

ESTIMATION OF A COVARIANCE MATRIX USING THE REFERENCE PRIOR¹

BY RUOYONG YANG AND JAMES O. BERGER

Purdue University

Estimation of a covariance matrix Σ is a notoriously difficult problem; the standard unbiased estimator can be substantially suboptimal. We approach the problem from a noninformative prior Bayesian perspective, developing the *reference* noninformative prior for a covariance matrix and obtaining expressions for the resulting Bayes estimators. These expressions involve the computation of high-dimensional posterior expectations, which is done using a recent Markov chain simulation tool, the *hit-and-run sampler*. Frequentist risk comparisons with previously suggested estimators are also given, and determination of the accuracy of the estimators is addressed.

1. Introduction. Suppose that $\mathbf{X}_1, \dots, \mathbf{X}_n$ are i.i.d. $N_p(0, \Sigma)$, and consider the problem of estimating the $p \times p$ positive-definite Σ under the losses

$$(1) \quad L_1(\widehat{\Sigma}, \Sigma) = \text{tr}(\widehat{\Sigma}\Sigma^{-1}) - \log |\widehat{\Sigma}\Sigma^{-1}| - p,$$

$$(2) \quad L_2(\widehat{\Sigma}, \Sigma) = \text{tr}(\widehat{\Sigma}\Sigma^{-1} - I)^2,$$

where $\widehat{\Sigma}$ denotes an arbitrary estimator. The first loss was advocated by Stein (1956) and is usually called entropy loss, while the second is typically called quadratic loss. The corresponding frequentist risk functions will be denoted by

$$(3) \quad R_i(\widehat{\Sigma}, \Sigma) = E_{\Sigma} L_i(\widehat{\Sigma}, \Sigma), \quad i = 1, 2.$$

Analogous losses and risks can be defined for the problem of estimating Σ^{-1} (see Section 3.1).

The usual (unbiased) estimator of Σ is the sample covariance matrix

$$(4) \quad \frac{1}{n}S = \frac{1}{n} \sum_{i=1}^n \mathbf{X}_i \mathbf{X}_i^t \sim \frac{1}{n} W_p(\Sigma, n),$$

where $W_p(\Sigma, n)$ is the Wishart distribution with scale matrix Σ and n degrees of freedom. This estimator and $S/(n+p+1)$ are the best scalar multiples of S for L_1 and L_2 , respectively [see, e.g., Haff (1980)]. It was, however, pointed out by Stein (1956, 1975) and Dempster (1969) that the eigenstructure of Σ tends to be systematically distorted by these estimators unless p/n is quite small. The problem is especially bad when $\Sigma \cong I$. Starting with Stein's Rietz lecture [Stein

Received February 1993; January 1994.

¹Supported by NSF Grant DMS-89-23071 at Purdue University.

AMS 1991 subject classifications. Primary 62C10; secondary 62F15, 62H12.

Key words and phrases. Jeffreys prior, reference prior, covariance matrix, information matrix, Markov chain simulation, hit-and-run sampler, entropy loss, quadratic loss, risk.

(1975)], several major efforts have been made to overcome this distortion. The literature includes Stein (1975, 1977a, b), Efron and Morris (1976), Haff (1977, 1979a, b, 1980, 1991), Olkin and Selliah (1977), Sharma (1980), Sugiura and Fujimoto (1982), Sharma and Krishnamoorthy (1983, 1985a, b), Takemura (1984), Dey and Srinivasan (1985, 1986), Lin and Perlman (1985), Dey (1988), Loh (1991a, b) and Perron (1992). Note that dramatic gains in risk are achievable.

Simulation studies [cf. Lin and Perlman (1985) and Haff (1991)] seem to suggest that the estimators of Stein (1975) and Haff (1991) are particularly successful in adequately "shrinking" the eigenstructure of S . Both estimators are approximately Bayes (especially that of Haff) but require incorporation of an isotonizing step in their computation to avoid overshrinkage of certain eigenvalues. Also, no approach to shrinkage estimation of Σ has produced reportable measures of the accuracy of $\hat{\Sigma}$. This is a serious limitation.

Because of the centrality in statistics of the covariance matrix estimation problem and because of the limitations of the existing estimation methods, it seemed desirable to attempt a fully Bayesian approach to the problem based on use of reference (noninformative) priors. These priors seem to be remarkably successful in many multivariate problems in producing estimators with simultaneously good Bayesian and frequentist properties [cf. Berger and Bernardo (1992a, b, c) and Ye and Berger (1991)]. Also, they tend to yield very satisfactory measures of accuracy, through the posterior covariance, or posterior expected loss.

Section 2 contains the development of the reference prior for this problem. Rather surprisingly, the reference prior turns out to be remarkably simple. Indeed, it is the prior proposed by Chang and Eaves (1990), which was based on the simpler (but less satisfactory) reference prior algorithm in Bernardo (1979). Not unexpectedly, however, computation with this prior is not possible in closed form; thus Section 3 develops an efficient computational scheme. Section 4 compares the reference prior Bayes estimator to the estimators of Stein (1975) and Haff (1991). Section 5 discusses determination of the accuracy of $\hat{\Sigma}$.

Note that there have been previous partial Bayesian approaches to estimation of Σ . These include the empirical Bayes analyses of Efron and Morris (1976) and Haff (1980). Conjugate priors have also been used [cf. Press (1982)], but these do not achieve the type of eigenvalue shrinkage that seems most desirable. A flexible and very appealing general class of prior distributions for Σ has recently been introduced by Léonard and Hsu (1992). Their approach allows for a wide variety of subjective shrinkage patterns, but it is not clear if the shrinkage pattern we seek can be reproduced in this way.

The common noninformative prior for the problem has been the Jeffreys prior

$$(5) \quad \pi(\Sigma) = (\det \Sigma)^{-(p+1)/2} d\Sigma.$$

This prior was developed by Jeffreys (1961), for $p = 1, 2$, and by Geisser and Cornfield (1963), Geisser (1965) and Villegas (1969), for arbitrary p . Use of the Jeffreys prior tends simply to reproduce classical answers, however, and hence also fails to shrink the eigenvalues appropriately.

The work most closely related to this study is that of Haff (1991), which proposes an estimator based on a variational form of the Bayes estimator. In the derivation of Haff’s estimator, however, a term in the expression for the Bayes estimator is (purposely) ignored, so that it is unclear if the result actually corresponds to a Bayes rule or what the implied prior distribution might be. We do, however, observe considerable similarity between our estimator and that of Haff.

2. The reference prior for a covariance matrix.

2.1. *The Fisher information for a covariance matrix.* We will use the following notation. The entries of a matrix A will be denoted by $A_{[i,j]}$, and A^t , $|A|$ and $\text{tr}(A)$ will denote the transpose, determinant and trace of a square matrix A , respectively. Denote the matrix operator which arranges the columns of a matrix into one long column as $\text{vec}(\)$. The Kronecker product of two matrices A and B will be denoted by $A \otimes B$. The covariance matrix Σ can be decomposed as $\Sigma = O^t D O$, with O an orthogonal matrix with positive elements for the first row and D a diagonal matrix, $D = \text{diag}(d_1, \dots, d_p)$, with $d_1 \geq d_2 \geq \dots \geq d_p \geq 0$. Write $O = (O_{12} O_{13} \dots O_{1p})(O_{23} \dots O_{2p}) \dots (O_{p-1,p}) D_\epsilon$, with O_{ij} being a simple orthogonal matrix such as

$$(6) \quad O_{ij} = O_{ij}(o_{ij}) = \begin{matrix} & & i & & j & & \\ & & & & & & \\ i & \left(\begin{array}{ccccc} I & 0 & 0 & 0 & 0 \\ 0 & \cos o_{ij} & 0 & -\sin o_{ij} & 0 \\ 0 & 0 & I & 0 & 0 \\ 0 & \sin o_{ij} & 0 & \cos o_{ij} & 0 \\ 0 & 0 & 0 & 0 & I \end{array} \right) & & \\ j & & & & & & \end{matrix}$$

where $-\pi/2 < o_{ij} \leq \pi/2$, and D_ϵ is a diagonal matrix with diagonal elements ± 1 [see Anderson, Olkin and Underhill (1987)]. Let $(d\Sigma)$ denote $\prod_{i < j} d\sigma_{ij}$, let (dD) denote $\prod_{i=1}^p dd_i$, let (dO) denote $\prod_{i < j} do_{ij}$ and let (dH) denote the *conditional invariant Haar measure* over the space of orthogonal matrices $\mathcal{O} = \{O: O^t O = I\}$ [see Anderson (1984) for definition].

LEMMA 1. *The Fisher information matrix for Σ , w.r.t. the reparametrization (D, O) , is of the form*

$$(7) \quad I(D, O) = \begin{pmatrix} I(D) & 0 \\ 0 & I(O) \end{pmatrix},$$

with $I(D) = \text{diag}(1/(2d_1^2), \dots, 1/(2d_p^2))$. [Note that the explicit form of $I(O)$ will not be needed.]

See Appendix A for the proof of Lemma 1.

LEMMA 2.

(i) *The determinant of the Fisher information matrix of Σ is*

$$(8) \quad |I(\Sigma)| \propto |\Sigma|^{-(p+1)}.$$

(ii) *The relationship between the Fisher information matrix w.r.t. the parameter Σ and (D, O) is*

$$(9) \quad I(D, O) = \left[\frac{\partial(\Sigma)}{\partial(D, O)} \right]^t I(\Sigma) \left[\frac{\partial(\Sigma)}{\partial(D, O)} \right]$$

and

$$(10) \quad (d\Sigma) \propto \left[\prod_{i < j} (d_i - d_j) \right] (dD)(dH)$$

$$(11) \quad \propto \left[\prod_{i=1}^{p-1} \prod_{j=i+1}^p \cos^{j-i-1} o_{ij} \right] \left[\prod_{i < j} (d_i - d_j) \right] (dD)(dO).$$

(iii) *The determinant of $I(O)$ in Lemma 1 is*

$$(12) \quad |I(O)| \propto \left[\prod_{i=1}^{p-1} \prod_{j=i+1}^p \cos^{j-i-1} o_{ij} \right]^2 \left[\prod_{i < j} (d_i - d_j) \right]^2 \prod_{i=1}^p d_i^{-(p-1)}.$$

PROOF. Equation (9) is trivial; for (8), see Press [(1982), page 79]; for (10), see Farrell [(1985), page 74)]; and for (11), see Anderson, Olkin and Underhill (1987). Part (iii) follows from the representation $\Sigma = O^t D O$. \square

2.2. *The reference prior.* Bernardo (1979) initiated an information-based approach to development of noninformative priors, called the reference prior approach. A review and discussion of the current status of the approach can be found in Berger and Bernardo (1992c). The motivation for developing the approach was the acknowledged problems of the Jeffreys prior in higher dimensions. Even Jeffreys would often alter the Jeffreys prior in multiparameter problems to remove perceived inadequacies. The reference prior approach seeks to overcome these difficulties by breaking up multiparameter problems into a series of conditional one-parameter problems, for which reasonable noninformative priors can be determined. The approach has proven to be remarkably successful in overcoming the inadequacies of the Jeffreys prior in multiparameter problems [cf. Berger and Bernardo (1989, 1992a, b, c) and Ye and Berger (1991)].

In the following theorem, the reference prior for Σ is given. The Jeffreys prior is also given for comparison. Note that the reference prior can depend on what

is called the group ordering, which is typically simply a listing of parameters according to perceived "importance." Note, also, that the reference prior here was first given in Chang and Eaves (1990), although their derivation utilized the early version of the reference prior algorithm in Bernardo (1979), which was improved in Berger and Bernardo (1992a, b, c).

THEOREM 1. *The reference prior for the parameter (D, O) is as follows, providing the group ordering used lists D before O and the $\{d_i\}$ are ordered monotonically (either increasing or decreasing):*

$$\begin{aligned} \pi_R(D, O)(dD)(dO) &\propto \frac{\left[\prod_{i=1}^{p-1} \prod_{j=i+1}^p \cos^{j-i-1} o_{ij} \right]}{|D|} (dD)(dO) \\ (13) \quad &\propto \frac{1}{|D|} (dD)(dH) \\ &\propto \frac{1}{|\Sigma| \prod_{i < j} (d_i - d_j)} (d\Sigma). \end{aligned}$$

The Jeffreys prior is

$$\begin{aligned} \pi_J(D, O)(dD)(dO) &\propto \frac{\prod_{i < j} (d_i - d_j)}{|D|^{(p+1)/2}} (dD)(dH) \\ (14) \quad &\propto |\Sigma|^{-(p+1)/2} (d\Sigma). \end{aligned}$$

See Appendix B for the proof of Theorem 1.

COROLLARY 1. *The resulting posterior distributions are*

$$\begin{aligned} \pi_R(\Sigma | S)(d\Sigma) &\propto \frac{\text{etr}\left(-\frac{1}{2}\Sigma^{-1}S\right)}{|\Sigma|^{n/2+1} \prod_{i < j} (d_i - d_j)} (d\Sigma) \\ (15) \quad &\propto \frac{\text{etr}\left(-\frac{1}{2}OD^{-1}O^tS\right)}{|D|^{n/2+1}} (dD)(dH); \end{aligned}$$

$$\begin{aligned} \pi_J(\Sigma | S)(d\Sigma) &\propto \frac{\text{etr}\left(-\frac{1}{2}\Sigma^{-1}S\right)}{|\Sigma|^{(n+p+1)/2}} (d\Sigma) \\ (16) \quad &\propto \prod_{i < j} (d_i - d_j) \frac{\text{etr}\left(-\frac{1}{2}OD^{-1}O^tS\right)}{|D|^{(n+p+1)/2}} (dD)(dH), \end{aligned}$$

where etr stands for $\exp(\text{tr}(\cdot))$.

PROOF. Simply multiply the prior and the likelihood. \square

Note that the posterior in (15) is proper, having all moments of order less than $n/2$ (including negative moments), because it is bounded by an inverse

Gamma distribution. Compared to the Jeffreys prior, note that the reference prior seems to put considerably more mass near the region of equality of the eigenvalues; thus it is intuitively plausible that the reference prior would produce a covariance matrix estimator with better eigenstructure shrinkage.

Sometimes Σ^{-1} , rather than Σ itself, is of interest. Note, however, that the reference prior for Σ^{-1} will be the same as that for Σ . This follows from the fact that $\Sigma^{-1} = O^t D^{-1} O$ and that the reference prior for the group ordering that lists first the ordered $\{d_i^{-1}\}$, and then O , is the same as that listing $\{d_i\}$ followed by O . (It can be shown that a one-to-one transformation of an element of the group ordering does not change the reference prior.) Similarly, if it is the eigenvalues of Σ that are of interest, the reference prior again turns out to be given by (13). It is methodologically pleasant that this same reference prior emerges for any of the usual quantities of interest.

3. Computation of the Bayes estimators.

3.1. *Bayes estimators for Σ and Σ^{-1} .* To find the Bayes estimators for Σ w.r.t. the loss functions L_1 and L_2 , one merely minimizes the associated posterior expected losses. The proofs of the following lemma and corollary are straightforward and are omitted.

LEMMA 3. *The Bayes estimator for Σ w.r.t. the posterior $\pi(\Sigma | S)$ and under L_1 is*

$$(17) \quad \delta_1^\pi = [E^{\pi(\Sigma | S)} \Sigma^{-1}]^{-1};$$

the Bayes estimator under L_2 is

$$(18) \quad \text{vec}(\delta_2^\pi) = [E^{\pi(\Sigma | S)} (\Sigma^{-1} \otimes \Sigma^{-1})]^{-1} \text{vec}[E^{\pi(\Sigma | S)} \Sigma^{-1}].$$

Both δ_1^π and δ_2^π are orthogonally invariant in the sense $\delta_i^{\pi(\Sigma | \Gamma S \Gamma^t)} = \Gamma \delta_i^{\pi(\Sigma | S)} \Gamma^t$, $i = 1, 2$, provided the prior is orthogonally invariant in the sense $\pi(\Gamma \Sigma \Gamma^t) = \pi(\Sigma)$, where Γ is an arbitrary orthogonal matrix. Also, for such priors, the Bayes estimators are diagonal when S is diagonal.

COROLLARY 2. *The Jeffreys prior Bayes estimator for the covariance matrix under L_1 is the usual unbiased estimator S/n .*

Often, estimation of Σ^{-1} , rather than Σ , is desired. The literature in this field includes Efron and Morris (1976), Haff (1977), Sharma and Krishnamoorthy (1985b), Sinha and Ghosh (1986), Krishnamoorthy and Gupta (1989) and Krishnamoorthy (1991). The commonly used loss functions are the natural analogues of (1) and (2), namely,

$$L_1(\widehat{\Sigma}^{-1}, \Sigma^{-1}) = \text{tr}(\widehat{\Sigma}^{-1} \Sigma) - \log |\widehat{\Sigma}^{-1} \Sigma| - p,$$

$$L_2(\widehat{\Sigma}^{-1}, \Sigma^{-1}) = \text{tr}(\widehat{\Sigma}^{-1} \Sigma - I)^2.$$

As in Lemma 3, these two loss functions result in Bayes estimators of Σ^{-1} given by, respectively,

$$[E^{\pi(\Sigma|S)}\Sigma]^{-1} \quad \text{and} \quad [E^{\pi(\Sigma|S)}(\Sigma \otimes \Sigma)]^{-1} \text{vec}[E^{\pi(\Sigma|S)}\Sigma].$$

Efron and Morris (1976) and Haff (1977) used a slightly different loss function,

$$L(\widehat{\Sigma}^{-1}, \Sigma^{-1}; S) = \frac{\text{tr}[(\widehat{\Sigma}^{-1} - \Sigma^{-1})^2 S]}{k \text{tr}(\Sigma^{-1})};$$

this would result in the Bayes estimator $E^{\pi(\Sigma|S)}[\Sigma^{-1}/\text{tr}(\Sigma^{-1})]$. Haff (1977) also considered the loss function

$$L(\widehat{\Sigma}^{-1}, \Sigma^{-1}) = \text{tr}[(\widehat{\Sigma}^{-1} - \Sigma^{-1})^2 Q],$$

where Q is an arbitrary positive-definite matrix; this would result in the simple Bayes estimator $E^{\pi(\Sigma|S)}[\Sigma^{-1}]$.

3.2. *Hit-and-run sampler.* For the reference posterior, analytical evaluation of the quantities in Lemma 3 appears to be quite difficult. Thus we turn to Monte Carlo integration to do the computation.

Recently, Monte Carlo methods for Bayesian integration have undergone extensive development. The methods that are commonly used are importance sampling, data augmentation and the Gibbs sampler. Attempts to apply these methods encountered difficulties, so we turned to the less common hit-and-run sampler, which is another Markov chain sampler.

The hit-and-run sampler was first proposed by Smith (1980, 1984) and later generalized by Belisle, Romeijn and Smith (1993). The algorithm we used is a version that was developed by Schmeiser and Chen (1991) and Chen and Schmeiser (1993) and is called the *Metropolized hit-and-run sampler*. This algorithm is particularly useful when the domain of the posterior along a random direction from a given point can be obtained without undue difficulty.

When sampling from the posterior (15), it will be convenient to transform from the space of positive-definite matrices to all of Euclidean space. We do this, as in Leonard and Hsu (1992), by defining $\Sigma^* = \log \Sigma$, or $\Sigma = \exp(\Sigma^*)$, in the sense that

$$(19) \quad \Sigma = \sum_{i=0}^{\infty} \frac{(\Sigma^*)^i}{i!}.$$

Using Lemma 2, it can be shown that the reference posterior for Σ^* is

$$(20) \quad \pi_R^*(\Sigma^* | S)(d\Sigma^*) \propto \frac{\text{etr}\{- (n/2)D^* - (1/2)O \exp(-D^*)O^t S\}}{\prod_{i < j} (d_i^* - d_j^*)} (d\Sigma^*),$$

with $\Sigma^* = OD^*O^t$, $D^* = \text{diag}(d_1^*, \dots, d_p^*)$, $d_1^* \geq d_2^* \geq \dots \geq d_p^*$ and O orthogonal.

Our actual sampling procedure thus proceeds as follows:

1. Select a starting positive-definite matrix Σ_0 , set $\Sigma_0^* = \log \Sigma_0$ and $k = 0$. Here we choose $\Sigma_0 = (1/n)S$.
2. Select a random direction symmetric $p \times p$ matrix T , with elements $t_{ij} = z_{ij}/\sqrt{\sum_{l \leq m} z_{lm}^2}$, where $z_{ij} \sim_{i.i.d} N(0, 1), i \leq j$. (The other elements of T are defined by symmetry.)
3. Generate $\lambda \sim N(0, 1)$.
4. Set $Y = \Sigma_k^* + \lambda T$. Then set

$$(21) \quad \Sigma_{k+1}^* = \begin{cases} Y, & \text{with probability } \min\left(1, \pi^*(Y|S)/\pi^*(\Sigma_k^*|S)\right), \\ \Sigma_k^*, & \text{otherwise.} \end{cases}$$

5. Set $k = k + 1$ and go back to step 2.

Finally, after a sufficiently large sample $\Sigma_1^*, \Sigma_2^*, \dots, \Sigma_N^*$ has been generated, one simply approximates a posterior expectation by $E^{\pi(\Sigma|S)}f(\Sigma) \approx (1/N)\sum_{k=1}^N f(e^{\Sigma_k^*})$, where f is the function of interest. As $N \rightarrow \infty$, the ergodic theorem asserts that the approximation converges to the true value [see Schmeiser and Chen (1991)]. Of course, one should simultaneously evaluate $E^{\pi(\Sigma|S)}[f(\Sigma)]$ for all f of interest. In the simulation in Section 4, we set $p = 5$ (so that the integrals are 15-dimensional) and $N = 50,000$. This gave simulation accuracies [$100 \times (\text{simulation error})/(\text{true value})$] of about 1.5% for the loss $L_1(\delta_1^\pi, \Sigma)$ and 0.75% for $L_2(\delta_2^\pi, \Sigma)$, the quantities needed in the risk evaluations. The individual elements of δ_1^π and δ_2^π were not quite so accurate, having simulation accuracies of about 5%.

4. Frequentist risk comparisons.

4.1. *Stein's and Haff's estimators.* Writing $S = VL V^t$, where V is an orthogonal matrix and $L = \text{diag}(l_1, \dots, l_p)$ with $l_1 \geq l_2 \geq \dots \geq l_p$, Stein (1975) considered the orthogonal invariant estimator

$$(22) \quad \hat{\Sigma} = V\Phi(L)V^t,$$

where $\Phi(L) = \text{diag}(\phi_1, \dots, \phi_p)$ with $\phi_i = l_i/\alpha_i$,

$$(23) \quad \alpha_i = (n - p + 1) + 2l_i \sum_{j \neq i} \frac{1}{l_i - l_j}, \quad i = 1, \dots, p.$$

This estimator has two problems. First, the intuitively compatible ordering $\phi_1 \geq \phi_2 \geq \dots \geq \phi_p$ is frequently violated. Second, and more serious, some of the ϕ_i may even be negative. Stein suggests an isotonizing algorithm to avoid these problems. The idea of the algorithm is to pool the adjacent pairs (l_i, α_i) . The resulting estimators of the eigenvalues are

$$(24) \quad \phi_i = \phi_{i+1} = \dots = \phi_{i+s} = \frac{l_i + l_{i+1} + \dots + l_{i+s}}{\alpha_i + \alpha_{i+1} + \dots + \alpha_{i+s}}.$$

The details of this isotonizing algorithm can be found in Lin and Perlman (1985).

Haff's estimator [Haff (1991)] is closely related to the above estimator. He minimizes the formal Bayes risk for an orthogonally invariant prior by a variational technique. Assuming the prior yields $1/|S|$ as the marginal distribution of S , this technique reproduces Stein's unconstrained estimator. By imposing the constraint $\phi_1 \geq \phi_2 \geq \dots \geq \phi_p$ in the minimization under L_1 , the formal Bayes estimator is of the form (23) with the eigenvalue estimators obtained by solving the equations

$$(25) \quad \varepsilon_i \sum_{j=1}^i \left[(\phi_j^s)^{-1} - (\phi_j)^{-1} \right] = 0, \quad i = 1, 2, \dots, p,$$

where $\phi_j^s = l_j/\alpha_j$, $j = 1, 2, \dots, p$, and $\varepsilon_1^2 = \phi_1 - \phi_2, \varepsilon_2^2 = \phi_2 - \phi_3, \dots, \varepsilon_p^2 = \phi_p$.

The two estimators discussed above are both obtained under L_1 . Stein's and Haff's methods are difficult to apply under L_2 . Thus it is common to take the L_1 estimators and simply rescale for L_2 [see Haff (1991) and Lin and Perlman (1985)]; that is, if $\widehat{\Sigma}$ is derived under L_1 , then one simply considers the estimator $n\widehat{\Sigma}/(n + p + 1)$ under L_2 . This corresponds to the optimal rescaling for the unbiased estimator under L_1 . Note that such ad hoc adjustments are not required for the Bayes estimators.

4.2. Risk simulations. The frequentist risks of the various estimators under L_1 and L_2 will be approximated by average losses in simulation. The simulation was designed as follows: Set $p = 5$ and $n = 10, 20, 40$. The test covariance matrices were chosen to be $\Sigma_1 = \text{diag}(1, 1, 1, 1, 1)$, $\Sigma_2 = \text{diag}(5, 4, 3, 2, 1)$ and $\Sigma_3 = \text{diag}(16, 8, 4, 2, 1)$. For fixed n and Σ_i , we do the following:

1. Generate 50 random $Y_k \sim W_p(I, n)$, $1 \leq k \leq 50$, using Bartlett's decomposition, and then transform them into $W_p(\Sigma_i, n)$ random variables S_k .
2. For each observation S_k , estimate the covariance matrix using the reference prior Bayes estimator, Stein's estimator and Haff's estimator, under L_1 and L_2 . Record the associated loss for each estimator.
3. Compute the mean and standard error of the *differences* in loss between the three different estimators.
4. Following the tradition of Lin and Perlman (1985), we also record the percentage reduction in average loss (PRIAL) of the three estimators relative to the usual estimator, defined by

$$(26) \quad \text{PRIAL} = \frac{R\left(\frac{1}{n}S, \Sigma\right) - R(\widehat{\Sigma}, \Sigma)}{R\left(\frac{1}{n}S, \Sigma\right)} \times 100 \quad \text{for } L_1;$$

$$(27) \quad \text{PRIAL} = \frac{R\left(\frac{1}{(n+p+1)}S, \Sigma\right) - R(\widehat{\Sigma}, \Sigma)}{R\left(\frac{1}{(n+p+1)}S, \Sigma\right)} \times 100 \quad \text{for } L_2.$$

TABLE 1
Risk differences of reference prior and Stein and Haff estimators

	<i>n</i>	<i>L</i> ₁		<i>L</i> ₂	
		<i>R</i> (Stein)– <i>R</i> (ref.)	<i>R</i> (Haff)– <i>R</i> (ref.)	<i>R</i> (Stein)– <i>R</i> (ref.)	<i>R</i> (Haff)– <i>R</i> (ref.)
Σ ₁	10	–0.14 (0.016)	–0.15 (0.029)	0.023 (0.029)	–0.13 (0.023)
	20	–0.056 (0.0063)	–0.059 (0.010)	0.0023 (0.015)	–0.061 (0.012)
	40	–0.031 (0.0033)	–0.029 (0.0049)	–0.025 (0.0093)	–0.036 (0.0071)
Σ ₂	10	0.049 (0.024)	0.045 (0.032)	0.089 (0.029)	0.035 (0.036)
	20	0.050 (0.011)	0.054 (0.011)	0.069 (0.019)	0.066 (0.019)
	40	0.011 (0.0037)	0.012 (0.0043)	0.014 (0.0070)	0.016 (0.0073)
Σ ₃	10	0.10 (0.026)	0.11 (0.032)	0.095 (0.031)	0.11 (0.042)
	20	0.044 (0.0089)	0.048 (0.0089)	0.045 (0.015)	0.051 (0.014)
	40	0.017 (0.0033)	0.017 (0.0034)	0.030 (0.0066)	0.030 (0.0070)

The simulation results for frequentist risk are given in Table 1, with the standard errors in parentheses. Table 2 presents the results for PRIAL.

The performance of the reference prior Bayes estimator is very comparable to that of the Stein and Haff estimators. It is somewhat worse when Σ = *I*, somewhat better otherwise. This behavior is indicative of an estimator that has more moderate shrinkage than that of the Stein or Haff estimator. This might well be desirable; indeed, for Σ far from the identity matrix (i.e., Σ₃), there is a suggestion in Table 2 that the Stein and Haff estimators slightly overshrink, at least for *L*₂, where the PRIAL can become negative.

To investigate this further, we considered alternative reference priors that happened to yield more shrinkage. For instance, if one applies the reference prior algorithm to the ordered group $\{(d_1, d_p), (d_2, \dots, d_{p-1}), (o_{12}, \dots, o_{p-1,p})\}$ [see Berger and Bernardo (1992c) for definition], the reference prior turns out to be

$$(28) \quad \pi_{R^*}(D, O)(dD)(dH) \propto |D|^{-1}[\log d_1 - \log d_p]^{-(p-2)}(dD)(dH),$$

TABLE 2
PRIAL relative to the usual estimator

	<i>n</i>	<i>L</i> ₁				<i>L</i> ₂			
		Ref.	Ref.*	Stein	Haff	Ref.	Ref.*	Stein	Haff
Σ ₁	10	59.02	80.03	66.98	67.34	35.02	59.17	33.76	41.92
	20	64.06	84.19	71.15	71.50	50.46	75.36	50.25	55.89
	40	63.61	80.95	71.63	71.06	55.78	75.31	59.55	61.23
Σ ₂	10	45.99	55.41	43.28	43.50	26.12	32.47	21.29	24.23
	20	34.71	31.58	28.35	27.84	23.69	16.50	17.52	17.82
	40	21.25	17.52	18.46	18.05	15.91	9.01	13.74	13.41
Σ ₃	10	31.92	33.71	26.28	25.55	16.08	16.32	10.89	10.33
	20	14.08	11.61	8.45	8.01	8.80	6.10	4.82	4.26
	40	6.09	6.02	1.73	1.75	4.39	3.88	–1.66	–2.24

TABLE 3
Risk differences of modified reference prior and Stein and Haff estimators

	<i>n</i>	L_1		L_2	
		$R(\text{Stein}) - R(\text{ref.}^*)$	$R(\text{Haff}) - R(\text{ref.}^*)$	$R(\text{Stein}) - R(\text{ref.}^*)$	$R(\text{Haff}) - R(\text{ref.}^*)$
Σ_1	10	0.23 (0.036)	0.23 (0.042)	0.47 (0.037)	0.32 (0.034)
	20	0.10 (0.015)	0.10 (0.016)	0.28 (0.024)	0.22 (0.023)
	40	0.036 (0.0076)	0.039 (0.0075)	0.10 (0.021)	0.092 (0.018)
Σ_2	10	0.22 (0.041)	0.21 (0.048)	0.21 (0.040)	0.15 (0.040)
	20	0.025 (0.019)	0.029 (0.018)	-0.011 (0.034)	-0.014 (0.032)
	40	-0.0036 (0.011)	-0.0021 (0.011)	-0.031 (0.032)	-0.029 (0.031)
Σ_3	10	0.13 (0.032)	0.15 (0.034)	0.10 (0.053)	0.11 (0.056)
	20	0.025 (0.017)	0.028 (0.016)	0.014 (0.026)	0.021 (0.024)
	40	0.017 (0.0051)	0.017 (0.0052)	0.027 (0.011)	0.027 (0.011)

which intuitively will induce more shrinkage than will (13). Risk differences and PRIAL's for the associated Bayes estimators are given in Table 3 and the "Ref.*" columns of Table 2, respectively. The apparently more aggressive shrinkage of these estimators results in dramatically improved risk when $\Sigma \cong I$. These high-shrinkage Bayes estimators seem somewhat superior to the Stein and Haff estimators.

Note that there is nothing particular compelling about π_{R^*} from a reference prior perspective. Indeed, we would still probably recommend π_R in (13); it is likely to have better confidence properties, and in general one should be wary of overshrinking. It must be admitted, however, that all these conclusions are quite tentative, being based on only a very limited study. Indeed, we feel that all four estimators considered here are comparably effective in practice.

5. Determination of accuracy. It is something of a bonus that the reference prior Bayes estimators may actually have superior risk properties, compared with existing shrinkage estimators. The primary motivation and value of the Bayesian approach to estimation problems such as this is, rather, that the Bayesian approach readily allows determination of accuracy of estimation and allows associated prediction and numerous other types of inference involving Σ . We illustrate this here by providing estimates of the loss of $\hat{\Sigma}$ under L_1 and L_2 . The proof of the following lemma is straightforward.

LEMMA 4. *The posterior expected loss of the Bayes estimator in Lemma 3 under L_1 is*

$$\rho_1(\pi(\Sigma | S), \delta_1^\pi) = E^{\pi(\Sigma | S)} \log |\Sigma| - \log |\delta_1^\pi|;$$

the posterior expected loss under L_2 is

$$\rho_2(\pi(\Sigma | S), \delta_2^\pi) = p - \text{tr}[\delta_2^\pi (\delta_1^\pi)^{-1}].$$

EXAMPLE 1. Consider, as data, the following matrix S/n , where S is generated from a $W_5(\Sigma, 10)$ distribution, with $\Sigma = \text{diag}(5, 4, 3, 2, 1)$:

$$\frac{S}{n} = \begin{pmatrix} 1.925 & 1.618 & 0.132 & -1.101 & 0.264 \\ 1.618 & 8.437 & 1.638 & -0.880 & -0.983 \\ 0.132 & 1.638 & 2.147 & -0.439 & -0.646 \\ -1.101 & -0.880 & -0.439 & 1.331 & -0.035 \\ 0.264 & -0.983 & -0.646 & -0.035 & 1.280 \end{pmatrix}.$$

The corresponding Bayes estimators are

$$\delta_1^\pi = \begin{pmatrix} 1.726 & 0.946 & 0.077 & -0.712 & 0.207 \\ 0.946 & 5.536 & 0.917 & -0.483 & -0.567 \\ 0.077 & 0.917 & 1.889 & -0.282 & -0.385 \\ -0.712 & -0.483 & -0.282 & 1.366 & -0.022 \\ 0.207 & -0.567 & -0.385 & -0.022 & 1.412 \end{pmatrix},$$

$$\delta_2^\pi = \begin{pmatrix} 1.233 & 0.723 & 0.057 & -0.567 & 0.165 \\ 0.723 & 4.120 & 0.682 & -0.364 & -0.413 \\ 0.057 & 0.682 & 1.371 & -0.215 & -0.307 \\ -0.567 & -0.364 & -0.215 & 0.936 & -0.018 \\ 0.165 & -0.413 & -0.307 & -0.018 & 0.993 \end{pmatrix}.$$

Using Lemma 4, the posterior expected losses of these estimators can be computed to be $\rho_1(\pi(\Sigma | S), \delta_1^\pi) = 1.152$ and $\rho_2(\pi(\Sigma | S), \delta_2^\pi) = 1.509$, respectively. Note that the actual losses (computable since we know Σ) are $L_1(\delta_1^\pi, \Sigma) = 1.270$ and $L_2(\delta_2^\pi, \Sigma) = 1.548$, respectively. [For comparison, observe that the actual losses for S/n and $S/(n+p+1)$ are $L_1(S/n, \Sigma) = 2.267$ and $L_2(S/(n+p+1), \Sigma) = 2.147$.]

Classical estimates of loss (the unbiased estimates of risk) are available here [cf. Haff (1991)], but they are unwieldy and potentially unreliable (e.g., they can even be negative).

Other estimates of accuracy could be found using the Bayesian approach, such as the posterior covariance matrix for Σ (a $[p(p+1)/2] \times [p(p+1)/2]$ matrix), or even credible intervals for components of Σ . This could also be done for functions of Σ that are of interest.

6. Comments and generalizations.

1. Although unquestionably computationally intensive, the rewards for adopting the Bayesian approach here are considerable. Resulting estimators have exceptional risk properties, and determination of accuracy and inference for functions of Σ is straightforward. Indeed, once the computation for $\hat{\Sigma}$ is set up, it is easy to compute the expectation of any function $g(\Sigma)$ that is of interest.

2. The situation in which the \mathbf{X}_i are i.i.d. $N_p(\mu, \Sigma)$ can be handled similarly; the reference prior will be constant over μ , so that μ can simply be integrated out to reduce the problem (in a Bayesian sense) to consideration of

$$S = \sum_{i=1}^n (\mathbf{X}_i - \bar{\mathbf{X}})(\mathbf{X}_i - \bar{\mathbf{X}})^t \sim W_p(\Sigma, n - 1).$$

Analysis then proceeds as before, with n replaced by $n - 1$. Note, however, that, from a frequentist or a hierarchical Bayesian perspective, there might be advantages in utilizing “shrinkage priors” for μ , rather than the constant prior.

3. As with all scale problems, choice of the loss is a rather perplexing question. From Example 1 it is clear that the effect can be substantial. We have no firm recommendation here, except to note that, when $p = 1$, the Jeffreys and reference priors both equal the usual invariant prior $1/\sigma^2$, and use of L_1 with this prior yields, as the Bayes estimator, the “standard” estimator S/n (see Corollary 2). Hence use of L_1 , and the corresponding δ_1^π , has some appeal.
4. It can be argued that the reference prior should depend on the loss function. For instance, when $p = 1$, the standard reference prior is $1/\sigma^2$, and this is completely satisfactory for invariant losses such as L_1 or L_2 , but it is not optimal for, say, squared error loss. Unfortunately, it is not easy, in general, to determine how the reference prior should depend on the loss [see Bernardo (1981) and Bernardo and Smith (1994) for discussion]. In our situation, the problem is probably not severe, since we only utilized invariant losses.
5. A related problem can arise if one is interested in some function $g(\Sigma)$ of Σ . Conceivably a better reference prior can be developed that recognizes the centrality of $g(\Sigma)$ [cf. Berger and Bernardo (1992c)]. Use of the given reference prior π_R is likely to be quite satisfactory, however, especially because it arises from so many different group orderings that it will be the reference prior for “most” $g(\Sigma)$.

APPENDIX A

PROOF OF LEMMA 1. Tracy and Jinadasa (1988) established that the Fisher information matrix for Σ is $I(\Sigma) = \frac{1}{2}G^t(\Sigma^{-1} \otimes \Sigma^{-1})G$, where G is defined as $G = \partial \text{vec}(\Sigma)/\partial \text{vecp}(\Sigma)$, with $\Sigma = O^t D O$ and $\text{vecp}(\Sigma) = (d_1, \dots, d_p, o_{12}, \dots, o_{1p}, o_{23}, \dots, o_{2p}, \dots, o_{p-1,p})$. Writing $\mathbf{a}_i = \text{vec} \partial(O^t D O)/\partial d_i$ and $\mathbf{b}_{ij} = \text{vec} \partial(O^t D O)/\partial o_{ij}$ yields $G = (\mathbf{a}_1, \dots, \mathbf{a}_p, \mathbf{b}_{12}, \dots, \mathbf{b}_{p-1,p})$, and thus the Fisher information matrix, w.r.t. the reparameterization (D, O) , is $I(D, O) = \frac{1}{2}G^t(\Sigma^{-1} \otimes \Sigma^{-1})G$, where G and Σ are to be considered functions of D and O . The elements of $I(D, O)$ are of three types: (i) $\mathbf{a}_i^t(\Sigma^{-1} \otimes \Sigma^{-1})\mathbf{a}_j$, (ii) $\mathbf{a}_i^t(\Sigma^{-1} \otimes \Sigma^{-1})\mathbf{b}_{rs}$ or $\mathbf{b}_{rs}^t(\Sigma^{-1} \otimes \Sigma^{-1})\mathbf{a}_i$; and (iii) $\mathbf{b}_{ij}^t(\Sigma^{-1} \otimes \Sigma^{-1})\mathbf{b}_{rs}$. To finish the proof of Lemma 1, we need only evaluate the first two types. We will utilize the following matrix equality [from, e.g., Magnus and Neudecker (1988)] to do so:

$$(29) \quad \text{tr}(ABCD) = (\text{vec } D^t)^t (C^t \otimes A)(\text{vec } B).$$

(i) For the first type,

$$\begin{aligned}
 \mathbf{a}_i^t(\Sigma^{-1} \otimes \Sigma^{-1})\mathbf{a}_j &= \left(\text{vec } \frac{\partial \Sigma}{\partial d_i}\right)(\Sigma^{-1} \otimes \Sigma^{-1})\left(\text{vec } \frac{\partial \Sigma}{\partial d_j}\right) \\
 &= \text{tr}\left(\Sigma^{-1} \frac{\partial \Sigma}{\partial d_j} \Sigma^{-1} \frac{\partial \Sigma}{\partial d_i}\right) \\
 (30) \qquad &= \text{tr}\left[D^{-1} \frac{\partial D}{\partial d_j} D^{-1} \frac{\partial D}{\partial d_i}\right] \\
 &= \begin{cases} 1 & \text{if } i = j, \\ d_i^2 & \\ 0 & \text{otherwise.} \end{cases}
 \end{aligned}$$

(ii) For the second type,

$$\begin{aligned}
 \mathbf{a}_i^t(\Sigma^{-1} \otimes \Sigma^{-1})\mathbf{b}_{rs} &= \text{tr}\left(\Sigma^{-1} \frac{\partial \Sigma}{\partial o_{rs}} \Sigma^{-1} \frac{\partial \Sigma}{\partial d_i}\right) \\
 (31) \qquad &= \text{tr}\left[D^{-1} \frac{\partial D}{\partial d_i} D^{-1} \cdot O \frac{\partial(O^t D O)}{\partial o_{rs}} O^t\right] \\
 &= \text{tr}\left[\frac{\partial D}{\partial d_i} D^{-1} \cdot O \frac{\partial O^t}{\partial o_{rs}} + D^{-1} \frac{\partial D}{\partial d_i} \cdot \frac{\partial O}{\partial o_{rs}} O^t\right] \\
 &= \text{tr}\left[\text{diag}\left(0, \dots, 0, \frac{1}{d_i}, 0, \dots, 0\right)(C + C^t)\right],
 \end{aligned}$$

where $C = O \partial O^t / \partial o_{rs}$.

To complete the proof, it suffices to show that C is skew symmetric. This is true, because

$$\begin{aligned}
 (32) \qquad C &= (O_{12} \cdots O_{1p})(O_{23} \cdots O_{2p}) \cdots (O_{p-1,p}) \frac{\partial [O_{p-1,p}^t \cdots (O_{1p}^t \cdots O_{12}^t)]}{\partial o_{rs}} \\
 &= O_{12} \cdots O_{rs} \frac{\partial O_{rs}^t}{\partial o_{rs}} \cdots O_{12}^t
 \end{aligned}$$

and $O_{rs} \partial O_{rs}^t / \partial o_{rs}$ is skew symmetric. \square

APPENDIX B

PROOF OF THEOREM 1. The general algorithm for computing ordered group reference priors can be found in Berger and Bernardo (1992c). In the following, we will use the notations $\Theta^l, h_i, \Theta^l(\theta_{[i]})$ and π_1^l from Berger and Bernardo (1992c).

We will take the ordered group to be $\{d_1, \dots, d_p, (o_{12}, \dots, o_{p-1,p})\}$ as an example of the computation of the reference prior for Σ ; all the other ordered

groups give the same answer, providing the $\{d_i\}$ are listed before the $\{o_{ij}\}$ and the $\{d_i\}$ are ordered monotonically (either increasing or decreasing). Define the compact subsets of the parameter space to be

$$\Theta^l = \{(D, O): 0 < a_l \leq d_p \leq \dots \leq d_1 \leq b_l < \infty, -\pi/2 < o_{ij} \leq \pi/2, \forall i \leq j\},$$

where $a_l \rightarrow 0$ and $b_l \rightarrow \infty$. Using Lemmas 1 and 2, note that

$$h_i = \frac{1}{2d_i^2}, \quad i = 1, \dots, p,$$

$$|h_{p+1}| \propto \left[\prod_{i=1}^{p-1} \prod_{j=i+1}^p \cos^{j-i-1} o_{ij} \right]^2 \left[\prod_{i < j} (d_i - d_j) \right]^2 \prod_{i=1}^p d_i^{-(p-1)}.$$

Also,

$$\Theta^l(\theta_{[i-1]}) = \{d_i: a_l \leq d_i \leq d_{i-1}\}, \quad i = 1, \dots, p,$$

$$\Theta^l(\theta_{\{(p+1)-1\}}) = \{0_{ij}: -\pi/2 < o_{ij} \leq \pi/2, \forall i \leq j\},$$

where d_0 is interpreted as b_l . Thus (2.3.6) in Berger and Bernardo (1992c) becomes

$$\pi_1^l(\theta) = \left(\prod_{i=1}^p \frac{1/d_i}{\int_{a_l \leq d_i \leq d_{i-1}} 1/d_i dd_i} \right) \left(\frac{|h_{p+1}|^{1/2}}{\int_{-\pi/2 < o_{ij} \leq \pi/2} |h_{p+1}|^{1/2} dO} \right) 1_{\Theta^l(\theta)}$$

$$\propto \frac{\prod_{i=1}^p 1/d_i}{(\log b_l - \log a_l) \prod_{i=2}^p (\log d_{i-1} - \log a_l)} \left[\prod_{i=1}^{p-1} \prod_{j=i+1}^p \cos^{j-i-1} o_{ij} \right].$$

The reference prior w.r.t. this group ordering is thus given by

$$\pi(D, O) \propto \lim_{l \rightarrow \infty} \frac{\pi_1^l(\theta)}{\pi_1^l(\theta^*)}$$

$$\propto \frac{[\prod_{i=1}^{p-1} \prod_{j=i+1}^p \cos^{j-i-1} o_{ij}]}{|D|} \quad \square$$

Acknowledgments. The authors would like to thank Ming-Hui Chen for his generous help with the computer calculations. We would also like to thank L. R. Haff for providing Fortran subroutines for computation of his estimators, and the referees for helpful suggestions.

REFERENCES

ANDERSON, T. W. (1984). *Introduction to Multivariate Statistical Analysis*. Wiley, New York.
 ANDERSON, T. W., OLKIN, I. and UNDERHILL, L. G. (1987). Generation of random orthogonal matrices. *SIAM J. Sci. Statist. Comput.* **8** 625-629.
 BELISLE, C. J. P., ROMELJIN, H. E. and SMITH, R. L. (1993). Hit-and-run algorithms for generating multivariate distributions. *Math. Oper. Res.* **18** 255-266.

- BERGER, J. and BERNARDO, J. M. (1989). Estimating a product of means: Bayesian analysis with reference priors. *J. Amer. Statist. Assoc.* **84** 200–207.
- BERGER, J. and BERNARDO, J. M. (1992a). Reference priors in a variance components problem. In *Bayesian Analysis in Statistics and Econometrics* (P. Goel and N. S. Iyengar, eds.) 177–194. Springer, New York.
- BERGER, J. and BERNARDO, J. M. (1992b). Ordered group reference priors with application to a multinomial problem. *Biometrika* **79** 25–37.
- BERGER, J. and BERNARDO, J. M. (1992c). On the development of reference priors (with discussion). In *Bayesian Statistics 4* (J. M. Bernardo, J. O. Berger, A. P. Dawid and A. F. M. Smith, eds.) 35–60. Oxford Univ. Press.
- BERNARDO, J. M. (1979). Reference posterior distributions for Bayes inference (with discussion). *J. Roy. Statist. Soc. Ser. B* **41** 113–147.
- BERNARDO, J. M. (1981). Reference decisions. *Sympos. Math.* **25** 85–94.
- BERNARDO, J. M. and SMITH, A. F. M. (1994). *Bayesian Theory*. Wiley, New York.
- CHANG, T. and EAVES, D. (1990). Reference priors for the orbit in a group model. *Ann. Statist.* **18** 1595–1614.
- CHEN, M. H. and SCHMEISER, B. W. (1993). Performance of the Gibbs, hit-and-run, and Metropolis samplers. *Journal of Computational and Graphical Statistics* **2** 251–272.
- DEMPSTER, A. P. (1969). *Elements of Continuous Multivariate Analysis*. Addison-Wesley, Reading, MA.
- DEY, D. K. (1988). Simultaneous estimation of eigenvalues. *Ann. Inst. Statist. Math.* **40** 137–147.
- DEY, D. K. and SRINIVASAN, C. (1985). Estimation of a covariance matrix under Stein's loss. *Ann. Statist.* **13** 1581–1591.
- DEY, D. K. and SRINIVASAN, C. (1986). Trimmed minimax estimator of a covariance matrix. *Ann. Inst. Statist. Math.* **38** 101–108.
- EFRON, B. and MORRIS, C. (1976). Multivariate empirical Bayes estimation of covariance matrices. *Ann. Statist.* **4** 22–32.
- FARRELL, R. H. (1985). *Multivariate Calculation*. Springer, New York.
- GEISSER, S. (1965). Bayesian estimation in multivariate analysis. *Ann. Math. Statist.* **36** 150–159.
- GEISSER, S. and CORNFIELD, J. (1963). Posterior distributions for multivariate normal parameter. *J. Roy. Statist. Soc. Ser. B* **25** 368–376.
- HAFF, L. R. (1977). Minimax estimators for a multinormal precision matrix. *J. Multivariate Anal.* **7** 374–385.
- HAFF, L. R. (1979a). Estimation of the inverse covariance matrix: random mixtures of the inverse Wishart matrix and the identity. *Ann. Statist.* **7** 1264–1276.
- HAFF, L. R. (1979b). An identity for the Wishart distribution with application. *J. Multivariate Anal.* **9** 531–542.
- HAFF, L. R. (1980). Empirical Bayes estimation of the multivariate normal covariance matrix. *Ann. Statist.* **8** 586–597.
- HAFF, L. R. (1991). The variational form of certain Bayes estimators. *Ann. Statist.* **19** 1163–1190.
- JEFFREYS, H. (1961). *Theory of Probability*. Oxford Univ. Press.
- KRISHNAMOORTHY, K. (1991). Estimation of normal covariance and precision matrices with incomplete data. *Comm. Statist. Theory Methods* **20** 757–770.
- KRISHNAMOORTHY, K. and GUPTA, A. K. (1989). Improved minimax estimation of a normal precision matrix. *Canad. J. Statist.* **17** 91–102.
- LEONARD, T. and HSU, J. S. J. (1992). Bayesian inference for a covariance matrix. *Ann. Statist.* **20** 1669–1696.
- LIN, S. P. and PERLMAN, M. D. (1985). A Monte Carlo comparison of four estimators for a covariance matrix. In *Multivariate Analysis 6* (P. R. Krishnaiah, ed.) 411–429. North-Holland, Amsterdam.
- LOH, W. L. (1991a). Estimating covariance matrices I. *Ann. Statist.* **19** 283–296.
- LOH, W. L. (1991b). Estimating covariance matrices II. *J. Multivariate Anal.* **36** 163–174.
- MAGNUS, J. R. and NEUDECKER, H. (1988). *Matrix Differential Calculus with Applications in Statistics and Econometrics*. Wiley, New York.

- OLKIN, I. and SELLIAH, J. B. (1977). Estimating covariance in a multivariate normal distribution. In *Statistical Decision Theory and Related Topics 2* (S. S. Gupta and D. Moore, eds.) 313–326. Academic, New York.
- PERRON, F. (1992). Minimax estimators of a covariance matrix. *J. Multivariate Anal.* **43** 16–28.
- PRESS, S. J. (1982). *Applied Multivariate Analysis: Using Bayesian and Frequentist Measures of Inference*, 2nd ed. Krieger, New York.
- SCHMEISER, B. W. and CHEN, M. H. (1991). On hit-and-run Monte Carlo sampling for evaluating multidimensional integrals. Technical Report 91-39, Dept. Statistics, Purdue Univ.
- SHARMA, D. (1980). An estimator of normal covariance matrix. *Calcutta Statist. Assoc. Bull.* **29** 161–167.
- SHARMA, D. and KRISHNAMOORTHY, K. (1983). Orthogonal equivariant estimators of bivariate normal covariance matrix and precision matrix. *Calcutta Statist. Assoc. Bull.* **32** 23–45.
- SHARMA, D. and KRISHNAMOORTHY, K. (1985a). Empirical Bayes estimators of normal covariance matrix. *Sankhyā Ser. B* **47** 247–254.
- SHARMA, D. and KRISHNAMOORTHY, K. (1985b). Improved minimax estimators of normal covariance and precision matrices from incomplete samples. *Calcutta Statist. Assoc. Bull.* **34** 23–42.
- SINHA, B. K. and GHOSH, M. (1986). Inadmissibility of the best equivariant estimators of the variance-covariance matrix, the precision matrix, and the generalized variance under entropy loss. *Statist. Decisions* **5** 201–227.
- SMITH, R. L. (1980). A Monte Carlo procedure for generating random feasible solutions to mathematical programs. A Bulletin of the ORSA/TIMS Joint National Meeting, Washington, D.C., 101.
- SMITH, R. L. (1984). Efficient Monte Carlo procedure for generating points uniformly distributed over bounded regions. *Oper. Res.* **32** 1297–1308.
- STEIN, C. (1956). Some problems in multivariate analysis. Part I. Technical Report 6, Dept. Statistics, Stanford Univ.
- STEIN, C. (1975). Estimation of a covariance matrix. Rietz Lecture, 39th Annual Meeting IMS, Atlanta, Georgia.
- STEIN, C. (1977a). Estimating the covariance matrix. Unpublished manuscript.
- STEIN, C. (1977b). Lectures on the theory of estimation of many parameters. In *Studies in the Statistical Theory of Estimation, Part I* (I. A. Ibragimov and M. S. Nikulin, eds.). *Proceedings of Scientific Seminars of the Steklov Institute, Leningrad Division* **74** 4–65. (In Russian.)
- SUGIURA, N. and FUJIMOTO, M. (1982). Asymptotic risk comparison of improved estimators for normal covariance matrix. *Tsukuba J. Math.* **6** 103–126.
- TAKEMURA, A. (1984). An orthogonally univariant minimax estimator of the covariance matrix of a multivariate normal population. *Tsukuba J. Math.* **8** 367–376.
- TRACY, D. S. and JINADASA, K. G. (1988). Patterned matrix derivatives. *Canad. J. Statist.* **16** 411–418.
- VILLEGAS, C. (1969). On the a priori distribution of the covariance matrix. *Ann. Math. Statist.* **40** 1098–1099.
- YE, K. Y. and BERGER, J. (1991). Noninformative priors for inference in exponential regression models. *Biometrika* **78** 645–656.

DEPARTMENT OF STATISTICS
PURDUE UNIVERSITY
WEST LAFAYETTE, INDIANA 47907