# Rgbp: An **R** Package for Gaussian, Poisson, and Binomial Random Effects Models with Frequency Coverage Evaluations

**Hyungsuk Tak**
Harvard University

**Joseph Kelly**
Google Inc.

**Carl Morris**
Harvard University

### Abstract

**Rgbp** is an R package that provides estimates and verifiable confidence intervals for random effects in two-level conjugate hierarchical models for overdispersed Gaussian, Poisson, and binomial data. **Rgbp** models aggregate data from $k$ independent groups summarized by observed sufficient statistics for each random effect, such as sample means, possibly with covariates. **Rgbp** uses approximate Bayesian machinery with unique improper priors for the hyper-parameters, which leads to good repeated sampling coverage properties for random effects. A special feature of **Rgbp** is an option that generates synthetic data sets to check whether the interval estimates for random effects actually meet the nominal confidence levels. Additionally, **Rgbp** provides inference statistics for the hyper-parameters, e.g., regression coefficients.

*Keywords*: overdispersion, hierarchical model, adjustment for density maximization, frequency method checking, R.

## 1. Introduction

Gaussian, Poisson, or binomial data from several independent groups sometimes have more variation than the assumed Gaussian, Poisson, or binomial distributions of the first-level observed data. To account for this extra variability, called overdispersion, a two-level conjugate hierarchical model regards first-level mean parameters as random effects that come from a population-level conjugate prior distribution. The main goal of our two-level conjugate modeling is to estimate these random effects for a comparison between groups. For example, this model can be used to estimate unknown true batting averages (random effects) of baseball players for a comparison among players based on their numbers of hits and at-bats possibly with their covariate information.

With an assumption of homogeneity within each group, the observed data are group-level aggregate data from $k$ independent groups, composed of sufficient statistics for their $k$ random effects (without the population-level data). Specifically, the data for the Gaussian model consist of each group's sample mean and its standard error, those for the Poisson model use each group's outcome count and an exposure measure, and those for the binomial model use the number of each group's successful outcomes together with the total number of trials. For example, the observed data for the binomial model can be the number of hits out of at-bats for each of $k$ baseball players, a sufficient statistic for the unknown true batting average of each player. The data analyzed by **Rgbp** may incorporate each group's covariate information, e.g., each player's position.

These types of data are common in various fields for estimating random effects. For example, biologists may be interested in unknown true tumor incidence rates in analyzing litter data composed of each litter's number of tumor-bearing animals out of a total number of animals at risk (Tamura and Young 1987). The unknown true mortality rates on myocardial infarction can be estimated based on the death rate data collected from several independent clinical studies via a meta analysis (Gelman, Carlin, Stern, and Rubin 2014). County-level or state-level summary data in small area estimation problems (Ghosh and Rao 1994; Rao 2003) can be used to estimate population parameters, such as unknown unemployment rates.

For such data, assuming homogeneity within each group, **Rgbp**'s two-level model may be viewed as a conjugate hierarchical generalized linear model (HGLM; Lee and Nelder 1996) where each random effect comes from a conjugate prior distribution. However, the HGLM focuses on estimating regression coefficients to explore associations between covariates and observed data. While **Rgbp** does that too, its emphasis concerns making valid point and interval estimates for the $k$ random effects for a comparison among groups.

A defining feature of **Rgbp** is to evaluate the repeated sampling coverage properties of the interval estimates for random effects (Christiansen and Morris 1997; Daniels 1999; Tang 2002; Morris and Tang 2011; Morris and Lysy 2012). This procedure distinguishes **Rgbp** from other R packages for similar hierarchical models, such as **hglm** (Alam, Rönnegård, and Shen 2015) for fitting conjugate hierarchical generalized models, because most software packages produce only estimation results without providing a quick way to evaluate their estimation procedures. The evaluation procedure which we call *frequency method checking* uses a parametric bootstrapping method that generates synthetic data sets given the fitted values of the estimated hyper-parameters. The frequency method checking procedure estimates the coverage rates of interval estimates based on the simulated data sets and checks whether the estimated coverage rates achieve (or exceed) the nominal confidence level.

**Rgbp** combines Bayesian modeling tools with our improper hyper-prior distributions on the second-level parameters. These hyper-prior distributions are known to produce good repeated sampling coverage rates for the Bayesian interval estimates for the $k$ random effects in two-level Gaussian models (Morris and Tang 2011; Morris and Lysy 2012; Kelly 2014). We extend these hyper-prior distributions for **Rgbp**'s Poisson and binomial hierarchical models.

For fitting the hierarchical model, **Rgbp** uses *adjustment for density maximization* (ADM; Morris 1988a; Christiansen and Morris 1997; Morris and Tang 2011). ADM approximates a posterior density or a likelihood function by fitting a selected (one dimensional) Pearson family, based on the first two derivatives of the given density function. For example, when the normal distribution is the chosen Pearson family, ADM reduces to a delta method via maxi-

mum likelihood estimation. Besides ADM, **Rgbp** provides an option for the binomial hierarchical model to draw independent posterior samples of random effects and hyper-parameters via an acceptance-rejection method (Robert and Casella 2013).

The rest of this paper is organized as follows. We specify the Bayesian hierarchical models and discuss their posterior propriety in Section 2. In Section 3, we explain the inferential models used to estimate the model parameters. We describe the estimation procedures including ADM and the acceptance-rejection method in Sections 4 and 5, respectively. We introduce frequency method checking techniques in Section 6. We explain the basic usages of **Rgbp**'s main functions with three examples in Section 7, and specify the functions further with various options in Section 8.

# 2. Conjugate hierarchical modeling structure

**Rgbp** allows users to choose one of three hierarchical models according to the type of data, namely normal-normal, Poisson-gamma, and binomial-beta models. Although there are more hierarchical models, we choose the three models because these are based on the most common types of data we may encounter in practice. Also, their conjugacy leads to linear posterior means simplifying computations.

Our parametrization of the three hierarchical models leads to an intuitive shrinkage interpretation in inference because the shrinkage factors under our parametrization are determined by the relative amount of information in the prior compared to the data (Morris 1983).

## 2.1. Normal-normal model for Gaussian data

The following normal-normal hierarchical model (hereafter the Gaussian model) assumed by **Rgbp** is useful when the group-level aggregate data from $k$ independent groups are continuous (or approximately continuous) variables with known standard errors. The subscript $j$ below indicates the $j$th group among $k$ groups in the data set. For $j = 1, 2, \ldots, k$,

$$y_j \mid \mu_j \overset{\text{indep.}}{\sim} \text{Normal}(\mu_j, V_j), \tag{1}$$

$$\mu_j \mid \boldsymbol{\beta}, A \overset{\text{indep.}}{\sim} \text{Normal}\left(\mu_j^E, A\right), \tag{2}$$

where $y_j$ is an observed unbiased estimate, e.g., the sample mean, for the random effect $\mu_j$, $V_j$ is a completely known standard error of $y_j$, $\mu_j^E$ is an expected random effect defined as $\mathsf{E}(\mu_j \mid \boldsymbol{\beta}, A) = \boldsymbol{x}_j^\top \boldsymbol{\beta} = \beta_1 x_{j,1} + \beta_2 x_{j,2} + \cdots + \beta_m x_{j,m}$, and $m$ is the number of unknown regression coefficients. It is assumed that the second-level variance $A$ is unknown and that the vector of $m$ regression coefficients $\boldsymbol{\beta}$ is also unknown unless otherwise specified. If no covariates are available, but with an unknown intercept term, then $\boldsymbol{x}_j^\top \boldsymbol{\beta} = \beta_1$ ($m = 1$) and thus $\mu_j^E = \mu^E = \beta_1$ for all $j$, resulting in an exchangeable conjugate prior distribution for the random effects. Based on these conjugate prior distributions for random effects, it is easy to derive the conditional posterior distribution of each random effect. For $j = 1, 2, \ldots, k$,

$$\mu_j \mid \boldsymbol{\beta}, A, \boldsymbol{y} \overset{\text{indep.}}{\sim} \text{Normal}\left((1 - B_j)y_j + B_j \mu_j^E, \ (1 - B_j)V_j\right), \tag{3}$$

where $B_j \equiv V_j/(V_j + A)$ is a shrinkage factor of group $j$ and $\boldsymbol{y} = (y_1, y_2, \ldots, y_k)^\top$. The conditional posterior mean $\mathsf{E}(\mu_j \mid \boldsymbol{\beta}, A, \boldsymbol{y})$, denoted by $\mu_j^*$, is a convex combination of the

observed sample mean $y_j$ and the expected random effect $\mu_j^E$ weighted by the shrinkage factor $B_j$. If the variance of the conjugate prior distribution, $A$, is smaller than the variance of the observed distribution, $V_j$, then we expect the posterior mean to borrow more information from the second-level conjugate prior distribution.

## 2.2. Poisson-gamma model for Poisson data

**Rgbp** can estimate a conjugate Poisson-gamma hierarchical model (hereafter the Poisson model) when the group-level aggregate data from $k$ independent groups consist of non-negative count data without upper limit. However, its usage is limited to the case where the expected random effect, $\lambda_j^E = \exp(\boldsymbol{x}_j^\top \boldsymbol{\beta})$, is known (or equivalently all the regression coefficients $\boldsymbol{\beta}$ are known ($m = 0$)); we may be able to obtain this information from the past studies or from experts. For $j = 1, 2, \ldots, k$,

$$y_j \mid \lambda_j \stackrel{\text{indep.}}{\sim} \text{Poisson}(n_j \lambda_j), \tag{4}$$

$$\lambda_j \mid r \stackrel{\text{indep.}}{\sim} \text{Gamma}\left(r\lambda_j^E, r\right), \tag{5}$$

where $y_j$ is the number of events happening, $n_j$ is the exposure of group $j$, which is not necessarily an integer, $\lambda_j^E = \mathsf{E}(\lambda_j \mid r)$ is the known expected random effect ($m = 0$), and $r$ is the unknown second-level variance component. The mean and variance of this conjugate gamma prior distribution are $\lambda^E$ and $\lambda^E/r$, respectively[1]. Albert (1988) interprets $r$ as the amount of prior information as $n_j$ represents the amount of observed information because the uncertainty of the conjugate prior distribution increases as $r$ decreases and vice versa. The conditional posterior distribution of the random effect $\lambda_j$ for this Poisson model is

$$\lambda_j \mid r, \boldsymbol{y} \stackrel{\text{indep.}}{\sim} \text{Gamma}\left(r\lambda_j^E + n_j \bar{y}_j, \ r + n_j\right), \tag{6}$$

where $\bar{y}_j \equiv y_j/n_j$. The mean and variance of the conditional posterior distribution are

$$\lambda_j^* \equiv \mathsf{E}(\lambda_j \mid r, \boldsymbol{y}) = (1 - B_j)\bar{y}_j + B_j \lambda_j^E \ \text{ and } \ \mathsf{VAR}(\lambda_j \mid r, \boldsymbol{y}) = \frac{\lambda_j^*}{r + n_j}. \tag{7}$$

where $B_j \equiv r/(r + n_j)$ is the shrinkage factor of group $j$, the relative amount of information in the prior compared to the data. The conditional posterior mean is a convex combination of $\bar{y}_j$ and $\lambda_j^E$ weighted by $B_j$. If the conjugate prior distribution contains more information than the observed data have, i.e., the ensemble sample size $r$ exceeds the individual sample size $n_j$, then the posterior mean shrinks towards the prior mean by more than 50%.

The conditional posterior variance in (7) is linear in the conditional posterior mean, whereas a slightly different parametrization for a Poisson-gamma model has been used elsewhere (Christiansen and Morris 1997) that makes the variances quadratic functions of means.

## 2.3. Binomial-beta model for binomial data

**Rgbp** can fit a conjugate binomial-beta hierarchical model (hereafter the binomial model) when the group-level aggregate data from $k$ independent groups are composed of each group's

---

[1]The density function of this gamma prior distribution in (5) is $f(\lambda_j \mid r) \propto \lambda_j^{r\lambda_j^E - 1} \exp(-r\lambda_j)$.

number of successes out of total number of trials. The expected random effect in the binomial model is either known ($m = 0$) or unknown ($m \geq 1$). For $j = 1, 2, \ldots, k$,

$$y_j \mid p_j \overset{\text{indep.}}{\sim} \text{Binomial}(n_j, \ p_j),$$

$$p_j \mid \boldsymbol{\beta}, r \overset{\text{indep.}}{\sim} \text{Beta}\left(rp_j^E, \ r(1 - p_j^E)\right),$$

where $y_j$ is the number of successes out of $n_j$ trials, $p_j^E$ is the expected random effect of group $j$ defined as $p_j^E \equiv \mathsf{E}(p_j \mid \boldsymbol{\beta}, r) = \exp(\boldsymbol{x}_j^\top \boldsymbol{\beta})/(1 + \exp(\boldsymbol{x}_j^\top \boldsymbol{\beta}))$. The vector of the $m$ logistic regression coefficients $\boldsymbol{\beta}$ and the second-level variance component $r$ are unknown. The mean and variance of the conjugate beta prior distribution for group $j$ are $p_j^E$ and $p_j^E(1 - p_j^E)/(r + 1)$, respectively. The resultant conditional posterior distribution of random effect $p_j$ is

$$p_j \mid \boldsymbol{\beta}, r, \boldsymbol{y} \overset{\text{indep.}}{\sim} \text{Beta}\left(n_j \bar{y}_j + rp_j^E, \ n_j(1 - \bar{y}_j) + r(1 - p_j^E)\right), \tag{8}$$

where $\bar{y}_j = y_j/n_j$ is the observed proportion of group $j$. The mean and variance of the conditional posterior distribution are

$$p_j^* \equiv \mathsf{E}(p_j \mid \boldsymbol{\beta}, r, \boldsymbol{y}) = (1 - B_j)\bar{y}_j + B_j p_j^E \quad \text{and} \quad \mathsf{VAR}(p_j \mid \boldsymbol{\beta}, r, \boldsymbol{y}) = \frac{p_j^*(1 - p_j^*)}{r + n_j + 1}.$$

The conditional posterior mean $p_j^*$ is a convex combination of $\bar{y}_j$ and $p_j^E$ weighted by $B_j \equiv r/(r + n_j)$ like for the Poisson model. If the conjugate prior distribution contains more information than the observed distribution does ($r > n_j$), then the resulting conditional posterior mean borrows more information from the conjugate beta prior distribution.

## 2.4. Hyper-prior distribution

Hyper-prior distributions are the distributions assigned to the second-level parameters called hyper-parameters. Our choices for the hyper-prior distributions are

$$\boldsymbol{\beta} \sim \text{Uniform on } \mathbf{R}^m \quad \text{and} \quad A \sim \text{Uniform}(0, \infty) \quad (\text{or } 1/r \sim \text{Uniform}(0, \infty)). \tag{9}$$

The improper flat hyper-prior distribution on $\boldsymbol{\beta}$ is a common non-informative choice. In the Gaussian case, the flat hyper-prior distribution on the second-level variance $A$ is known to produce good repeated sampling coverage properties of the Bayesian interval estimates for the random effects (Morris and Tang 2011; Morris and Lysy 2012; Kelly 2014). The resulting full posterior distribution of the random effects and hyper-parameters is proper if $k \geq m + 3$ (Morris and Tang 2011; Kelly 2014).

In the other two cases, Poisson and binomial, the flat prior distribution on $1/r$ induces the same improper prior distribution on shrinkages ($dB_j/B_j^2$) as does $A$ with the Uniform$(0, \infty)$ for the Gaussian case. The Poisson model with this hyper-prior distribution on $r$, i.e., $dr/r^2$, provides posterior propriety if there are at least two groups whose observed values $y_j$ are non-zero and the expected random effects, $\lambda_j^E$, are known ($m = 0$); see Appendix A for its proof. If $\lambda_j^E$ is unknown, the reliability of **Rgbp** cannot be guaranteed as posterior propriety has not been verified in this case. If the Poisson is being used as an approximation to the binomial and the exposures are known integer values, then we recommend using the binomial model with the same hyper-prior distributions.

As for posterior propriety of the binomial model, let us define an *interior group* as the group whose number of successes $y_j$ are neither 0 nor $n_j$, and $k_y$ as the number of interior groups among $k$ groups. The full posterior distribution of random effects and hyper-parameters is proper if and only if there are at least two interior groups in the data and the $k_y \times m$ covariate matrix of the interior groups is of full rank $m$ ($k_y \geq m$) (Tak and Morris 2017).

## 3. The inferential model

The likelihood function of hyper-parameters, $A$ and $\boldsymbol{\beta}$, for the Gaussian model is derived from the independent normal distributions of the observed data with random effects integrated out.

$$L(A, \boldsymbol{\beta}) = \prod_{j=1}^{k} f(y_j \mid A, \boldsymbol{\beta}) = \prod_{j=1}^{k} \frac{1}{\sqrt{2\pi(A + V_j)}} \exp\left(-\frac{(y_j - \mu_j^E)^2}{2(A + V_j)}\right). \tag{10}$$

The joint posterior density function of hyper-parameters for the Gaussian model is proportional to their likelihood function in (10) because we use flat improper hyper-prior density functions for $A$ and $\boldsymbol{\beta}$:

$$f(A, \boldsymbol{\beta} \mid \boldsymbol{y}) \propto L(A, \boldsymbol{\beta}). \tag{11}$$

The likelihood function of $r$ for the Poisson model comes from the independent negative-binomial distributions of the observed data with the random effects integrated out.

$$L(r) = \prod_{j=1}^{k} f(y_j \mid r) = \prod_{j=1}^{k} \frac{\Gamma(r\lambda_j^E + y_j)}{\Gamma(r\lambda_j^E)(y_j!)} (1 - B_j)^{y_i} B_j^{r\lambda_j^E}, \tag{12}$$

where $\Gamma(a)$ is the gamma function defined as $\int_0^\infty x^{a-1} \exp(x) dx$ for a positive constant $a$. The posterior density function of $r$ for the Poisson model is the likelihood function in (12) times the hyper-prior density function of $r$, i.e., $dr/r^2$:

$$f(r \mid \boldsymbol{y}) \propto L(r)/r^2. \tag{13}$$

The likelihood function of hyper-parameters $r$ and $\boldsymbol{\beta}$ for the binomial model is derived from the independent beta-binomial distributions of the observed data with random effects integrated out (Skellam 1948).

$$L(r, \boldsymbol{\beta}) = \prod_{j=1}^{k} f(y_j \mid r, \boldsymbol{\beta}) = \prod_{j=1}^{k} \binom{n_j}{y_j} \frac{B(y_j + rp_j^E, \ n_j - y_j + r(1 - p_j^E))}{B(rp_j^E, \ r(1 - p_j^E))}, \tag{14}$$

where the notation $B(a, b)$ ($\equiv \int_0^1 v^{a-1}(1 - v)^{b-1} dv$) indicates the beta function for positive constants $a$ and $b$. The joint posterior density function of hyper-parameters $f(r, \boldsymbol{\beta} \mid \boldsymbol{y})$ for the binomial model is proportional to their likelihood function in (14) multiplied by the hyper-prior density functions of $r$ and $\boldsymbol{\beta}$ based on distributions in (9):

$$f(r, \boldsymbol{\beta} \mid \boldsymbol{y}) \propto L(r, \boldsymbol{\beta})/r^2. \tag{15}$$

Our goal is to obtain the point and interval estimates of the random effects from their joint unconditional posterior density which, for the Gaussian model, can be expressed as

$$f(\boldsymbol{\mu} \mid \boldsymbol{y}) = \int f(\boldsymbol{\mu} \mid A, \boldsymbol{\beta}, \boldsymbol{y}) f(A, \boldsymbol{\beta} \mid \boldsymbol{y}) dA d\boldsymbol{\beta},$$

where $\boldsymbol{\mu} = (\mu_1, \mu_2, \ldots, \mu_k)^\top$ and the distributions in the integrand are given in (3) and (11). For the Poisson model, the joint unconditional posterior density for the random effects is

$$f(\boldsymbol{\lambda} \mid \boldsymbol{y}) = \int f(\boldsymbol{\lambda} \mid r, \boldsymbol{y}) f(r \mid \boldsymbol{y}) dr,$$

where $\boldsymbol{\lambda} = (\lambda_1, \lambda_2, \ldots, \lambda_k)^\top$ and the distributions in the integrand are given in (6) and (13). For the binomial model, the joint unconditional posterior density for the random effects is

$$f(\boldsymbol{p} \mid \boldsymbol{y}) = \int f(\boldsymbol{p} \mid r, \boldsymbol{\beta}, \boldsymbol{y}) f(r, \boldsymbol{\beta} \mid \boldsymbol{y}) dr d\boldsymbol{\beta}, \tag{16}$$

where $\boldsymbol{p} = (p_1, p_2, \ldots, p_k)^\top$ and the distributions in the integrand are given in (8) and (15).

# 4. Estimation via the adjustment for density maximization

We illustrate our estimation procedure which utilizes adjustment for density maximization (ADM; Morris 1988a; Christiansen and Morris 1997; Morris and Tang 2011). ADM is a method to approximate a distribution by a member of the Pearson family of distributions and obtain moment estimates via maximization. The ADM procedure for the Gaussian model adopted in **Rgbp** is well documented in Kelly (2014) and thus we describe the ADM procedure using the Poisson and binomial model in this section.

## 4.1. Estimation for shrinkage factors and expected random effects

Our goal here is to estimate the unconditional posterior moments of the shrinkage factors and the expected random effects because they are used to estimate the unconditional posterior moments of the random effects.

*Unconditional posterior moments of shrinkage factors*

It is noted that the shrinkage factors are a function of $r$, i.e., $B_j = B_j(r) = r/(r + n_j)$ (or a function of $A$ for the Gaussian model). A common method of estimation of $B_j$ is to approximate the likelihood of $r$ with two derivatives and use a delta method for an asymptotic normal distribution of $\hat{B}_j(\hat{r}_{\mathrm{MLE}})$. This normal approximation, however, is defined on $(-\infty, \infty)$ whereas $B_j$ lies on the unit interval between 0 and 1, and hence in small sample sizes the delta method can result in point estimates lying on the boundary of the parameter space, from which the restricted MLE procedure sometimes suffers (Morris and Tang 2011).

To continue with a maximization-based estimation procedure but to steer clear of aforementioned boundary issues we make use of ADM. The ADM approximates the distribution of the function of the parameter of interest by one of the Pearson family distributions using the first two derivatives as the delta method does; the delta method is a special case of the ADM based on the normal distribution.

The ADM procedure specified in Morris and Tang (2011) assumes that the unconditional posterior distribution of a shrinkage factor follows a beta distribution; for $j = 1, 2, \ldots, k$,

$$B_j \mid \boldsymbol{y} \sim \mathrm{Beta}(a_{1j},\ a_{0j}). \tag{17}$$

The mean of the beta distribution $a_{1j}/(a_{1j} + a_{0j})$ is not the same as its mode $(a_{j1} - 1)/(a_{j1} + a_{j0} - 2)$. The ADM works on an adjusted posterior distribution $f^A(B_j \mid \boldsymbol{y}) \propto B_j(1 -$

$B_j)f(B_j|\boldsymbol{y})$ so that the mode of $f^A(B_j \mid \boldsymbol{y})$ is the same as the mean of the original beta distribution. The assumed posterior mean and variance of the $j$th shrinkage factor are

$$\mathsf{E}(B_j \mid \boldsymbol{y}) = \frac{a_{1j}}{a_{1j} + a_{0j}} = \underset{B_j}{\operatorname{argmax}} \ f^A(B_j \mid \boldsymbol{y}) \equiv B_j', \tag{18}$$

$$\mathsf{VAR}(B_j \mid \boldsymbol{y}) = \frac{B_j'(1 - B_j')}{a_{1j} + a_{0j} + 1} = \frac{B_j'(1 - B_j')}{B_j'(1 - B_j')\left[ -\frac{d^2}{dB_j^2} \log\left(f^A(B_j \mid \boldsymbol{y})\right) \Big|_{B_j = B_j'} \right] + 1}. \tag{19}$$

The ADM estimates these mean and variance using the marginal posterior distribution of $r$, $f(r \mid \boldsymbol{y}) \propto L(r)/r^2$. The marginal likelihood, $L(r) = \int L(\boldsymbol{\beta}, r)d\boldsymbol{\beta}$, for the binomial model is obtained via the Laplace approximation with a Lebesgue measure on $\boldsymbol{\beta}$ and that for the Poisson model is specified in (12).

Considering that (18) involves maximization and (19) involves calculating the second derivative of $\log(f^A(B_j \mid \boldsymbol{y}))$, we work on a logarithmic scale of $r$, i.e., $\alpha = -\log(r)$ (or $\alpha = \log(A)$ for the Gaussian model). This is because the distribution of $\alpha$ is more symmetric than that of $r$ and $\alpha$ is defined on a real line without any boundary issues. Because $f^A(B_j \mid \boldsymbol{y})$ is proportional to the marginal posterior density $f(\alpha \mid \boldsymbol{y}) \propto \exp(\alpha)L(\alpha)$ as shown in Morris and Tang (2011), the posterior mean in (18) is estimated by

$$\hat{B}_j' = \frac{\exp(-\hat{\alpha})}{n_j + \exp(-\hat{\alpha})}, \tag{20}$$

where $\hat{\alpha}$ is the mode of $f(\alpha \mid \boldsymbol{y})$, i.e., $\operatorname{argmax}_\alpha\{\alpha + \log(L(\alpha))\}$ .

To estimate the variance in (19), Morris and Tang (2011) introduced the invariant information defined as

$$\begin{aligned} I_{\text{inv}} &\equiv -\frac{d^2 \log(f^A(B_j \mid \boldsymbol{y}))}{d[\operatorname{logit}(B_j)]^2} \Big|_{B_j = \hat{B}_j'} = -\frac{d^2 \log(f^A(B_j(r) \mid \boldsymbol{y}))}{d[\log(r)]^2} \Big|_{r = \hat{r}} \\ &= -\frac{d^2 \log(f^A(B_j(r(\alpha)) \mid \boldsymbol{y}))}{d\alpha^2} \Big|_{\alpha = \hat{\alpha}}. \end{aligned} \tag{21}$$

This invariant information is the negative second-derivative of $\alpha + \log(L(\alpha))$ evaluated at the mode $\hat{\alpha}$. Using the invariant information, we estimate the unconditional posterior variance of the shrinkage factor in (19) by

$$\widehat{\mathsf{VAR}}(B_j \mid \boldsymbol{y}) = \frac{(\hat{B}_j')^2(1 - \hat{B}_j')^2}{I_{\text{inv}} + \hat{B}_j'(1 - \hat{B}_j')}. \tag{22}$$

We obtain the estimates of $a_{1j}$ and $a_{0j}$, the two parameters of the beta distribution in (17), by matching them to the estimated unconditional posterior mean and variance of $B_j$ specified in (20) and (22) as follows.

$$\hat{a}_{1j} = \frac{I_{\text{inv}}}{1 - \hat{B}_j'} \ \text{ and } \ \hat{a}_{0j} = \frac{I_{\text{inv}}}{\hat{B}_j'}. \tag{23}$$

The moments of the beta distribution are well defined as a function of $a_{1j}$ and $a_{0j}$, i.e., $\mathsf{E}(B_j^c \mid \boldsymbol{y}) = B(a_{1j} + c, a_{0j})/B(a_{1j}, a_{0j})$ for $c \geq 0$. Their estimates are

$$\widehat{\mathsf{E}}(B_j^c \mid \boldsymbol{y}) = \frac{B(\hat{a}_{1j} + c, \ \hat{a}_{0j})}{B(\hat{a}_{1j}, \ \hat{a}_{0j})}. \tag{24}$$

The ADM approximation to the shrinkage factors via beta distributions is empirically proven to be more accurate than a Laplace approximation (Morris 1988a; Christiansen and Morris 1997; Morris and Tang 2011; Morris and Lysy 2012).

*Unconditional posterior moments of expected random effects*

We estimate the unconditional posterior moments of expected random effects using their relationship to the conditional posterior moments. For a non-negative constant $c$, the unconditional posterior moments are

$$\mathsf{E}\left(\left(p_j^E\right)^c \mid \boldsymbol{y}\right) = \mathsf{E}\left[\mathsf{E}\left(\left(p_j^E\right)^c \mid \alpha, \boldsymbol{y}\right) \mid \boldsymbol{y}\right].$$

We approximate the unconditional posterior moments on the left hand side by the conditional posterior moments with $\hat{\alpha}$ inserted (Kass and Steffey 1989), i.e., by $\mathsf{E}((p_j^E)^c \mid \hat{\alpha}, \boldsymbol{y})$.

However, calculating conditional posterior moments of each expected random effect involves an intractable integration. For example, the first conditional posterior moment of $p_j^E$ is

$$\mathsf{E}(p_j^E \mid \hat{\alpha}, \boldsymbol{y}) = E\left(\frac{\exp(\boldsymbol{x}_j^\top \boldsymbol{\beta})}{1 + \exp(\boldsymbol{x}_j^\top \boldsymbol{\beta})} \bigg| \hat{\alpha}, \boldsymbol{y}\right) = \int_{\mathbf{R}^m} \frac{\exp(\boldsymbol{x}_j^\top \boldsymbol{\beta})}{1 + \exp(\boldsymbol{x}_j^\top \boldsymbol{\beta})} f(\beta \mid \hat{\alpha}, \boldsymbol{y}) d\boldsymbol{\beta}.$$

Thus, we use another ADM, assuming the conditional posterior distribution of each expected random effect is a beta distribution as follows.

$$p_j^E \mid \hat{\alpha}, \boldsymbol{y} = \frac{\exp(\boldsymbol{x}_j^\top \boldsymbol{\beta})}{1 + \exp(\boldsymbol{x}_j^\top \boldsymbol{\beta})} \bigg| \hat{\alpha}, \boldsymbol{y} \sim \mathrm{Beta}(b_{1j},\ b_{0j}) \sim \frac{G_1}{G_1 + G_0}, \tag{25}$$

where $G_1$ is a random variable following a $\mathrm{Gamma}(b_{1j}, 1)$ distribution and independently $G_0$ has a $\mathrm{Gamma}(b_{0j}, 1)$ distribution. The representation in (25) is equivalent to $\exp(\boldsymbol{x}_j^\top \boldsymbol{\beta}) \mid \hat{\alpha}, \boldsymbol{y} \sim G_1/G_0$, a ratio of two independent gamma random variables. Its mean and variance are

$$\mathsf{E}(\exp(\boldsymbol{x}_j^\top \boldsymbol{\beta}) \mid \hat{\alpha}, \boldsymbol{y}) = \mathsf{E}\left(\frac{G_1}{G_0}\right) = \frac{b_{1j}}{b_{0j} - 1} \equiv \eta_j, \tag{26}$$

$$\mathsf{VAR}(\exp(\boldsymbol{x}_j^\top \boldsymbol{\beta}) \mid \hat{\alpha}, \boldsymbol{y}) = \mathsf{VAR}\left(\frac{G_1}{G_0}\right) = \frac{\eta_j(1 + \eta_j)}{b_{0j} - 2}. \tag{27}$$

To estimate $b_{1j}$ and $b_{0j}$, we assume that the conditional posterior distribution of $\boldsymbol{\beta}$ given $\hat{\alpha}$ and $\boldsymbol{y}$ follows a normal distribution with mean $\hat{\boldsymbol{\beta}}$ and variance-covariance matrix $\hat{\Sigma}$, where $\hat{\boldsymbol{\beta}}$ is the mode of $f(\boldsymbol{\beta} \mid \hat{\alpha}, \boldsymbol{y})$ and $\hat{\Sigma}$ is an inverse of the negative Hessian matrix at the mode. Thus, the posterior distribution of $\boldsymbol{x}_j^\top \boldsymbol{\beta}$ is also normal with mean $\boldsymbol{x}_j^\top \hat{\boldsymbol{\beta}}$ and variance $\boldsymbol{x}_j^\top \hat{\Sigma} \boldsymbol{x}_j$.

Using the property of the log-normal distribution for $\exp(\boldsymbol{x}_j^\top \boldsymbol{\beta})$, we estimate the posterior mean and variance in (26) and (27) as

$$\hat{\mathsf{E}}(\exp(\boldsymbol{x}_j^\top \boldsymbol{\beta}) \mid \hat{\alpha}, \boldsymbol{y}) = \exp\left(\boldsymbol{x}_j^\top \hat{\boldsymbol{\beta}} + \boldsymbol{x}_j^\top \hat{\Sigma} \boldsymbol{x}_j / 2\right) = \hat{\eta}_j, \tag{28}$$

$$\widehat{\mathsf{VAR}}(\exp(\boldsymbol{x}_j^\top \boldsymbol{\beta}) \mid \boldsymbol{y}) = \hat{\eta}_j^2\left(\exp(\boldsymbol{x}_j^\top \hat{\Sigma} \boldsymbol{x}_j) - 1\right). \tag{29}$$

We estimate the values of $b_{1j}$ and $b_{0j}$ by matching them to the estimated unconditional posterior mean and variance of $\exp(\boldsymbol{x}_j^\top \boldsymbol{\beta})$ in (28) and (29), that is,

$$\hat{b}_{1j} = \hat{\eta}_j(\hat{b}_{0j} - 1) \ \ \text{and} \ \ \hat{b}_{0j} = \frac{1 + \hat{\eta}_j}{\hat{\eta}_j\left(\exp(\boldsymbol{x}_j^\top \hat{\Sigma} \boldsymbol{x}_j) - 1\right)} + 2.$$

Finally, we estimate the unconditional posterior moments of the expected random effects by

$$\hat{\mathsf{E}}\left(\left(p_j^E\right)^c \mid \hat{\alpha}, \boldsymbol{y}\right) = \frac{B\left(\hat{b}_{1j} + c, \hat{b}_{0j}\right)}{B\left(\hat{b}_{1j}, \hat{b}_{0j}\right)} \text{ for } c \geq 0. \tag{30}$$

The ADM approximation to a log-normal density via a *F*-distribution (represented by a ratio of two independent gamma random variables) is known to be more accurate than the Laplace approximation (Morris 1988a).

For the Gaussian model (Morris and Tang 2011), the conditional posterior distribution of $\boldsymbol{\beta}$ given $\hat{A}$ and $\boldsymbol{y}$ is normal whose mean and variance-covariance matrix are

$$\left(X^\top D_{V+\hat{A}}^{-1} X\right)^{-1} X^\top D_{V+\hat{A}}^{-1} \boldsymbol{y} \text{ and } \left(X^\top D_{V+\hat{A}}^{-1} X\right)^{-1},$$

respectively, where $X \equiv (\boldsymbol{x}_1, \boldsymbol{x}_2, \ldots, \boldsymbol{x}_k)^\top$ is a $k \times m$ covariate matrix and $D_{V+\hat{A}}$ is a $k \times k$ diagonal matrix with the $j$th diagonal element equal to $V_j + \hat{A}$. Because $\boldsymbol{x}^\top \boldsymbol{\beta}$ given $\hat{A}$ and $\boldsymbol{y}$ is also normally distributed, we easily obtain the conditional posterior moments of $\mu_j^E = \boldsymbol{x}^\top \boldsymbol{\beta}$ given $\hat{A}$ and use them to estimate unconditional posterior moments of $\mu_j^E$.

## 4.2. Estimation for random effects

We illustrate how we obtain approximate unconditional posterior distributions of random effects using the estimated unconditional posterior moments of shrinkage factors and those of expected random effects. It is intractable to derive analytically the unconditional posterior distribution of each random effect for the three models. Thus, we approximate the distributions by matching the estimated posterior moments with a skewed-normal distribution (Azzalini 1985) for the Gaussian model, a gamma distribution for the Poisson model, and a beta distribution for the binomial model; for $j = 1, 2, \ldots, k$,

$$\mu_j \mid \boldsymbol{y} \sim \text{Skewed-Normal}(\phi, \ \omega, \ \delta), \tag{31}$$

$$\lambda_j \mid \boldsymbol{y} \sim \text{Gamma}(s_{1j}, \ s_{0j}), \tag{32}$$

$$p_j \mid \boldsymbol{y} \sim \text{Beta}(t_{1j}, \ t_{0j}), \tag{33}$$

where $(\phi, \omega, \delta)$ of the skewed-normal distribution are location, scale, and skewness parameters, respectively.

Morris and Lysy (2012) first noted that the unconditional posterior distribution of the random effect in a two-level conjugate Gaussian model might be skewed. Kelly (2014) shows that the skewed-normal approximation to the unconditional posterior distribution of the random effect is better than a normal approximation ($\mu_j \mid \boldsymbol{y} \sim \text{Normal}$) in terms of the repeated sampling coverage properties of random effects. Kelly (2014) estimates the first three moments of the random effects by noting that $\mu_j$ is normally distributed given $A$ and $\boldsymbol{y}$, and thus estimates the moments by using the ADM approximation of the shrinkage factors, $B_j$, and the law of third cumulants (Brillinger 1969). The three estimated moments are then matched to the first three moments of the skewed-normal distribution, i.e., $\mathsf{E}(\mu_j \mid \boldsymbol{y}) = \phi + \omega\delta\sqrt{2/\pi}$, $\mathsf{VAR}(\mu_j \mid \boldsymbol{y}) = \omega^2(1 - 2\delta^2/\pi)$, and $\text{Skewness}(\mu_j \mid \boldsymbol{y}) = (4 - \pi)\delta^3/[2(\pi/2 - \delta^2)^{3/2}]$ (Azzalini 1985). The full derivation can be found in Kelly (2014).

The unconditional posterior mean and variance of random effect $\lambda_j$ in the Poisson model are

$$\mathsf{E}(\lambda_j \mid \boldsymbol{y}) = \mathsf{E}(\mathsf{E}(\lambda_j \mid r, \boldsymbol{y}) \mid \boldsymbol{y}) = (1 - \mathsf{E}(B_j \mid \boldsymbol{y}))\bar{y}_j + \mathsf{E}(B_j \mid \boldsymbol{y})\lambda_j^E, \tag{34}$$

$$\mathsf{VAR}(\lambda_j \mid \boldsymbol{y}) = \mathsf{E}(\mathsf{VAR}(\lambda_j \mid r, \boldsymbol{y}) \mid \boldsymbol{y}) + \mathsf{VAR}(\mathsf{E}(\lambda_j \mid r, \boldsymbol{y}) \mid \boldsymbol{y}) \tag{35}$$

$$= \mathsf{E}(\lambda_j^*/(r + n_j) \mid \boldsymbol{y}) + \mathsf{VAR}(B_j(\bar{y}_j - \lambda_j^E) \mid \boldsymbol{y}) \tag{36}$$

$$= \left[\bar{y}_j\mathsf{E}\left((1 - B_j)^2|\boldsymbol{y}\right) + \lambda_j^E\mathsf{E}\left((1 - B_j)B_j|\boldsymbol{y}\right)\right]/n_j + (\bar{y}_j - \lambda_j^E)^2\mathsf{VAR}(B_j|\boldsymbol{y}). \tag{37}$$

To estimate these, we insert the estimated unconditional posterior moments of shrinkage factors in (24) into both (34) and (37). Let $\hat{\mu}_{\lambda_j}$ and $\hat{\sigma}_{\lambda_j}^2$ denote the estimated unconditional posterior mean and variance, respectively. The estimates of the two parameters $s_{1j}$ and $s_{0j}$ in (32) are

$$\hat{s}_{1j} = \frac{\hat{\mu}_{\lambda_j}^2}{\hat{\sigma}_{\lambda_j}^2}, \text{ and } \hat{s}_{0j} = \frac{\hat{\mu}_{\lambda_j}}{\hat{\sigma}_{\lambda_j}^2}.$$

To estimate the unconditional posterior moments of random effects in the binomial model, we assume that hyper-parameters $r$ and $\boldsymbol{\beta}$ are independent *a posteriori*. With this assumption, the unconditional posterior mean and variance of random effect $p_j$ are

$$\mathsf{E}(p_j \mid \boldsymbol{y}) = \mathsf{E}(\mathsf{E}(p_j \mid r, \boldsymbol{\beta}, \boldsymbol{y}) \mid \boldsymbol{y}) = (1 - \mathsf{E}(B_j \mid \boldsymbol{y}))\bar{y}_j + \mathsf{E}(B_j \mid \boldsymbol{y})\mathsf{E}(p_j^E \mid \boldsymbol{y}), \tag{38}$$

$$\mathsf{VAR}(p_j \mid \boldsymbol{y}) = \mathsf{E}(\mathsf{VAR}(p_j \mid r, \boldsymbol{\beta}, \boldsymbol{y}) \mid \boldsymbol{y}) + \mathsf{VAR}(\mathsf{E}(p_j \mid r, \boldsymbol{\beta}, \boldsymbol{y}) \mid \boldsymbol{y}) \tag{39}$$

$$= \mathsf{E}(p_j^*(1 - p_j^*)/(r + n_j + 1) \mid \boldsymbol{y}) + \mathsf{VAR}(B_j(\bar{y}_j - p_j^E) \mid \boldsymbol{y}) \tag{40}$$

$$\approx \mathsf{E}(p_j^*(1 - p_j^*)(1 - B_j)/n_i \mid \boldsymbol{y}) + \mathsf{VAR}(B_j(\bar{y}_j - p_j^E) \mid \boldsymbol{y}) \tag{41}$$

$$= \{(1 - \bar{y}_j)\bar{y}_j[1 - \mathsf{E}(B_j \mid \boldsymbol{y})] + (2\bar{y}_j - 1)\mathsf{E}(B_j(1 - B_j) \mid \boldsymbol{y})(\bar{y}_j - \mathsf{E}(p_j^E \mid \boldsymbol{y}))$$
$$+ \mathsf{E}(B_j^2(1 - B_j) \mid \boldsymbol{y})\mathsf{E}((\bar{y}_j - p_j^E)^2 \mid \boldsymbol{y})\}/n_j + \mathsf{VAR}(B_j(\bar{y}_j - p_j^E)|\boldsymbol{y}), \tag{42}$$

where the approximation in (41) is a first-order Taylor approximation. By inserting the estimated unconditional posterior moments of shrinkage factors in (24) and those of expected random effect in (30) into both (38) and (42), we obtain the estimates of the unconditional posterior mean and variance of each random effect, denoted by $\hat{\mu}_{p_j}$ and $\hat{\sigma}_{p_j}^2$, respectively. We thus obtain the estimates of two parameters $t_{1j}$ and $t_{0j}$ in (33) as follows.

$$\hat{t}_{1j} = \left(\frac{\hat{\mu}_{p_j}(1 - \hat{\mu}_{p_j})}{\hat{\sigma}_{p_j}^2} - 1\right)\hat{\mu}_{p_j}, \text{ and } \hat{t}_{0j} = \left(\frac{\hat{\mu}_{p_j}(1 - \hat{\mu}_{p_j})}{\hat{\sigma}_{p_j}^2} - 1\right)(1 - \hat{\mu}_{p_j}).$$

Finally, the assumed unconditional posterior distribution of the random effect for the Gaussian model is

$$\mu_j \mid \boldsymbol{y} \sim \text{Skewed-Normal}(\hat{\phi}, \hat{\omega}, \hat{\delta}), \tag{43}$$

that for the Poisson model is

$$\lambda_j \mid \boldsymbol{y} \sim \text{Gamma}(\hat{s}_{1j}, \hat{s}_{0j}), \tag{44}$$

and that for the binomial model is

$$p_j \mid \boldsymbol{y} \sim \text{Beta}(\hat{t}_{1j}, \hat{t}_{0j}). \tag{45}$$

Our point and interval estimates of each random effect are the mean and (2.5%, 97.5%) quantiles (if we assign 95% confidence level) of the assumed unconditional posterior distribution in (43), (44), or (45).

For the binomial model, **Rgbp** provides a fully Bayesian approach for drawing posterior samples of random effects and hyper-parameters, which we illustrate in the next section.

## 5. The acceptance-rejection method for the binomial model

In this section, we illustrate an option of **Rgbp** that provides a way to draw posterior samples of random effects and hyper-parameters via the acceptance-rejection (A-R) method (Robert and Casella 2013) for the binomial model. Unlike the approximate Bayesian machinery specified in the previous section, this method does not assume that hyper-parameters are independent *a posteriori*. The joint posterior density function of $\alpha = -\log(r)$ and $\boldsymbol{\beta}$ based on their joint hyper-prior density function in (9) is

$$f(\alpha, \boldsymbol{\beta} \mid \boldsymbol{y}) \propto f(\alpha, \boldsymbol{\beta})L(\alpha, \boldsymbol{\beta}) \propto \exp(\alpha)L(\alpha, \boldsymbol{\beta}). \tag{46}$$

The A-R method is useful when it is difficult to sample a parameter of interest $\theta$ directly from its target probability density $f(\theta)$, which is known up to a normalizing constant, but an easy-to-sample envelope function $g(\theta)$ is available. The A-R method samples $\theta$ from the envelope $g(\theta)$ and accepts it with a probability $f(\theta)/(Mg(\theta))$, where $M$ is a constant making $f(\theta)/g(\theta) \leq M$ for all $\theta$. The distribution of the accepted $\theta$ exactly follows $f(\theta)$. The A-R method is stable as long as the tails of the envelope function are thicker than those of the target density function.

The goal of the A-R method for the binomial model is to draw posterior samples of hyper-parameters from (46), using an easy-to-sample envelope function $g(\alpha, \boldsymbol{\beta})$ that has thicker tails than the target density function.

We factor the envelope function into two parts, $g(\alpha, \boldsymbol{\beta}) = g_1(\alpha)g_2(\boldsymbol{\beta})$ to model the tails of each function separately. We consider the tail behavior of the conditional posterior density function $f(\alpha \mid \boldsymbol{\beta}, \boldsymbol{y})$ to establish $g_1(\alpha)$; $f(\alpha \mid \boldsymbol{\beta}, \boldsymbol{y})$ behaves as $\exp(-\alpha(k-1))$ when $\alpha$ goes to $\infty$ and as $\exp(\alpha)$ when $\alpha$ goes to $-\infty$. This indicates that $f(\alpha \mid \boldsymbol{\beta}, \boldsymbol{y})$ is skewed to the left because the right tail touches the $x$-axis faster than the left tail does as long as $k > 1$. A skewed $t$-distribution is a good candidate for $g_1(\alpha)$ because it behaves as a power law on both tails, leading to thicker tails than those of $f(\alpha \mid \boldsymbol{\beta}, \boldsymbol{y})$.

It is too complicated to figure out the tail behaviors of $f(\boldsymbol{\beta} \mid \alpha, \boldsymbol{y})$. However, because $f(\boldsymbol{\beta} \mid \alpha, \boldsymbol{y})$ in the Gaussian model (as an approximation) has a multivariate Gaussian density function (Morris and Tang 2011; Kelly 2014), we consider a multivariate $t$-distribution with four degrees of freedom as a good candidate for $g_2(\boldsymbol{\beta})$.

Specifically, we assume

$$
\begin{aligned}
g_1(\alpha) &= g_1(\alpha; l, \sigma, a, b) &\equiv \text{Skewed-}t(\alpha \mid l, \sigma, a, b), \tag{47} \\
g_2(\boldsymbol{\beta}) &= g_2(\boldsymbol{\beta}; \boldsymbol{\xi}, S_{(m \times m)}) &\equiv t_4(\boldsymbol{\beta} \mid \boldsymbol{\xi}, S), \tag{48}
\end{aligned}
$$

where Skewed-$t(\alpha \mid l, \sigma, a, b)$ represents a density function of a skewed $t$-distribution of $\alpha$ with location $l$, scale $\sigma$, degree of freedom $a + b$, and skewness $a - b$ for any positive constants $a$ and $b$ (Jones and Faddy 2003). Jones and Faddy (2003) derive the mode of $g_1(\alpha)$ as

$$l + \frac{(a-b)\sqrt{a+b}}{\sqrt{(2a+1)(2b+1)}}, \tag{49}$$

and provide a representation to generate random variables that follows Skewed-$t(\alpha \mid l, \sigma, a, b)$;

$$\alpha \sim l + \sigma \frac{\sqrt{a+b}(2T-1)}{2\sqrt{T(1-T)}}, \text{ where } T \sim \text{Beta}(a, b). \tag{50}$$

They also show that the tails of the skewed-$t$ density function follow a power law with $\alpha^{-(2a+1)}$ on the left and $\alpha^{-(2b+1)}$ on the right when $b > a$.

The notation $t_4(\boldsymbol{\beta} \mid \boldsymbol{\xi}, S)$ in (48) indicates a density function of a multivariate $t$-distribution of $\boldsymbol{\beta}$ with four degrees of freedom, a location vector $\boldsymbol{\xi}$, and a $m \times m$ scale matrix $S$ that leads to the variance-covariance matrix $2S$.

**Rgbp** determines the parameters of $g_1(\alpha)$ and $g_2(\boldsymbol{\beta})$, i.e., $l$, $\sigma$, $a$, $b$, $\boldsymbol{\xi}$, and $S$, to make the product of $g_1(\alpha)$ and $g_2(\boldsymbol{\beta})$ similar to the target joint posterior density $f(\alpha, \boldsymbol{\beta} \mid \boldsymbol{y})$. First, **Rgbp** obtains the mode of $f(\alpha, \boldsymbol{\beta} \mid \boldsymbol{y})$, $(\hat{\alpha}, \hat{\boldsymbol{\beta}})$, and the inverse of the negative Hessian matrix at the mode. We define $-H_{\hat{\alpha}}^{-1}$ to indicate the $(1, 1)$th element of the negative Hessian matrix and $-H_{\hat{\boldsymbol{\beta}}}^{-1}$ to represent the negative Hessian matrix without the first row and the first column.

For $g_1(\alpha)$, **Rgbp** sets $(a, b)$ to $(k, 2k)$ if $k$ is less than 10 (or to $(\log(k), 2\log(k))$ otherwise) for left-skewness and these small values of $a$ and $b$ lead to thick tails. **Rgbp** matches the mode of $g_1(\alpha)$ specified in (49) to $\hat{\alpha}$ by setting the location parameter $l$ to $\hat{\alpha} - (a - b)\sqrt{a+b}/\sqrt{(2a+1)(2b+1)}$. **Rgbp** sets the scale parameter $\sigma$ to $(-H_{\hat{\alpha}}^{-1})^{0.5}\psi$, where $\psi$ is a tuning parameter; when the A-R method produces extreme weights defined in (51) below, we need to enlarge the value of $\psi$.

For $g_2(\boldsymbol{\beta})$, **Rgbp** sets the location vector $\boldsymbol{\xi}$ to the mode $\hat{\boldsymbol{\beta}}$ and the scale matrix $S$ to $-H_{\hat{\boldsymbol{\beta}}}^{-1}/2$ so that the variance-covariance matrix becomes $-H_{\hat{\boldsymbol{\beta}}}^{-1}$.

For implementation of the acceptance-rejection method, **Rgbp** draws four times more trial samples than the desired number of samples, denoted by $N$, independently from $g_1(\alpha)$ and $g_2(\boldsymbol{\beta})$. **Rgbp** calculates $4N$ weights, each of which is defined as

$$w_i \equiv w(\alpha^{(i)}, \boldsymbol{\beta}^{(i)}) = \frac{f(\alpha^{(i)}, \boldsymbol{\beta}^{(i)} \mid \boldsymbol{y})}{g_1(\alpha^{(i)})g_2(\boldsymbol{\beta}^{(i)})}, \text{ for } i = 1, 2, \ldots, 4N. \tag{51}$$

**Rgbp** accepts each pair of $(\alpha^{(i)}, \boldsymbol{\beta}^{(i)})$ with a probability $w_i/M$ where $M$ is set to the maximum of all the $4N$ weights. When **Rgbp** accepts more than $N$ pairs, it discards the redundant. If **Rgbp** accepts less than $N$ pairs, then it additionally draws $N'$ (six times the shortage) pairs and calculates a new maximum $M'$ from all the previous and new weights; **Rgbp** accepts or rejects the entire pairs again with new probabilities $w_j/M'$, $j = 1, 2, \ldots, 4N + N'$.

After obtaining posterior samples of hyper-parameters, **Rgbp** draws posterior samples of random effects from $f(\boldsymbol{p} \mid \boldsymbol{y})$ in (16). The integration on the right hand side of (16) can be done by sampling $\boldsymbol{p}$ from $f(p_j \mid \boldsymbol{\beta}, r, \boldsymbol{y})$ in (8) for $j = 1, 2, \ldots, k$, given $r = \exp(-\alpha)$ and $\boldsymbol{\beta}$ that are already sampled from $f(\alpha, \boldsymbol{\beta} \mid \boldsymbol{y})$ via the A-R method.

# 6. Frequency method checking

The question as to whether the interval estimates of random effects for a given confidence level obtained by a specific model achieve the nominal coverage rate for any true parameter values is one of the key model evaluation criteria. Unlike standard model checking methods

$$\mathbf{A}, \boldsymbol{\beta} \overset{\nearrow}{\underset{\searrow}{\rightarrow}} \begin{matrix} \boldsymbol{\mu}^{(1)} & \rightarrow & \boldsymbol{y}^{(1)} \\ \vdots & \rightarrow & \vdots \\ \boldsymbol{\mu}^{(N_{sim})} & \rightarrow & \boldsymbol{y}^{(N_{sim})} \end{matrix}$$
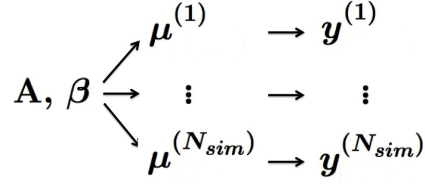
Figure 1: Pseudo-data generating process.

that test whether a two-level model is appropriate for data (Dean 1992; Christiansen and Morris 1996), frequency method checking is a procedure to evaluate the coverage properties of the model. Conditioning on that the two-level model is appropriate, the frequency method checking procedure generates pseudo-data sets given specific values of hyper-parameters and estimates unknown coverage probabilities based on these mock data sets (i.e., parametric bootstrapping). We describe the frequency method checking based on the Gaussian model because the idea can be easily applied to the other two models.

### 6.1. Pseudo-data generation

Figure 1 displays the process of generating pseudo-data sets. It is noted that the conjugate prior distribution of each random effect in (2) is completely determined by two hyper-parameters, $A$ and $\boldsymbol{\beta}$. Fixing these hyper-parameters at specific values, we generate $N_{\text{sim}}$ sets of random effects from the conjugate prior distribution, i.e., $\{\boldsymbol{\mu}^{(i)}, \ i = 1, \ldots, N_{\text{sim}}\}$, where the superscript $(i)$ indicates the $i$th simulation. Next, using the distribution of observed data in (1), we generate $N_{\text{sim}}$ sets of observed data sets $\{\boldsymbol{y}^{(i)}, \ i = 1, \ldots, N_{\text{sim}}\}$ given each $\boldsymbol{\mu}^{(i)}$.

### 6.2. Coverage probability estimation

After fitting the Gaussian model for each simulated data set, we obtain interval estimates of the random effects $\boldsymbol{\mu}^{(i)}$. Let $(\hat{\mu}_{j,\ \text{low}}^{(i)}, \ \hat{\mu}_{j,\ \text{upp}}^{(i)})$ represent the lower and upper bounds of the interval estimate of random effect $j$ based on the $i$th simulation given a specific confidence level. We define the coverage indicator of random effect $j$ on the $i$th mock data set as

$$I\left(\mu_j^{(i)}\right) = \begin{cases} 1, & \text{if } \mu_j^{(i)} \in \left(\hat{\mu}_{j,\ \text{low}}^{(i)}, \ \hat{\mu}_{j,\ \text{upp}}^{(i)}\right), \\ 0, & \text{otherwise.} \end{cases} \tag{52}$$

This shrinkage indicator is equal to the value one if the random effect $j$ in simulation $i$ is between its interval estimates and zero otherwise.

*Simple unbiased coverage estimator*

When the confidence level is 95%, the proportion of 95% interval estimates that contain random effect $j$ is an intuitive choice for the coverage rate estimator for random effect $j$. This estimator implicitly assumes that there exist $k$ unknown coverage probabilities of random effects, denoted by $C_{A,\beta}(\mu_j)$ for $j = 1, 2, \ldots, k$, depending on the values of the hyper-parameters that generate random effects and mock data sets. The coverage indicators for random effect $j$ in (52) is assumed to follow an independent and identically distributed Bernoulli distribution given the unknown coverage rate $C_{A,\beta}(\mu_j)$. The sample mean of these coverage indicators is

a simple unbiased coverage estimator for $C_{A,\beta}(\mu_j)$; for $j = 1, 2, \ldots, k$,

$$\bar{I}(\mu_j) = \frac{1}{N_{\text{sim}}} \sum_{i=1}^{N_{\text{sim}}} I\left(\mu_j^{(i)}\right).$$

The unbiased variance estimator of $\mathsf{VAR}(\bar{I}(\mu_j))$ is, for $j = 1, 2, \ldots, k$,

$$\widehat{\mathsf{VAR}}\left(\bar{I}(\mu_j)\right) = \frac{1}{N_{\text{sim}}(N_{\text{sim}} - 1)} \sum_{i=1}^{N_{\text{sim}}} \left(I(\mu_j^{(i)}) - \bar{I}(\mu_j)\right)^2. \tag{53}$$

*Rao-Blackwellized unbiased coverage estimator*

Frequency method checking is computationally expensive in nature because it fits a model on every mock data set. The situation deteriorates if the number of simulations or the size of data is large, or the estimation method is computationally demanding. Christiansen and Morris (1997) and Tang (2002) use a Rao-Blackwellized (RB) unbiased coverage estimator for the unknown coverage rate of each random effect, which is more efficient than the simple unbiased coverage estimator. For $j = 1, 2, \ldots, k$,

$$C_{A,\beta}(\mu_j) = \mathsf{E}\left(\bar{I}(\mu_j) \mid A, \beta\right) = \mathsf{E}\left[\frac{1}{N_{\text{sim}}} \sum_{i=1}^{N_{\text{sim}}} \mathsf{E}\left(I(\mu_j^{(i)}) \mid A, \beta, \boldsymbol{y}^{(i)}\right) \,\middle|\, A, \beta\right], \tag{54}$$

where the sample mean of the interior conditional expectations in (54) is the RB unbiased coverage estimator. Specifically,

$$\bar{I}^{RB}(\mu_j) = \frac{1}{N_{\text{sim}}} \sum_{i=1}^{N_{\text{sim}}} \mathsf{E}\left(I(\mu_j^{(i)}) \mid A, \beta, \boldsymbol{y}^{(i)}\right) \tag{55}$$

$$= \frac{1}{N_{\text{sim}}} \sum_{i=1}^{N_{\text{sim}}} \mathsf{P}\left(\mu_j^{(i)} \in (\hat{\mu}_{j,\,\text{low}}^{(i)}, \; \hat{\mu}_{j,\,upp}^{(i)}) \mid A, \beta, \boldsymbol{y}^{(i)}\right). \tag{56}$$

We can easily compute the conditional posterior probabilities in (56) using the cumulative density function of the Gaussian conditional posterior distribution of each random effect in (3). The variance of $\bar{I}^{RB}(\mu_j)$ does not exceed the variance of a simple unbiased coverage estimator, $\bar{I}(\mu_j)$ (Rao 1945; Blackwell 1947).

If one data set $\boldsymbol{y}^{(i)}$ is simulated for each set of random effects $\boldsymbol{\mu}^{(i)}$, the variance estimator below is an unbiased estimator of $\mathsf{VAR}(\bar{I}^{RB}(\mu_j))$. For $j = 1, 2, \ldots, k$,

$$\widehat{\mathsf{VAR}}(\bar{I}^{RB}(\mu_j)) \equiv \frac{1}{N_{\text{sim}}(N_{\text{sim}} - 1)} \sum_{i=1}^{N_{\text{sim}}} \left(\mathsf{E}(I(\mu_j^{(i)}) \mid A, \beta, \boldsymbol{y}^{(i)}) - \bar{I}^{RB}(\mu_j)\right)^2. \tag{57}$$

*Overall unbiased coverage estimator*

To summarize the frequency method checking, we report the overall unbiased coverage estimate and its variance estimate,

$$\bar{\bar{I}}^{RB} = \frac{1}{k} \sum_{j=1}^{k} \bar{I}^{RB}(\mu_j) \quad \text{and} \quad \widehat{\mathsf{VAR}}(\bar{\bar{I}}_{RB}) = \frac{1}{k^2} \sum_{j=1}^{k} \widehat{\mathsf{VAR}}(\bar{I}^{RB}(\mu_j)). \tag{58}$$

# 7. Examples

In this section, we demonstrate how **Rgbp** can be used to analyze three realistic data sets: medical profiling of 31 hospitals with Poisson distributed fatality counts; educational assessment of eight schools with normally distributed data; and evaluation of 18 baseball hitters with binomial success rates and one covariate. For each example, we construct 95% confidence intervals. Additional usages and options of the functions in **Rgbp** can be found in Section 8.

## 7.1. Poisson data: Known expected random effect

We analyze a data set of 31 hospitals in New York State consisting of the outcomes of the coronary artery bypass graft (CABG) surgery (Morris and Lysy 2012). The data set contains the number of deaths, $y$, for a specified period after CABG surgeries out of the total number of patients, $n$, receiving CABG surgeries in each hospital. A goal would be to obtain the point and interval estimates for the unknown true fatality rates (random effects) of 31 hospitals to evaluate each hospital's reliability on the CABG surgery (Morris and Christiansen 1995 use a similar Poisson model to handle these hospital profile data). We interpret the caseloads, $n$, as exposures and assume that the state-level fatality rate per exposure of this surgery is known, $\lambda_j^E = 0.03$ ($m = 0$). The following code can be used to load these data into R.

```
R> library("Rgbp")
R> data("hospital", package = "Rgbp")
R> y <- hospital$d
R> n <- hospital$n
```

The function `gbp` can then be used to fit a Poisson-gamma to the fatality rates in New York States with the expected random effect, $\lambda_j^E$, equal to 0.03.

```
R> p.output <- gbp(y, n, mean.PriorDist = 0.03, model = "poisson")
R> p.output
```

Summary for each unit (sorted by n):

|     | obs.mean | n | prior.mean | shrinkage | low.intv | post.mean | upp.intv | post.sd |
|-----|----------|-----|------------|-----------|----------|-----------|----------|---------|
| 1   | 0.0448   | 67  | 0.03       | 0.911     | 0.0199   | 0.0313    | 0.0454   | 0.00653 |
| 2   | 0.0294   | 68  | 0.03       | 0.910     | 0.0189   | 0.0299    | 0.0435   | 0.00631 |
| 3   | 0.0238   | 210 | 0.03       | 0.765     | 0.0185   | 0.0285    | 0.0407   | 0.00566 |
| 4   | 0.0430   | 256 | 0.03       | 0.728     | 0.0225   | 0.0335    | 0.0467   | 0.00619 |
| 5   | 0.0335   | 269 | 0.03       | 0.718     | 0.0208   | 0.0310    | 0.0432   | 0.00573 |
| 6   | 0.0438   | 274 | 0.03       | 0.714     | 0.0229   | 0.0339    | 0.0472   | 0.00621 |
| 7   | 0.0432   | 278 | 0.03       | 0.711     | 0.0228   | 0.0338    | 0.0469   | 0.00617 |
| 8   | 0.0136   | 295 | 0.03       | 0.699     | 0.0157   | 0.0250    | 0.0366   | 0.00534 |
| 9   | 0.0288   | 347 | 0.03       | 0.663     | 0.0200   | 0.0296    | 0.0410   | 0.00536 |
| 10  | 0.0372   | 349 | 0.03       | 0.662     | 0.0222   | 0.0325    | 0.0446   | 0.00571 |
| 11  | 0.0391   | 358 | 0.03       | 0.656     | 0.0228   | 0.0331    | 0.0454   | 0.00579 |
| 12  | 0.0177   | 396 | 0.03       | 0.633     | 0.0165   | 0.0255    | 0.0363   | 0.00506 |
| 13  | 0.0278   | 431 | 0.03       | 0.613     | 0.0200   | 0.0292    | 0.0400   | 0.00511 |
| 14  | 0.0249   | 441 | 0.03       | 0.608     | 0.0191   | 0.0280    | 0.0387   | 0.00502 |

| | | | | | | | | |
|---|---|---|---|---|---|---|---|---|
| 15 | 0.0273 | 477 | 0.03 | 0.589 | 0.0199 | 0.0289 | 0.0394 | 0.00499 |
| 16 | 0.0455 | 484 | 0.03 | 0.585 | 0.0256 | 0.0364 | 0.0491 | 0.00601 |
| 17 | 0.0304 | 494 | 0.03 | 0.580 | 0.0211 | 0.0302 | 0.0409 | 0.00506 |
| 18 | 0.0220 | 501 | 0.03 | 0.577 | 0.0180 | 0.0266 | 0.0369 | 0.00483 |
| 19 | 0.0277 | 505 | 0.03 | 0.575 | 0.0202 | 0.0290 | 0.0395 | 0.00494 |
| 20 | 0.0204 | 540 | 0.03 | 0.559 | 0.0173 | 0.0258 | 0.0358 | 0.00474 |
| 21 | 0.0284 | 563 | 0.03 | 0.548 | 0.0206 | 0.0293 | 0.0395 | 0.00485 |
| 22 | 0.0236 | 593 | 0.03 | 0.535 | 0.0187 | 0.0270 | 0.0369 | 0.00466 |
| 23 | 0.0150 | 602 | 0.03 | 0.532 | 0.0147 | 0.0230 | 0.0329 | 0.00466 |
| 24 | 0.0238 | 629 | 0.03 | 0.521 | 0.0188 | 0.0271 | 0.0368 | 0.00460 |
| 25 | 0.0204 | 636 | 0.03 | 0.518 | 0.0173 | 0.0254 | 0.0351 | 0.00455 |
| 26 | 0.0480 | 729 | 0.03 | 0.484 | 0.0286 | 0.0393 | 0.0516 | 0.00587 |
| 27 | 0.0306 | 849 | 0.03 | 0.446 | 0.0223 | 0.0303 | 0.0397 | 0.00445 |
| 28 | 0.0274 | 914 | 0.03 | 0.428 | 0.0208 | 0.0285 | 0.0374 | 0.00423 |
| 29 | 0.0213 | 940 | 0.03 | 0.421 | 0.0176 | 0.0249 | 0.0335 | 0.00407 |
| 30 | 0.0293 | 1193 | 0.03 | 0.364 | 0.0223 | 0.0296 | 0.0379 | 0.00397 |
| 31 | 0.0201 | 1340 | 0.03 | 0.338 | 0.0170 | 0.0235 | 0.0310 | 0.00360 |
| Mean | | 517 | 0.03 | 0.600 | 0.0201 | 0.0293 | 0.0403 | 0.00517 |

The output contains information about (from the left) the observed fatality rates $\bar{y}_j$, caseloads $n_j$, known expected random effect $\lambda_j^E$, shrinkage estimates $\hat{B}'_j$, lower bounds (2.5%) of posterior interval estimates $\hat{\lambda}_{j,\text{low}}$, posterior means $\hat{\mathsf{E}}(\lambda_j|\boldsymbol{y})$, upper bounds (97.5%) of posterior interval estimates $\hat{\lambda}_{j,\text{upp}}$, and posterior standard deviations $\widehat{\text{SD}}(\lambda_j|\boldsymbol{y})$ for random effects based on the assumed unconditional gamma posterior distributions in (44).

A function `summary` shows selective information about hospitals with minimum, median, and maximum exposures and the estimation result of the hyper-parameter $\alpha = -\log(r)$.

```
R> summary(p.output)
```

Main summary:

| | obs.mean | n | prior.mean | shrinkage | low.intv | post.mean |
|---|---|---|---|---|---|---|
| Unit with min(n) | 0.0448 | 67 | 0.03 | 0.911 | 0.0199 | 0.0313 |
| Unit with median(n) | 0.0455 | 484 | 0.03 | 0.585 | 0.0256 | 0.0364 |
| Unit with max(n) | 0.0201 | 1340 | 0.03 | 0.338 | 0.0170 | 0.0235 |
| Overall Mean | | 517 | 0.03 | 0.600 | 0.0201 | 0.0293 |

| upp.intv | post.sd |
|---|---|
| 0.0454 | 0.00653 |
| 0.0491 | 0.00601 |
| 0.0310 | 0.00360 |
| 0.0403 | 0.00517 |

Second-level Variance Component Estimation Summary:
alpha=log(A) for Gaussian or alpha=log(1/r) for binomial and Poisson data:

**Shrinkage plot**



**95 % Interval plot**

Groups sorted by the ascending order of n
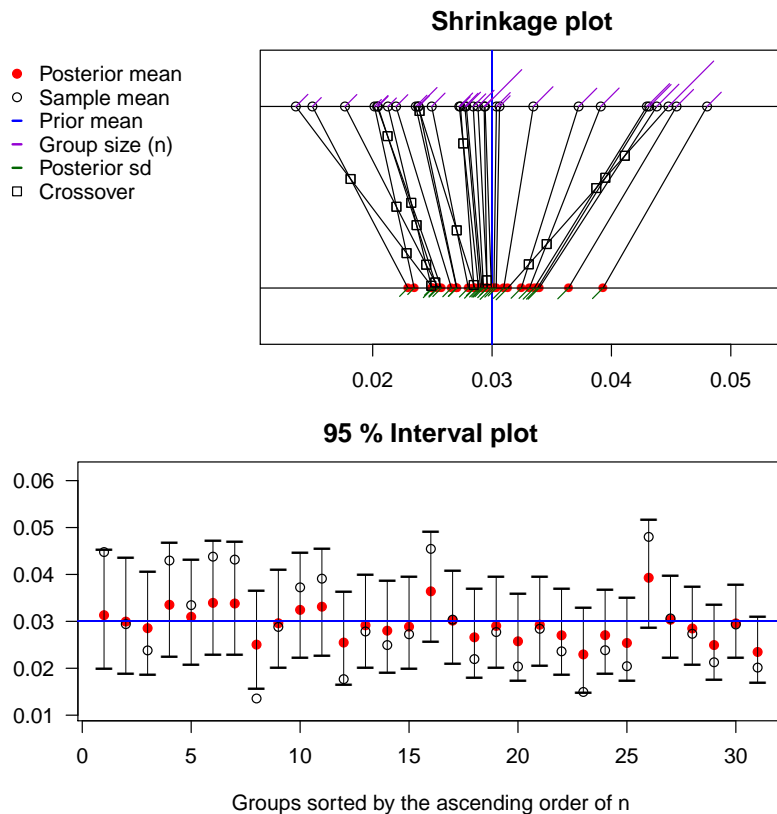
Figure 2: Shrinkage plot and 95% interval plot for fatality rates at 31 hospitals.

```
post.mode.alpha post.sd.alpha post.mode.r
         -6.53          0.576          684
```

The output of `summary` shows that $\hat{r} = \exp(6.53) = 684$, which is an indicator of how valuable and informative the second-level hierarchy is. This means that the 25 hospitals with caseload less than 684 patients shrink their sample means towards the prior mean (0.03) more than 50%. For example, the shrinkage estimate of the first hospital ($\hat{B}_1 = 0.911$) was calculated by 684 / (684 + 67), where 67 is its caseload ($n_1$). For this hospital, using more information from the conjugate prior distribution is an appropriate choice because the amount of observed information (67) is much less than the amount of state-level information (684).

To obtain a graphical summary, we use the function `plot`.

```
R> plot(p.output)
```

The shrinkage plot (Efron and Morris 1975; Morris and Lysy 2012) in the first panel of Figure 2 shows the regression towards the mean; the observed fatality rates, denoted by empty dots on the upper horizontal line, are shrinking towards the known expected random effect, denoted by a blue vertical line at 0.03, to the different extents. The red dots on the bottom line denote the estimated posterior means. Some hospitals' ranks have changed by shrinking more sharply towards 0.03 than the others. For example, an empty square at the crossing point of the two left-most lines (8th and 23rd hospitals on the list above) indicates

that the seemingly safest hospital in terms of the observed mortality rate is probably not the safest in terms of the estimated posterior mean accounting for the different caseloads of these two hospitals.

To be specific, their observed fatality rates ($y_j$, $j = 8, 23$) are 0.0136 and 0.0150 and caseloads ($n_j$, $j = 8, 23$) are 295 and 602, respectively. Considering solely the observed fatality rates may lead to an unfair comparison because the latter hospital handled twice the caseload. **Rgbp** accounts for this caseload difference, making the posterior mean for the random effect of the former hospital shrink toward the state-level mean ($\lambda_j^E = 0.03$) more strongly than that for the latter hospital.

The point estimates are not enough to evaluate hospital reliability because one hospital may have a lower point estimate but larger uncertainty (variance) than the other. The second plot of Figure 2 displays the 95% interval estimates. Each posterior mean (red dot) is between the sample mean (empty dot) and the known expected random effect (a blue horizontal line).

This 95% interval plot reveals that the 31st hospital has the lowest upper bound even though its point estimate ($\hat{\lambda}_{31} = 0.0235$) is slightly larger than that of the 23rd hospital ($\hat{\lambda}_{23} = 0.0230$). The observed mortality rates for these two hospitals ($y_j, j = 23, 31$) are 0.0150 and 0.0201 and the caseloads ($n_j, j = 23, 31$) are 602 and 1340 each. The 31st hospital has twice the caseload, which leads to borrowing less information from the New York State-level hierarchy (or shrinking less toward the state-level mean, 0.03) with smaller variance. Based on the point and interval estimates, the 31st hospital seems a better choice than the 23rd hospital. (Note that this conclusion is based on the data, assuming no covariate information about the overall case difficulties in each hospital. A more reliable analysis must take into account all the possible covariate information and instead of our Poisson model, we recommend using our binomial model to account for covariate information.)

Next, we perform frequency method checking to question how reliable the estimation procedure is, assuming $r$ equals its estimated value, $\hat{r} = 683.53$. The function `coverage` generates synthetic data sets starting with the estimated value of $r$ as a generative value. For reference, we could designate other generative values of $r$ and $\lambda_j^E$ by adding two arguments, `A.or.r` and `mean.PriorDist`, into the code below, see Section 8.1 for details.

```
R> p.coverage <- coverage(p.output, nsim = 1000)
```

In Figure 3, the black horizontal line at 0.95 represents the nominal confidence level and the red circles indicate RB unbiased coverage estimates, $\bar{I}^{RB}(\lambda_j)$ for $j = 1, 2, \ldots, 31$. The overall unbiased coverage estimate across all the hospitals, $\bar{\bar{I}}^{RB}$ in (58), is 0.955. None of the RB unbiased coverage estimates for the 31 hospitals are less than 0.95 regardless of their caseloads, which range from 67 for hospital 1 to 1,340 for hospital 31. This result shows that the interval estimate for this particular data set adequately achieves a 95% confidence level if $r = \hat{r}$.

The following code provides 31 RB unbiased coverage estimates and their standard errors (the output is omitted for space reasons).

```
R> p.coverage$coverageRB
R> p.coverage$se.coverageRB
```

The code below produces 31 simple unbiased coverage estimates and their standard errors.
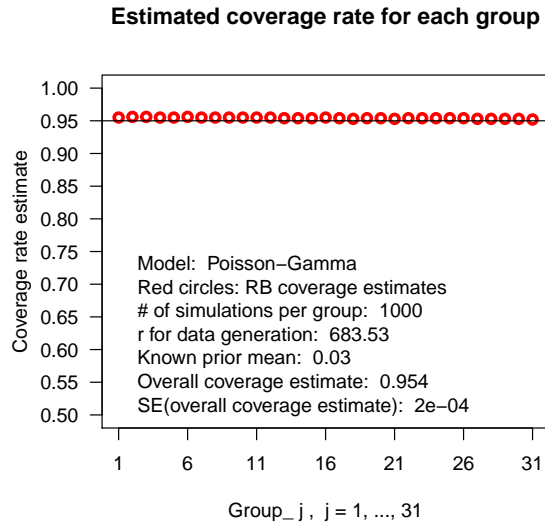
**Estimated coverage rate for each group**



Figure 3: Coverage plot via frequency method checking for 31 hospitals.

```
R> p.coverage$coverageS
R> p.coverage$se.coverageS
```

It turns out that the variance estimate of the RB unbiased coverage estimate for the first hospital $(0.0016^2)$ is about 19 times smaller than that of the simple one $(0.0070^2)$. This means that the RB unbiased coverage estimates based on 1,000 simulations $(N_{\text{sim}})$ are as precise as the simple unbiased coverage estimates based on 19,000 simulations in terms of estimating the coverage probability for the first hospital, $C_{r,\lambda^{\text{E}}}(\lambda_1)$.

### 7.2. Gaussian data: Unknown expected random effect and no covariates

The Education Testing Service conducted randomized experiments in eight separate schools (groups) to test whether students' (units') SAT scores are affected by coaching. The data set contains the estimated coaching effects on SAT scores ($y_j$, $j = 1, \ldots, 8$) and standard errors ($V_j^{0.5}$, $j = 1, \ldots, 8$) of the eight schools (Rubin 1981). These data are contained in the package and can be loaded into R as follows.

```
R> data("schools", package = "Rgbp")
R> y  <- schools$y
R> se <- schools$se
```

Due to the nature of the test each school's coaching effect has an approximately normal sampling distribution with approximately known sampling variances, based on large sample consideration. At the second hierarchy, the mean for each school is assumed to be drawn from a common normal distribution ($m = 1$).

```
R> g.output <- gbp(y, se, model = "gaussian")
R> g.output
```

Summary for each group (sorted by the descending order of se):

|      | obs.mean | se   | prior.mean | shrinkage | low.intv | post.mean | upp.intv | post.sd |
|------|----------|------|------------|-----------|----------|-----------|----------|---------|
| 8    | 12.00    | 18.0 | 8.17       | 0.734     | -10.21   | 9.19      | 29.9     | 10.23   |
| 3    | -3.00    | 16.0 | 8.17       | 0.685     | -17.13   | 4.65      | 22.5     | 10.10   |
| 1    | 28.00    | 15.0 | 8.17       | 0.657     | -2.32    | 14.98     | 38.8     | 10.56   |
| 4    | 7.00     | 11.0 | 8.17       | 0.507     | -8.78    | 7.59      | 23.6     | 8.26    |
| 6    | 1.00     | 11.0 | 8.17       | 0.507     | -13.03   | 4.63      | 20.1     | 8.44    |
| 2    | 8.00     | 10.0 | 8.17       | 0.459     | -7.25    | 8.08      | 23.4     | 7.81    |
| 7    | 18.00    | 10.0 | 8.17       | 0.459     | -1.29    | 13.48     | 30.8     | 8.18    |
| 5    | -1.00    | 9.0  | 8.17       | 0.408     | -13.30   | 2.74      | 16.7     | 7.63    |
| Mean |          | 12.5 | 8.17       | 0.552     | -9.16    | 8.17      | 25.7     | 8.90    |

This output from `gbp` summarizes the results. In this Gaussian model the amount of shrinkage for each unit is governed by the shrinkage factor, $B_j = V_j/(V_j + A)$. As such, schools whose variation within the school ($V_j$) is less than the between-school variation ($A$) will shrink greater than 50%. The results provided by `gpb` suggests that there is little evidence that the training provided much added benefit due to the fact that every school's 95% posterior interval contains zero. In the case where the number of groups is large **Rgbp** provides a summary feature:

```
R> summary(g.output)
```

Main summary:

|                       | obs.mean | se   | prior.mean | shrinkage | low.intv | post.mean |
|-----------------------|----------|------|------------|-----------|----------|-----------|
| Unit with min(se)     | -1.00    | 9.0  | 8.17       | 0.408     | -13.30   | 2.74      |
| Unit with median(se)1 | 1.00     | 11.0 | 8.17       | 0.507     | -13.03   | 4.63      |
| Unit with median(se)2 | 7.00     | 11.0 | 8.17       | 0.507     | -8.78    | 7.59      |
| Unit with max(se)     | 12.00    | 18.0 | 8.17       | 0.734     | -10.21   | 9.19      |
| Overall Mean          |          | 12.5 | 8.17       | 0.552     | -9.16    | 8.17      |

| upp.intv | post.sd |
|----------|---------|
| 16.7     | 7.63    |
| 20.1     | 8.44    |
| 23.6     | 8.26    |
| 29.9     | 10.23   |
| 25.7     | 8.90    |

Second-level Variance Component Estimation Summary:
alpha=log(A) for Gaussian or alpha=log(1/r) for Binomial and Poisson data:

| post.mode.alpha | post.sd.alpha | post.mode.A |
|-----------------|---------------|-------------|
| 4.77            | 1.14          | 118         |

Regression Summary:

|       | estimate | se   | z.val | p.val |
|-------|----------|------|-------|-------|
| beta1 | 8.168    | 5.73 | 1.425 | 0.154 |

**Shrinkage plot**



**95 % Interval plot**
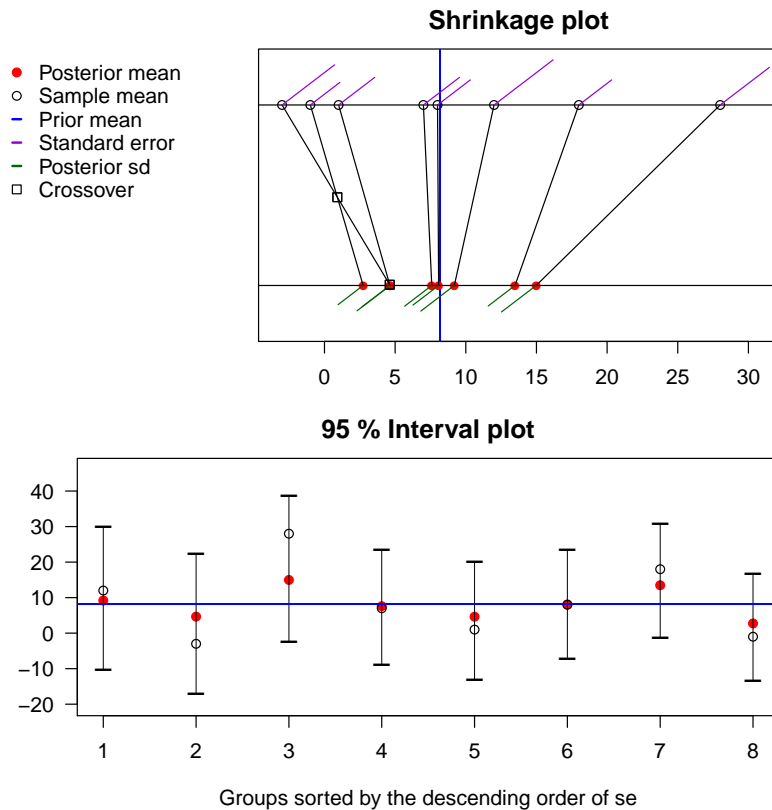
Groups sorted by the descending order of se

Figure 4: Shrinkage plot and 95% interval plot for eight schools.

The summary provides results regarding the second level hierarchy parameters. It can be seen that the estimate of the expected random effect, $\mu^E = \beta_1$ (beta1), is not significantly different from zero suggesting that there is no effect of the coaching program on SAT math scores.

**Rgbp** also provides functionality to plot the results of the analysis as seen in Figure 4. Plotting the results provides a visual aid to understanding but is only largely beneficial when the number of groups ($k$) is small.

```
R> plot(g.output)
```

The frequency method checking procedures generates new pseudo-data from our assumed model. Unless otherwise specified, the procedure fixes the hyper-parameter values at their estimates ($\hat{A}$ and $\hat{\beta}_1$ in this example) and then simulates random effects $\mu_j$ for each group $j$. The model is then estimated and this is repeated an $N_{\text{sim}}$ (nsim) number of times to estimate the coverage probabilities of the procedure.

```
R> g.coverage <- coverage(g.output, nsim = 1000)
```

As seen in Figure 5 the desired 95% confidence level, denoted by a black horizontal line at 0.95, is achieved for each school in this example. All the coverage estimates depend on the chosen generative values of $A$ and $\beta_1$, and the assumption that the model is valid.

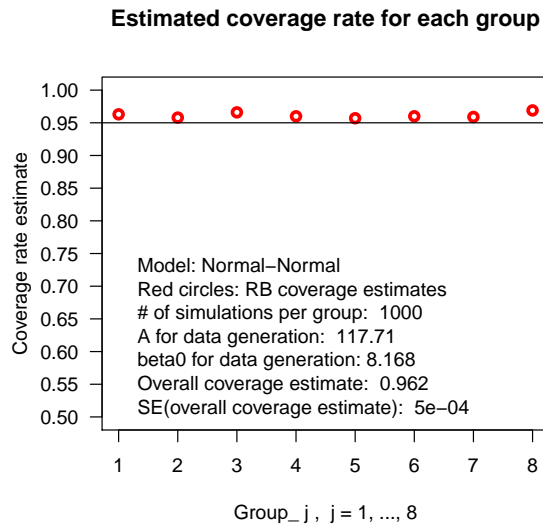**Estimated coverage rate for each group**



Figure 5: Coverage plot via frequency method checking for 8 schools.

In addition, the RB unbiased coverage estimate and its standard error for each school can be calculated with the command below.

```
R> g.coverage$coverageRB
```

```
[1] 0.963 0.958 0.966 0.960 0.957 0.960 0.959 0.969
```

```
R> g.coverage$se.coverageRB
```

```
[1] 0.0015 0.0013 0.0015 0.0013 0.0011 0.0013 0.0013 0.0013
```

### 7.3. Binomial data: Unknown expected random effects and one covariate

The data of 18 major league baseball players contain the batting averages through their first 45 official at-bats of the 1970 season (Efron and Morris 1975). A binary covariate is created that is equal to the value one if a player is an outfielder and zero otherwise. The data can be loaded into R with the following code.

```
R> data("baseball", package = "Rgbp")
R> y <- baseball$Hits
R> n <- baseball$At.Bats
R> x <- ifelse(baseball$Position == "fielder", 1, 0)
```

Conditional on the unknown true batting average (random effect) of each player it is assumed that the at-bats are independent and therefore, $y_j \mid p_j \sim \text{Binomial}(45, p_j)$ independently for $j = 1, \ldots, 18$. Our goal is to obtain point and interval estimates of each random effect whilst considering the additional information on whether the player is an outfielder or not. The function `gbp` provides a way to incorporate such covariate information seamlessly into the model so that the regression towards the mean occurs within outfielders and non-outfielders separately.

```
R> b.output <- gbp(y, n, x, model = "binomial")
R> b.output
```

Summary for each unit (sorted by n):

|      | obs.mean | n  | X1  | prior.mean | shrinkage | low.intv | post.mean | upp.intv | post.sd |
|------|----------|----|-----|------------|-----------|----------|-----------|----------|---------|
| 1    | 0.400    | 45 | 1.0 | 0.310      | 0.715     | 0.248    | 0.335     | 0.429    | 0.0462  |
| 2    | 0.378    | 45 | 1.0 | 0.310      | 0.715     | 0.244    | 0.329     | 0.420    | 0.0448  |
| 3    | 0.356    | 45 | 1.0 | 0.310      | 0.715     | 0.240    | 0.323     | 0.411    | 0.0437  |
| 4    | 0.333    | 45 | 1.0 | 0.310      | 0.715     | 0.236    | 0.316     | 0.403    | 0.0429  |
| 5    | 0.311    | 45 | 1.0 | 0.310      | 0.715     | 0.230    | 0.310     | 0.396    | 0.0424  |
| 6    | 0.311    | 45 | 0.0 | 0.233      | 0.715     | 0.179    | 0.256     | 0.341    | 0.0415  |
| 7    | 0.289    | 45 | 0.0 | 0.233      | 0.715     | 0.175    | 0.249     | 0.331    | 0.0400  |
| 8    | 0.267    | 45 | 0.0 | 0.233      | 0.715     | 0.171    | 0.243     | 0.323    | 0.0388  |
| 9    | 0.244    | 45 | 0.0 | 0.233      | 0.715     | 0.166    | 0.237     | 0.315    | 0.0380  |
| 10   | 0.244    | 45 | 1.0 | 0.310      | 0.715     | 0.210    | 0.291     | 0.379    | 0.0432  |
| 11   | 0.222    | 45 | 0.0 | 0.233      | 0.715     | 0.161    | 0.230     | 0.308    | 0.0377  |
| 12   | 0.222    | 45 | 0.0 | 0.233      | 0.715     | 0.161    | 0.230     | 0.308    | 0.0377  |
| 13   | 0.222    | 45 | 0.0 | 0.233      | 0.715     | 0.161    | 0.230     | 0.308    | 0.0377  |
| 14   | 0.222    | 45 | 1.0 | 0.310      | 0.715     | 0.202    | 0.285     | 0.375    | 0.0441  |
| 15   | 0.222    | 45 | 1.0 | 0.310      | 0.715     | 0.202    | 0.285     | 0.375    | 0.0441  |
| 16   | 0.200    | 45 | 0.0 | 0.233      | 0.715     | 0.155    | 0.224     | 0.302    | 0.0377  |
| 17   | 0.178    | 45 | 0.0 | 0.233      | 0.715     | 0.148    | 0.218     | 0.297    | 0.0381  |
| 18   | 0.156    | 45 | 0.0 | 0.233      | 0.715     | 0.140    | 0.211     | 0.292    | 0.0389  |
| Mean |          | 45 | 0.4 | 0.267      | 0.715     | 0.191    | 0.267     | 0.351    | 0.0410  |

The shrinkage estimates are the same for all players because all players have the same 45 at-bats, i.e., the same amount of the observed information.

```
R> summary(b.output)
```

Main summary:

|                            | obs.mean | n  | X1    | prior.mean | shrinkage | low.intv |
|----------------------------|----------|----|-------|------------|-----------|----------|
| Unit with min(obs.mean)    | 0.156    | 45 | 0.000 | 0.233      | 0.715     | 0.140    |
| Unit with median(obs.mean)1 | 0.244   | 45 | 0.000 | 0.233      | 0.715     | 0.166    |
| Unit with median(obs.mean)2 | 0.244   | 45 | 1.000 | 0.310      | 0.715     | 0.210    |
| Unit with max(obs.mean)    | 0.400    | 45 | 1.000 | 0.310      | 0.715     | 0.248    |
| Overall Mean               |          | 45 | 0.444 | 0.267      | 0.715     | 0.191    |

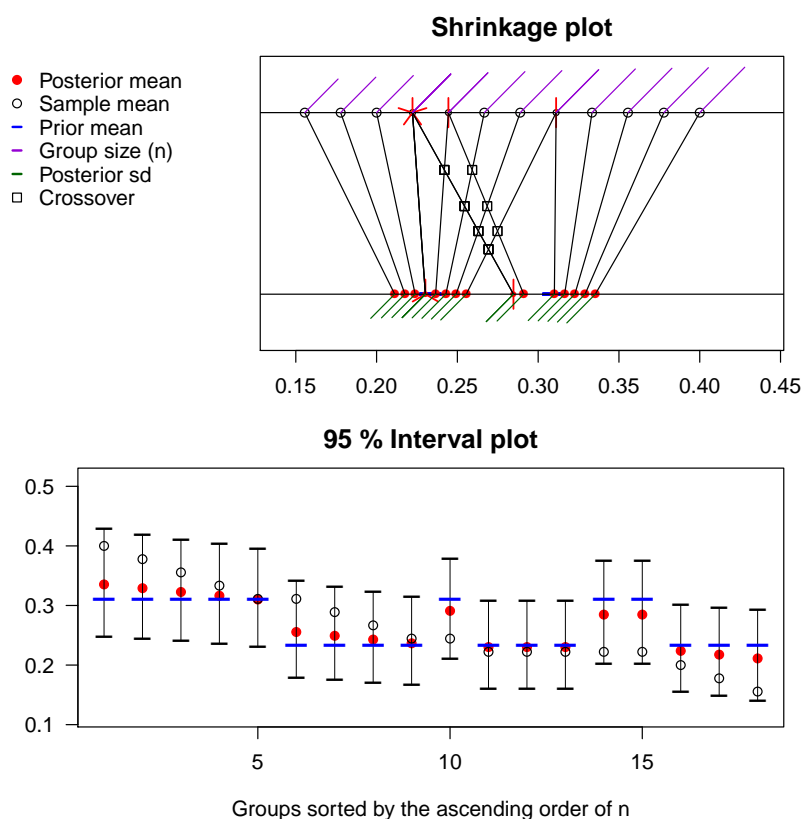| post.mean | upp.intv | post.sd |
|-----------|----------|---------|
| 0.211     | 0.292    | 0.0389  |
| 0.237     | 0.315    | 0.0380  |
| 0.291     | 0.379    | 0.0432  |
| 0.335     | 0.429    | 0.0462  |
| 0.267     | 0.351    | 0.0410  |

Figure 6: Shrinkage plot and 95% interval plot for 18 baseball players.

```
Second-level Variance Component Estimation Summary:
alpha=log(A) for Gaussian or alpha=log(1/r) for Binomial and Poisson data:

post.mode.alpha post.sd.alpha post.mode.r
          -4.73         0.957           113


Regression Summary:

      estimate    se  z.val p.val
beta1   -1.194 0.131 -9.129 0.000
beta2    0.389 0.187  2.074 0.038
```

The regression coefficient for the outfielder indicator is significant, considering that the $p$ value for $\hat{\beta}_2$ (`beta2`) is 0.038. It means that the two estimates for the expected random effects for the outfielders and infielders are significantly different. Also, the positive sign of $\hat{\beta}_2$ indicates that the population batting average for outfielders tends to be higher than that for infielders. The estimated odds ratio is $\exp(0.389) = 1.48$.

```
R> plot(b.output)
```

**Estimated coverage rate for each group**
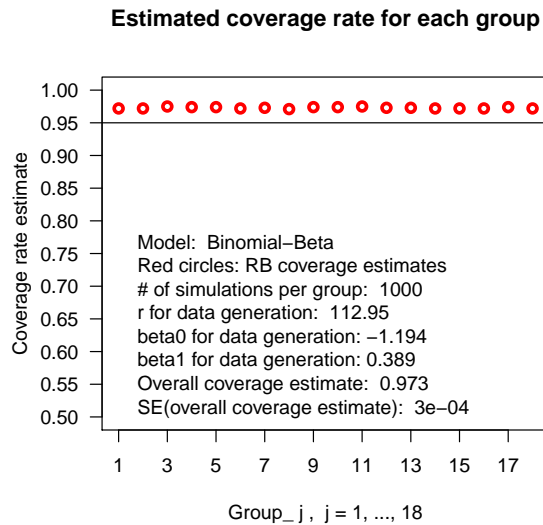


Figure 7: Coverage plot via frequency method checking for 18 players.

The shrinkage plot in Figure 6 shows that the observed batting averages (empty dots) on the upper horizontal line shrink towards the two expected random effects, 0.233 and 0.310. The short red line symbols near some empty dots are for when two or more points have the same mean and are plotted over each other. For example, five players (from the 11th player to the 15th) have the same batting average, 0.222, and at this point on the upper horizontal line, there are short red lines toward five directions.

The 95% interval plot in Figure 6 shows the range of true batting average for each player, which clarifies the regression towards the mean within two groups. The 10th, 14th, and 15th players, for example, are outfielders but their observed batting averages are far lower than the first five outfielders. This can be attributed to their bad luck because their observed batting averages are close to the lower bounds of their interval estimates. The regression towards the mean indicates that their batting averages shrink towards the expected random effect of outfielders (0.310) in the long run.

To check the level of trust in these interval estimates, we proceed to frequency method checking by assuming the estimates, 112.95 for $\hat{r}$ and $(-1.194, 0.389)$ for $\hat{\boldsymbol{\beta}}$, are the generative values.

```
R> b.coverage <- coverage(b.output, nsim = 1000)
```

In Figure 7, the estimated coverage probabilities for random effects are beyond 0.95, conservatively satisfying the 95% confidence level if $r = \hat{r}$ and $\boldsymbol{\beta} = \hat{\boldsymbol{\beta}}$. The overall unbiased coverage estimate across all the players is 0.972.

We can check the RB unbiased coverage estimates and their standard errors for each player;

```
R> b.coverage$coverageRB
```

```
 [1] 0.972 0.972 0.975 0.974 0.974 0.972 0.973 0.971 0.974 0.974 0.975 0.973
[13] 0.973 0.972 0.972 0.972 0.974 0.972
```

```
R> b.coverage$se.coverageRB
```

```
 [1] 0.0012 0.0012 0.0009 0.0010 0.0011 0.0011 0.0010 0.0013 0.0009 0.0011
[11] 0.0010 0.0010 0.0010 0.0012 0.0011 0.0012 0.0010 0.0011
```

If we want to draw 2,000 posterior samples of random effects and hyper-parameters from their full posterior distribution via the A-R method, we use the following R code.

```
R> b.output <- gbp(y, n, x, model = "binomial", n.AR = 2000)
```

The sampling result saved in `b.output` consists of 8,000 weights (`b.output$weight`), 2,000 posterior samples of $\alpha$ (`b.output$alpha`), a $2,000 \times 2$ matrix of $\boldsymbol{\beta}$ (`b.output$beta`) each column of which corresponds to 2,000 posterior samples of each regression coefficient, and a k $\times$ 2,000 matrix of random effects (`b.output$p`) each row of which has posterior samples of each random effect.

If we run the frequency method checking using this sampling result, `b.output`, obtained via the A-R method, the $N_{\text{sim}}$ simulations also run the A-R method each time.

# 8. Usage of functions in Rgbp

In this section, we describe more specific usage with various options of the two main functions of **Rgbp**, i.e., `gbp` for model fitting and `coverage` for frequency method checking.

## 8.1. Model fitting

The function `gbp` creates an S3 object 'gbp' for which three generic functions `plot`, `print`, and `summary` are defined.

There are two cases according to whether covariates are available or not. When no covariates are available, the function `gbp` requires fitting an intercept term or designating known values of the expected random effects, i.e., the intercept term must be either estimated or known. The default of `gbp` is to fit an intercept term. The value(s) of the known expected random effect(s) can be assigned through an optional argument `mean.PriorDist`. Note that `gbp` can fit the Poisson model only when the values of expected random effects, $\lambda_j^E$, are known. The usage of `gbp` to fit each model without any covariates is

```
R> g.output <- gbp(y, se.or.n, model = "gaussian")
R> b.output <- gbp(y, se.or.n, model = "binomial")
R> p.output <- gbp(y, se.or.n, mean.PriorDist, model = "poisson")
```

The argument `y` is a vector of $k$ observed sample means for the Gaussian model, $k$ observed numbers of successful outcomes for the binomial model, and $k$ observed outcome counts for the Poisson model. The argument `se.or.n` is a vector of $k$ standard errors of each sample mean for the Gaussian model, $k$ numbers of trials for the binomial model, and $k$ exposures for the Poisson model. The argument `mean.PriorDist` is either a constant (if all the known expected random effects are the same) or a vector of $k$ known expected random effects.

If covariate information for each group is available, users can fit the Gaussian and binomial models, using the following codes.

```
R> g.output <- gbp(y, se.or.n, X, model = "gaussian")
R> b.output <- gbp(y, se.or.n, X, model = "binomial")
```

The argument X is a matrix of covariate(s) each column of which corresponds to one covariate for $k$ groups. For example, if users have two covariates for each group, the argument X must be a k $\times$ 2 matrix to estimate three regression coefficients $\boldsymbol{\beta} = (\beta_1, \beta_2, \beta_3)$ including an intercept term, $\beta_1$, as a default. If users do not want to include an intercept term ($\beta_1 = 0$), estimating two regression coefficients for the two covariates, users can add an optional argument `intercept` as follows.

```
R> g.output <- gbp(y, se.or.n, X, model = "gaussian", intercept = FALSE)
```

The function `gbp` contains more optional arguments. The argument `confidence.lvl`, whose default value is 0.95, sets the confidence level, producing $100 \times$ `confidence.lvl`% interval estimates for the random effects. For the Gaussian model, setting the argument `normal.CI` to `TRUE` lets `gbp` use a normal approximation to the unconditional posterior distribution of the random effect (Morris and Tang 2011). The default value of `normal.CI` is `FALSE` for the skewed-normal approximation (Kelly 2014).

The function `gbp` uses the A-R method to fit the binomial model if users assign the desired number of posterior samples ($N$ in (51)) through the argument `n.AR`; its default value is zero. There are several arguments related to the A-R method. The argument `n.AR.factor` determines how many trial samples the method draws; its default value is four, meaning that the function `gbp` draws $4 \times$ `n.AR` trial samples and accepts or rejects them. The argument `trial.scale` is $\psi$ determining the scale parameter $\sigma$ of the skewed-$t$ distribution; its default value is 1.3. The argument `save.result` indicates whether `gbp` saves the whole posterior samples of the random effects and hyper-parameters; its default value is `TRUE`. The two arguments `t` and `u`, taking on non-negative and positive values, respectively, allow users to choose the joint hyper-prior density function, $f(r, \boldsymbol{\beta}) \propto 1/(t+r)^{u+1}$; the default values for `t` and `u` are 0 and 1, respectively, for the joint hyper-prior density function specified in (9).

For example, when there are two covariates, the following code produces 2,000 posterior samples of random effects and those of hyper-parameters, $r$ and $\boldsymbol{\beta}_{(3\times 1)}$ including an intercept term, via the A-R method with 8,000 trial samples.

```
R> b.output <- gbp(y, se.or.n, X, model = "binomial", n.AR = 2000)
```

The object `b.output` above contains 8,000 weights (`b.output$weight`), 2,000 posterior samples of $\alpha$ (`b.output$alpha`), a 2,000 $\times$ 3 matrix of $\beta$ (`b.output$beta`) each column of which corresponds to 2,000 posterior samples of each regression coefficient, and a k $\times$ 2,000 matrix of random effects (`b.output$p`) each row of which has posterior samples of each random effect.

The S3 object 'gbp' benefits from three generic functions, `print`, `summary`, and `plot`. The estimation result for all the random effects appears if users type the 'gbp' object in the R console, which plays the same role of the function `print` with its default argument `sort = TRUE`. When the argument `sort` is set to `TRUE`, the function `print` displays the estimation result for all the groups in the ascending order of $n$ for the binomial and Poisson model and in the descending order of standard errors for the Gaussian model. When the argument `sort` is `FALSE`, the estimation result is returned in the order of data input.

```
R> b.output
R> print(b.output, sort = FALSE)
```

The function `summary` prints a detailed estimation result, including the estimation result for the hyper-parameters, $A$ (or $r$) and $\beta$.

```
R> summary(b.output)
```

The function `plot` draws a shrinkage plot and $100 \times$ `confidence.lvl`% interval plot for random effects, see Figures 2, 4, or 6 for example. Its default argument `sort = TRUE` displays the $100 \times$ `confidence.lvl`% interval plot in the ascending order of $n$ for the binomial and Poisson model and the descending order of standard errors for the Gaussian model. When the argument `sort` is set to `FALSE` the $100 \times$ `confidence.lvl`% interval plot is displayed in the order of data input.

```
R> plot(b.output)
R> plot(b.output, sort = FALSE)
```

## 8.2. Frequency method checking

The function `coverage` conducts the frequency method checking. It estimates the coverage properties for our estimators of the random effects at a particular value of the hyper-parameters by averaging the coverage over many simulated data sets. The basic usage of `coverage` needs a 'gbp' object, such as `b.output` above, as the first argument;

```
R> cov <- coverage(b.output, nsim = 1000)
```

The argument `nsim` sets the number of simulations, $N_{\text{sim}}$, defined in Section 6.1. If users do not assign values of the hyper-parameters through the arguments `A.or.r` and `reg.coef`, then the function `coverage` automatically sets the estimated posterior modes of hyper-parameters saved in the 'gbp' object (or their posterior medians if the acceptance-rejection method for the binomial model is used) as the generative values of hyper-parameters. If users want to conduct the frequency method checking with different generated values of hyper-parameters, for example, $r = 100$ and $\beta = (2, 5)^\top$ when one covariate was used with an intercept term, then users can specify them via the arguments `A.or.r` and `reg.coef`;

```
R> cov <- coverage(b.output, A.or.r = 100, reg.coef = c(2, 5), nsim = 1000)
```

When users fit a model via `gbp` with known expected random effects, e.g., a Poisson model with known values of $\{\lambda_1, \lambda_2, \dots, \lambda_k\}$, `coverage` conducts the frequency method checking based on these known values as a default. However, users may want to conduct the frequency method checking with different known values of the expected random effects. For example, if users want to try a different value of the expected random effect, e.g., $\lambda_j = 30$ (or can also be a vector of different values), the argument `mean.PriorDist` is added as follows.

```
R> cov <- coverage(p.output, mean.PriorDist = 30, nsim = 1000)
```

The resulting frequency method checking is based on the estimated posterior mode of $r$ (because it is not specified through `A.or.r`) and the newly specified value of the expected random effect, 30, as its known value.

Though the function `coverage` does not produce an S3 object, the result of `coverage` contains various numerical details; $k$ RB coverage estimates (`cov$coverageRB`) and their standard errors (`cov$se.coverageRB`), overall unbiased coverage estimate (`cov$overall.coverageRB`) and its standard error (`cov$se.overall.coverageRB`), etc.

A coverage plot summarizing the result of `coverage` automatically appears. If the result is saved in a variable such as `cov` above, then users can recall the coverage plot, using the function `coverage.plot`.

```
R> coverage.plot(cov)
```

# 9. Discussion

**Rgbp** is an R package for estimating and validating two-level Gaussian, Poisson, and binomial hierarchical models. The package aims to provide a procedure that is computationally efficient with good frequency properties and includes frequency method checking functionality to examine repeated sampling properties and to test that the method is valid at specified hyper-parameter values.

As an alternative to other maximization based estimation methods such as MLE and REML, **Rgbp** provides approximate point and interval estimates of parameters via ADM. Using the ADM approach, with our specified choice of priors, protects from cases of overshrinkage and undercoverage from which the aforementioned methods suffer (Morris 1988b).

A benefit of **Rgbp** is that it produces non-random output (except the A-R method for the binomial model) and so results are easily reproduced and compared across studies. In addition to being a stand-alone analysis tool the package can be used as an aid in a broader estimation procedure. For example, by checking the similarity of output of **Rgbp** and that of another estimation procedure such as MCMC (Markov chain Monte Carlo), the package can be used as a confirmatory tool to check whether the alternative procedure has been programmed correctly. In addition, the parameter estimates obtained via **Rgbp** can be used to initialize an MCMC sampling procedure thus decreasing time to convergence. Lastly, due to its speed and ease of use, **Rgbp** can be used as a method of preliminary data analysis. Such results may tell statisticians and practitioners alike whether a more intensive method in terms of implementation and computational time, such as MCMC, is needed.

# Acknowledgments

# References

Alam M, Rönnegård L, Shen X (2015). **hglm**: *Hierarchical Generalized Linear Models*. URL https://CRAN.R-project.org/package=hglm.

Albert JH (1988). "Computational Methods Using a Bayesian Hierarchical Generalized Linear Model." *Journal of the American Statistical Association*, **83**(404), 1037–1044. `doi:10.1080/01621459.1988.10478698`.

Azzalini A (1985). "A Class of Distributions Which Includes the Normal Ones." *Scandinavian Journal of Statistics*, **12**(2), 171–178.

Blackwell D (1947). "Conditional Expectation and Unbiased Sequential Estimation." *The Annals of Mathematical Statistics*, pp. 105–110. `doi:10.1214/aoms/1177730497`.

Brillinger DR (1969). "The Calculation of Cumulants via Conditioning." *The Annals of the Institute of Statistical Mathematics*, **21**(1), 215–218. `doi:10.1007/bf02532246`.

Christiansen C, Morris C (1996). "Fitting and Checking a Two-Level Poisson Model: Modeling Patient Mortality Rates in Heart Transplant Patients." In D Berry, D Stangl (eds.), *Bayesian Biostatistics*, pp. 467–501. Chapman & Hall/CRC.

Christiansen C, Morris C (1997). "Hierarchical Poisson Regression Modeling." *Journal of the American Statistical Association*, **92**(438), 618–632. `doi:10.2307/2965709`.

Daniels MJ (1999). "A Prior for the Variance in Hierarchical Models." *Canadian Journal of Statistics*, **27**(3), 567–578. `doi:10.2307/3316112`.

Dean CB (1992). "Testing for Overdispersion in Poisson and Binomial Regression Models." *Journal of the American Statistical Association*, **87**(418), 451–457. `doi:10.1080/01621459.1992.10475225`.

Efron B, Morris C (1975). "Data Analysis Using Stein's Estimator and its Generalizations." *Journal of the American Statistical Association*, **70**(350), 311–319. `doi:10.1080/01621459.1975.10479864`.

Gelman A, Carlin JB, Stern HS, Rubin DB (2014). *Bayesian Data Analysis.* 2nd edition. Taylor & Francis.

Ghosh M, Rao JNK (1994). "Small Area Estimation: An Appraisal." *Statistical Science*, **9**(1), 55–76. `doi:10.1214/ss/1177010647`.

Jones MC, Faddy MJ (2003). "A Skew Extension of the $t$-Distribution, with Applications." *Journal of the Royal Statistical Society B*, **65**(1), 159–174. `doi:10.1111/1467-9868.00378`.

Kass RE, Steffey D (1989). "Approximate Bayesian Inference in Conditionally Independent Hierarchical Models (Parametric Empirical Bayes Models)." *Journal of the American Statistical Association*, **84**(407), 717–726. `doi:10.1080/01621459.1989.10478825`.

Kelly J (2014). *Advances in the Normal-Normal Hierarchical Model.* Ph.D. thesis, Harvard University.

Lee Y, Nelder JA (1996). "Hierarchical Generalized Linear Models." *Journal of the Royal Statistical Society B*, **58**(4), 619–678.

Morris C (1988a). "Approximating Posterior Distributions and Posterior Moments." In JM Bernardo, MH DeGroot, DV Lindley, AFM Smith (eds.), *Bayesian Statistics 3*, pp. 327–344. Oxford University Press.

Morris C (1988b). "Determining the Accuracy of Bayesian Empirical Bayes Estimates in the Familiar Exponential Families." In SS Gupta, JO Berger (eds.), *Statistical Decision Theory and Related Topics IV*, pp. 251–263. Springer-Verlag.

Morris C, Christiansen C (1995). "Hierarchical Models for Ranking and for Identifying Extremes, with Application." In J Bernardo, J Berger, A Dawid, A Smith (eds.), *Bayesian Statistics 5*, pp. 227–296. Oxford University Press, New York.

Morris C, Lysy M (2012). "Shrinkage Estimation in Multilevel Normal Models." *Statistical Science*, **27**(1), 115–134. `doi:10.1214/11-sts363`.

Morris C, Tang R (2011). "Estimating Random Effects via Adjustment for Density Maximization." *Statistical Science*, **26**(2), 271–287. `doi:10.1214/10-sts349`.

Morris CN (1983). "Natural Exponential Families with Quadratic Variance Functions: Statistical Theory." *The Annals of Statistics*, **11**(2), 515–529. `doi:10.1214/aos/1176346158`.

Rao CR (1945). "Information and Accuracy Attainable in the Estimation of Statistical Parameters." *Bulletin of the Calcutta Mathematical Society*, **37**(3), 81–91.

Rao JNK (2003). *Small Area Estimation.* John Wiley & Sons.

Robert C, Casella G (2013). *Monte Carlo Statistical Methods.* Springer-Verlag.

Rubin DB (1981). "Estimation in Parallel Randomized Experiments." *Journal of Educational Statistics*, **6**(4), 377–401. `doi:10.3102/10769986006004377`.

Skellam JG (1948). "A Probability Distribution Derived from the Binomial Distribution by Regarding the Probability of Success as Variable between the Sets of Trials." *Journal of the Royal Statistical Society B*, **10**(2), 257–261.

Tak H, Morris C (2017). "Data-Dependent Posterior Propriety of a Bayesian Beta-Binomial-Logit Model." *Bayesian Analysis*, **12**(2), 533–555. `doi:10.1214/16-ba1012`.

Tamura RN, Young SS (1987). "A Stabilized Moment Estimator for the Beta-Binomial Distribution." *Biometrics*, **43**(4), 813–824. `doi:10.2307/2531535`.

Tang R (2002). *Fitting and Evaluating Certain Two-Level Hierarchical Models.* Ph.D. thesis, Harvard University.

# A. Posterior propriety of the Poisson model

If the posterior distribution of $r$ is proper, then the full posterior distribution of random effects and $r$ is also proper because

$$f(\boldsymbol{\lambda}, r \mid \boldsymbol{y}) = f(\boldsymbol{\lambda} \mid \boldsymbol{y}) \times f(r \mid \boldsymbol{y}),$$

where $f(\boldsymbol{\lambda} \mid \boldsymbol{y})$ is a product of $k$ proper conditional posterior density function in (32). Thus, our goal is to show that $\int_0^\infty f(r \mid \boldsymbol{y}) dr < \infty$:

$$f(r \mid \boldsymbol{y}) \propto \frac{1}{r^2} L(r) \propto \frac{1}{r^2} \prod_{j=1}^k \frac{\Gamma(r\lambda_j^E + y_j)}{\Gamma(r\lambda_j^E)} (1 - B_j)^{y_i} B_j^{r\lambda_j^E} \tag{59}$$

$$= \frac{1}{r^2} \left[ r^{\sum_{j=1}^k y_j} + \cdots + a_k r^k \right] \exp\left( -r \sum_{j=1}^k \lambda_j^E \log(1 + n_j/r) \right) \prod_{j=1}^k \left( \frac{n_j}{n_j + r} \right)^{y_j}, \tag{60}$$

where the polynomial function of $r$ in the bracket has constant coefficients.

If there are at least two groups whose observed values $y_j$ are non-zero, then $f(r \mid \boldsymbol{y})$ goes to zero as $r$ approaches zero due to the polynomial function of $r$ in (60); the following two factors in (60) approach one. As $r$ becomes infinite, $f(r \mid \boldsymbol{y})$ touches zero exponentially fast due to the exponential term in the middle of (60). Thus, the integration of $f(r \mid \boldsymbol{y})$ must be finite.

**Affiliation:**

Hyungsuk Tak, Carl Morris
Department of Statistics
Harvard University
1 Oxford Street
Cambridge, MA, United States of America
E-mail: hyungsuk.tak@gmail.com, morris@stat.harvard.edu

Joseph Kelly
Google, Inc.
76 Ninth Avenue
New York, NY, United States of America
E-mail: josephkelly@google.com