituted for λ_A and λ_B , the general mean μ is typically estimated by generalized least squares, producing an empirical version of what is known as the Best Linear Unbiased Predictor (BLUP). There is extensive literature on the estimation of the variance components. For a textbook account, see, for example, Chapters 5 and 6 of Searle, Casella and McCulloch (1992); and Section 7.3.2 in Fahrmeir et al. (2013), who also discuss Bayesian inference in Section 7.4.2. Computational issues are discussed in Harville (1977). The main methods for estimation of variance components in linear mixed models are maximum-likelihood (ML), restricted maximum-likelihood (REML) and the ANOVA methods (Method-of-Moments), including the three original ANOVA methods of Henderson [Henderson (1984)]. Here, we concentrate on the commonly used maximum-likelihood estimates, which are implemented in the popular R package lme4 [Bates (2010)]. If $\mathcal{L}(\mu, \lambda_A, \lambda_B; \mathbf{y})$ denotes the marginal likelihood of \mathbf{y} according to (3), then the maximum-likelihood (ML) estimates are

$$(7) \quad (\hat{\mu}^{ML}, \hat{\lambda}_A^{ML}, \hat{\lambda}_B^{ML}) = \underset{\mu \in [\hat{a}_\tau, \hat{b}_\tau], \lambda_A \geq 0, \lambda_B \geq 0}{\arg \max} \mathcal{L}(\mu, \lambda_A, \lambda_B; \mathbf{y}).$$

The corresponding empirical Bayes estimator is $\hat{\eta}^{ML} = \hat{\eta}(\hat{\mu}^{ML}, \hat{\lambda}_A^{ML}, \hat{\lambda}_B^{ML})$ and will be referred to as EBMLE (for Empirical Bayes Maximum-Likelihood).

LEMMA 2.2. *ML estimates defined in (7) satisfy the following equations:*

$$(8) \quad \hat{\mu} = \hat{\mu}_1 \cdot I\{\hat{\mu}_1 \in [\hat{a}_\tau, \hat{b}_\tau]\} + \hat{a}_\tau \cdot I\{\hat{\mu}_1 < \hat{a}_\tau\} + \hat{b}_\tau \cdot I\{\hat{\mu}_1 > \hat{b}_\tau\},$$

where $\hat{\mu}_1 = (\mathbf{1}^\top \hat{\Sigma}^{-1} \mathbf{y}) / (\mathbf{1}^\top \hat{\Sigma}^{-1} \mathbf{1})$; if $\hat{\mu}_1 \in [\hat{a}_\tau, \hat{b}_\tau]$ and if $\hat{\lambda}_a, \hat{\lambda}_b$ are both strictly positive, they satisfy

$$(9) \quad \begin{aligned} & \text{tr}(\hat{\Sigma}^{-1} Z_A Z_A^\top) - \sigma^{-2} \mathbf{y}^\top (I - \hat{P})^\top \hat{\Sigma}^{-1} Z_A Z_A^\top \hat{\Sigma}^{-1} (I - \hat{P}) \mathbf{y} = 0, \\ & \text{tr}(\hat{\Sigma}^{-1} Z_B Z_B^\top) - \sigma^{-2} \mathbf{y}^\top (I - \hat{P})^\top \hat{\Sigma}^{-1} Z_B Z_B^\top \hat{\Sigma}^{-1} (I - \hat{P}) \mathbf{y} = 0, \end{aligned}$$

where $\hat{P} = \mathbf{1}(\mathbf{1}^\top \hat{\Sigma}^{-1} \mathbf{1})^{-1} \mathbf{1}^\top \hat{\Sigma}^{-1}$.

The derivation is standard and provided in Section S.2.1 of the supplement, which also contains the estimating equation for the case when $\hat{\mu}_1 \notin [\hat{a}_\tau, \hat{b}_\tau]$. If the solution to the estimating equations (9) includes a negative component, adjustments are needed in order to produce the maximum-likelihood estimates of the scale hyperparameters [see McCulloch and Searle (2001), Section 2.2b-iii for a discussion of the one-way case].

We propose an alternative method for estimating the shrinkage parameters. Following the approach of Xie, Kou and Brown (2012), for fixed $\tau \in (0, 1]$ we choose

the shrinkage parameters by minimizing unbiased risk estimate (URE) over estimators $\hat{\eta}(\mu, \lambda_A, \lambda_B)$ in $\mathcal{S}(\tau)$. By Lemma S.2.2 of the supplement, an unbiased estimate of the risk, $R_{r,c}(\eta, \hat{\eta}(\mu, \lambda_A, \lambda_B)) \triangleq (rc)^{-1} \mathbb{E} \|\hat{\eta}(\mu, \lambda_A, \lambda_B) - \eta\|_2^2$ is given by $\text{URE}(\mu, \lambda_A, \lambda_B)$, which equals

$$(10) \quad \frac{1}{rc} \{ \sigma^2 \text{tr}(M) - 2\sigma^2 \text{tr}(\Sigma^{-1}M^2) + (\mathbf{y} - \mathbf{1}\mu)^\top [\Sigma^{-1}M^2\Sigma^{-1}](\mathbf{y} - \mathbf{1}\mu) \}.$$

Hence we propose to estimate the tuning parameters of the class $\mathcal{S}(\tau)$ by

$$(11) \quad (\hat{\mu}^U, \hat{\lambda}_A^U, \hat{\lambda}_B^U) = \underset{\mu \in [\hat{a}_\tau, \hat{b}_\tau], \lambda_A \geq 0, \lambda_B \geq 0}{\text{arg min}} \quad \text{URE}(\mu, \lambda_A, \lambda_B).$$

The corresponding empirical Bayes estimator is $\hat{\eta}^{\text{URE}} = \hat{\eta}(\hat{\mu}^U, \hat{\lambda}_A^U, \hat{\lambda}_B^U)$. As in the case of maximum likelihood estimation, there is no closed-form solution to (11), but we can characterize the solutions by the corresponding estimating equations.

LEMMA 2.3. *URE based estimates of (11) satisfy the following estimating equations:*

$$(12) \quad \hat{\mu} = \hat{\mu}_1 \cdot I\{\hat{\mu}_1 \in [\hat{a}_\tau, \hat{b}_\tau]\} + \hat{a}_\tau \cdot I\{\hat{\mu}_1 < \hat{a}_\tau\} + \hat{b}_\tau \cdot I\{\hat{\mu}_1 > \hat{b}_\tau\},$$

where $\hat{\mu}_1 = (\mathbf{1}^\top [\hat{\Sigma}^{-1}M^2\hat{\Sigma}^{-1}]\mathbf{y}) / (\mathbf{1}^\top [\hat{\Sigma}^{-1}M^2\hat{\Sigma}^{-1}]\mathbf{1})$; if $\hat{\mu}_1 \in [\hat{a}_\tau, \hat{b}_\tau]$ and if $\hat{\lambda}_a, \hat{\lambda}_b$ are both strictly positive, they satisfy

$$(13) \quad \begin{aligned} & \text{tr}(\hat{\Sigma}^{-1}Z_A Z_A^\top \hat{\Sigma}^{-1}M^2) \\ & - \sigma^{-2} \mathbf{y}^\top (I - \hat{P})^\top \hat{\Sigma}^{-1}Z_A Z_A^\top \hat{\Sigma}^{-1}M^2\hat{\Sigma}^{-1}(I - \hat{P})\mathbf{y} = 0, \\ & \text{tr}(\hat{\Sigma}^{-1}Z_B Z_B^\top \hat{\Sigma}^{-1}M^2) \\ & - \sigma^{-2} \mathbf{y}^\top (I - \hat{P})^\top \hat{\Sigma}^{-1}Z_B Z_B^\top \hat{\Sigma}^{-1}M^2\hat{\Sigma}^{-1}(I - \hat{P})\mathbf{y} = 0, \end{aligned}$$

where $\hat{P} = \mathbf{1}(\mathbf{1}^\top [\hat{\Sigma}^{-1}M^2\hat{\Sigma}^{-1}]\mathbf{1})^{-1} \mathbf{1}^\top \hat{\Sigma}^{-1}M^2\hat{\Sigma}^{-1}$.

The derivation is provided in Section S.2.1 of the supplementary materials. Comparing the two systems of equations (9) and (13) without substituting the value of μ , it can be seen that the URE equation involves an extra term $\hat{\Sigma}^{-1}M^2$ in both summands of the left-hand side, as compared to the ML equation. The estimating equations therefore imply that the ML and URE solutions may differ when the design is unbalanced (see Section S.5 of the supplement for a discussion of the balanced case). In Section 3, we show that the URE estimate $\hat{\eta}^{\text{URE}}$ is asymptotically optimal as $r, c \rightarrow \infty$, and the numerical simulations in Section 4 demonstrate that in certain situations EBMLE performs significantly worse.

REMARK 1. To compute the hyperparameter estimates by the URE method, one could attempt to solve the estimating equations in (13), which have no closed-form solution in general. For example, one could fix the value of λ_A to some initial

positive value and solve the first equation in λ_B . Then plug the solution into the second equation and solve for λ_A , and keep iterating between the two equations until convergence. If this approach is taken, a nontrivial issue to overcome will be obtaining the actual minimizing values λ_A and λ_B when one of the solutions to (13) is negative. Another issue will be ascertaining the global optimality of the solutions, as URE is not necessarily convex in $(\mu, \lambda_A, \lambda_B)$. To bypass these issues, we minimize URE by conducting a grid search on (λ_A, λ_B) , and μ is subsequently estimated by (12). For efficiency in handling large data sets, our implementation adopts some of the key computational elements from the `lme4` package [Section 5.4, Bates (2010)]; Details are provided in Section A.3 of the Appendix.

REMARK 2. A comment is in order regarding shrinkage estimators for the general homoscedastic linear model. Note that model (1) could be written for individual, homoscedastic observations (with an additional subscript k) instead of for the cell averages. With the corresponding design matrix, the two-way additive model is therefore a special case of the homoscedastic Gaussian linear model, $\mathbf{y} \sim N_n(X\boldsymbol{\gamma}, \sigma^2 I)$, where $X \in \mathbb{R}^{n \times p}$ a known matrix and $\boldsymbol{\gamma} \in \mathbb{R}^p$ is the unknown parameter. Thus, the various so-called Stein-type shrinkage methods that have been proposed for estimating $\boldsymbol{\gamma}$ can also be applied to our problem. Specifically, a popular approach is to reduce the problem of estimating $\boldsymbol{\gamma}$ to the problem of estimating the mean of a p -dimensional *heteroscedastic* normal vector with known variances [see, e.g., Johnstone (2011), Section 2.9] by applying orthogonal transformations to the parameter $\boldsymbol{\gamma}$ and data \mathbf{y} . Thereafter, Stein-type shrinkage estimators can be constructed as empirical Bayes rules by putting a prior which is either i.i.d. on the transformed coordinates or i.i.d. on the original coordinates of the parameter [Rolph (1976), referred to priors of the first type as *proportional* priors and to those of the second kind as *constant* priors]. In the case of factorial designs, however, neither of these choices is very sensible, because they do not capture the (within-batch) symmetries of cross-classified models. Hence, procedures relying on models that take exchangeability into account can potentially achieve a significant and meaningful reduction in estimation risk. The methodology we develop here incorporates the exchangeable structure in (3).

2.3. *A model with missing cells.* A more general model than (1) allows some cells to be empty. Hence, consider

$$(14) \quad y_{ij} = \eta_{ij} + \varepsilon_{ij}, \quad \eta_{ij} = \mu + \alpha_i + \beta_j, \varepsilon_{ij} \sim N(0, \sigma^2 K_{ij}^{-1}), (i, j) \in \mathcal{E},$$

where $\mathcal{E} = \{(i, j) : K_{ij} \geq 1\} \subseteq \{1, \dots, r\} \times \{1, \dots, c\}$ is the set of indices corresponding to the nonempty cells. As before, $\sigma^2 > 0$ is assumed to be known. Our goal is in general to estimate all cell means that are *estimable* under (14) rather than only the means of observed cells. For ease of presentation and without loss

of generality, from here on we assume that \mathcal{E} is a *connected* design⁴ so that all rc cell means are estimable.

We will need some new notation to distinguish between $\mathbb{E}[\mathbf{y}] \in \mathbb{R}^{|\mathcal{E}|}$ and the rc vector consisting of all cell means. In general, the notation in (3) is reserved for quantities associated with the observed variables. As before, $\boldsymbol{\theta} = (\boldsymbol{\alpha}^\top, \boldsymbol{\beta}^\top)^\top$. The matrix $M = \text{diag}(K_{ij}^{-1} : (i, j) \in \mathcal{E})$, where the indices of diagonal elements are in lexicographical order. Let $Z_c = [\mathbf{1}_{rc} I_R \otimes \mathbf{1}_C \mathbf{1}_R \otimes I_C]$ be the $rc \times (r + c + 1)$ design matrix associated with the unobserved complete model. The $|\mathcal{E}| \times (r + c + 1)$ “observed” design matrix Z is obtained from Z_c by deleting the subset of rows corresponding to \mathcal{E}^c . With the new definitions for Z and M , we define Σ by (4). Finally, let $\boldsymbol{\eta}_c = Z_c \boldsymbol{\theta} \in \mathbb{R}^{rc}$ be the vector of all estimable cell means and $\boldsymbol{\eta} = Z \boldsymbol{\theta} \in \mathbb{R}^{|\mathcal{E}|}$ be the vector of cell means for only the observed cells of (14). Hence, assuming \mathcal{E} corresponds to a connected design, we consider estimating $\boldsymbol{\eta}_c$ under the normalized sum-of-squares loss. Note that since $\boldsymbol{\eta}_c$ is estimable, it must be a linear function of $\boldsymbol{\eta}$. The following lemma is an application of the basic theory of estimable functions and is proved in the Section S.2 of the supplementary materials.

LEMMA 2.4. *If $\boldsymbol{\eta}_c$ is estimable, then $\boldsymbol{\eta}_c = Z_c(Z^\top Z)^- Z^\top \boldsymbol{\eta}$, where $(Z^\top Z)^-$ is any generalized inverse of $Z^\top Z$.*

In particular, writing Z^\dagger for the Moore–Penrose pseudo-inverse of Z , we therefore have $\boldsymbol{\eta}_c = Z_c Z^\dagger \boldsymbol{\eta}$. Thus, we can rewrite the loss function as

$$(15) \quad L_{r,c}(\boldsymbol{\eta}_c, \hat{\boldsymbol{\eta}}_c) \triangleq \frac{1}{rc} \|\hat{\boldsymbol{\eta}}_c - \boldsymbol{\eta}_c\|^2 = \frac{1}{rc} (\hat{\boldsymbol{\eta}} - \boldsymbol{\eta})^\top Q (\hat{\boldsymbol{\eta}} - \boldsymbol{\eta}) = L_{r,c}^Q(\boldsymbol{\eta}, \hat{\boldsymbol{\eta}}),$$

where

$$(16) \quad Q = (Z_c Z^\dagger)^\top Z_c Z^\dagger.$$

In other words, the problem of estimating $\boldsymbol{\eta}_c$ under sum-of-squares loss can be recast as the problem of estimating $\boldsymbol{\eta} = \mathbb{E}[\mathbf{y}]$ under appropriate quadratic loss. This allows us to build on the techniques developed in the previous section and extend their applicability to the loss in (15). The standard unbiased estimator of $\boldsymbol{\eta}_c$ is the weighted least squares estimator. The form of the Bayes estimator for $\boldsymbol{\eta}$ under (3) is not affected by the generalized quadratic loss $L_{r,c}^Q$ and is still given by (5), with M, Σ^{-1} as defined in the current section. As before, for any pre-specified $\tau \in (0, 1]$ we consider the class of estimators $\mathcal{S}[\tau] := \mathcal{S}(\hat{a}_\tau, \hat{b}_\tau)$, with $\mathcal{S}(a, b)$ defined in (6). The EBMLE estimates the hyperparameters $\mu, \lambda_A, \lambda_B$ based on the marginal likelihood \mathbf{y} according to (3), where M, Σ^{-1} are as defined in the current

⁴A design is disconnected if there is a partition of the row effects into nonempty subsets such that any two members of distinct subsets never appear with the same column effect. A design is called connected if it is not disconnected.

section. As shown in Lemma S.2.3 of the supplement, an unbiased estimator of the point risk corresponding to (15),

$$R_{r,c}^Q(\boldsymbol{\eta}, \widehat{\boldsymbol{\eta}}(\mu, \lambda_A, \lambda_B)) \triangleq \mathbb{E}\{L_{r,c}^Q(\boldsymbol{\eta}, \widehat{\boldsymbol{\eta}}(\mu, \lambda_A, \lambda_B))\},$$

is given by $\text{URE}^Q(\mu, \lambda_A, \lambda_B)$ which is evaluated as

$$(17) \quad \begin{aligned} \text{URE}^Q(\mu, \lambda_A, \lambda_B) &= (rc)^{-1}[\sigma^2 \text{tr}(QM) - 2\sigma^2 \text{tr}(\Sigma^{-1}MQM) \\ &\quad + (\mathbf{y} - \mu\mathbf{1})^\top [\Sigma^{-1}MQM\Sigma^{-1}](\mathbf{y} - \mu\mathbf{1})]. \end{aligned}$$

The URE estimates of the tuning parameters are

$$(18) \quad (\widehat{\mu}^{UQ}, \widehat{\lambda}_A^{UQ}, \widehat{\lambda}_B^{UQ}) = \underset{\mu \in [\widehat{a}_\tau, \widehat{b}_\tau], \lambda_A \geq 0, \lambda_B \geq 0}{\text{arg min}} \quad \text{URE}^Q(\mu, \lambda_A, \lambda_B),$$

and the corresponding EB estimate is $\widehat{\boldsymbol{\eta}}^{\text{URE}} = \widehat{\boldsymbol{\eta}}(\widehat{\mu}^{UQ}, \widehat{\lambda}_A^{UQ}, \widehat{\lambda}_B^{UQ})$. Equivalently, the estimate for $\boldsymbol{\eta}_c$ is $\widehat{\boldsymbol{\eta}}_c^{\text{URE}} = Z_c Z_c^\dagger \widehat{\boldsymbol{\eta}}(\widehat{\mu}^{UQ}, \widehat{\lambda}_A^{UQ}, \widehat{\lambda}_B^{UQ})$. The estimating equations for the URE as well as ML estimates of $\mu, \lambda_A, \lambda_B$ can be derived similar to those in the all-cells-filled model.

3. Risk properties and asymptotic optimality of the URE based estimator.

We now present the results that establish the optimality properties of our proposed URE-based estimator. We present the result for the general quadratic loss $L_{r,c}^Q$ of the previous section with the matrix Q defined in (16). Replacing Q with I_{rc} will give the results for the all-cells-filled model (1), which are explained separately. In proving our theoretical results, we make the following assumptions:

A1. *On the parameter space:* We assume that the parameter $\boldsymbol{\eta}_c$ in the complete model is estimable and satisfies the following second-order moment condition:

$$(A1) \quad \lim_{r,c \rightarrow \infty} \frac{1}{rc} \sum_{i=1}^r \sum_{j=1}^c \eta_{i,j}^2 < \infty.$$

This assumption is very mild, and similar versions are widely used in the EB literature [see Assumption C' of Xie, Kou and Brown (2012)]. It mainly facilitates a shorter technical proof and can be avoided by considering separate analyses of the extreme cases.

A2. *On the design matrix:* Denoting the largest eigenvalue of a matrix A by $\lambda_1(A)$, the matrix Q in (16) is assumed to satisfy

$$(A2) \quad \lim_{r,c \rightarrow \infty} (rc)^{-1/8} (\log(rc))^2 v_{r,c} \lambda_1(Q) = 0,$$

where $v_{r,c} = \max\{K_{ij} : (i, j) \in \mathcal{E}\} / \min\{K_{ij} : (i, j) \in \mathcal{E}\}$. As shown in Lemma S.3.1 of the supplement, $\lambda_1(Q)$ equals the largest eigenvalue of $(Z_c^\top Z_c)(Z^\top Z)^\dagger$. Intuitively, it represents the difference in information between

the observed data matrix and the complete data matrix Z_c . If there are many empty cells, $\lambda_1((Z_c^T Z_c)(Z^T Z)^\dagger)$ will be large and may violate the above condition. On the contrary, in the case of the completely observed data we have $\lambda_1(Q) = 1$ (see Lemma S.3.1). Thus, in that case the assumption reduces to $\lim_{r,c \rightarrow \infty} (rc)^{-1/8} (\log(rc))^2 \nu_{r,c} = 0$. This condition amounts to controlling in some sense the extent of imbalance in the number of observations procured per cell. Here, we are allowing the imbalance in the design to asymptotically grow to infinity but at a slower rate than $(rc)^{1/8} / (\log(rc))^2$. This assumption on the design matrix is essential for our asymptotic optimality proofs. Section A.1 of the Appendix shows its role in our proofs and a detailed discussion about it is provided in the supplementary materials.

Asymptotic results. Our decision theoretic optimality results depend on the following pointwise approximation of the true loss of estimators in $\mathcal{S}[\tau]$ by the URE methodology. Consider the following design-dependent quantities:

$$(19) \quad d_{r,c} = m_{r,c}^7 \nu_{r,c}^3 \lambda_1^3(Q), \quad m_{r,c} = \log(rc).$$

Note that, as $r, c \rightarrow \infty$, by Assumption A2, $d_{r,c} \nu_{r,c} \lambda_1(Q) = o(\sqrt{rc})$ which will be used afterward in the proofs for bounding the aforementioned pointwise loss approximation error. The following theorem shows that the unbiased risk estimator approximates the true loss at an asymptotic L_1 error rate smaller than $d_{r,c}^{-1}$ for every hyperparameter value in the set where λ_A, λ_B is nonnegative and the location hyperparameter μ is restricted to the interval $[-m_{r,c}, m_{r,c}]$, which grows as r, c increase.

THEOREM 3.1. *Under Assumptions A1–A2, we have*

$$\lim_{\substack{r \rightarrow \infty \\ c \rightarrow \infty}} d_{r,c} \cdot \left\{ \sup_{\substack{|\mu| \leq m_{r,c} \\ \lambda_A, \lambda_B \geq 0}} \mathbb{E} | \text{URE}_{r,c}^Q(\mu, \lambda_A, \lambda_B) - L_{r,c}^Q(\eta, \hat{\eta}(\mu, \lambda_A, \lambda_B)) | \right\} = 0.$$

The set of all hyperparameters considered in $\mathcal{S}[\tau]$ differs from hyper-parameter set considered in the above theorem, as there μ was restricted to be in the data-dependent set $[\hat{a}_\tau, \hat{b}_\tau]$. However, as $r, c \rightarrow \infty$, $[\hat{a}_\tau, \hat{b}_\tau]$ is asymptotically contained in $[-m_{r,c}, m_{r,c}]$ (see Lemma A.4), so Theorem 3.1 asymptotically covers all hyperparameters considered in $\mathcal{S}[\tau]$ for any $\tau \in (0, 1]$. This explains intuitively why in choosing the hyperparameters by minimizing an unbiased risk estimate as in (17), we can expect the resulting estimate $\hat{\eta}(\hat{\mu}^{U_Q}, \hat{\lambda}_A^{U_Q}, \hat{\lambda}_B^{U_Q})$ to have competitive performance.

Decision theoretic optimality. To compare the asymptotic performance of our proposed estimate, we define the oracle loss (OL) hyperparameter as

$$(\tilde{\mu}^{\text{OL}}, \tilde{\lambda}_A^{\text{OL}}, \tilde{\lambda}_B^{\text{OL}}) = \arg \min_{\mu \in [\hat{a}_\tau, \hat{b}_\tau]; \lambda_A, \lambda_B \geq 0} L^Q(\eta, \hat{\eta}(\mu, \lambda_A, \lambda_B))$$

and the corresponding oracle rule

$$(20) \quad \tilde{\eta}_c^{OL} = Z_c Z_c^\dagger \hat{\eta}(\tilde{\mu}^{OL}, \tilde{\lambda}_A^{OL}, \tilde{\lambda}_B^{OL}).$$

Note that the oracle rule depends on the unknown cell means η_c and is therefore not a “legal” estimator. It serves as the theoretical benchmark for the minimum attainable error by any possible estimator: by its definition, no EB estimator in our class can have smaller risk than η_c^{OL} . The following two theorems show that our proposed URE-based estimator performs asymptotically nearly as well as the oracle loss estimator. The results hold for any class $\mathcal{S}[\tau]$ where $\tau \in (0, 1]$. These results are in terms of the usual quadratic loss on the vector of all cell-means. Note that, based on our formulation of the problem in Sections 2.1 and 2.3, both Theorems 3.2 and 3.3 simultaneously cover the missing and all-cells-filled model.

THEOREM 3.2. *Under Assumptions A1–A2, for any $\varepsilon > 0$ we have*

$$\lim_{\substack{r \rightarrow \infty \\ c \rightarrow \infty}} P\{L_{r,c}(\eta_c, \hat{\eta}_c^{URE}) \geq L_{r,c}(\eta_c, \tilde{\eta}_c^{OL}) + \varepsilon\} = 0.$$

The next theorem asserts that under the same conditions, the URE-based estimator is asymptotically as good as the oracle estimator in terms of risk.

THEOREM 3.3. *Under Assumptions A1–A2, the following holds:*

$$\lim_{\substack{r \rightarrow \infty \\ c \rightarrow \infty}} R_{r,c}(\eta_c, \hat{\eta}_c^{URE}) - \mathbb{E}[L_{r,c}(\eta_c, \tilde{\eta}_c^{OL})] = 0.$$

Finally, as the oracle performs better than any empirical Bayes estimator associated with $\mathcal{S}[\tau]$, a consequence of the above two theorems is that the URE-based estimator cannot be improved by any other such empirical Bayes estimator.

COROLLARY 3.1. *Under Assumptions A1–A2, it holds that for any estimator $\hat{\eta}(\hat{\mu}, \hat{\lambda}_A, \hat{\lambda}_B)$ corresponding to the class $\mathcal{S}[\tau]$ we have:*

- (a) $\lim_{r \rightarrow \infty, c \rightarrow \infty} P\{L_{r,c}^Q(\eta, \hat{\eta}^{URE}) \geq L_{r,c}^Q(\eta, \hat{\eta}(\hat{\mu}, \hat{\lambda}_A, \hat{\lambda}_B)) + \varepsilon\} = 0.$
- (b) $\limsup_{r \rightarrow \infty, c \rightarrow \infty} R_{r,c}^Q(\eta, \hat{\eta}^{URE}) - R_{r,c}^Q(\eta, \hat{\eta}(\hat{\mu}, \hat{\lambda}_A, \hat{\lambda}_B)) \leq 0.$

Unlike the above two theorems, this corollary is based on the quadratic loss L^Q . It emphasizes the nature of our estimating class $\mathcal{S}[\tau]$. In Section 2, we saw that the EBMLE and URE generally produce different solutions in unbalanced designs; combined with Corollary (3.1), this implies that, asymptotically EBMLE generally does not achieve the optimal risk of an EB estimator corresponding to the class $\mathcal{S}[\tau]$ (otherwise the EBML estimate for η would have to be very close to the URE estimate).

Proof Overview. The pointwise loss estimation result of Theorem 3.1 is proved by a moment-based concentration approach, which translates the problem into bounding moments of Gaussian quadratic forms involving matrices with possibly dependent rows and columns. The following two lemmas, which are used in proving Theorem 3.1, display our moment-based convergence approach, where the concentration of relevant quantities about their respective mean is proved. To prove Theorem 3.1, we first show (Lemma 3.1) that for each estimator in $\mathcal{S}[\tau]$ that shrinks toward the origin (i.e., with μ set to 0) the URE methodology estimates the risk in L_2 norm below the desired error rate. Thereafter, in Lemma 3.2 we prove that the loss of those estimators concentrate around their expected values (risk) when we have a large number of row and column effects.

LEMMA 3.1. *Under Assumptions A1–A2,*

$$\lim_{\substack{r \rightarrow \infty \\ c \rightarrow \infty}} d_{r,c}^2 \cdot \left\{ \sup_{\lambda_A, \lambda_B \geq 0} \mathbb{E}[\text{URE}_{r,c}^Q(0, \lambda_A, \lambda_B) - R_{r,c}^Q(\boldsymbol{\eta}, \hat{\boldsymbol{\eta}}(0, \lambda_A, \lambda_B))]^2 \right\} = 0.$$

LEMMA 3.2. *Under Assumptions A1–A2,*

$$\lim_{\substack{r \rightarrow \infty \\ c \rightarrow \infty}} d_{r,c}^2 \cdot \left\{ \sup_{\lambda_A, \lambda_B \geq 0} \mathbb{E}[L_{r,c}^Q(\boldsymbol{\eta}, \hat{\boldsymbol{\eta}}(0, \lambda_A, \lambda_B)) - R_{r,c}^Q(\boldsymbol{\eta}, \hat{\boldsymbol{\eta}}(0, \lambda_A, \lambda_B))]^2 \right\} = 0.$$

If we restrict ourselves to only estimators in $\mathcal{S}[\tau]$ that shrink toward the origin, then Theorem 3.1 follows directly from the above two lemmas. As such, for this subset of estimators, the lemmas prove a stronger version of the theorem with convergence in L_2 norm. The proof is extended to general shrinkage estimators by controlling the L_1 deviation between the true loss and its URE-based approximation through the nontrivial use of the location invariance structure of the problem. The proofs of all of these results are provided in Section A.1 of the Appendix. The results for the weighted loss $L_{r,c}^{\text{wgt}}$ (defined in Section 2) are discussed in Section S.3.1 of the supplement.

We would like to point out the qualitative differences between the results and the proof techniques employed here and those usually used for the one-way Normal mean estimation problem exemplified in Xie, Kou and Brown (2012). For the cell mean estimation problem in two-way layouts with unbalanced designs, there is no indexing on the parametric space under which the “row” effects and the “column” effects can be decoupled. Thus, in unbalanced designs (see Section S.5 of the supplement for discussion on balanced designs) the two-way cell mean estimation problem cannot be reduced to the one-way setting. The approach of Xie, Kou and Brown (2012), which would require showing uniform convergence of the difference between the URE and the loss over the entire set $\mathcal{H} = \{\mu \in [\hat{a}_\tau, \hat{b}_\tau]; \lambda_A, \lambda_B \geq 0\}$ of possible hyperparameter values, that is, showing L_1 convergence of $\sup_{(\mu, \lambda_A, \lambda_B) \in \mathcal{H}} |\text{URE}_{r,c}^Q(\mu, \lambda_A, \lambda_B) - L_{r,c}^Q(\boldsymbol{\eta}, \hat{\boldsymbol{\eta}}(\mu, \lambda_A, \lambda_B))|$ to 0, cannot be trivially adapted to the two-way layout.

Instead, in Theorem 3.1 we show pointwise control on the expected absolute difference between the URE and the loss. The pointwise control was established through computations involving the variance of Gaussian quadratic forms, which greatly helps in tackling the difficulties encountered when passing to the two-way problem, where computations involving matrices are generally unavoidable. Specifically, we show that the expected absolute difference between the URE and the loss is asymptotically controlled at $o(d_{r,c}^{-1})$ over the essential support of the hyperparameters. Thereafter, for establishing oracle optimality of the hyperparameters produced by our URE method, we leverage the differentiability of the loss as function of the hyperparameters. This allows us to concentrate on controlling the expected absolute difference uniformly over a discrete subset of \mathcal{H} whose cardinality diverges to infinity as $r, c \rightarrow \infty$ but is of the order of $O(d_{r,c})$ (see Appendix A.2 for details). The oracle optimality results follow by applying the pointwise control of $o(d_{r,c}^{-1})$ on the expected absolute difference between the URE and the loss at every point on this discrete subset which has only $O(d_{r,c})$ hyperparameter values; thus, ensuring at most $o(d_{r,c}^{-1}) \cdot O(d_{r,c}) = o(1)$ cumulative error over the concerned discrete subset. Section A.2 contains the detailed proofs of the oracle optimality results of Theorems 3.2, 3.3 and that of Corollary 3.1.

4. Empirical studies. We carry out numerical experiments to compare the performance of the URE-based estimator to that of different estimators discussed in the previous sections. As the standard technique, we consider the weighted least squares estimator $\hat{\eta}^{LS} = \hat{\mu}^{LS}\mathbf{1} + Z\hat{\theta}^{LS}$, where $(\hat{\mu}^{LS}, \hat{\theta}^{LS})$ is any pair that minimizes

$$(\mathbf{y} - \mu \cdot \mathbf{1} - Z\theta)^\top M^{-1}(\mathbf{y} - \mu \cdot \mathbf{1} - Z\theta).$$

The two-way shrinkage estimators reported are the maximum-likelihood empirical Bayes (EBML) estimator $\hat{\eta}^{ML}$ and the URE-based estimator $\hat{\eta}^{URE}$, as well as versions of these two estimators which shrink toward the origin (i.e., with μ fixed at 0); these are designated in Table 1 as “EBMLE (origin)” and “URE (origin).”

TABLE 1
Estimation errors relative to the Least Squares (LS) estimator. The columns in the table correspond to the six simulation examples described in Section 4

	(a)	(b)	(c)	(d)	(e)	(f)
LS	1.00	1.00	1.00	1.00	1.00	1.00
EBMLE	0.31	1.79	0.48	1.37	0.21	0.96
URE	0.31	0.45	0.19	0.21	0.18	0.58
EBMLE (origin)	0.31	0.69	0.45	1.42	0.58	0.95
URE (origin)	0.31	0.46	0.20	0.53	0.57	0.63
XKB	0.31	0.58	0.28	0.44	0.20	–
Oracle	0.30	0.42	0.16	0.20	0.17	0.56

We also consider an adaptation of the one-way estimator of Xie, Kou and Brown (2012) to the two-way layout, which estimates the scale hyperparameters based on two independent one-way shrinkage problems and shrinks toward a general data-driven location; details for this estimator, which we label “XKB” in Table 1, appear in Section S.5 of the supplement. As a benchmark, we consider the oracle rule $\tilde{\boldsymbol{\eta}}^{\text{OL}} = \hat{\boldsymbol{\eta}}(\tilde{\mu}^{\text{OL}}, \tilde{\lambda}_A^{\text{OL}}, \tilde{\lambda}_B^{\text{OL}})$, where

$$(21) \quad (\tilde{\mu}^{\text{OL}}, \tilde{\lambda}_A^{\text{OL}}, \tilde{\lambda}_B^{\text{OL}}) = \arg \min_{\mu, \lambda_A \geq 0, \lambda_B \geq 0} \|\mathbf{y} - M \Sigma^{-1}(\mathbf{y} - \mu \cdot \mathbf{1}) - \boldsymbol{\eta}\|^2.$$

Simulation Experiments. We report results across 6 simulation scenarios. For each of these, we draw $(\boldsymbol{\alpha}, \boldsymbol{\beta}, M^{-1} = \text{diag}(K_{11}, K_{12}, \dots, K_{rc}))$ jointly from some distribution such that the cell counts K_{ij} are i.i.d. and $(\boldsymbol{\alpha}, \boldsymbol{\beta})$ are drawn from some conditional distribution given the K_{ij} . We then draw $y_{ij} \sim N(\mu + \alpha_i + \beta_j, \sigma^2 K_{ij}^{-1})$ independently, fixing $\mu = 0$ throughout and setting σ^2 to some (known) constant value. This process is repeated for $N = 100$ times for each pair (r, c) in a range of values, and the average squared loss over the N rounds is computed for each of the estimators mentioned above. With 100 repetitions, the standard error of the average loss for each estimator is at least one order-of-magnitude smaller than the estimated differences between the risks; hence, the differences can be safely considered significant. The URE estimate is computed using the implementation described in Section A.3 of the Appendix, and the oracle “estimate” is computed employing a similar technique. The EBMLE estimate is computed using the R package `lme4` [Bates (2010)].

Table 1 shows the mean square error (MSE) of different estimators as a fraction of the estimated risk of the Least Squares (LS) estimator. We have equal number of row and column levels for all experiments except for scenario (c). Figure 1 displays the MSE of the URE, EBMLE, LS, XKB and the Oracle loss (OL) estimators across the six experiments as the number of levels in the design varies. The figure shows how the estimation errors of the different estimators compare with the minimum achievable (oracle) error rates as the number of levels in the design increases. An additional simulation, examining the performance of the estimators under model misspecification, is included in the supplement.

The general pattern reflected in the subplots shows an initial sharp decline with a gradual flattening-out of the error rates as the number of levels exceeds 100, suggesting a setting within the asymptotic regime. In all the examples, the performance of the URE-based method is close to that of the oracle when the number of levels is large; when the number of levels is bigger than 60, there is no other estimator which is much better at any instance than the URE. On the contrary, in all examples except scenario (a) the EBMLE performs quite poorly, and is outperformed even by the “one-way” XKB estimator. In cases with dependency be-

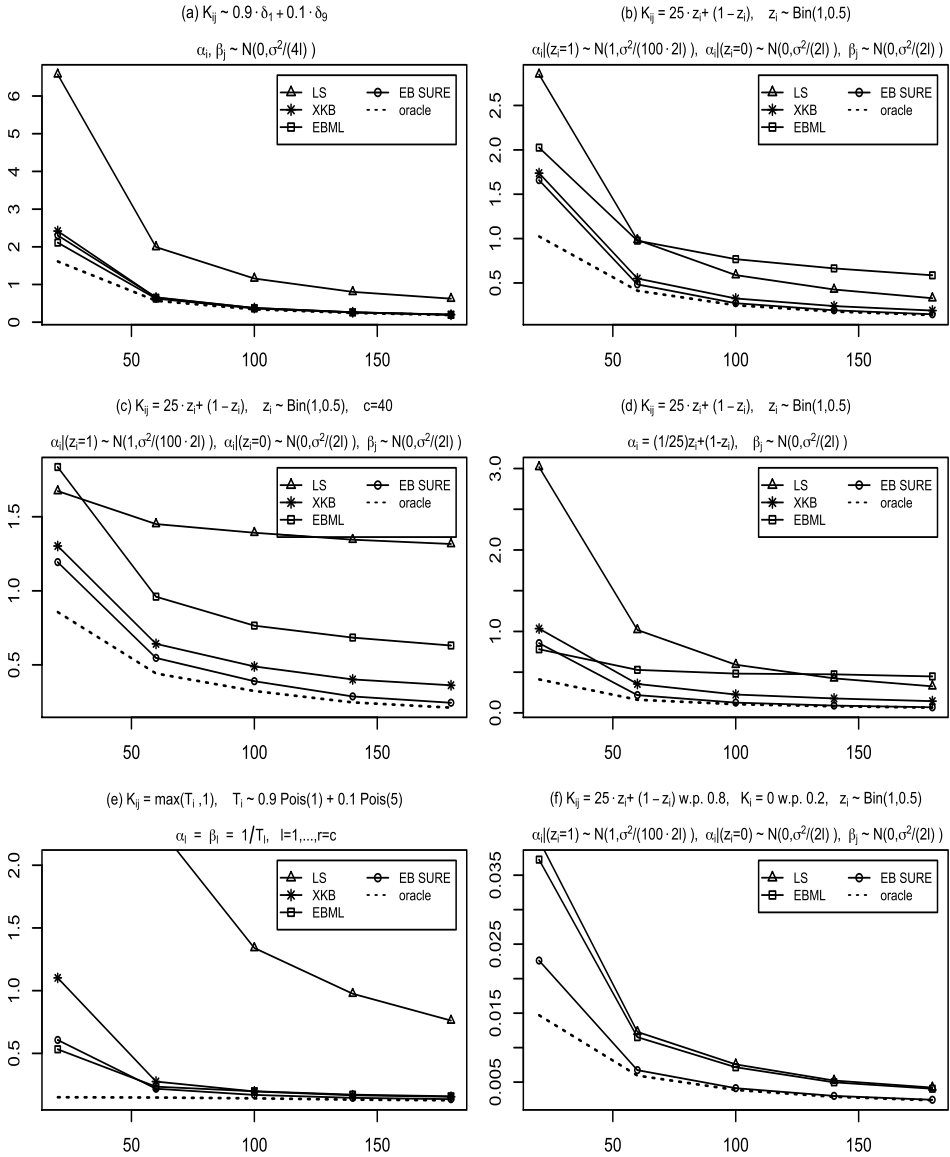


FIG. 1. Risk of the various estimators in the six simulation scenarios of Table 1. The ordinate shows the risk of the estimators while we vary L along the abscissa. In all experiments $r = c = L$, except (d) where $L = r$ and c was fixed at 40. The sampling variance is $\sigma^2 = 10$ in all experiments except (f), where $\sigma^2 = 1$.

tween the effects and the cell counts, even the LS estimator can be preferable to the EBMLE [experiments (b) and (d)]. A detailed description of each of the simulation examples is provided in the supplement to this article.

TABLE 2
Estimated fixed effect (μ) and components λ_{county} and λ_{hour} , which determine shrinkage

	μ	County	Hour
EBMLE	1.10	0.57	0.05
URE	0.78	0.07	0.80

Real Data Analysis. In addition to the simulation study, we analyzed real data collected on Nitrate levels measured in water sources across the US. The data was obtained from the Water Quality Portal cooperative (<http://waterqualitydata.us/>). We estimated the (log transformed) average Nitrate level based on the location of the water resource and time when the measurement was taken. Our data set includes a total of 858 observations categorized into 108 counties and 9 different hour-slots (8–16). The data is highly unbalanced: 57% of the cells are empty, and the cell counts among the nonempty cells vary between 1 to 12. Both the EBMLE and the URE estimators shrink the LS estimates in this example, but the shrinkage factors are quite different. There is a marked difference between the methods in the estimates of the two variance components, as Table 2 shows. Note specifically that EBMLE applies much more shrinkage to “hour”; this is in accord with the fact that the p-value for “hour” in a two-way ANOVA (using Type-II sums-of-squares) is not significant (p-value = 0.25). In terms of the *fitted* values, EBMLE applies more shrinkage, although the differences are not very big. To compare the performance of the different estimators, we carried out two separate analyses. In the first one, we split the data evenly and used the first portion for estimation and the second portion for validation. The second analysis is a data-informed simulation intended to compare performance of the estimators when the two-way additive model (1) is correctly specified. In the analysis with validation, EBMLE had a smaller estimated total squared error (0.42 vs. 0.5, presented as fraction of TSE of LS), but in the data-informed simulation, the URE achieved smaller squared error (0.72 vs. 0.81). As both EBMLE and the URE estimators are designed for the additive model, the results from the latter analysis might be considered a better basis for comparison between the methods. The two experiments are described in details in the supplement.

5. Discussion. We considered estimation under sum-of-squares loss of the cell means in a two-way linear model with additive fixed effects, where the focus was on the unbalanced case. Minimax shrinkage estimators exist which differ from, and hence dominate, the Least Squares estimator for the more general linear regression setup [Rolph (1976)]. However, such estimators do not exploit the special structure of the two-factor additive model, and might lead to undesirable shrinkage patterns which are difficult to interpret. Instead, we considered a

parametric class of Bayes estimators corresponding to a prior motivated from exchangeability considerations. The resulting estimates exhibit meaningful shrinkage patterns and, when appropriately calibrated, achieve significant risk reduction as compared to the least squares estimator in practical situations.

To calibrate the Bayes estimator, we considered substituting the hyperparameters governing the prior with data-dependent values, and proposed a method, which chooses these values in an asymptotically optimal way. We contrasted the proposed estimator with the traditional likelihood-based empirical BLUP estimator, which was shown to generally produce asymptotically suboptimal estimates of the cell means. Since it relies on the postulated two-level model, the likelihood-based empirical BLUP estimator might be led astray when there is dependency between the cell counts and the true cell means; this was clearly shown in our simulation examples.

The theory developed here employs proof techniques that differ in fundamental aspects from those commonly used to prove asymptotic optimality in the one-way Normal mean estimation problem. We offered a flexible approach for proving asymptotic optimality via efficient pointwise risk estimation which reduces to uniformly controlling the variance of the associated quadratic forms. Our proof techniques can be extended to k -way additive models, although computational difficulty of our proposed method might become a problem when k is even moderately large. It would be interesting to investigate whether our proof techniques can be extended to tackle URE-based shrinkage estimation in more complex models (non-Gaussian or misspecified) by using the general quadratic form concentration results established in [Dicker and Erdogdu \(2017\)](#).

APPENDIX

A.1. Pointwise loss estimation by URE: Proof details. Here, we present the proofs of [Theorem 3.1](#), [Lemma 3.1](#) and [Lemma 3.2](#). Henceforth, we assume $\sigma = 1$. It is done mainly for the ease of presentation, and the proofs can easily be modified for any known value of σ . We fix the following notation. Denote by $\sigma_k(A)$ the k th largest singular value of a matrix A . Denote by $\lambda_k(B)$ the k th largest eigenvalue of a symmetric matrix B . Also, $G \doteq M\Sigma^{-1}$ and $H \doteq G^T QG = \Sigma^{-1} M Q M \Sigma^{-1}$ where M , Σ^{-1} and Q are defined in [Section 2.3](#). Let $W = M^{1/2} \Sigma^{-1} M^{1/2}$. As $0 < M \preceq \Sigma$ we have $W \preceq I$, and also $W^2 \preceq I$. We will use the following result of [Searle, Casella and McCulloch \(1992\)](#) ([Theorem S4](#), page 467).

LEMMA A.1 (Central moments of Gaussian Quadratic Forms).

If $\mathbf{y} \sim N(\boldsymbol{\eta}, V)$,

then $\mathbb{E}(\mathbf{y}^T A \mathbf{y}) = 2\text{tr}[AV] + \boldsymbol{\eta}^T A \boldsymbol{\eta}$, and

$\text{Var}(\mathbf{y}^T A \mathbf{y}) = 2\text{tr}[(AV)^2] + 4\boldsymbol{\eta}^T A V A \boldsymbol{\eta}$.

The proof of Theorem 3.1 for the case when the general effect $\mu = 0$ follows directly from the results of Lemma 3.1 and Lemma 3.2 as $\mathbb{E}\{\text{URE}^Q_{r,c}(0, \lambda_A, \lambda_B) - L^Q_{r,c}(\boldsymbol{\eta}, \hat{\boldsymbol{\eta}}(0, \lambda_A, \lambda_B))\}^2$ is bounded above by

$$2\mathbb{E}\{\text{URE}^Q_{r,c}(0, \lambda_A, \lambda_B) - R^Q_{r,c}(\boldsymbol{\eta}, \hat{\boldsymbol{\eta}}(0, \lambda_A, \lambda_B))\}^2 + 2\mathbb{E}\{L^Q_{r,c}(\boldsymbol{\eta}, \hat{\boldsymbol{\eta}}(0, \lambda_A, \lambda_B)) - R^Q_{r,c}(\boldsymbol{\eta}, \hat{\boldsymbol{\eta}}(0, \lambda_A, \lambda_B))\}^2.$$

In fact, this proves Theorem 3.1 in the stronger L_2 norm (with its correspondingly adjusted rate). We now concentrate on proving the lemmas and will thereafter prove the theorem for the general ($\mu \neq 0$) case.

PROOF OF LEMMA 3.1. As the URE is an unbiased estimator of the risk of an estimator in \mathcal{S} , for any fixed $\lambda_A, \lambda_B \geq 0$, we have

$$(22) \quad \mathbb{E}[\text{URE}^Q(0, \lambda_A, \lambda_B) - R^Q_{r,c}(\boldsymbol{\eta}; \hat{\boldsymbol{\eta}}(0, \lambda_A, \lambda_B))]^2 = \text{Var}[\text{URE}^Q(0, \lambda_A, \lambda_B)].$$

Based on the expression of the URE estimator in (17) we know that

$$\text{URE}^Q(0, \lambda_A, \lambda_B) = (rc)^{-2}\{\sigma^2\text{tr}(QM) - 2\sigma^2\text{tr}(\Sigma^{-1}MQM) + \mathbf{y}^\top H\mathbf{y}\},$$

and so the RHS of (22) reduces to $(rc)^{-2} \text{Var}(\mathbf{y}^\top H\mathbf{y})$ which, being the variance of a quadratic form of the Gaussian random vector \mathbf{y} , can in turn be evaluated by using Lemma A.1 to give

$$(23) \quad \text{Var}[\text{URE}^Q(0, \lambda_A, \lambda_B)] = (rc)^{-2}\{2\text{tr}(HMHM) + 4\boldsymbol{\eta}^\top HMMH\boldsymbol{\eta}\}.$$

Our goal now is to show that each of the terms on the RHS, after being multiplied by $d^2_{r,c}$, uniformly converges to 0 for all choices of λ_A and λ_B . For this purpose, we concentrate on the second term of the RHS first. As H is p.s.d. by R2 (see Section S.6 of the supplement), HMH is also p.s.d. Thus, $\boldsymbol{\eta}^\top HMMH\boldsymbol{\eta} \leq \lambda_1(HMH)\|\boldsymbol{\eta}\|^2$. Now, using the bound $\lambda_1(HMH) \leq \nu_{r,c}\lambda^2_1(Q)$ on the largest eigenvalue of HMH (for detailed derivation, see Lemma S.3.2 in the supplements) we arrive at the following upper bound:

$$(rc)^{-2}d^2_{r,c} \sup_{\lambda_A, \lambda_B \geq 0} \boldsymbol{\eta}^\top HMMH\boldsymbol{\eta} \leq (rc)^{-2}d^2_{r,c}\nu_{r,c}\lambda^2_1(Q)\|\boldsymbol{\eta}\|^2$$

which, under Assumptions A1 and A2, converges to 0 as $r, c \rightarrow \infty$. From Lemma S.3.2, we have $\text{tr}(HMHM) \leq rc \cdot \lambda^2_1(Q)$. Hence, as $r, c \rightarrow \infty$, by Assumption A2 the first term in (23) scaled by $d^2_{r,c}$ also converges to 0 uniformly over the ranges of λ_A and λ_B . This completes the proof of the lemma. \square

PROOF OF LEMMA 3.2. As the risk is the expectation of the loss, to prove the lemma we need to show

$$d^2_{r,c} \sup_{\lambda_A, \lambda_B \geq 0} \text{Var}[L^Q(\boldsymbol{\eta}, \hat{\boldsymbol{\eta}}(0, \lambda_A, \lambda_B))] \rightarrow 0 \quad \text{as } r, c \rightarrow \infty.$$

Again, the loss of the estimator $\hat{\boldsymbol{\eta}}_0 = \hat{\boldsymbol{\eta}}(0, \lambda_A, \lambda_B)$ can be decomposed as

$$\begin{aligned} L^Q(\boldsymbol{\eta}, \hat{\boldsymbol{\eta}}_0) &= (rc)^{-1}(\hat{\boldsymbol{\eta}}_0 - \boldsymbol{\eta})^\top Q(\hat{\boldsymbol{\eta}}_0 - \boldsymbol{\eta}) = (rc)^{-1}(\mathbf{y} - \boldsymbol{\eta} - G\mathbf{y})^\top Q(\mathbf{y} - \boldsymbol{\eta} - G\mathbf{y}) \\ &= (rc)^{-1}\{(\mathbf{y} - \boldsymbol{\eta})^\top Q(\mathbf{y} - \boldsymbol{\eta}) + \mathbf{y}^\top H\mathbf{y} - 2(\mathbf{y} - \boldsymbol{\eta})^\top QG\mathbf{y}\} \\ &= (rc)^{-1}\{L_1 + L_2 - L_3 + L_4\}, \end{aligned}$$

where $L_1 = (\mathbf{y} - \boldsymbol{\eta})^\top Q(\mathbf{y} - \boldsymbol{\eta})$, $L_2 = \mathbf{y}^\top H\mathbf{y}$, $L_3 = 2\mathbf{y}^\top QG\mathbf{y}$, $L_4 = 2\boldsymbol{\eta}^\top QG\mathbf{y}$.

Hence, it suffices to show that $d_{r,c}^2 \sup_{\lambda_A, \lambda_B} \text{Var}((rc)^{-1}L_i) \rightarrow 0$ as $r, c \rightarrow \infty$ for all $i = 1, \dots, 4$. Uniform convergence of the desired scaled variance of L_2 was already shown in the proof of Lemma 3.1.

For the first term L_1 , we have

$$\begin{aligned} \text{Var}[(rc)^{-1}L_1] &= (rc)^{-2} \text{Var}[(\mathbf{y} - \boldsymbol{\eta})^\top Q(\mathbf{y} - \boldsymbol{\eta})] = 2(rc)^{-2} \text{tr}(QMQM) \\ &= 2(rc)^{-2} \text{tr}\{(M^{\frac{1}{2}}QM^{\frac{1}{2}})^2\} \leq 2(rc)^{-1} \lambda_1^2(M^{\frac{1}{2}}QM^{\frac{1}{2}}) \end{aligned}$$

which by Assumption A2 is $o(d_{r,c}^{-2})$ as $r, c \rightarrow \infty$ for any value of the hyperparameter. As \mathbf{y} is normally distributed, the fourth term can be explicitly evaluated as

$$\begin{aligned} 4^{-1} \text{Var}(L_4) &= \text{Var}(\boldsymbol{\eta}^\top QG\mathbf{y}) = \boldsymbol{\eta}^\top QGMG^\top Q\boldsymbol{\eta} \leq \lambda_1(QGMG^\top Q)\|\boldsymbol{\eta}\|^2 \\ &\leq \nu_{r,c} \lambda_1^2(Q)\|\boldsymbol{\eta}\|^2, \end{aligned}$$

where the last inequality follows from Lemma S.3.2. By Assumptions A1–A2, the above is $o(r^2c^2d_{r,c}^{-2})$.

The third term requires detailed analysis. It breaks into two components:

$$\begin{aligned} \text{Var}[(rc)^{-1}L_3] &= 4(rc)^{-2} \text{Var}(\mathbf{y}^\top QG\mathbf{y}) \\ (24) \quad &= 8(rc)^{-2} \text{tr}(\tilde{G}M\tilde{G}M) + 16(rc)^{-2} \boldsymbol{\eta}^\top \tilde{G}M\tilde{G}\boldsymbol{\eta}, \end{aligned}$$

where $\tilde{G} = QG + G^\top Q$ is a symmetric matrix. By Lemma S.3.2, the first term on the RHS of (24) is bounded above by $32(rc)^{-1} \lambda_1^2(Q)$ and hence, by Assumption A2 is uniformly bounded by $o(d_{r,c}^{-2})$ as $r, c \rightarrow \infty$. We now concentrate on the second term of the RHS of (24). Note that $(rc)^{-2} \boldsymbol{\eta}^\top \tilde{G}M\tilde{G}\boldsymbol{\eta} \leq (rc)^{-2} \boldsymbol{\eta}^\top \boldsymbol{\eta} \sigma_1(\tilde{G}M\tilde{G})$. If we can uniformly bound the largest eigenvalue of $\tilde{G}M\tilde{G}$ as $o(rcd_{r,c}^{-2})$ then by Assumption A1 the second term of the RHS of (24) is $o(d_{r,c}^{-2})$ as $r, c \rightarrow \infty$. As $\tilde{G} = QG + G^\top Q = QM\Sigma^{-1} + \Sigma^{-1}MQ$, we have

$$\tilde{G}M\tilde{G} = H_1 + H_1^\top + H_2 + H_3,$$

where

$$\begin{aligned} H_1 &= QM\Sigma^{-1}MQM\Sigma^{-1}, \\ H_2 &= QM\Sigma^{-1}M\Sigma^{-1}MQ, \quad H_3 = \Sigma^{-1}MQMQM\Sigma^{-1}. \end{aligned}$$

To uniformly bound the eigenvalues of $\tilde{G}M\tilde{G}$ as desired before we just show that for each of $i = 1, \dots, 3$, $(rc)^{-1}d_{r,c}^{-2}\sigma_1(H_i) \rightarrow 0$ as $r, c \rightarrow \infty$.

Note that, H_2, H_3 is symmetric but H_1 is not. By Lemma S.3.2, we have $\max\{\sigma_1(H_1), \lambda_1(H_3)\} \leq \nu_{r,c}\lambda_1^2(Q)$, from which the desired asymptotic control on the variance follows. For H_2 using $W^2 \preceq I$, we have

$$\lambda_1(H_2) = \lambda_1(QM^{\frac{1}{2}}W^2M^{\frac{1}{2}}Q) \leq \lambda_1(QM^{\frac{1}{2}}M^{\frac{1}{2}}Q) \leq \lambda_1(M^{-1})\lambda_1^2(M^{\frac{1}{2}}QM^{\frac{1}{2}})$$

which is uniformly controlled at $o(rcd_{r,c}^{-2})$ by Assumption A2. This completes the proof of the lemma. \square

PROOF OF THEOREM 3.1. As shown in the beginning of this section, by the above two lemmas the theorem easily follows for the case when $\mu = 0$. We now prove the theorem for the general case. First, note that for arbitrary fixed $\mu \in \mathbb{R}$, the loss $(\hat{\eta}(\mu, \lambda_A, \lambda_B) - \eta)^\top Q(\hat{\eta}(\mu, \lambda_A, \lambda_B) - \eta)$ decomposes into the following components:

$$\begin{aligned} & (\hat{\eta}(0, \lambda_A, \lambda_B) - \eta)^\top Q(\hat{\eta}(0, \lambda_A, \lambda_B) - \eta) + \mu^2 \mathbf{1}^\top H \mathbf{1} \\ & - 2\mu \mathbf{1}^\top H \mathbf{y} + 2\mu \mathbf{1}^\top G^\top Q(\mathbf{y} - \eta). \end{aligned}$$

Comparing it with the definition of URE, we have

$$\begin{aligned} \text{URE}_{r,c}^Q(\mu, \lambda_A, \lambda_B) - L_{r,c}^Q(\eta, \hat{\eta}(0, \lambda_A, \lambda_B)) \\ = \text{URE}_{r,c}^Q(0, \lambda_A, \lambda_B) - L_{r,c}^Q(\eta, \hat{\eta}(0, \lambda_A, \lambda_B)) + 2(rc)^{-1}\mu \mathbf{1}^\top G^\top Q(\mathbf{y} - \eta). \end{aligned}$$

We have already proved the theorem for the case of $\mu = 0$; hence, in light of the above identity, the proof of the general case will follow if we show:

$$(25) \quad \lim_{\substack{r \rightarrow \infty \\ c \rightarrow \infty}} \sup_{\substack{|\mu| \leq m_{r,c} \\ \lambda_A, \lambda_B \geq 0}} d_{r,c} \cdot (rc)^{-1} \cdot \mathbb{E}|\mu \mathbf{1}^\top G^\top Q(\mathbf{y} - \eta)| = 0.$$

Noting that for any fixed η the random variable $F = \mathbf{1}^\top G^\top Q(\mathbf{y} - \eta)$ follows a univariate normal distribution with mean 0 and variance $\mathbf{1}^\top G^\top Q M Q G \mathbf{1}$, the above holds if we show

$$(26) \quad \lim_{r \rightarrow \infty, c \rightarrow \infty} m_{r,c} \cdot d_{r,c} \cdot (rc)^{-1} \cdot \sup_{\lambda_A, \lambda_B \geq 0} \{\text{Var}(|\mathbf{1}^\top G^\top Q(\mathbf{y} - \eta)|)\}^{1/2} = 0.$$

The above is true based on our assumptions as we can upper bound the variance of $|F|$ by

$$\text{Var}(F) \leq \mathbf{1}^\top G^\top Q M Q G \mathbf{1} \leq rc \lambda_1(G^\top Q M Q G) \leq rc \nu_{r,c} \lambda_1^2(Q),$$

where the last inequality follows from Lemma S.3.2. This based on Assumption A2 establishes (26) and completes the proof of the theorem. \square

A.2. Proof of the oracle optimality results. The proofs of Theorems 3.2, 3.3 and Corollary 3.1 are presented here. For that purpose, we first construct a discrete subset of the set of hyperparameters and define analogous versions of the URE and oracle estimators over it.

Discretization. In (18) and (20), the URE and oracle estimators involve minimizing the hyperparameters $(\mu, \lambda_A, \lambda_B)$ simultaneously over $\hat{T}_{r,c} = [\hat{a}_\tau, \hat{b}_\tau] \times [0, \infty] \times [0, \infty]$ where the range of the location hyperparameter μ depends on the data. We define a discrete product grid $\Theta_{r,c} = \Theta_{r,c}^{[1]} \times \Theta_{r,c}^{[2]} \times \Theta_{r,c}^{[3]}$ which only depends on r, c and not on the data. Details for the construction of $\Theta_{r,c}$ is provided afterward. It contains countably infinite grid points as $r, c \rightarrow \infty$. We define the discretized version of the oracle estimator where the minimization is conducted over all the points in the discrete grid $\Theta_{r,c}$ that are contained in $\hat{T}_{r,c}$. We define the discretized oracle loss hyper-parameters as

$$(\tilde{\mu}^{\text{OD}}, \tilde{\lambda}_A^{\text{OD}}, \tilde{\lambda}_B^{\text{OD}}) = \arg \min_{(\mu, \lambda_A, \lambda_B) \in \Theta_{r,c} \cap \hat{T}_{r,c}} L_{r,c}^Q(\eta, \hat{\eta}(\mu, \lambda_A, \lambda_B)),$$

and the corresponding oracle rule by $\tilde{\eta}_c^{\text{OD}} = Z_c Z^\dagger \hat{\eta}(\tilde{\mu}^{\text{OD}}, \tilde{\lambda}_A^{\text{OD}}, \tilde{\lambda}_B^{\text{OD}})$. We define the URE estimators over the discrete grid by projecting the URE estimates of equation (18) in $\Theta_{r,c} \cap \hat{T}_{r,c}$: if the URE hyperparameters given by equation (18) are such that:

$$\mu_1 \leq \hat{\mu}^{\text{UQ}} \leq \mu_2, \quad \lambda_1 \leq \hat{\lambda}_A^{\text{UQ}} \leq \lambda_2, \quad \text{and} \quad \lambda_3 \leq \hat{\lambda}_B^{\text{UQ}} \leq \lambda_4,$$

where μ_1, μ_2 are neighboring points in $\Theta_{r,c}^{[1]} \cap [\hat{a}_\tau, \hat{b}_\tau]$, λ_1, λ_2 are neighboring points in $\Theta_{r,c}^{[2]}$ and λ_3, λ_4 are neighboring points in $\Theta_{r,c}^{[3]}$, then the URE estimates of the tuning parameters over the discrete grid is defined as the minima over the nearest 8-point subset of the grid:

$$(27) \quad (\hat{\mu}^{\text{UD}}, \hat{\lambda}_A^{\text{UD}}, \hat{\lambda}_B^{\text{UD}}) = \arg \min_{(\mu, \lambda_A, \lambda_B) \in \{\mu_1, \mu_2\} \times \{\lambda_1, \lambda_2\} \times \{\lambda_3, \lambda_4\}} \text{URE}^Q(\mu, \lambda_A, \lambda_B).$$

The corresponding discretized EB estimate is $\hat{\eta}^{\text{UD}} = \hat{\eta}(\hat{\mu}^{\text{UD}}, \hat{\lambda}_A^{\text{UD}}, \hat{\lambda}_B^{\text{UD}})$. The corresponding estimate for η_c is $\hat{\eta}_c^{\text{UD}} = Z_c Z^\dagger \hat{\eta}(\hat{\mu}^{\text{UD}}, \hat{\lambda}_A^{\text{UD}}, \hat{\lambda}_B^{\text{UD}})$. If the URE estimators for any of the three hyperparameters are outside the grid, then the nearest boundary of the grid is taken as the UD estimate for that hyper-paramter. We will show afterward that the probability of such events is negligible. Also, by construction, $L(\eta_c, \hat{\eta}_c^{\text{UD}}) \geq L(\eta_c, \tilde{\eta}_c^{\text{OD}}) \geq L(\eta_c, \tilde{\eta}_c^{\text{OL}})$.

Construction of the grid $\Theta_{r,c}$. The grid $\Theta_{r,c}$ is a product grid. The grid $\Theta_{r,c}^{[1]}$ on the location hyperparameter μ is an equispaced discrete set $\{-m_{r,c} = \mu[1] < \mu[2] < \dots < \mu[n_1] \leq m_{r,c}\}$, which covers $[-m_{r,c}, m_{r,c}]$ at a spacing of $\delta_{r,c}^{[1]}$. Thus, the cardinality of $\Theta_{r,c}^{[1]}$, $n_1 = \lceil 2m_{r,c} \{\delta_{r,c}^{[1]}\}^{-1} \rceil$. We choose the spacing as

$$(28) \quad \delta_{r,c}^{[1]} = \{m_{r,c}^{4/3} \cdot \nu_{r,c} \cdot \lambda_1(Q)\}^{-1}.$$

For constructing the grid $\Theta_{r,c}^{[2]}$ on the scale hyperparameter, we consider the following transformation $\tilde{\lambda}_A = (1 + \lambda_A)^{-1/2}$. Note that $\tilde{\lambda}_A \in [0, 1]$ as λ_A varies over $[0, \infty]$. Let $\tilde{\lambda}_A[k] = (k - 1)\delta_{r,c}^{[2]}, n_2 = \lceil \{\delta_{r,c}^{[2]}\}^{-1} \rceil$. Consider the equispaced grid on $\tilde{\lambda}_A$ between 0 and 1 at a spacing of $\delta_{r,c}^{[2]}$: $\{0 = \tilde{\lambda}_A[1] < \tilde{\lambda}_A[2] < \dots < \tilde{\lambda}_A[n_2] \leq 1\}$. The grid on $\tilde{\lambda}_A$ is then retransformed to produce the grid $\Theta_{r,c}^{[2]}$ on the scale hyperparameter λ_A in the domain $[0, \infty]$. The grid $\Theta_{r,c}^{[3]}$ on λ_B is similarly constructed with $\delta_{r,c}^{[3]}$ distances between two corresponding grid points in $\tilde{\lambda}_B$ scale. The spaces were chosen as

$$(29) \quad \delta_{r,c}^{[2]} = \delta_{r,c}^{[3]} = \{m_{r,c}^{7/3} \cdot v_{r,c} \cdot \lambda_1(Q)\}^{-1}.$$

Now, as $r, c \rightarrow \infty, n_1 = O(m_{r,c}^{7/3} \cdot v_{r,c} \cdot \lambda_1(Q)), n_2 = O(m_{r,c}^{7/3} \cdot v_{r,c} \cdot \lambda_1(Q))$, and thus the cardinality of $\Theta_{r,c}$ is $|\Theta_{r,c}| = O(m_{r,c}^7 v_{r,c}^3 \lambda_1^3(Q)) = O(d_{r,c})$.

The following two lemmas enable us to work with the more tractable, discretized versions of the URE and oracle estimators when proving our decision theoretic results.

LEMMA A.2. *Under Assumptions A1–A2 for any fixed $\varepsilon > 0$ as $r, c \rightarrow \infty$,*

- A. $P\{L_{r,c}(\eta_c, \tilde{\eta}_c^{OD}) - L_{r,c}(\eta_c, \tilde{\eta}_c^{OL}) > \varepsilon\} \rightarrow 0,$
- B. $\mathbb{E}|L_{r,c}(\eta_c, \tilde{\eta}_c^{OD}) - L_{r,c}(\eta_c, \tilde{\eta}_c^{OL})| \rightarrow 0,$
- C. $P\{|L_{r,c}(\eta_c, \hat{\eta}_c^{UD}) - L_{r,c}(\eta_c, \hat{\eta}_c^{URE})| > \varepsilon\} \rightarrow 0,$
- D. $\mathbb{E}|L_{r,c}(\eta_c, \hat{\eta}_c^{UD}) - L_{r,c}(\eta_c, \hat{\eta}_c^{URE})| \rightarrow 0.$

LEMMA A.3. *Under Assumptions A1–A2 for any fixed $\varepsilon > 0$ as $r, c \rightarrow \infty$,*

- A. $P\{\text{URE}^Q(\hat{\mu}^{UD}, \hat{\lambda}_A^{UD}, \hat{\lambda}_B^{UD}) - \text{URE}^Q(\hat{\mu}^{UQ}, \hat{\lambda}_A^{UQ}, \hat{\lambda}_B^{UQ}) > \varepsilon\} \rightarrow 0,$
- B. $\mathbb{E}[\text{URE}^Q(\hat{\mu}^{UD}, \hat{\lambda}_A^{UD}, \hat{\lambda}_B^{UD}) - \text{URE}^Q(\hat{\mu}^{UQ}, \hat{\lambda}_A^{UQ}, \hat{\lambda}_B^{UQ})] \rightarrow 0.$

Lemma A.2 above shows that the difference in the loss between the true estimators and their discretized versions is asymptotically controlled at any prefixed level. It uses the following two lemmas. Lemma A.3 above shows that the URE values for the estimator is also asymptotically close for the discretized version. The proofs of all these lemmas (A.2, A.3, A.4 and A.5) are provided in the supplementary materials.

LEMMA A.4. *Under Assumption A1 in model (14), for any fixed $\tau \in (0, 1]$ and $m_{r,c}$ of (19), the event $A_{r,c}(Y) = \{[\hat{a}_\tau(Y), \hat{b}_\tau(Y)] \subseteq [-m_{r,c}, m_{r,c}]\}$ satisfies $P\{A_{r,c}\} \rightarrow 1$ as $r, c \rightarrow \infty$.*

LEMMA A.5. Under Assumptions A1–A2, for any fixed $\tau \in (0, 1]$ and $m_{r,c}$ of (19), the event $A_{r,c}(\mathbf{Y}) = \{[\hat{a}_\tau(\mathbf{Y}), \hat{b}_\tau(\mathbf{Y})] \subseteq [-m_{r,c}, m_{r,c}]\}$ satisfies:

- A. $\mathbb{E}\{|L_{r,c}(\boldsymbol{\eta}_c, \hat{\boldsymbol{\eta}}_c^{\text{OD}}) - L_{r,c}(\boldsymbol{\eta}_c, \hat{\boldsymbol{\eta}}_c^{\text{OL}})| \cdot I\{A_{r,c}(\mathbf{Y})\}\} \rightarrow 0$ as $r, c \rightarrow \infty$.
- B. $\mathbb{E}\{|L_{r,c}(\boldsymbol{\eta}_c, \hat{\boldsymbol{\eta}}_c^{\text{UD}}) - L_{r,c}(\boldsymbol{\eta}_c, \hat{\boldsymbol{\eta}}_c^{\text{URE}})| \cdot I\{A_{r,c}(\mathbf{Y})\}\} \rightarrow 0$ as $r, c \rightarrow \infty$.

We next present the proof of the decision theoretic properties where Lemmas A.2, A.3 will be repeatedly used.

PROOF OF THEOREM 3.2. We know that

$$P\{L(\boldsymbol{\eta}_c, \hat{\boldsymbol{\eta}}_c^{\text{URE}}) \geq L(\boldsymbol{\eta}_c, \hat{\boldsymbol{\eta}}_c^{\text{OL}}) + \varepsilon\} \leq P\{L(\boldsymbol{\eta}_c, \hat{\boldsymbol{\eta}}_c^{\text{URE}}) \geq L(\boldsymbol{\eta}_c, \hat{\boldsymbol{\eta}}_c^{\text{OD}}) + \varepsilon/2\} + P\{L(\boldsymbol{\eta}_c, \hat{\boldsymbol{\eta}}_c^{\text{OD}}) \geq L(\boldsymbol{\eta}_c, \hat{\boldsymbol{\eta}}_c^{\text{OL}}) + \varepsilon/2\}.$$

By Lemma A.2, the second term converges to 0. The first term is less than

$$P\{L(\boldsymbol{\eta}_c, \hat{\boldsymbol{\eta}}_c^{\text{UD}}) \geq L(\boldsymbol{\eta}, \hat{\boldsymbol{\eta}}_c^{\text{OD}}) + \varepsilon/4\} + P\{|L(\boldsymbol{\eta}_c, \hat{\boldsymbol{\eta}}_c^{\text{URE}}) - L(\boldsymbol{\eta}, \hat{\boldsymbol{\eta}}_c^{\text{UD}})| \leq \varepsilon/4\}.$$

As $r, c \rightarrow \infty$ the second term in the RHS above converges to 0 by Lemma A.2. For the first term note that, by definition, $\text{URE}^{\text{Q}}(\hat{\mu}^{\text{UQ}}, \hat{\lambda}_A^{\text{UQ}}, \hat{\lambda}_B^{\text{UQ}}) \leq \text{URE}^{\text{Q}}(\tilde{\mu}^{\text{OD}}, \tilde{\lambda}_A^{\text{OD}}, \tilde{\lambda}_B^{\text{OD}})$ which, combined with Lemma A.3, suggests that

$$(30) \quad P\{\text{URE}^{\text{Q}}(\hat{\mu}^{\text{UD}}, \hat{\lambda}_A^{\text{UD}}, \hat{\lambda}_B^{\text{UD}}) \leq \text{URE}^{\text{Q}}(\tilde{\mu}^{\text{OD}}, \tilde{\lambda}_A^{\text{OD}}, \tilde{\lambda}_B^{\text{OD}}) + \varepsilon/8\} \rightarrow 0$$

as $r, c \rightarrow \infty$.

Thus, showing $P\{L(\boldsymbol{\eta}_c, \hat{\boldsymbol{\eta}}_c^{\text{UD}}) \geq L(\boldsymbol{\eta}, \hat{\boldsymbol{\eta}}_c^{\text{OD}}) + \varepsilon/4\} \rightarrow 0$ as $r, c \rightarrow \infty$ can be reduced to showing $\lim_{r,c \rightarrow \infty} P\{A(\mathbf{y}; \boldsymbol{\eta}_c) \geq B(\mathbf{y}; \boldsymbol{\eta}_c) + \varepsilon/8\} = 0$ where

$$A(\mathbf{y}; \boldsymbol{\eta}_c) = L(\boldsymbol{\eta}_c, \hat{\boldsymbol{\eta}}_c^{\text{UD}}) - \text{URE}^{\text{Q}}(\hat{\mu}^{\text{UD}}, \hat{\lambda}_A^{\text{UD}}, \hat{\lambda}_B^{\text{UD}}),$$

and

$$B(\mathbf{y}; \boldsymbol{\eta}_c) = L(\boldsymbol{\eta}_c, \hat{\boldsymbol{\eta}}_c^{\text{OD}}) - \text{URE}^{\text{Q}}(\tilde{\mu}^{\text{OD}}, \tilde{\lambda}_A^{\text{OD}}, \tilde{\lambda}_B^{\text{OD}}).$$

As $L(\boldsymbol{\eta}_c, \hat{\boldsymbol{\eta}}_c^{\text{UD}}) = L^{\text{Q}}(\boldsymbol{\eta}, \hat{\boldsymbol{\eta}}^{\text{UD}})$ and $L(\boldsymbol{\eta}_c, \hat{\boldsymbol{\eta}}_c^{\text{OD}}) = L^{\text{Q}}(\boldsymbol{\eta}, \hat{\boldsymbol{\eta}}^{\text{OD}})$, by using Markov’s inequality we get

$$P\{A(\mathbf{y}; \boldsymbol{\eta}_c) \geq B(\mathbf{y}; \boldsymbol{\eta}_c) + \varepsilon/8\} \leq 8^{-1} \varepsilon^{-1} \mathbb{E}\{|A(\mathbf{y}; \boldsymbol{\eta}_c) - B(\mathbf{y}; \boldsymbol{\eta}_c)|\}.$$

Now, by triangle inequality the RHS above is upper bounded by

$$\begin{aligned} & 16\varepsilon^{-1} \mathbb{E}\left\{ \sup_{(\mu, \lambda_A, \lambda_B) \in \Theta_{r,c}} |L^{\text{Q}}(\boldsymbol{\eta}, \hat{\boldsymbol{\eta}}(\mu, \lambda_A, \lambda_B)) - \text{URE}^{\text{Q}}(\mu, \lambda_A, \lambda_B)| \right\} \\ & \leq 16\varepsilon^{-1} \mathbb{E}\left\{ \sum_{(\mu, \lambda_A, \lambda_B) \in \Theta_{r,c}} |L^{\text{Q}}(\boldsymbol{\eta}, \hat{\boldsymbol{\eta}}(\mu, \lambda_A, \lambda_B)) - \text{URE}^{\text{Q}}(\mu, \lambda_A, \lambda_B)| \right\} \\ & \leq 16\varepsilon^{-1} |\Theta_{r,c}| \sup_{\substack{|\mu| \leq m_{r,c} \\ \lambda_A, \lambda_B \geq 0}} \mathbb{E}\{|L^{\text{Q}}(\boldsymbol{\eta}, \hat{\boldsymbol{\eta}}(\mu, \lambda_A, \lambda_B)) - \text{URE}^{\text{Q}}(\mu, \lambda_A, \lambda_B)|\}. \end{aligned}$$

As $|\Theta_{r,c}| = O(d_{r,c})$ by Theorem 3.1, the above expression converges to zero when $r, c \rightarrow \infty$. This completes the proof of the theorem. \square

PROOF OF THEOREM 3.3. Decompose the loss as the sum of the following components: $L(\eta_c, \hat{\eta}_c^{\text{URE}}) - L(\eta_c, \hat{\eta}_c^{\text{UD}})$, $L(\eta_c, \tilde{\eta}_c^{\text{OD}}) - L(\eta_c, \tilde{\eta}_c^{\text{OL}})$ and $L(\eta_c, \hat{\eta}_c^{\text{UD}}) - L(\eta_c, \tilde{\eta}_c^{\text{OD}})$. By Lemma A.2, expectations of the absolute values of the first two terms converge to 0 as $r, c \rightarrow \infty$. Write the third term as

$$\begin{aligned} & \{L(\eta_c, \hat{\eta}_c^{\text{UD}}) - \text{URE}^Q(\hat{\mu}^{\text{UD}}, \hat{\lambda}_A^{\text{UD}}, \hat{\lambda}_B^{\text{UD}})\} \\ & - \{L(\eta_c, \tilde{\eta}_c^{\text{OD}}) - \text{URE}^Q(\tilde{\mu}^{\text{OD}}, \tilde{\lambda}_A^{\text{OD}}, \tilde{\lambda}_B^{\text{OD}})\} \\ & + \{\text{URE}^Q(\hat{\mu}^{\text{UD}}, \hat{\lambda}_A^{\text{UD}}, \hat{\lambda}_B^{\text{UD}}) - \text{URE}^Q(\tilde{\mu}^{\text{OD}}, \tilde{\lambda}_A^{\text{OD}}, \tilde{\lambda}_B^{\text{OD}})\}. \end{aligned}$$

Now, by definition $\text{URE}^Q(\hat{\mu}^{\text{UQ}}, \hat{\lambda}_A^{\text{UQ}}, \hat{\lambda}_B^{\text{UQ}}) \leq \text{URE}^Q(\tilde{\mu}^{\text{OD}}, \tilde{\lambda}_A^{\text{OD}}, \tilde{\lambda}_B^{\text{OD}})$ which, combined with Lemma A.3, suggests that the last term above has asymptotically non-positive expectation. Thus, there exist constants $\kappa_{r,c}$ with $\kappa_{r,c} \rightarrow 0$ as $r, c \rightarrow \infty$ such that $\mathbb{E}\{L(\eta_c, \hat{\eta}_c^{\text{UD}}) - L(\eta_c, \tilde{\eta}_c^{\text{OD}})\}$ is

$$\begin{aligned} & \leq 2\mathbb{E}\left\{\sup_{(\mu, \lambda_A, \lambda_B) \in \Theta_{r,c}} |L^Q(\eta, \hat{\eta}(\mu, \lambda_A, \lambda_B)) - \text{URE}^Q(\mu, \lambda_A, \lambda_B)|\right\} + \kappa_{r,c} \\ & \leq 2\mathbb{E}\left\{\sum_{(\mu, \lambda_A, \lambda_B) \in \Theta_{r,c}} |L^Q(\eta, \hat{\eta}(\mu, \lambda_A, \lambda_B)) - \text{URE}^Q(\mu, \lambda_A, \lambda_B)|\right\} + \kappa_{r,c} \\ & \leq 2|\Theta_{r,c}| \sup_{|\mu| \in m_{r,c}; \lambda_A, \lambda_B \geq 0} \mathbb{E}\{|L^Q(\eta, \hat{\eta}(\mu, \lambda_A, \lambda_B)) - \text{URE}^Q(\mu, \lambda_A, \lambda_B)|\} + \kappa_{r,c}. \end{aligned}$$

As $|\Theta_{r,c}| = O(d_{r,c})$, the above expression tends to zero when $r, c \rightarrow \infty$ by Theorem 3.1. This completes the proof of Theorem 3.3. \square

PROOF OF COROLLARY 3.1. (a) and (b) are direct consequences, respectively, of Theorems 3.2 and 3.3, since $L^Q(\eta, \hat{\eta}(\hat{\mu}, \hat{\lambda}_A, \hat{\lambda}_B)) \geq L^Q(\eta, \eta^{\text{OL}})$, and hence, also $\mathbb{E}\{L^Q(\eta, \hat{\eta}(\hat{\mu}, \hat{\lambda}_A, \hat{\lambda}_B))\} \geq \mathbb{E}\{L^Q(\eta, \eta^{\text{OL}})\}$. Unlike in the above two theorems, here we only have optimality over the loss L^Q defined over the observed cells with Q in (16). As explained in Section 2.3, the loss L^Q for the observed cells is the same as the (normalized) sum-of-squares loss over all (observed and missing) rc cell means for an estimator of the form $Z_c Z^\top \hat{\eta}(\hat{\mu}, \hat{\lambda}_A, \hat{\lambda}_B)$, where $\hat{\mu}, \hat{\lambda}_A, \hat{\lambda}_B$ are any estimates of the hyper-parameters. \square

A.3. Computations for implementing the URE method. By definition, $\Sigma = Z\Lambda\Lambda^\top Z^\top + M$. By the matrix inverse identity, we have $\Sigma^{-1} = M^{-1} - M^{-1}Z\Lambda(\Lambda^\top Z^\top M^{-1}Z\Lambda + I_q)^{-1}\Lambda^\top Z^\top M^{-1}$. Hence, we get

$$\begin{aligned} M\Sigma^{-1} &= I_{rc} - Z\Lambda(\Lambda^\top Z^\top M^{-1}Z\Lambda + I_q)^{-1}\Lambda^\top Z^\top M^{-1}, \\ M\Sigma^{-1}M &= M - Z\Lambda(\Lambda^\top Z^\top M^{-1}Z\Lambda + I_q)^{-1}\Lambda^\top Z^\top. \end{aligned}$$

Using the above, we get

$$\text{tr}(\Sigma^{-1}M^2) = \text{tr}(M\Sigma^{-1}M) = \text{tr}(M) - \text{tr}(Z\Lambda(\Lambda^\top Z^\top M^{-1}Z\Lambda + I_q)^{-1}\Lambda^\top Z^\top).$$

Therefore, the URE expression in (10) multiplied by rc can be written as

$$-\sigma^2\text{tr}(M) + 2\sigma^2\text{tr}\{(\Lambda^\top Z^\top M^{-1}Z\Lambda + I_q)^{-1}(\Lambda^\top Z^\top Z\Lambda)\} + \|M\Sigma^{-1}(\mathbf{y} - \mathbf{1}\mu)\|^2.$$

In computing the above expression:

1. The middle term is computed as the sum of the *elementwise* product of $(\Lambda^\top Z^\top M^{-1}Z\Lambda + I_q)^{-1}$ and $\Lambda^\top Z^\top Z\Lambda$, using the property $\text{tr}(A^\top B) = \sum_{i,j} A_{ij}B_{ij}$.

2. $(\Lambda^\top Z^\top M^{-1}Z\Lambda + I_q)^{-1}$ is computed efficiently employing a sparse Cholesky factorization of $\Lambda^\top Z^\top M^{-1}Z\Lambda + I_q$ similar to the implementation in the lme4 package in R.

3. The quantity $\min_{\mu} \|M\Sigma^{-1}(\mathbf{y} - \mathbf{1}\mu)\|^2$ is computed by regressing $M\Sigma^{-1}\mathbf{y}$ on $M\Sigma^{-1}\mathbf{1}_{rc}$ using the `lm` function in R. In doing that, the vector $M\Sigma^{-1}\mathbf{x}$ (for $\mathbf{x} = \mathbf{y}$ and $\mathbf{x} = \mathbf{1}_{rc}$) is computed as

$$(31) \quad M\Sigma^{-1}\mathbf{x} = \mathbf{x} - Z\Lambda(\Lambda^\top Z^\top M^{-1}Z\Lambda + I_q)^{-1}\Lambda^\top Z^\top(M^{-1}\mathbf{x}),$$

where (31) is implemented proceeding “from right to left” to always compute a product of a matrix and a *vector*, instead of two matrices: First, find $M^{-1}\mathbf{x}$, then find $(\Lambda^\top Z^\top)(M^{-1}\mathbf{x})$, and so on.

Acknowledgement. We thank Tony Cai, Samuel Kou and Art Owen for helpful discussions. We would also like to thank the Associate Editor and three referees for constructive suggestions to shorten and improve the paper.

SUPPLEMENTARY MATERIAL

Supplement to “Empirical Bayes estimates for a two-way cross-classified model” (DOI: [10.1214/17-AOS1599SUPP](https://doi.org/10.1214/17-AOS1599SUPP); .pdf). The supplement [Brown, Mukherjee and Weinstein (2018)] contains detailed proofs of the lemmas that were used in the [Appendix](#) for proving the results in Section 3; and derivations and further discussions on the results of Sections 2 and 4.

REFERENCES

- BATES, D. M. (2010). lme4: Mixed-effects modeling with R. <http://lme4.r-forge.r-project.org/book>.
 BROWN, L. D., MUKHERJEE, G. and WEINSTEIN, A. (2018). Supplement to “Empirical Bayes estimates for a two-way cross-classified model.” DOI:[10.1214/17-AOS1599SUPP](https://doi.org/10.1214/17-AOS1599SUPP).
 CANDÈS, E. J., SING-LONG, C. A. and TRZASKO, J. D. (2013). Unbiased risk estimates for singular value thresholding and spectral estimators. *IEEE Trans. Signal Process.* **61** 4643–4657. MR3105401
 DICKER, L. H. (2013). Optimal equivariant prediction for high-dimensional linear models with arbitrary predictor covariance. *Electron. J. Stat.* **7** 1806–1834. MR3084672

- DICKER, L. H. and ERDOGDU, M. A. (2017). Flexible results for quadratic forms with applications to variance components estimation. *Ann. Statist.* **45** 386–414. [MR3611496](#)
- DONOHO, D. L., JOHNSTONE, I. M., KERKYACHARIAN, G. and PICARD, D. (1995). Wavelet shrinkage: Asymptopia? *J. R. Stat. Soc., B* **57** 301–369. [MR1323344](#)
- DRAPER, N. R. and VAN NOSTRAND, R. C. (1979). Ridge regression and James–Stein estimation: Review and comments. *Technometrics* **21** 451–466. [MR0555086](#)
- 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](#)
- FAHRMEIR, L., KNEIB, T., LANG, S. and MARX, B. (2013). *Regression: Models, Methods and Applications*. Springer, Heidelberg. [MR3075546](#)
- GELMAN, A. (2005). Analysis of variance—why it is more important than ever. *Ann. Statist.* **33** 1–53. [MR2157795](#)
- GELMAN, A., CARLIN, J. B., STERN, H. S. and RUBIN, D. B. (2004). *Bayesian Data Analysis*, 2nd ed. Chapman & Hall/CRC, Boca Raton, FL. [MR2027492](#)
- GHOSH, M., NICKERSON, D. M. and SEN, P. K. (1987). Sequential shrinkage estimation. *Ann. Statist.* **15** 817–829. [MR0888442](#)
- GOLDSTEIN, H., BROWNE, W. and RASBASH, J. (2002). Multilevel modelling of medical data. *Stat. Med.* **21** 3291–3315.
- HARVILLE, D. A. (1977). Maximum likelihood approaches to variance component estimation and to related problems. *J. Amer. Statist. Assoc.* **72** 320–340. [MR0451550](#)
- HENDERSON, C. (1984). ANOVA, MIVQUE, REML, and ML algorithms for estimation of variances and covariances. In *Statistics: An Appraisal: Proceedings 50th Anniversary Conference* (H. A. David and H. T. David, eds.) 257–280. The Iowa State University Press, Ames, IA.
- JAMES, W. and STEIN, C. (1961). Estimation with quadratic loss. In *Proc. 4th Berkeley Sympos. Math. Statist. and Prob., Vol. I* 361–379. Univ. California Press, Berkeley, CA. [MR0133191](#)
- JIANG, J., NGUYEN, T. and RAO, J. S. (2011). Best predictive small area estimation. *J. Amer. Statist. Assoc.* **106** 732–745. [MR2847987](#)
- JOHNSTONE, I. M. (2011). Gaussian estimation: Sequence and wavelet models. *Unpublished Manuscript*.
- JOHNSTONE, I. M. and SILVERMAN, B. W. (2004). Needles and straw in haystacks: Empirical Bayes estimates of possibly sparse sequences. *Ann. Statist.* **32** 1594–1649. [MR2089135](#)
- KOU, S. C. and YANG, J. J. (2017). Optimal shrinkage estimation in heteroscedastic hierarchical linear models. In *Big and Complex Data Analysis. Contrib. Stat.* 249–284. Springer, Cham. [MR3674720](#)
- 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](#)
- LINDLEY, D. (1962). Discussion of the paper by Stein. *J. R. Stat. Soc., B* **24** 265–296.
- LINDLEY, D. V. and SMITH, A. F. M. (1972). Bayes estimates for the linear model. *J. R. Stat. Soc., B* **34** 1–41. [MR0415861](#)
- MASON, W. M., WONG, G. Y. and ENTWISLE, B. (1983). Contextual analysis through the multi-level linear model. *Sociol. Method.* **1984** 72–103.
- MCCULLOCH, C. E. and SEARLE, S. R. (2001). *Generalized, Linear, and Mixed Models*. Wiley, New York. [MR1884506](#)
- OMAN, S. D. (1982). Shrinking towards subspaces in multiple linear regression. *Technometrics* **24** 307–311. [MR0687188](#)
- RASBASH, J. and GOLDSTEIN, H. (1994). Efficient analysis of mixed hierarchical and cross-classified random structures using a multilevel model. *J. Educ. Behav. Stat.* **19** 337–350.
- ROLPH, J. E. (1976). Choosing shrinkage estimators for regression problems. *Comm. Statist. Theory Methods* **A5** 789–802. [MR0436481](#)

- SCLOVE, S. L. (1968). Improved estimators for coefficients in linear regression. *J. Amer. Statist. Assoc.* **63** 596–606. [MR0237057](#)
- SEARLE, S. R., CASELLA, G. and MCCULLOCH, C. E. (1992). *Variance Components*. Wiley, New York. [MR1190470](#)
- STEIN, C. M. (1962). Confidence sets for the mean of a multivariate normal distribution. *J. R. Stat. Soc., B* **24** 265–296. [MR0148184](#)
- TAN, Z. (2016). Steinized empirical Bayes estimation for heteroscedastic data. *Statist. Sinica* **26** 1219–1248. [MR3559950](#)
- 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](#)
- ZACCARIN, S. and RIVELLINI, G. (2002). Multilevel analysis in social research: An application of a cross-classified model. *Stat. Methods Appl.* **11** 95–108.

L. D. BROWN
DEPARTMENT OF STATISTICS
UNIVERSITY OF PENNSYLVANIA
400 JON M. HUNTSMAN HALL
3730 WALNUT STREET
PHILADELPHIA, PENNSYLVANIA 19104
USA
E-MAIL: lbrown@wharton.upenn.edu

G. MUKHERJEE
DEPARTMENT OF DATA SCIENCES AND OPERATIONS
MARSHALL SCHOOL OF BUSINESS
UNIVERSITY OF SOUTHERN CALIFORNIA
LOS ANGELES, CALIFORNIA 90089-0809
USA
E-MAIL: gourab@usc.edu

A. WEINSTEIN
DEPARTMENT OF STATISTICS
SEQUOIA HALL, 390 SERRA MALL
STANFORD UNIVERSITY
STANFORD, CALIFORNIA 94305-4065
USA
E-MAIL: asafw@stanford.edu