

EMPIRICAL BAYES ESTIMATION OF PROPORTIONS WITH APPLICATION TO COWBIRD PARASITISM RATES¹

WILLIAM A. LINK AND D. CALDWELL HAHN

Patuxent Wildlife Research Center, Laurel, Maryland 20708 USA

Abstract. Bayesian models provide a structure for studying collections of parameters such as are considered in the investigation of communities, ecosystems, and landscapes. This structure allows for improved estimation of individual parameters by considering them in the context of a group of related parameters. Individual estimates are differentially adjusted toward an overall mean, with the magnitude of their adjustment based on their precision. Consequently, Bayesian estimation allows for a more reliable ranking of parameters and, in particular, a more credible identification of extreme values from a collection of estimates.

In Bayesian models, individual parameters are regarded as values sampled from a specified probability distribution, called a *prior*. The requirement that the prior be known is often regarded as an unattractive feature of Bayesian analysis and may be the reason Bayesian analyses are not frequently applied in ecological studies. *Empirical Bayes* methods provide an alternative approach that incorporates the structural advantages of Bayesian models while requiring a less stringent specification of prior knowledge. Empirical Bayes methods require only that the prior be in a certain family of distributions, indexed by hyperparameters that can be estimated from the available data. This structure is of interest per se, in addition to its value in allowing for improved estimation of individual parameters; for example, hypotheses regarding the existence of distinct subgroups in a collection of parameters can be considered under the empirical Bayes framework by allowing the hyperparameters to vary among subgroups.

We describe the empirical Bayes approach in application to estimation of proportions, using data obtained in a community-wide study of Brown-headed Cowbird parasitism rates for illustration. Empirical Bayes estimates identify those species for which there is the greatest evidence of extreme parasitism rates.

Subgroup analysis of our data on cowbird parasitism rates indicates that parasitism rates for neotropical migrants as a group are no greater than those of resident/short-distance migrant species in this forest community. Our data and analyses demonstrate that the parasitism rates for certain neotropical migrant species (Wood Thrush and Rose-breasted Grosbeak) are remarkably low while those for others (Ovenbird and Red-eyed Vireo) are remarkably high.

Key words: Bayes; Brown-headed Cowbirds; empirical Bayes; *Molothrus ater*; neotropical migrants; parasitism; Wood Thrush.

INTRODUCTION

Empirical Bayes methods have received considerable attention in the statistical literature since their introduction four decades ago (Robbins 1956). These procedures yield improved estimates of individual parameters by considering them in the context of a group of related parameters; these estimates have “borrowed strength from the ensemble” (Morris 1983). Empirical Bayes analysis also provides a more dependable ranking of parameters and aids in the identification of extreme values in the group.

Morris (1983) documents application of empirical Bayes methodology by the U.S. Census Bureau, in industry, and in epidemiology. More recently, empirical Bayes methods have been applied to geographic mapping problems in the social and medical sciences (sui-

cide rates, Lui et al. 1990; sudden infant death rates, Cressie 1992; rates for several causes of death in the Netherlands, Heisterkamp et al. 1993), to problems in forest science (Burk and Ek 1982, Green and Strawderman 1985), to monitoring of air pollution (Suggs and Curran 1983), and in a variety of medical applications (Stijnen and Van Houwelingen 1990, Comenges and Etcheverry 1993). Nevertheless, there have been only a few applications of empirical Bayes methodology in nonmedical biological settings. Johnson (1981), in a report of an application of empirical Bayes methods to the estimation of avian population sizes, noted that “despite the theoretical justification of empirical Bayes methods, their use has not been widespread,” a statement that remains largely true with regard to the ecological literature. Some recent exceptions are applications to estimating population sizes from survey data (Johnson 1989), to estimating numbers of species (Mingoti and Meeden 1992), to capture-

¹ Manuscript received 30 May 1995; revised 29 January 1996; accepted 2 February 1996.

recapture data (Smith 1991), to toxicity data (Piegorisch 1994), to summary analyses of avian trends (Peterjohn et al. 1994, Link and Sauer 1995), and to identification of extremes in collections of parameter estimates (Link and Sauer 1996). Given current interest in assessing the relative effects of environmental insults, ecological competitors, and disturbances on collections of species within a community or landscape, empirical Bayes methods seem well adapted to fill a niche in the toolbox of ecological methods.

In this paper, we apply empirical Bayes methods to the estimation of proportions. Many fundamental descriptions of biological systems involve proportions: survival rates, nest predation rates, breeding rates, and parasitism rates are a few examples. To illustrate the application of empirical Bayes modeling to proportions, we use data obtained in a community-wide study of parasitism by Brown-headed Cowbirds (*Molothrus ater*).

Conservation biologists have been concerned that cowbird parasitism is a major factor reducing breeding success of neotropical migrants in North America (Brittingham and Temple 1983, Robbins et al. 1989, Terborgh 1989, Robinson 1993). This concern has prompted a number of studies comparing the incidence and effect of brood parasitism on different host species (Thompson 1993, Rothstein and Robinson 1994, Cook et al. 1996). The data we examine consist of observed cowbird parasitism rates for 26 host species in a forest habitat (Hahn and Hatfield 1995, 1996). Host species considered include 18 neotropical migrant species and 8 resident/short-distance migrant species; we tested for differences in average parasitism rates between these two groups.

In the data that we consider, and typically in similar studies, there is considerable variation in the sample sizes of nests of host species. This creates difficulties in ranking parasitism rates. To give an extreme example, suppose that only one nest is observed for some potential host species. The observed parasitism rate will be either 0 or 100%; clearly, it would be inappropriate to say that such a species was the least or most severely affected. The empirical Bayes methods described in this paper account for differences in sample sizes among host species, providing an improved ranking of parasitism rates and allowing for a more appropriate identification of the host species for which there is the greatest evidence of extreme parasitism rates.

THE NEED FOR ENSEMBLE ESTIMATION

Suppose that we wish to create a list of host species, ranking them by their parasitism rates. For each of a number of species, data available are observed parasitism rates in random samplings of nests. In particular, three of five nests for Species A (60%), and 19 of 50 nests for Species B (38%), are parasitized.

We would place little confidence in the statement that the actual rate for Species A is higher than that

for Species B (in fact, Fisher's test for equality of proportions yields a P value of 0.32, from which the hypothesis of equal rates is not rejected). Yet Species A would be placed higher than Species B in the ordered list of parasitism rate estimates. Is this ordering reasonable?

The answer we give to this question will depend on our general knowledge of parasitism rates across species. If we knew that parasitism rates were typically close to, say, 55%, then the 38% observed for Species B might be taken as evidence of an exceptionally low rate for that individual species (P value = 0.0232 in testing H_0 : rate for Species B = 55%), while that for Species A would be taken as nothing extraordinary; it would be reasonable to assume that the actual rate for Species A is greater than that for Species B.

On the other hand, suppose that the typical parasitism rate across species were known to be 20%. Then we would have evidence that Species B had an unusually high rate (P value = 0.0050 in testing H_0 : rate for Species B = 20%), but we would not have evidence that Species A had an unusually high rate, despite the high observed rate of 60% (P value = 0.1024 in testing H_0 : rate for Species A = 20%). In this case we might suspect that the actual rate for Species B is higher than that for Species A.

Thus, our impression as to whether the ordering of estimates for Species A and B is correct depends on our prior knowledge of the rates of parasitism among host species in general. Even with a less specific knowledge of typical parasitism rates than we have supposed, our impression regarding the ranking of estimates for Species A and B may be based on more information than the observed rates for these two species alone. It may be based on observed rates for a collection of species, of which the data for Species A and B are a subset. Then, our thinking would be guided by an unstated hypothesis that there is a degree of similarity of rates within the group of host species. Thus, we examine the data available for each species in light of an informal model of similarity across species.

Clearly, there is a need to formalize the unstated hypotheses and informal models. Bayesian modeling and inference offers one such formalization. Bayesian methods are designed to incorporate prior knowledge of group structure in producing individual parameter estimates. A Bayesian analysis of cowbird parasitism rates for Species A and B would begin by specifying typical parasitism rates for a group of related species, summarized, perhaps, by their mean and variance. It then adjusts the estimates for Species A and B in light of this prior knowledge.

Many biologists would be unwilling to make a priori claims regarding the mean and variance of parameters for a collection of species under investigation, and to them the Bayesian approach is not appealing. Empirical Bayes methods can be thought of as a step in the direction of Bayesian analysis, but requiring a less strin-

gent specification of prior knowledge. To describe empirical Bayes methods, however, we must begin with the Bayesian paradigm.

A HEURISTIC EXPLANATION OF BAYES AND EMPIRICAL BAYES METHODS

We begin with a heuristic explanation of Bayes estimation. Suppose that before having collected the data for Species A, we were required to provide an educated guess of its parasitism rate. If we could assume that the rate for Species A is not too different from those in a collection of species for which data are available, our best guess would be the average parasitism rate for those host species; the standard deviation among those rates would provide a measure of our uncertainty in this educated guess. To give a concrete example, suppose that the average parasitism rate for the collection of host species is known to be $\pi = 20\%$, and that the standard deviation among these rates is $\tau = 10\%$. We treat 20% as an "estimate" of the rate for Species A, and 10% as the standard deviation of this "estimate."

Upon collecting data for Species A, we obtain a second estimate of the parasitism rate for Species A. Since five nests were encountered, of which three were parasitized, this second estimate is $\hat{p} = 60\%$ with standard error:

$$\sigma(\hat{p}) = \sqrt{\hat{p}(1 - \hat{p})/5} = 21.9\%;$$

thus, the data tell us that the rate for Species A is likely to be near 60%, give or take 21.9%. We now have two estimates of the same quantity, one based on prior knowledge (π is referred to as the *prior mean*) and the other based on new observations. The usual expedient when we have two estimates of the same quantity is to combine estimates using a precision-weighted average. Under certain conditions (see *The beta-binomial model*) the Bayes estimate for Species A is obtained in just this way. The Bayes estimate for Species A is:

$$\begin{aligned}\hat{p}_{A,\text{Bayes}} &= w\pi + (1 - w)\hat{p}_A \\ &= w(20\%) + (1 - w)(60\%),\end{aligned}$$

where

$$w = \frac{\frac{1}{(10\%)^2}}{\frac{1}{(10\%)^2} + \frac{1}{(21.9\%)^2}} = 0.83$$

is the weight placed on the prior mean. The Bayes estimate for Species A is therefore 26.9%. We note that the rate for Species A was very heavily adjusted, moving 83% of the way from the raw estimate (60%) to the overall mean (20%).

Carrying out similar calculations for Species B (for which 19 of 50 nests were found to be parasitized, so that $\hat{p} = 38\%$, $\sigma(\hat{p}) = 6.9\%$, and $w = 0.32$), the Bayes estimate for Species B is:

$$\hat{p}_{B,\text{Bayes}} = 0.32(20\%) + 0.68(38\%) = 32.2\%.$$

The rate for Species B was less severely adjusted than that for Species A, moving only 32% of the way from the raw estimate (38%) to the overall mean (20%). The estimate for Species A (based on a small sample size, $N = 5$) was heavily adjusted, while the more reliable estimate for Species B (based on $N = 50$) received a more moderate adjustment.

A Bayesian analysis can be thought of as a combining of existing knowledge with new knowledge (sample data), the two being combined in such a way as to account for the amount of confidence that is placed in each source of knowledge. The existing knowledge base and updated knowledge base are summarized by probability distributions describing the likely range of values for each unknown parameter; these are referred to as the *prior* and *posterior distributions*, respectively. Informally, the prior distribution can be thought of as an approximation to the histogram of the true, unknown values of the parameters under investigation. There is a posterior distribution for each of the unknown parameters, and it is to this that we would turn, as informed gamblers, for setting odds in betting on the value of the parameter given our updated knowledge base.

Two other probability distributions encountered in describing Bayesian modeling are the *sampling* and *marginal distributions* of the data used to update the knowledge base. The sampling distribution describes the variability of an observed statistic about an unobserved parameter value. Thus, the sampling distribution describes a second source of variation, the first source being the variability among species as described by the prior distribution. The marginal distribution describes the totality of the variation in the data, incorporating the variability described by the prior and sampling distributions. For a detailed accounting of the Bayesian approach to statistical modeling, the reader is referred to the text by Box and Tiao (1992).

The empirical Bayes methods we describe use the same models as Bayesian methods, and similar analytical procedures. Raw estimates are adjusted toward an overall mean by a factor based on the precision of the raw estimate relative to the variability among parameters in the group. The difference between and Bayes and empirical Bayes analyses is that in the latter the overall mean and the variability among parameters in the group are now *estimated* from the data, rather than postulated on the basis of prior knowledge. Rather than requiring a complete specification of the prior distribution, we require only a partial specification of the prior distribution; namely, that it be a member of a known family of distributions. We might not know specifically what the prior distribution looks like, but we are willing to make a general statement about its form.

Thus the prior distribution is assumed to be a member of a specified family of distributions, but characterized by unknown parameters. (We might, for example, specify that the prior is a normal distribution,

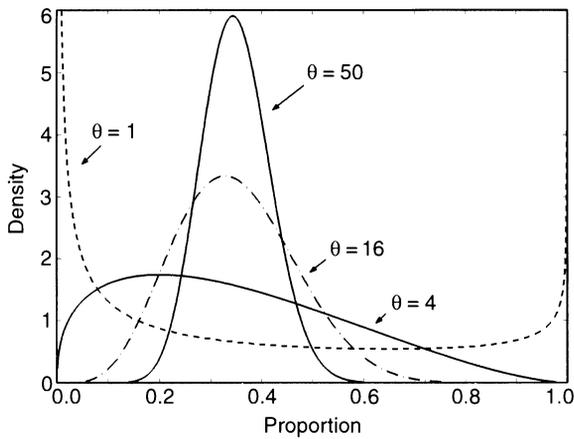


FIG. 1. Four beta densities, with common mean $\pi = 0.35$, and spread parameter $\theta = 1, 4, 16$, and 50 . The larger the value of θ , the greater the variability in the distribution. The numbers on the y-axis scale (density) are unitless.

with unknown mean and variance.) To distinguish the parameters of the prior from the parameters of primary interest (which are a sampling from the prior), we refer to the parameters of the prior as *hyperparameters*. In our example, cowbird parasitism rates are the parameters, and hyperparameters describe the distribution of these rates across species. The first stage of an empirical Bayes analysis is to estimate the unknown hyperparameters and thereby to obtain an estimate of the prior distribution.

THE BETA-BINOMIAL MODEL

We now turn to an explicit description of a model that is useful for empirical Bayes estimation of proportions. We begin by assuming that the prior distribution is in the class of *beta distributions*, i.e., those having density functions of the form:

$$f(p) = \frac{\Gamma(\alpha + \beta)}{\Gamma(\alpha)\Gamma(\beta)} p^{\alpha-1}(1 - p)^{\beta-1}, \quad 0 < p < 1,$$

where α and β are positive numbers and $\Gamma(\cdot)$ is the gamma function. A Bayesian analysis requires a complete specification of the prior, in which the parameters α and β are assumed to be known; in the empirical Bayes framework, α and β are unknowns to be estimated.

For the present application, it is useful to re-parameterize the prior distributions by setting $\pi = \alpha/(\alpha + \beta)$ and $\theta = (\alpha + \beta)$. Then, if P is a random variable to be sampled from a beta distribution, its mean is $E(P) = \pi$ and its variance is $Var(P) = \pi(1 - \pi)/(1 + \theta)$. Expressed in terms of our cowbird parasitism data, the hyperparameters π and θ are descriptive of the collection of true parasitism rates for host species. Roughly speaking, the parameter π is the average parasitism rate across species and the parameter θ controls the range of values of parasitism rates among species. The pa-

rameter θ , which we refer to as the *spread parameter*, is more useful for comparison of distributions than is the variance, since the variance depends on the mean.

The choice of a beta prior is made primarily for reasons of technical convenience, but should not be regarded as an overly restrictive assumption; the beta family of distributions is very flexible. Figs. 1 and 2 illustrate beta distributions with various values of π and θ .

For an individual host species, we let X denote the number of parasitized nests in a random sample of N nests. Given that the true parasitism rate is P , X is assumed to be a binomial random variable with mean $N \cdot P$ (the observed parasitism rate is $\hat{p} = X/N$). Taken together, the assumptions of a beta prior and binomial sampling distribution comprise the *beta-binomial model*.

Under the beta-binomial model, the number of parasitized nests for a randomly selected species has distribution function given by:

$$p(x; N) = \binom{N}{x} \frac{\Gamma(\theta)\Gamma(x + \pi\theta)\Gamma(N - x + (1 - \pi)\theta)}{\Gamma(\pi\theta)\Gamma((1 - \pi)\theta)\Gamma(N + \theta)},$$

$$x = 0, 1, \dots, N;$$

this is the marginal distribution. In an empirical Bayes analysis, the marginal distributions are used to estimate the hyperparameters. Given data for n species, (X_i, N_i) , $i = 1, 2, \dots, n$, maximum likelihood estimates of the hyperparameters π and θ are obtained by maximizing the loglikelihood:

$$l(\theta, \pi) = \sum_i \log(p(x_i; N_i))$$

as a function of π and θ . Maximum likelihood estimates are not available in closed form; they must be obtained using numerical optimization methods such as the

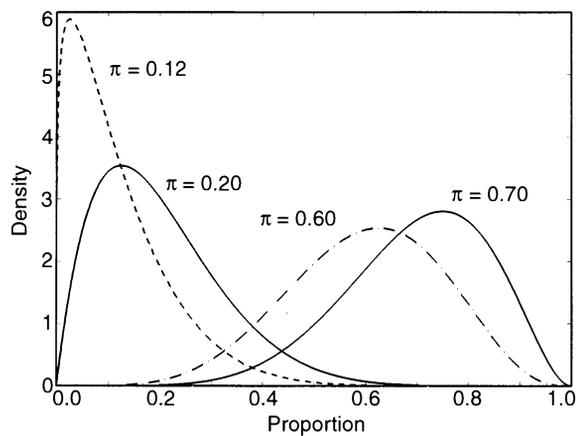


FIG. 2. Four beta densities, with common spread parameter $\theta = 10$, and means $\pi = 0.12, 0.20, 0.60$, and 0.70 . Range of possible values of spread parameter, unlike that of variance, is not restricted by the values of the mean and hence is more useful in characterizing dispersion of beta distributions. The numbers on the y-axis scale (density) are unitless.

Newton-Raphson method. As a byproduct of Newton-Raphson optimization of the loglikelihood, one can obtain information matrix-based estimates of variances and covariances of the estimated hyperparameters (for details, see Burnham et al. 1987).

Alternatively, π and θ can be estimated by moment estimators as described in Appendix 1. These estimators are slightly less efficient than the maximum likelihood estimators, but they are much more easily calculated.

Given the beta-binomial model, we may address the question of how new knowledge (X and N) can be used to update our knowledge base (the prior distribution, $f(p)$). The Bayesian response is given by the posterior distribution of P , given \hat{p} . Given the model structure specified, the posterior distribution of P is also of the beta family of distributions, with the parameters α and β replaced by $\alpha + X$ and $\beta + N - X$, respectively. Expressed in terms of the (π, θ) parameterization, the posterior distribution of P given \hat{p} is therefore:

$$f_{\pi,\theta}(P|\hat{p}) = \frac{\Gamma(\theta + N)}{\Gamma(\pi\theta + X)\Gamma((1 - \pi)\theta + N - X)} \times P^{(\pi\theta + X - 1)}(1 - P)^{((1 - \pi)\theta + N - X - 1)}$$

The mean of the posterior distribution (also described as the expected value of P given \hat{p}) is:

$$E(P|\hat{p}) = \frac{\alpha + X}{\alpha + \beta + N} = \pi \left(\frac{\theta}{\theta + N} \right) + \hat{p} \left(\frac{N}{\theta + N} \right)$$

This quantity is the Bayes estimator of P . It is a weighted average of the raw estimate \hat{p} and the overall mean π ; the weights on π and \hat{p} are inversely proportional to $\text{Var}(P)$ and the sampling variance of \hat{p} .

To obtain empirical Bayes estimates, we need only substitute estimates of the hyperparameters π and θ in the formula for the Bayes estimator. We thus define an empirical Bayes estimate of a proportion P by:

$$\hat{P}_{\text{EBayes}} = \hat{\pi} \left(\frac{\hat{\theta}}{\hat{\theta} + N} \right) + \hat{p} \left(\frac{N}{\hat{\theta} + N} \right)$$

the circumflexes “^” on π and θ indicating that these quantities are estimates.

We note also that the estimated hyperparameters can be used to estimate the posterior distribution for P , from which we can make probability statements regarding the likely range of values for P . Substituting the hyperparameter estimates for the true values in the formula for the posterior distributions, we obtain the estimated posterior distribution $f_{\pi,\theta}(P|\hat{p})$. We describe the use of the estimated posterior distribution in our application of these methods to cowbird parasitism data.

APPLICATION OF EMPIRICAL BAYES ANALYSIS TO COWBIRD PARASITISM DATA

Table 1 provides sample sizes (N) and numbers parasitized (X) for a sampling of nests of 26 host species in an eastern forest community. The collection of these data is described elsewhere (Hahn and Hatfield 1995, 1996); here we note only that the data represent totals over a 3-yr study and were pooled after finding no significant year effects in a loglinear analysis.

Our empirical Bayes analysis treats X_i (the data for species i) as a binomial random variable with parameter P_i , with P_i sampled from a beta distribution. The binomial assumption requires independence of sampled nests: the parasitism status of a given nest is not related to the parasitism status of other nests. Violations of this assumption could result from pooling data across sites or date of observation. The assumption that the P_i are a sampling from a beta distribution is not terribly restrictive, but could be violated by the existence of distinct subgroups of species. We address the evaluation of this assumption in the next section.

The moment (Mom) estimators of the hyperparameters are not substantially different from the maximum likelihood estimators. We have:

$$\begin{aligned} \hat{\theta}_{\text{Mom}} &= 5.898 & \hat{\pi}_{\text{Mom}} &= 0.297 \\ \hat{\theta}_{\text{MLE}} &= 6.959(3.183) & \hat{\pi}_{\text{MLE}} &= 0.293(0.043); \end{aligned}$$

the bracketed quantities are information matrix-based estimates of the standard errors of the maximum likelihood estimators (MLEs). The choice of estimators had little effect in creating the empirical Bayes estimates for the 26 species: the absolute difference between empirical Bayes estimates calculated with the two sets of estimated hyperparameters averaged 0.7% and never exceeded 1.7%.

Fig. 3 gives plots of the estimated posterior distributions for four species evidencing distinct levels of parasitism and of various sample sizes: Wood Thrush, Worm-eating Warbler, Hooded Warbler, and Red-eyed Vireo. We note that the posterior distribution for Wood Thrush is considerably less dispersed than that of the other species (Fig. 3). This is a consequence of the larger sample of Wood Thrush nests ($N = 105$); the posterior distribution for Hooded Warbler ($N = 1$) is only slightly different from the estimated prior.

The estimated posterior distributions can be used to create intervals similar to confidence intervals for the parasitism rates (Table 1). For example, the central 95% of the estimated posterior distribution describing rates for Worm-eating Warbler rates covers the range (0.1085, 0.4560). Thus, supposing that the prior distribution for parasitism rates (across species) is beta with parameters $\pi = 0.297$ and $\theta = 5.898$; given the observed data, we would place 95% confidence on the statement that the rate for Worm-eating Warbler is in the range (0.1085, 0.4560). This interval, sometimes called a “95% credibility interval” so as to distinguish

TABLE 1. Cowbird parasitism data for 26 host species in a forest community.

Host	X	N	\hat{p}	\hat{p}_{EB}	Bayesian 95% C.I.	Rank ₁	Rank ₂
Wood Thrush (<i>Hylocichla mustelina</i>)	10	105	0.095	0.108	(0.057, 0.171)	8	1
Rose-breasted Grosbeak (<i>Pheucticus ludovicianus</i>)	1	20	0.050	0.113	(0.025, 0.253)	7	2
Blue-gray Gnatcatcher (<i>Poliopitila caerulea</i>)	0	7	0.000	0.146	(0.020, 0.365)	3.5	3
Eastern Phoebe (<i>Sayornis phoebe</i>)	11	78	0.141	0.153	(0.085, 0.234)	11	4
Northern Cardinal (<i>Cardinalis cardinalis</i>)	1	10	0.100	0.179	(0.042, 0.387)	9	5
Rufous-sided Towhee (<i>Pipilo erythrophthalmus</i>)	2	15	0.133	0.184	(0.055, 0.366)	10	6
Eastern Wood-pewee (<i>Contopus virens</i>)	0	4	0.000	0.186	(0.027, 0.451)	3.5	7
Great-crested Flycatcher (<i>Myiarchus crinitus</i>)	0	2	0.000	0.227	(0.034, 0.533)	3.5	8.5
Warbling Vireo (<i>Vireo gilvus</i>)	0	2	0.000	0.227	(0.034, 0.533)	3.5	8.5
Veery (<i>Catharus fuscescens</i>)	18	74	0.243	0.247	(0.160, 0.346)	12	10
Black-capped Chickadee (<i>Parus atricapillus</i>)	0	1	0.000	0.256	(0.039, 0.586)	3.5	11.5
Yellow-rumped Warbler (<i>Dendroica coronata</i>)	0	1	0.000	0.256	(0.039, 0.586)	3.5	11.5
Worm-eating Warbler (<i>Helmitheros vermivorus</i>)	4	16	0.250	0.263	(0.109, 0.456)	13.5	13
Carolina Wren (<i>Thryothorus ludovicianus</i>)	1	4	0.250	0.277	(0.069, 0.561)	13.5	14
Chestnut-sided Warbler (<i>Dendroica pensylvanica</i>)	1	3	0.333	0.305	(0.077, 0.606)	16.5	15.5
Solitary Vireo (<i>Vireo solitarius</i>)	1	3	0.333	0.305	(0.077, 0.606)	16.5	15.5
American Redstart (<i>Setophaga ruticilla</i>)	26	81	0.321	0.319	(0.226, 0.419)	15	17
Hooded Warbler (<i>Wilsonia citrina</i>)	1	1	1.000	0.382	(0.102, 0.716)	25.5	18.5
Yellow-throated Vireo (<i>Vireo flavifrons</i>)	1	1	1.000	0.382	(0.102, 0.716)	25.5	18.5
Louisiana Waterthrush (<i>Seiurus motacilla</i>)	6	14	0.429	0.383	(0.193, 0.595)	18	20
Least Flycatcher (<i>Empidonax minimus</i>)	3	6	0.500	0.389	(0.154, 0.655)	21.0	21
Hermit Thrush (<i>Catharus guttatus</i>)	8	18	0.444	0.402	(0.223, 0.596)	19	22
Scarlet Tanager (<i>Piranga olivacea</i>)	4	8	0.500	0.404	(0.179, 0.652)	21.0	23
Black-and-White Warbler (<i>Mniotilta varia</i>)	5	10	0.500	0.415	(0.200, 0.649)	21.0	24
Ovenbird (<i>Seiurus aurocapillus</i>)	21	37	0.568	0.524	(0.378, 0.668)	23	25
Red-eyed vireo (<i>Vireo olivaceus</i>)	16	24	0.667	0.583	(0.408, 0.747)	24	26

Note: X = number parasitized nests; N = sample size; $\hat{p} = X/N$; \hat{p}_{EB} = empirical Bayes estimate, using maximum likelihood estimators of hyperparameters. Bayesian 95% credibility intervals (C.I.) are lower and upper endpoint of central 95% interval in estimated posterior distribution. Rank₁ = rank based on raw estimates \hat{p} ; Rank₂ = rank based on empirical Bayes estimates \hat{p}_{EB} . (Tied ranks replaced by averages.)

it from confidence intervals, is shorter than the usual approximate confidence interval:

$$\hat{p} \pm 1.96 \sigma(\hat{p})$$

$$(0.0378, 0.4620).$$

It is also shorter than the exact confidence interval (0.0727, 0.5238) obtained by inverting the binomial distribution. The shorter range of values expresses the increased knowledge brought to bear on parameter es-

timation by considering the parameter in the context of the group.

It may be desirable to account for the uncertainty in the estimation of the posterior distribution, perhaps by a lengthening of the credibility intervals. Two computationally intensive procedures for doing so are described in Appendix 2. For the present example, and in general when the prior parameters are reasonably well estimated, the effects of such procedures are slight.

TESTING FOR SUBGROUPS

The assumption of a common prior distribution for all of the proportions being investigated may be untenable. There may be reason to believe that proportions within distinct subgroups are more similar to each other than to the collection as a whole. For example, the host species in the cowbird data set under consideration can be classified as to migratory status (neotropical migrants vs. residents/short-distance migrants). In such a case, we might consider allowing distinct prior distributions for distinct subgroups. The prior distributions might be completely unrelated, or partially related, in the sense of having some common values of hyperparameters.

To illustrate the use of distinct priors for distinct subgroups, we consider models relating the hyperparameters π and θ to the migratory status of the species. There are four possible models describing the two distinct priors, each of which is described by a vector $(\theta_1,$

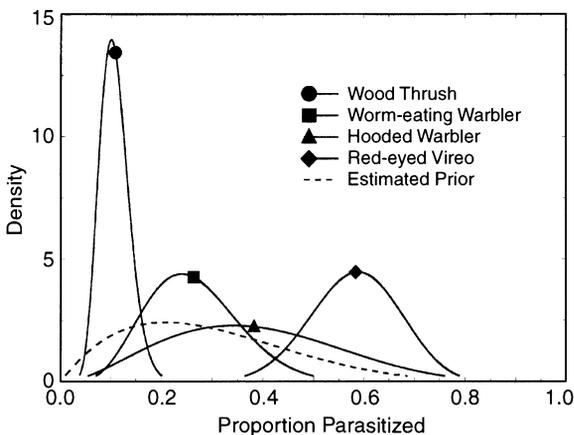


FIG. 3. Distributions of cowbird parasitism rates. Estimated prior distribution and corresponding estimates of posterior distribution rates for four host species. Symbols are aligned with posterior means.

TABLE 2. Maximum likelihood comparison of subgroup models. Values in parentheses are information matrix-based estimates of standard errors of parameter estimates.

Parameter estimates (estimated standard errors)					
Model	θ_1	θ_2	π_1	π_2	MLL
1. ($\theta_1, \theta_2, \pi_1, \pi_2$)	5.63 (3.00)	28.34 (42.02)	0.324 (0.056)	0.205 (0.054)	-52.0262
2. ($\theta, \theta, \pi_1, \pi_2$)	7.59 (3.56)	7.59 (3.56)	0.319 (0.051)	0.227 (0.071)	-52.7823
3. ($\theta_1, \theta_2, \pi, \pi$)	5.80 (3.01)	13.19 (14.21)	0.281 (0.046)	0.281 (0.046)	-53.0435
4. (θ, θ, π, π)	6.96 (3.18)	6.96 (3.18)	0.293 (0.043)	0.293 (0.043)	-53.3048

Summary of likelihood ratio tests				
Models compared	H_0	χ^2	df	P value
2 vs. 1	$\theta_1 = \theta_2$	1.5122	1	0.22
3 vs. 1	$\pi_1 = \pi_2$	2.0346	1	0.15
4 vs. 2	$\pi_2 = \pi_2$, given $\theta_1 = \theta_2$	1.0450	1	0.31
4 vs. 3	$\theta_1 = \theta_2$, given $\pi_1 = \pi_2$	0.5226	1	0.47
4 vs. 1	$\theta_1 = \theta_2$ and $\pi_1 = \pi_2$	2.5572	2	0.28

Note: Subscripts "1" on parameters correspond to priors for neotropical migrants; subscripts "2" correspond to resident/short-distant migrants. MLL = maximum log-likelihood.

θ_2, π_1, π_2). The model ($\theta, \theta, \pi_1, \pi_2$), for example, requires that the two prior distributions have the same relative spread, but allows the means to differ.

Maximum likelihood estimation (based on the marginal distributions of the observed numbers of nests parasitized) can be used to obtain parameter estimates, and the models compared by likelihood ratio tests. Parameter estimates and likelihood ratio test results for the parasitism rates we have investigated here are given in Table 2.

The parasitism rates for the subset of neotropical migrants appear to be slightly higher and less spread than those for short-distance migrant/resident species, but the differences between groups were not statistically significant (P value = 0.28; Table 2). We thus conclude that the best estimates of parasitism rates for these species are those that have been adjusted assuming a common prior distribution. For simplicity, we refer to these estimates as \hat{p}_{EB4} , with reference to model 4 of Table 2.

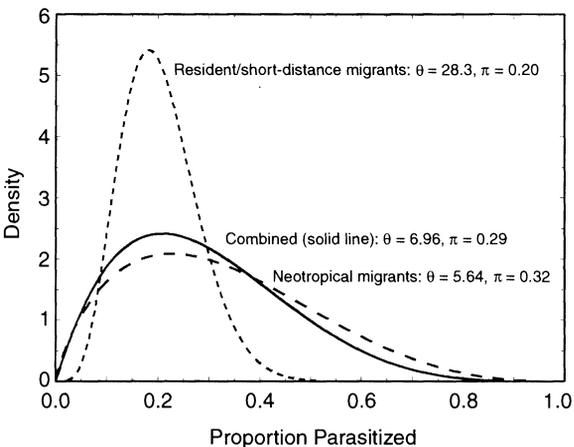


FIG. 4. Estimates of prior distributions in subgroup analysis. Estimated prior for resident and short-distance migrants is more peaked and has lower mean than that for neotropical migrant species, but not significantly different.

It is instructive to compare these empirical Bayes estimates to those obtained assuming that distinct subgroups exist. We denote these latter as \hat{p}_{EB1} , with reference to model 1 of Table 2. The estimated prior distributions for the two subgroups can be compared by inspection of Fig. 4; Table 3 displays both sets of empirical Bayes estimates. As would be anticipated, there are generally only minor differences between the two sets of empirical Bayes estimates, and these are most notable for the species with the smallest sample sizes. The empirical Bayes procedure heavily adjusts estimates based on small sample sizes toward their group means. Thus, the difference between the two sets of estimates is greatest for those species that are most heavily adjusted (i.e., those with small N).

One species for which \hat{p}_{EB1} differs substantially from \hat{p}_{EB4} is the Hermit Thrush. Considered among the entire group of host species, our best estimate of the rate for Hermit Thrush was 40%, while considered among the resident/short-distance migrants our best estimate of the rate was 30%. This difference is simply explained by noting that the raw estimate for Hermit Thrush is exceptionally high relative to the prior distribution for short-distance migrants/residents; it is high relative to the combined group as well, but not as exceptionally so. Thus, the rate is more heavily adjusted in the former context than in the latter. This underscores the point that empirical Bayes estimates are context specific; they are based on our understanding of an individual species as a member of a group. Since the likelihood ratio test did not reject the model (θ, θ, π, π) in favor of the model ($\theta_1, \theta_2, \pi_1, \pi_2$), we are inclined to place more credence in the estimate for Hermit Thrush in the context of the entire group, i.e., the rate of 40%.

DISCUSSION

Empirical Bayes analyses incorporate the structural advantages of Bayesian models without requiring the complete specification of a prior distribution. Like Bayesian analyses, empirical Bayes analyses allow for improved

TABLE 3. Comparison of estimates of parasitism rates.

Host	X	N	\hat{p}	\hat{p}_{EB4}	\hat{p}_{EB1}
A) Neotropical migrants					
Wood Thrush (<i>Hylocichla mustelina</i>)	10	105	0.095	0.108	0.107
Rose-breasted Grosbeak (<i>Pheucticus ludovicianus</i>)	1	20	0.050	0.113	0.110
Blue-Gray Gnatcatcher (<i>Poliophtila caerulea</i>)	0	7	0.000	0.146	0.145
Eastern Wood-pewee (<i>Contopus virens</i>)	0	4	0.000	0.186	0.190
Great-crested Flycatcher (<i>Myiarchus crinitus</i>)	0	2	0.000	0.227	0.239
Warbling Vireo (<i>Vireo gilvus</i>)	0	2	0.000	0.227	0.239
Veery (<i>Catharus fuscescens</i>)	18	74	0.243	0.247	0.249
Worm-eating Warbler (<i>Helmitheros vermivorus</i>)	4	16	0.250	0.263	0.269
American Redstart (<i>Setophaga ruticilla</i>)	26	81	0.321	0.319	0.321
Chestnut-sided Warbler (<i>Dendroica pensylvanica</i>)	1	3	0.333	0.305	0.327
Louisiana Waterthrush (<i>Seiurus motacilla</i>)	6	14	0.429	0.383	0.399
Least Flycatcher (<i>Empidonax minimus</i>)	3	6	0.500	0.389	0.415
Yellow-throated Vireo (<i>Vireo flavifrons</i>)	1	1	1.000	0.382	0.426
Hooded Warbler (<i>Wilsonia citrina</i>)	1	1	1.000	0.382	0.426
Scarlet Tanager (<i>Piranga olivacea</i>)	4	8	0.500	0.404	0.427
Black-and-White Warbler (<i>Mniotilta varia</i>)	5	10	0.500	0.415	0.437
Ovenbird (<i>Seiurus aurocapillus</i>)	21	37	0.568	0.524	0.535
Red-eyed Vireo (<i>Vireo olivaceus</i>)	16	24	0.667	0.583	0.601
B) Residents/short-distance migrants					
Eastern Phoebe (<i>Sayornis phoebe</i>)	11	78	0.141	0.153	0.158
Northern Cardinal (<i>Cardinalis cardinalis</i>)	1	10	0.100	0.179	0.177
Rufous-sided Towhee (<i>Pipilo erythrophthalmus</i>)	2	15	0.133	0.184	0.180
Yellow-rumped Warbler (<i>Dendroica coronata</i>)	0	1	0.000	0.256	0.198
Black-capped Chickadee (<i>Parus atricapillus</i>)	0	1	0.000	0.256	0.198
Carolina Wren (<i>Thryothorus ludovicianus</i>)	1	4	0.250	0.277	0.210
Solitary Vireo (<i>Vireo solitarius</i>)	1	3	0.333	0.305	0.217
Hermit Thrush (<i>Catharus guttatus</i>)	8	18	0.444	0.402	0.298

Note: X = number nests parasitized, N = number of nests sampled, $\hat{p} = X/N$, \hat{p}_{EB4} = empirical Bayes estimates under model 4 of Table 2 (i.e., assuming common priors across subgroups), \hat{p}_{EB1} = empirical Bayes estimates under model 1 of Table 2 (i.e., allowing prior distributions to differ between subgroups).

estimation of individual parameters by considering them in the context of a group of related parameters; similarly, empirical Bayes analyses allow for a more credible identification of extremes in the group of parameters.

We have described the empirical Bayes analysis of proportions assuming an underlying beta-binomial model. Our empirical Bayes analysis of cowbird parasitism rates leads us to identify Wood Thrush as the least frequently parasitized species in the forest community we studied. Seven other species had lower observed parasitism rates, but these rates were based on small sample sizes (five of the seven species had sample sizes <5). Our finding of a low parasitism rate for Wood Thrush in this community provides an interesting contrast to Robinson's (1992) results, which indicated very high parasitism rates on Wood Thrush in Illinois, leading him to speculate that this neotropical migrant was targeted by cowbirds or was in some sense more vulnerable to parasitism.

The ranking of other species in the community on the basis of empirical Bayes estimates indicates that the four most frequently parasitized species were Red-eyed Vireo ($\hat{p} = 67\%$, $N = 24$; $\hat{p}_{EB} = 58\%$), Ovenbird ($\hat{p} = 57\%$, $N = 57$, $\hat{p}_{EB} = 52\%$), Black-and-White Warbler ($\hat{p} = 50\%$, $N = 10$, $\hat{p}_{EB} = 42\%$), and Scarlet Tanager ($\hat{p} = 50\%$, $N = 8$, $\hat{p}_{EB} = 40\%$), and in this agrees with the ranking based on raw parasitism rates. It is interesting to note that in the ranking of species based on raw estimates, Black-and-White Warbler and Scarlet Tanager are

tied, having the same observed rates. By contrast, the empirical Bayes estimates rank Black-and-White Warbler higher because a larger sample size provides stronger evidence of a distinction from the overall group mean.

Empirical Bayes estimates discount the extreme values of the two species with $N = 1$ and 100% observed parasitism rates (Hooded Warbler and Yellow-throated Vireo) and moved them toward the middle, lowering them in rank from 25.5 to the rank of 18.5; the two species with zero parasitism and $N = 1$ (Black-capped Chickadee and Yellow-rumped Warbler) were similarly moved from ranks of 3.5 to ranks of 11.5.

We believe that the empirical Bayes approach will be valuable in community studies using proportions like parasitism, predation, or survival rates in multi-species rankings or in comparisons of subgroups. Empirical Bayes analysis will also be a useful tool in analyses at the landscape scale that compare parasitism or predation rates, for example, in different habitat types (Martin 1992), different vegetation types, or at different distances to ecological features like powerlines, roads, clearcuts, or other edges (Brittingham and Temple 1983).

ACKNOWLEDGMENTS

The authors thank the associate editor and two referees for helpful suggestions made during the review of this paper. They also thank J. S. Hatfield, J. D. Nichols, and J. R. Sauer for their reviews of an early version of this paper. Special

thanks are extended to Joe R. Hill for his insightful comments and patient assistance with technical details.

LITERATURE CITED

- Berger, J. O. 1985. Statistical decision theory and Bayesian analysis. Second edition. Springer-Verlag, New York, New York, USA.
- Box, G. E. P., and G. C. Tiao. 1992. Bayesian inference in statistical science. John Wiley and Sons, New York, New York, USA.
- Brittingham, M. C., and S. A. Temple. 1983. Have cowbirds caused forest songbirds to decline? *BioScience* **33**:31–35.
- Burk, T. E., and A. R. Ek. 1982. Application of empirical Bayes James Stein procedures to simultaneous estimation problems in forest inventory. *Forest Science* **28**:753–771.
- Burnham, K. P., D. R. Anderson, G. C. White, C. Brownie, and K. H. Pollock. 1987. Design and analysis methods for fish survival experiments based on release–recapture. American Fisheries Society Monographs, Monograph 5. Bethesda, Maryland, USA.
- Commenges, D., and B. Etcheverry. 1993. An empirical Bayes approach to the estimation of the incidence curve of HIV infection. *Statistics in Medicine* **12**:1317–1324.
- Cook, T. L., S. Robinson, S. Rothstein, J. N. M. Smith, and S. Sealy. 1996. The ecology and management of cowbirds. University of Texas Press, Austin, Texas, USA.
- Cressie, N. 1992. Smoothing regional maps using empirical Bayes predictors. *Geographical Analysis* **24**:75–95.
- Green, E. J., and W. E. Strawderman. 1985. The use of Bayes-empirical Bayes in individual tree volume development. *Forest Science* **31**:975–990.
- Hahn, D. C., and J. S. Hatfield. 1995. Parasitism at the landscape scale: cowbirds prefer forest. *Conservation Biology* **9**:1415–1424.
- Hahn, D. C., and J. Hatfield. 1996. Host selection in the forest interior: cowbirds target ground-nesting species. In T. L. Cook, S. Robinson, S. Rothstein, S. Sealy, and J. N. M. Smith, editors. The ecology and management of cowbirds. University of Texas Press, Austin, Texas, USA, *in press*.
- Heisterkamp, S. H., G. Doornbos, and M. Gankema. 1993. Disease mapping using empirical Bayes and Bayes methods on mortality statistics in the Netherlands. *Statistics in Medicine* **12**:1895–1913.
- Johnson, D. H. 1981. Improved population estimates through the use of auxiliary information. Pages 436–440 in C. J. Ralph and J. M. Scott, editors. Estimating numbers of terrestrial birds. Studies in avian biology Number 6. Cooper Ornithological Society, Lawrence, Kansas, USA.
- . 1989. An empirical Bayes approach to analyzing recurring animal surveys. *Ecology* **70**:945–952.
- Laird, N. M., and T. A. Louis. 1987. Empirical Bayes confidence intervals based on bootstrap samples. *Journal of the American Statistical Association* **82**:739–757.
- Link, W. A., and J. R. Sauer. 1995. Estimation and confidence intervals for empirical mixing distributions. *Biometrics* **51**:799–810.
- Link, W. A., and J. R. Sauer. 1996. Extremes in ecology: on avoiding the misleading effects of sampling variation in summary analyses. *Ecology* **77**:1633–1640.
- Louis, T. A. 1991. Using empirical Bayes methods in biopharmaceutical research. *Statistics in Medicine* **10**:811–829.
- Lui, K.-J., B. Martinez, and J. Mercy. 1990. An application of the empirical Bayes approach to directly adjusted rates: a note on suicide mapping in California, USA. *Suicide and Life-Threatening Behavior* **20**:240–253.
- Mingoti, S. A., and G. Meeden. 1992. Estimating the total number of distinct species using presence and absence data. *Biometrics* **48**:863–875.
- Morris, C. N. 1983. Parametric empirical Bayes inference: theory and applications (with comments). *Journal of the American Statistical Association* **78**:47–65.
- Peterjohn, B. G., J. R. Sauer, and W. A. Link. 1994. The 1992–1993 summary of the North American Breeding Bird Survey. *Bird Populations* **2**:46–61.
- Piegorsch, W. W. 1994. Empirical Bayes calculations of concordance between endpoints in environmental toxicity experiments. *Environmental and Ecological Statistics* **1**:153–164.
- Robbins, C. S., D. K. Dawson, and B. A. Dowell. 1989. Habitat area requirements of breeding forest birds of the middle Atlantic states. *Wildlife Monographs* **103**:1–34.
- Robbins, H. 1956. An empirical Bayes approach to statistics. *Proceedings of the Third Berkeley Symposium* **1**:157–163.
- Robinson, S. K. 1992. Population dynamics of breeding Neotropical migrants in a fragmented Illinois landscape. Pages 455–471 in J. M. Hagan and D. W. Johnston, editors. Ecology and conservation of Neotropical migrant landbirds. Smithsonian Institution Press, Washington, D.C., USA.
- Rothstein, S. I., and S. R. Robinson. 1994. Conservation and coevolutionary implications of brood parasitism by cowbirds. *Trends in Ecology and Evolution* **9**:162–164.
- Smith, P. J. 1991. Bayesian analyses for a multiple capture–recapture model. *Biometrika* **78**:399–407.
- Stijnen, T., and J. C. Van Houwelingen. 1990. Empirical Bayes methods in clinical trials meta-analysis. *Biometrical Journal* **32**:335–346.
- Suggs, J. C., and T. C. Curran. 1983. An empirical Bayes method for comparing air pollution data to air quality standards. *Atmospheric Environment* **17**:837–842.
- Terborgh, J. 1989. Where have the birds gone? Princeton University Press, Princeton, New Jersey, USA.
- Thompson, F. R. 1993. Behavior, movements, and habitat selection of Brown-Headed Cowbirds in Midwest landscapes. Report to North Central Forest Experiment Station, U.S. Department of Agriculture, U.S. Forest Service, St. Paul, Minnesota, USA.

APPENDIX 1

MOMENT ESTIMATORS OF HYPERPARAMETERS

Let n denote the number of host species under consideration. The data for species i are N_i , the number of nests sampled, and X_i , the number of these that are parasitized. The observed rate for species i is $\hat{p}_i = X_i/N_i$, with standard error estimated by:

$$\sigma_i = \sqrt{\hat{p}_i(1 - \hat{p}_i)/N_i}.$$

The mean of the prior distribution can be estimated by:

$$\hat{\pi} = \frac{1}{n} \sum_{i=1}^n \hat{p}_i.$$

An estimate of the spread parameter, θ , can be obtained by evaluating the expected values of the statistics:

$$\hat{\sigma}^2 = \frac{1}{n} \sum_{i=1}^n \sigma_i^2,$$

and

$$V = \frac{1}{n-1} \sum_{i=1}^n (\hat{p}_i - \hat{\pi})^2.$$

Using these, we estimate θ by:

$$\hat{\theta} = \frac{\hat{\pi}(1 - \hat{\pi})}{V - \hat{\sigma}^2} - 1.$$

These estimators are slightly less efficient than the maximum likelihood estimators, but they are much more easily calculated.

APPENDIX 2

ACCOUNTING FOR UNCERTAINTY IN ESTIMATION OF POSTERIOR DISTRIBUTIONS

Empirical Bayes analyses that do not somehow account for the uncertainty associated with the estimation of the hyperparameters have been criticized by some authors (Berger 1985, Laird and Louis 1987, Louis 1991). In particular, concern focuses on whether posterior intervals—the “credibility intervals” described above (see *Application of empirical Bayes analysis*)—should be made longer so as to reflect this uncertainty. A simple and familiar case is the use of a *t* multiplier rather than a *z* multiplier in the construction of interval estimates for a normal mean, when the variance is unknown. (From the Bayesian perspective, these intervals are credibility intervals for the normal mean, supposing a non-informative prior distribution for the mean. The prior distributions for the variance are a point mass on the true variance for the *z* interval, and a noninformative prior for the *t* interval. A detailed discussion of the Bayesian interpretation of these intervals is given by Box and Tiao 1992.)

In this Appendix we describe two methods for adjusting the credibility intervals so as to reflect our imperfect knowledge of the prior parameters. These methods are applicable in many settings; we describe them in the context of the beta-binomial model.

The first method is via an “empirical Bayes bootstrap” (Laird and Louis 1987). The idea is to reproduce, as nearly as possible, the processes that generated the data in hand, with the goal of obtaining “replicate” collections of hyperparameter estimates. First, generate a bootstrap sample of proportions from the estimated prior beta distribution with parameters $\hat{\pi}$ and $\hat{\theta}$. The *i*th of these, p_i^* , can be thought of as a replication of the “true” parasitism rate for species *i*. Next, for each *i*, generate a binomial random variable X_i^* , the number of successes in *N* Bernoulli trials with success rate p_i^* . The collection of pairs (X_i^*, N) is a bootstrap replication of the original data set, from which bootstrap replicate estimates of the hyperparameters can be obtained. These replicate estimates of the hyperparameters are then used with the original

data (X_i, N) to obtain a bootstrap replicate estimate of the posterior distribution for species *i*. This procedure is repeated a large number of times (say, 200); the average of these bootstrap replicate posterior distributions is then used as the estimate of the posterior distribution, instead of the estimate based on the original hyperparameter estimates. The credibility interval is taken from this new estimate of the posterior distribution.

Another approach is to use “hierarchical Bayes” models in which the hyperparameters themselves are regarded as sampled from a “hyper-prior” distribution. To illustrate, we treated the pair (π, θ) as having been sampled from a uniform distribution over the rectangle $[0.16, 0.42] \times [0, 16.5]$. (These ranges were selected on the basis of the marginal likelihood function for (π, θ) ; the likelihood of values outside of this rectangle is very small.) Next, we calculated 2500 distinct posterior distributions for species *i*, one for each pair (π, θ) on a 50×50 grid covering the rectangle. We then took a weighted average of these, with weights given by the marginal likelihood of the pair (π, θ) . Under the hierarchical Bayes model, this weighted average is a numerical integration approximating the posterior distribution for species *i*. For details on the hierarchical Bayes approach, the reader is referred to Berger (1985).

One could consider a variety of prior distributions for (π, θ) : for example, we carried out similar calculations to those described, but under the assumption that $(\pi, \ln(\theta))$ has a uniform hyperprior distribution over the rectangle $[0.15, 0.45] \times [0.5, 3.5]$. For the present data set, however, there was little difference in the results. The 95% credibility intervals are given in Fig. A1, along with the “naive intervals,” which do not account for the uncertainty in the hyperparameters. The effects of the uncertainty in hyperparameters are seen to be slight. This is probably a consequence of the beta-binomial model and the precision with which the hyperparameters are estimated, and may not hold in general.

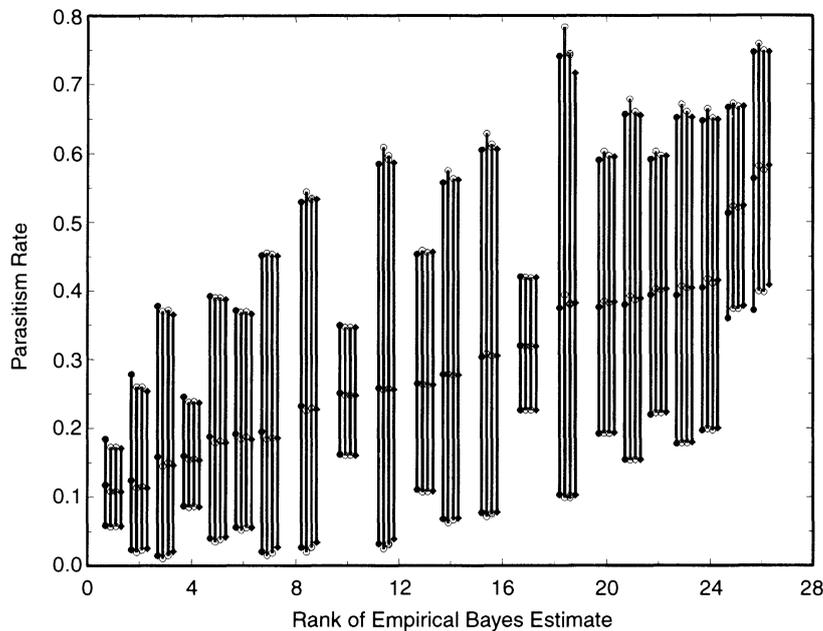


FIG. A1. Graphical comparison of sets of four credibility intervals and posterior means for cowbird parasitism rates of Table 1. The x axis identifies species by rank of parasitism rate based on empirical Bayes estimates \hat{p}_{EB} (Rank₂ in Table 1). Each group of four intervals consists of results obtained using (from left to right) Laird and Louis bootstrap, direct substitution of estimated hyperparameters in formula for posterior distribution, hierarchical Bayes modeling with noninformative priors on π and θ , and hierarchical Bayes modeling with noninformative priors on π and $\ln(\theta)$. There is little difference among the results.