

(wileyonlinelibrary.com) DOI: 10.1002/jrsm.1117

A Bayesian nonparametric meta-analysis model

George Karabatsos,^{a*} Elizabeth Talbott^b and Stephen G. Walker^c

In a meta-analysis, it is important to specify a model that adequately describes the effect-size distribution of the underlying population of studies. The conventional normal fixed-effect and normal random-effects models assume a normal effect-size population distribution, conditionally on parameters and covariates. For estimating the mean overall effect size, such models may be adequate, but for prediction, they surely are not if the effect-size distribution exhibits non-normal behavior. To address this issue, we propose a Bayesian nonparametric meta-analysis model, which can describe a wider range of effect-size distributions, including unimodal symmetric distributions, as well as skewed and more multimodal distributions. We demonstrate our model through the analysis of real meta-analytic data arising from behavioral-genetic research. We compare the predictive performance of the Bayesian nonparametric model against various conventional and more modern normal fixed-effects and random-effects models. Copyright © 2014 John Wiley & Sons, Ltd.

Keywords: meta-analysis; Bayesian nonparametric regression; meta-regression; effect sizes; publication bias

1. Introduction

A research synthesis aims to integrate results from empirical research so as to produce generalizations (Cooper and Hedges, 2009). Meta-analysis, also referred to as the analysis of analyses (Glass, 1976), provides a quantitative synthesis of statistics that are reported by multiple research studies. Specifically, each study reports an effect-size statistic, and provides information about its sampling variance, while the features of the study may be described by one or more covariates. Important examples of effect-size statistics include the unbiased standardized mean difference between two independent groups (Hedges, 1981), among others (Konstantopoulos, 2007). Given a sample of effect-size data, the primary aim of a meta-analysis is to infer the overall effect-size distribution from the given study population, as well as to infer the heterogeneity of the effect sizes. Conventional summaries of the overall effect-size distribution include the mean, which is often referred to as the 'overall effect-size,' and the variance which describes the heterogeneity of reported effect sizes in the overall effect-size distribution. Heterogeneity can be further investigated in a meta-regression analysis, in order to investigate how the mean effect-size relates to key study-level covariates (e.g., Berkey *et al.*, 1995; Thompson and Sharp, 1999; Thompson and Higgins, 2002; Higgins and Thompson, 2004). Also, meta-regression analysis can be used to investigate and test for publication bias in the data, by relating effect sizes with their standard errors (precisions; the square root of the effect-size sampling variances) (Thompson and Sharp, 1999). This actually provides a regression analysis for the funnel plot (Egger *et al.*, 1997).

The normal fixed-effects model and the normal random-effects model provide two traditional and alternative approaches to meta-analysis (Hedges and Vevea, 1998; Konstantopoulos, 2007; Borenstein *et al.*, 2010). Each model is a weighted linear regression model, which treats study-reported effect-size as the dependent variable, weights each reported effect-size by the inverse of its sampling variance, assumes normally distributed regression errors, and represents the overall (mean) effect size by the intercept parameter. In other words, each of these

^aDepartment of Educational Psychology, Program in Measurement, Evaluation Statistics, and Assessments, College of Education, University of Illinois – Chicago, 1040 W. Harrison St. (MC 147), Chicago, IL 60607, USA

^bDepartment of Special Education, College of Education, University of Illinois – Chicago, Chicago, IL, USA

^cDivision of Statistics and Scientific Computation, The University of Texas at Austin, Austin, TX, USA

*Correspondence to: Professor George Karabatsos, Department of Educational Psychology, Program in Measurement, Evaluation Statistics, and Assessments, College of Education, University of Illinois-Chicago, 1040 W. Harrison St. (MC 147), Chicago, IL 60607, USA.
E-mail: georgek@uic.edu; gkarabatsos1@gmail.com

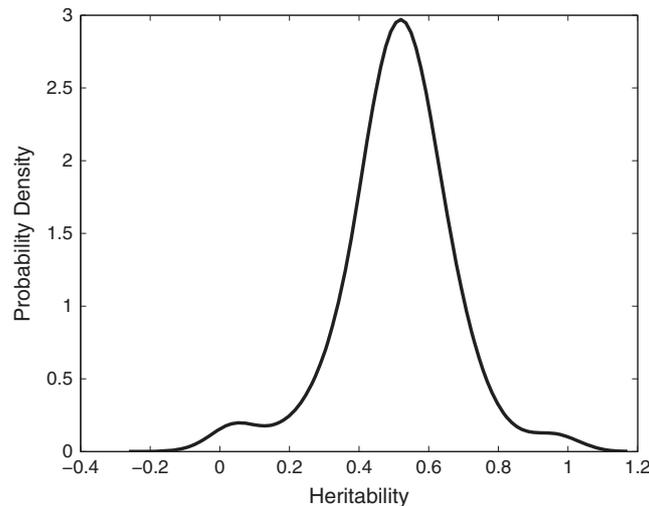


Figure 1. Kernel density estimate of the effect-size (heritability).

models assumes that the effect-size distribution is a unimodal and symmetric, normal distribution, conditionally on all model parameters, and conditionally on any set of chosen values of covariates in the model. The fixed-effects model is an ordinary (weighted) linear regression model. The normal random-effects model is a two-level model that extends the fixed-effects model, by allowing for between-study variance in the effect sizes, through the addition of random intercept parameters that are assumed to have a normal distribution over the given study population. Specifically, effect sizes at the first level are nested within studies at the second level. A three-level meta-analysis model can accommodate data structures where, for example, studies (level 2) are themselves nested within units, such as school districts (level 3) (for more details, see Konstantopoulos, 2011). Given a sample set of meta-analytic data, the parameters of a normal fixed-effects or normal random-effects model can be estimated via maximum likelihood. Alternatively, a Bayesian inference approach can be taken, which involves the specification of a prior distribution on all parameters of the given meta-analytic model (e.g., Higgins *et al.*, 2009). Then, Bayesian inference of the data is based on the posterior distribution of the parameters, formed by combining the data (likelihood) information of the model, with the information from the prior distribution.

Our motivation for the current paper involves the meta-analysis of 71 effect sizes, which are heritability estimates that were reported by a collection of behavioral-genetic studies of the heritability of antisocial behavior (Talbot *et al.*, 2012). Figure 1 presents a simple kernel probability density estimate of the effect sizes. This estimate shows skewness and at least two modes in the effect-size distribution, thereby violating assumptions of a normal fixed-effects or a normal random-effects model, namely, that this distribution is symmetric and unimodal.[†]

Arguably, a normal model (fixed-effects or random-effects) can provide an adequate analysis of these data, specifically for the purposes of estimating the overall mean of the effect size, and possibly also the overall variance. However, such a model may not be adequate for predictive purposes, given the lack of distributional symmetry and unimodality in the data. The importance of predicting from meta-analytic data has been detailed by Higgins *et al.* (2009). As they state, 'Predictions are one of the most important outcomes of a meta-analysis, since the purpose of reviewing research is generally to put knowledge gained into future application. Predictions also offer a convenient format for expressing the full uncertainty around inferences, since both magnitude and consistency of effects may be considered.' Accurate prediction with non-normal data requires flexible models that support a wider range of distributions, beyond the normal distribution. This has already been recognized by Burr and Doss (2005) and Branscum and Hanson (2008), who proposed Bayesian nonparametric models for meta-analysis. More generally speaking, Bayesian nonparametric models are 'nonparametric' in the sense that they avoid the more restrictive assumptions of 'parametric' models, namely, that the data distribution can be fully described by a few finite number of parameters. For example, a normal model assumes that the data distribution is unimodal and symmetric, and therefore can be entirely described by a mean parameter and a variance parameter. A Bayesian nonparametric model assigns a prior distribution to an infinite number (or a very large number) of model parameters. This is carried out for the purposes of specifying a highly flexible model that describes a wide range of data distributions, including unimodal distributions, and more multimodal distributions (Müller and Quintana, 2004).

[†]This is because either normal model assumes that effect sizes have a common mean but different variances, respectively, and then the marginal distribution over effect sizes (studies) is a scale mixture of normal distributions (Andrews and Mallows, 1974; West, 1984).

Consequently, we propose a Bayesian nonparametric model for meta-analysis, which is more flexible than the normal, fixed-effects, and random-effects models. Our model is a special case of the general model introduced by Karabatsos and Walker (2012), which was studied in more general regression settings. The new model specifies the effect-size distribution by an infinite random intercept mixture of normal distributions, conditional on any covariate(s) of interest, with covariate-dependent mixture weights. Therefore, the model is flexible enough to describe a very wide range of effect-size distributions, including all normal distributions, as well as all (smooth) distributions that are more skewed and/or multimodal. Also, the model avoids the empirically falsifiable assumption that effect sizes arise strictly from symmetric unimodal distributions, such as normal distributions. Furthermore, the model's high flexibility encourages a rich and graphical inference of the whole effect-size distribution, as previously recommended for meta-analytic practice (Higgins *et al.*, 2009).

Also, in the spirit of meta-regression analysis, our Bayesian nonparametric meta-analysis model allows the whole effect-size distribution to change flexibly and non-linearly as a function of key study-level covariates. This feature permits a rich and flexible meta-regression analysis, whereas the previous Bayesian nonparametric regression models do not account for covariate information (Burr and Doss, 2005; Branscum and Hanson, 2008). Moreover, for a given meta-analysis, our Bayesian model can automatically identify the subset of covariates that significantly predict changes in the mean effect-size, in a model-based and non-ad-hoc fashion. Specifically, the model makes use of spike-and-slab priors for the regression coefficients that allow for automatic covariate (predictor) selection in the posterior distribution. Such priors were developed for Bayesian normal linear regression models (George and McCulloch, 1997). Moreover, under either a normal fixed-effects or normal random-effects model inferred under a non-Bayesian (Frequentist) framework of maximum-likelihood estimation, the identification or selection of significant study-level covariates (predictors) is challenging because it deals with the standard issues of multiple hypothesis testing over predictors (e.g., Thompson and Higgins, 2002). The often-used stepwise procedures of covariate selection are known to be ad-hoc and sub-optimal.

We now describe the layout of the rest of the paper. As the paper covers various key statistical concepts, it is necessary to first give them a brief review in Section 2. In Section 2.1, we review the basic data framework of meta-analysis, including effect sizes. In Section 2.2, we review of the traditional normal fixed-effects and normal random-effects models, and in Section 2.3, we briefly review the Bayesian statistical inference framework. In Section 2.4, we review the traditional Bayesian meta-analytic models, including conventional and more modern versions of normal fixed-effects and normal random-effects models. In Section 3, we describe our new Bayesian meta-analysis model. In Section 4, we review a standard criterion for comparing the predictive performance between different Bayesian models that are fit to a common data set, for the purposes of identifying the single model that has best predictive-fit, that is, of identifying the single model that best describes the underlying population distribution of the sample data. Section 5 illustrates our Bayesian meta-analytic model through the analysis of the large meta-analytic data set of behavioral-genetic studies, which was briefly described previously, and which involves 24 covariates. In that section, we also compare the predictive accuracy of our Bayesian nonparametric model, against the predictive accuracy of conventional and more modern normal fixed-effects and normal random-effects models. Section 6 ends with conclusions.

2. Review of meta-analytic modeling concepts

Before we review the key concepts underlying the various approaches to meta-analysis, we describe some notation that we use in the remainder of this paper. Following the standard notation of statistics, \sim will mean 'distributed as'; $n(\cdot|\mu, \nu)$ denotes the ('bell-shaped') PDF of the normal distribution having mean and variance (μ, ν) ; the PDF of the n -variate normal distribution with mean vector μ and (symmetric and positive-definite) variance-covariance matrix Σ is denoted by $n_n(\cdot|\mu, \Sigma)$; the PDF of a gamma distribution with shape and rate parameters (a, b) is denoted by $ga(\cdot|a, b)$; and the PDF of a uniform distribution with minimum and maximum parameters (a, b) is denoted by $un(\cdot|a, b)$. Also, we denote a CDF by a capital letter, such as G , and $\Phi(\cdot)$ denotes the Normal(0,1) CDF. Finally, $\delta_\theta(\cdot)$ denotes the degenerate distribution that assigns probability 1 (full support) to the number θ .

2.1. Data framework of meta-analysis

In a typical meta-analysis context, n' ($\leq n$) studies provide data on n study reports $\{y_i\}_{i=1}^n$ of a common effect size (y) of interest. Each effect-size report y_i is based on n_i observations, has a sampling variance $(\hat{\sigma}_i^2)$, and corresponds to p covariates $\mathbf{x}_i = (1, x_{1i}, \dots, x_{pi})^T$ describing study characteristics, including a constant (1) term for future notational convenience. A full meta-analytic data set is denoted by $\mathcal{D}_n = \{(y_i, \mathbf{x}_i, \hat{\sigma}_i^2)\}_{i=1}^n$.

Table 1 presents some typical examples of effect-size statistics that are often used in meta-analysis, along with their sampling variances $(\hat{\sigma}_i^2)$ (Konstantopoulos, 2007; Borenstein, 2009; Fleiss and Berlin, 2009). Textbook examples of these effect sizes, from the social or medical sciences, includes the standardized between-gender difference in conformity, the correlation between student ratings and student achievement (Cooper *et al.*, 2009,

Effect-size description	Effect-size (y_i)	Variance ($\hat{\sigma}_i^2$)
Unbiased standardized mean difference, two independent groups (Hedges, 1981).	$\frac{\hat{\mu}_{1i} - \hat{\mu}_{2i}}{\sqrt{\frac{(n_{1i}-1)\hat{\sigma}_{1i} + (n_{2i}-1)\hat{\sigma}_{2i}}{n_{1i} + n_{2i} - 2}}} c^*$	$\left(\frac{n_{1i} + n_{2i}}{n_{1i}n_{2i}} + \frac{y_i^2}{2(n_{1i} + n_{2i})} \right) c^{*2};$ $c^* = 1 - \frac{3}{4(n_{1i} + n_{2i} - 2) - 1}$
Fisher z transformation of the correlation $\hat{\rho}_i$.	$\frac{1}{2} \log \frac{1 + \hat{\rho}_i}{1 - \hat{\rho}_i}$	$\frac{1}{n_i + 3}$
Log odds ratio for two binary (0–1) variables.	$\log \left(\frac{n_{11i}/n_{10i}}{n_{01i}/n_{00i}} \right)$	$\frac{1}{n_{11i}} + \frac{1}{n_{10i}} + \frac{1}{n_{01i}} + \frac{1}{n_{00i}}$

Appendix A), and the log odds death rate of patients receiving versus patients not receiving beta-blocker treatment (Gelman *et al.*, 2004, Section 5.6). As mentioned, in the current manuscript, we consider the analysis of a large meta-analytic data set of behavioral genetic studies, which together reported estimates of the heritability of antisocial behavior. See Section 5 for more details. Heritability provides a standard measure of effect size in the behavioral-genetics field, which is concerned with the study of genetic and environmental influences on behaviors (e.g., Bazzett, 2008). This field is highly interdisciplinary and cross-cuts with many fields of the social, physical, and medical sciences, including psychology, biology, genetics, epigenetics, ethology, and statistics.

Finally, we may more generally consider a sampling covariance matrix for the n study reports, $\hat{\Sigma}_n = (\hat{\sigma}_{ij})_{n \times n}$, having diagonal elements $\hat{\sigma}_{ii} = \hat{\sigma}_i^2$, where each off-diagonal element is the sampling covariance for a given effect-size pair (y_k, y_l), with $k \neq l$ (Gleser and Olkin, 2009). To maintain notational simplicity throughout the paper, we will present the meta-analytic models under the common assumption that $\hat{\Sigma}_n$ is a diagonal matrix (i.e., $\hat{\Sigma}_n = \text{diag}(\sigma_1^2, \dots, \sigma_n^2)$), implying the assumption of zero sampling covariances. Though, as we discuss in Section 3, this diagonal matrix assumption can be made for the Bayesian nonparametric meta-analytic model, without loss of generality in terms of being able to model covariances between distinct pairs of study effect-size reports.

2.2. Traditional meta-analytic models

A traditional meta-analysis model assumes that, for a given set of data $\mathcal{D}_n = \{(y_i, \mathbf{x}_i, \hat{\sigma}_i^2)\}_{i=1}^n$, the effect-size distribution follows the general form:

$$f(y_i | \mathbf{x}_i, \hat{\sigma}_i^2; \zeta) = n(y_i | \mathbf{x}_i^T \boldsymbol{\beta} + \mu_{0i} + \mu_{00t(i)}, \hat{\sigma}_i^2), \quad i = 1, \dots, n; \quad (1a)$$

$$\mathbf{x}_i^T \boldsymbol{\beta} = \beta_0 + \beta_1 x_{i1} + \dots + \beta_p x_{ip}; \quad (1b)$$

$$(\mu_{01}, \dots, \mu_{0n}) | \Sigma_0 \sim n_n(0, \sigma_0^2 \mathbf{I}_n + \psi \mathbf{M}_n); \quad (1c)$$

$$\mu_{00t} | \sigma_{00}^2 \sim n(0, \sigma_{00}^2), \quad t = 1, \dots, T. \quad (1d)$$

Typical meta-analytic models are special cases of the general normal model shown in equation (1a–d). In all, the general traditional model (1a–d) has likelihood density $f(y | \mathbf{x}, \hat{\sigma}^2; \zeta)$ with parameters $\zeta = (\boldsymbol{\beta}, \mu_0, \mu_{00}, \sigma_0^2, \sigma_{00}^2, \psi)^T$. These parameters are explained as follows.

The intercept parameter β_0 is interpreted as the *mean effect-size* over the given population of studies (Louis and Zelterman, 1994). This interpretation holds true, provided that each of the p covariates has data observations $(x_{k1}, \dots, x_{kn})^T$ that have already been centered to have mean zero, as we assume throughout. Also, the p covariates are, respectively, parameterized by linear slope coefficients $(\beta_1, \dots, \beta_p)$.

Also, $\boldsymbol{\mu}_0 = (\mu_{01}, \dots, \mu_{0n})^T$ are the level-2 random intercept parameters. Similarly, $\boldsymbol{\mu}_{00} = (\mu_{001}, \dots, \mu_{00T})^T$ are the level-3 random intercept parameters, with each of the n study reports being nested within exactly one of $T \leq n$ study reports ($t = 1, \dots, T$), and with $\mu_{00t(i)}$ meaning that the level-3 intercept μ_{00t} is assigned to the i th study report. As shown in the model equations (1c) and (1d), the random intercepts $\boldsymbol{\mu}_0$ and $\boldsymbol{\mu}_{00}$ are each assumed to have a multivariate normal distribution (probability density).

A normal fixed-effects model assumes that all the random intercept parameters ($\boldsymbol{\mu}_0, \boldsymbol{\mu}_{00}$) are zero. In terms of the general normal meta-analytic model (1a–d), this corresponds to the assumption of zero variances, that is, $\sigma_0^2 = \sigma_{00}^2 = \psi = 0$ (Hedges and Vevea, 1998). A 2-level normal random-effects model allows for non-zero random intercepts $\boldsymbol{\mu}_0$, by allowing for non-zero variances (σ_0^2, ψ) , as shown in equation (1c). Typical normal random-effects models assume that the random intercepts $\boldsymbol{\mu}_0$ are uncorrelated, with n -variate normal density $n_n(\boldsymbol{\mu}_0 | 0, \sigma_0^2 \mathbf{I}_n)$ (Hedges and Vevea, 1998). However, it is possible to model correlations between the level-2 random intercepts $\boldsymbol{\mu}_0$. For example, Stevens and Taylor (2009) consider a 2-level normal random-effects model, which assumes that the level-2 random intercepts have a n -variate normal density $n_n(\boldsymbol{\mu}_0 | 0, \Sigma_0)$, with a more general covariance structure $\Sigma_0 = \sigma_0^2 \mathbf{I}_n + \psi \mathbf{M}_n$. Here, the

parameter ψ represents the covariance between pairs of study reports. Also, \mathbf{M}_n is a fixed $n \times n$ indicator (0–1) matrix, with 1s specified in the off-diagonal to reflect a-priori beliefs as to which pairs of the n study reports have correlated level-2 random intercepts, and with zeros specified for all the other entries of \mathbf{M}_n . A 3-level normal random-effects model allows for non-zero random intercepts μ_{0i} , by allowing for a positive variance σ_{00}^2 , as shown in equation (1d) of the general normal model (Konstantopoulos, 2011).

The general normal model described in equation (1) can be written explicitly as a normal mixture of multivariate normal model:

$$f(\mathbf{y}|\mathbf{X}, \hat{\Sigma}_n; \beta, \sigma_0^2, \psi, \sigma_{00}^2) = \int n_n(\mathbf{y}|\mathbf{X}\beta + \mu_0 + \mu_{00}^*, \hat{\Sigma}_n) dG_2(\mu_0) dG_3(\mu_{00}),$$

for effect-size data $\mathbf{y} = (y_1, \dots, y_n)^\top$, given the n -by- $(p+1)$ matrix \mathbf{X} of row vectors \mathbf{x}_i^\top ($i=1, \dots, n$), $\mu_{00}^* = (\mu_{00t(1)}, \dots, \mu_{00t(n)})^\top$, and $\hat{\Sigma}_n = \text{diag}(\sigma_1^2, \dots, \sigma_n^2)$. Specifically, the mixture model presented previously assumes that the mixture distribution $G_2(\mu_0)$ is a multivariate distribution with probability density $n_n(\mu_0|0, \sigma_0^2 \mathbf{I}_n + \psi \mathbf{M}_n)$, and that $G_3(\mu_{00})$ is a multivariate normal distribution with probability density $n_T(\mu_{00}|0, \sigma_{00}^2 \mathbf{I}_n)$.

For any of the models described in this section, full maximum-likelihood methods can be used to estimate the parameters $\zeta = (\beta, \mu_0, \mu_{00}, \sigma_0^2, \sigma_{00}^2, \psi)^\top$, from a given data set \mathcal{D}_n . Alternatively, the parameters of a random-effects model can also be estimated by the restricted maximum-likelihood method, which focuses estimation on the variance parameters (Harville, 1977; Raudenbush and Bryk, 2002, Ch. 3, 13–14; Stevens and Taylor, 2009). For any one of the models that is described in this section, the estimate of the effect-size distribution of the underlying population is given by the density $f_n(y|\mathbf{x}_0, \hat{\sigma}^2; \hat{\zeta}) = n(y|\hat{\beta}_0, \hat{\sigma}^2 + \hat{\sigma}_0^2 + \hat{\sigma}_{00}^2)$ of the normal distribution, given a maximum-likelihood estimate $\hat{\zeta}$ obtained from a sample data set \mathcal{D}_n , and after conditioning on all p covariates with values $\mathbf{x} = \mathbf{x}_0 = (1, 0, \dots, 0)^\top$.

As an alternative to the maximum-likelihood approach, parameter estimation can be performed in a fully Bayesian framework (e.g., Higgins *et al.*, 2009), which is described next.

2.3. Review of Bayesian inference

For a general meta-analytic model, let $f(y|\mathbf{x}, \hat{\sigma}^2; \zeta)$ denote the likelihood density of the effect-size data point y , conditionally on covariates \mathbf{x} and model parameter ζ , which has space Ω_ζ . A given meta-analytic data set $\mathcal{D}_n = \{(y_i, \mathbf{x}_i, \hat{\sigma}_i^2)\}_{i=1}^n$ has likelihood $L(\mathcal{D}_n; \zeta) = \prod_{i=1}^n f(y_i|\mathbf{x}_i, \hat{\sigma}_i^2; \zeta)$ under the model. In the Bayesian approach to statistical inference, the model parameter ζ is assigned a prior probability density $\pi(\zeta)$ over the parameter space Ω_ζ . This density reflects pre-experimental beliefs about the plausible values of the parameter, for the meta-analytic data set \mathcal{D}_n at hand. Then according to Bayes' theorem, the data \mathcal{D}_n , via the model's likelihood $L(\mathcal{D}_n; \zeta)$, combines with the prior density $\pi(\zeta)$, to yield a posterior density for ζ , defined by

$$\pi(\zeta|\mathcal{D}_n) = \frac{L(\mathcal{D}_n; \zeta)\pi(\zeta)}{\int_{\Omega_\zeta} L(\mathcal{D}_n; \zeta)d\Pi(\zeta)} = \frac{\prod_{i=1}^n f(y_i|\mathbf{x}_i, \hat{\sigma}_i^2; \zeta)\pi(\zeta)}{\int_{\Omega_\zeta} \prod_{i=1}^n f(y_i|\mathbf{x}_i, \hat{\sigma}_i^2; \zeta)d\Pi(\zeta)}, \quad (2)$$

where $\Pi(\zeta)$ denotes the CDF of the prior density $\pi(\zeta)$. The posterior density describes the plausible values of the model parameters ζ , given prior beliefs about ζ and data \mathcal{D}_n .

The specification of the prior density $\pi(\zeta)$ is an important step in a Bayesian analysis. The posterior $\pi(\zeta|\mathcal{D}_n)$ can be quite sensitive to the choice of prior, especially when the sample size (n) is not large. Also, often in practice, there is a lack of prior information about the parameters ζ of the given model. This lack of prior information is reflected by a 'diffuse' prior probability density $\pi(\zeta)$ that has high variance, and assigns rather-equal but broad support over the parameter space Ω_ζ . Such priors are often referred to as 'non-informative', even though technically speaking, a prior cannot be fully non-informative. As a consequence of specifying a diffuse prior $\pi(\zeta)$, the posterior density $\pi(\zeta|\mathcal{D}_n)$ becomes mostly determined by the data \mathcal{D}_n likelihood $L(\mathcal{D}_n; \zeta)$, relative to the prior.

Prediction is a basic function of statistical modeling. A Bayesian meta-analytic model makes predictions of Y , given a chosen \mathbf{x} , on the basis of the posterior predictive density:

$$f_n(y|\mathbf{x}, \hat{\sigma}^2) = \int f(y|\mathbf{x}, \hat{\sigma}^2; \zeta)\pi(\zeta|\mathcal{D}_n)d\zeta, \quad (3)$$

and this density has posterior predictive mean (expectation) $E_n(Y|\mathbf{x}, \hat{\sigma}^2) = \int yf_n(y|\mathbf{x}, \hat{\sigma}^2)dy$, and posterior predictive variance $\text{Var}_n(Y|\mathbf{x}, \hat{\sigma}^2) = \int \{y - E_n(Y|\mathbf{x}, \hat{\sigma}^2)\}^2 f_n(y|\mathbf{x}, \hat{\sigma}^2)dy$. The posterior predictive density $f_n(y|\mathbf{x}, \hat{\sigma}^2)$ provides an estimate of the true effect-size density (distribution) for the underlying study population, given sample data \mathcal{D}_n and covariates \mathbf{x} of interest, under squared-error loss (Aitchison, 1975).

In most Bayesian meta-analytic models, ζ is a high-dimensional parameter vector. Then the direct evaluation of the posterior equations (2) and (3) requires computationally prohibitive high-dimensional integrations. In such

situations, one may use Markov chain Monte Carlo (MCMC) sampling methods to estimate such posterior densities of the given model. Such methods include the Gibbs sampler (Gelfand and Smith, 1990), Metropolis–Hastings algorithm (Chib and Greenberg, 1995), and other sampling algorithms (see for e.g., Brooks *et al.*, 2011).

2.4. Review of Bayesian normal meta-analytic models

In reference to the Bayesian inference framework described in Section 2.3, consider the general normal meta-analytic model (1a–d), which has likelihood density given by (1a), and which has parameters $\zeta = (\beta, \gamma, \mu_0, \mu_{00}, \sigma_0^2, \sigma_{00}^2, \psi)^\top$. Here, $\gamma = (\gamma_1, \dots, \gamma_p)^\top$ are included as parameters, which, respectively, indicate (0–1) whether or not the p covariates are included ($\gamma = 1$) or excluded ($\gamma = 0$) from the model. In typical practice involving such a model (including special cases), the prior density has the general form:

$$\pi(\zeta) = n(\beta|0, \mathbf{V}_\gamma)\pi(\gamma)n_n(\mu_0|0, \sigma_0^2\mathbf{I}_n + \psi\mathbf{M}_n)n_T(\mu_{00}|0, \sigma_{00}^2\mathbf{I}_T)\pi(\sigma_0^2, \psi)\pi(\sigma_{00}^2). \quad (4)$$

In Bayesian meta-analytic modeling, it is a common practice to specify a diffuse normal prior density for the coefficients β by taking $\mathbf{V}_\gamma = v\mathbf{I}_{p+1}$, with $v \rightarrow \infty$ (e.g., DuMouchel and Normand, 2000), along with the implicit prior assumption that $\pi(\gamma = \mathbf{1}) = 1$.

Also, it is a common practice in Bayesian normal random-effects modeling to attempt to assign a non-informative prior for $(\sigma_0^2, \sigma_{00}^2)$ via the specification of inverse-gamma priors $\pi(\sigma_0^2) = \text{ga}(\sigma_0^{-2}|\epsilon, \epsilon)$ and $\pi(\sigma_{00}^2) = \text{ga}(\sigma_{00}^{-2}|\epsilon, \epsilon)$, for small choice of constant $\epsilon > 0$ (Gelman, 2006), implying a prior density of the form $\pi(\sigma_0^2, \psi) = \pi(\sigma_0^2)\delta_0(\psi)$. Though, recall from Section 2.2 that a more general multivariate normal $n_n(\mu_0|0, \sigma_0^2\mathbf{I}_n + \psi\mathbf{M}_n)$ mixture distribution can be specified (Stevens and Taylor, 2009). This mixture distribution prior allows for correlated level-2 random intercepts μ_0 via a parameter ψ that measures the covariance between specific pairs of the total n study reports, and where $\mathbf{M}_n = (m_{ij})_{n \times n}$ is a fixed matrix, which indicates (0–1) which pairs of the study reports are expected to yield correlated level-2 random intercepts μ_0 , with zeros in the diagonal. For the parameters (σ_0^2, ψ) , Stevens and Taylor (2009) propose the rather non-informative prior density

$$\pi(\sigma_0^2, \psi) = \left(c_0 / (c_0 + \sigma_0^2)\right) \text{un}(\psi | -\sigma_0^2 / (K - 1), \sigma_0^2), \quad (5)$$

where the first term in the product gives a log-logistic prior density for σ_0^2 , where $c_0 = \{n / \text{tr}([\text{diag}(\hat{\Sigma}_n)] - 1)\}^{1/2}$ is the harmonic mean of the sampling variances $\hat{\sigma}_i^2$ ($i = 1, \dots, n$), and $K = \max_i \left\{ \sum_{l=1}^n m_{il} \right\}$ is the largest group of related study reports. This log-logistic prior is right-skewed, highly dispersed, with quartiles $(c_0/3, c_0, 3c_0)$.

Simpler versions of the general normal random-effects model (1a–d) can be specified via appropriate straightforward modifications of the prior density (4). A Bayesian 2-level normal random-effects model, which assumes $\sigma_{00}^2 = 0$ (i.e., $\mu_{00} = 0$) but allows for correlated level-2 random intercepts μ_0 assigns the prior density $\pi(\sigma_{00}^2) = \delta_0(\sigma_{00}^2)$; a 2-level normal random-effects model, which assumes $\sigma_{00}^2 = \psi = 0$ (i.e., $\mu_{00} = 0$) and assumes independent level-2 random intercepts μ_0 assigns the prior density $\pi(\sigma_{00}^2, \psi) = \delta_0(\sigma_{00}^2)\delta_0(\psi)$; and a fixed-effects model, which assumes $\sigma_0^2 = \sigma_{00}^2 = \psi = 0$ (i.e., $\mu_0 = 0$ and $\mu_{00} = 0$) assigns the prior density $\pi(\sigma_0^2, \sigma_{00}^2, \psi) = \delta_0(\sigma_0^2)\delta_0(\sigma_{00}^2)\delta_0(\psi)$.

Other simple modifications of the prior density (4) can be used to specify other important versions of the general normal random-effects model (1a–d), which consider different priors for the random intercepts (μ_0, μ_{00}) . For example, Gelman (2006) notes that posterior inference of the parameter σ_0^2 , under the often-used ‘non-informative’ gamma $\text{ga}(\sigma_0^{-2}|\epsilon, \epsilon)$ prior, is very sensitive to the choice of small ϵ , especially when the data support small values of σ_0^2 . The same is true for the gamma $\text{ga}(\sigma_{00}^{-2}|\epsilon, \epsilon)$ prior for the parameter σ_{00}^2 . Therefore, in a situation where there is little prior information available about the parameters $(\sigma_0^2, \sigma_{00}^2)$, he alternatively recommends the specification of uniform prior densities $\pi(\sigma_0) = \text{un}(\sigma_0|0, b_0)$ and $\pi(\sigma_{00}) = \text{un}(\sigma_{00}|0, b_{00})$ for reasonably large values (b_0, b_{00}) , whenever n and T are both at least 5. When more prior information is desired, say when n is less than 5, he recommends the half- t prior density of the general form $\pi(\sigma_0^2) \propto \left(1 + a_0^{-1}(\sigma_0/b_0)^2\right)^{-(a+1)/2}$, and similarly for the level-3 variance parameter σ_{00} .

Finally, while it is a common practice to assume a diffuse prior for the regression coefficients $\beta = (\beta_0, \beta_1, \dots, \beta_p)^\top$, along with $\pi(\gamma = \mathbf{1}) = 1$, in principle, one may specify spike-and-slab priors for the slope parameters $(\beta_1, \dots, \beta_p)$ in order to enable automatic variable (covariate) selection via posterior inference (George and McCulloch, 1997), along with a diffuse prior $\beta_0 \sim n(0, v \rightarrow \infty)$. These spike-and-slab priors are defined by independent normal and Bernoulli prior densities, so that the $n(\beta|0, \mathbf{V}(\gamma))\pi(\gamma)$ prior in (4) is based on

$$\mathbf{V}_\gamma = \text{diag}\left(v \rightarrow \infty, v_1^{\gamma_1} v_0^{1-\gamma_1}, \dots, v_1^{\gamma_p} v_0^{1-\gamma_p}\right) \quad (6a)$$

$$\pi(\gamma) = \prod_{k=1}^p \text{Pr}(\gamma_k = 1)^{\gamma_k} [1 - \text{Pr}(\gamma_k = 1)]^{1-\gamma_k}, \quad (6b)$$

where v_0 is a small prior variance (e.g., $v_0 = 0.001$), v_1 is a large prior variance (e.g., $v_1 = 10$), and $\text{Pr}(\gamma_k = 1)$ is the Bernoulli probability parameter that is often set to .5 in practice (George and McCulloch, 1997). So on the one

hand, with prior probability $Pr[\gamma_k = 1] = .5$, the k th covariate is included in the model as a ‘significant’ predictor of the effect-size, by assigning its regression coefficient β_k a normal $n(\beta_k|0, v_1)$ prior density that supports a large range of β_k values. On the other hand, with prior probability $Pr[\gamma_k = 0] = .5$, that covariate is excluded from the model, by assigning its regression coefficient β_k a normal $n(\beta_k|0, v_0)$ prior that places all its support on values $\beta_k \approx 0$. Given MCMC samples from the posterior, $\pi(\zeta|\mathcal{D}_n)$, the k th covariate can be viewed as a ‘significant predictor,’ when the posterior inclusion probability of the covariate, $Pr[\gamma_k = 1|\mathcal{D}_n]$, is at least .5 (Barbieri and Berger, 2004). We assume throughout that the spike-and-slab prior specifications are given by $v_0 = .001$ and $v_1 = 10$. These specifications are consistent with the previous recommendation that the ratio v_1/v_0 be no greater than 10,000, for the purposes of reliably estimating the parameters (β, γ) via MCMC methods (George and McCulloch, 1997, p. 368).

A given Bayesian regression model, assigned spike-and-slab priors, can handle a large number of covariates relative to the number of observations, as these priors control the parametric dimensionality of the model. For example, if there is not enough information in the given data set to support the specification of many covariates, then these priors will eliminate most of the irrelevant covariates by setting their coefficients to zero with high posterior probability. See George and McCulloch (1997) for more details.

For any one of the models described in this section, the posterior predictive density $f_n(y|\mathbf{x}_0, \hat{\sigma}^2)$ provides an estimate of the ‘overall’ effect-size distribution of the underlying population, and is a symmetric and unimodal distribution, conditionally on covariates $\mathbf{x} = \mathbf{x}_0 = (1, 0, \dots, 0)^T$. For example, under a Bayesian normal fixed-effects model, the posterior predictive density $f_n(y|\mathbf{x}_0, \hat{\sigma}^2)$ estimate of the overall population effect-size distribution is a normal density (distribution). Under a Bayesian normal 2-level random-effects model with $ga(\sigma_0^{-2}|a, b)$ prior for independent level-2 random intercepts μ_0 , the posterior predictive estimate $f_n(y|\mathbf{x}_0, \hat{\sigma}^2)$ of the population effect-size distribution is given by a student density (distribution) (e.g., Denison et al., 2002, Appendix). The Student’s t distribution is very similar to a normal distribution, except that the student distribution has thicker tails.

For a given Bayesian normal meta-analytic model with prior density having the general form (4), the posterior density $\pi(\beta, \gamma, \mu_0, \mu_{00}, \sigma_0^2, \sigma_{00}^2, \psi|\mathcal{D}_n)$, the posterior predictive density $f_n(y|\mathbf{x}, \hat{\sigma}^2)$, and any functional of these densities, can be estimated through the use of standard MCMC Gibbs and Metropolis sampling algorithms for normal linear and random-effects models (e.g., Gilks et al., 1993; Denison et al., 2002).

3. The Bayesian nonparametric meta-analysis model

For effect-size data, our Bayesian nonparametric meta-analysis model is defined by an infinite random intercepts mixture of regressions. The model assumes data likelihood that is given by $\prod_{i=1}^n f(y_i|\mathbf{x}_i, \hat{\sigma}_i^2; \zeta)$, with

$$\begin{aligned} f(y_i|\mathbf{x}_i, \hat{\sigma}_i^2; \zeta) &= \int n(y_i|\mu_0 + \mathbf{x}_i^T \beta, \phi \hat{\sigma}_i^2) dG_{\mathbf{x}}(\mu_0) \\ &= \sum_{j=-\infty}^{\infty} n(y_i|\mu_{0j} + \mathbf{x}_i^T \beta, \phi \hat{\sigma}_i^2) \omega_j(\mathbf{x}_i^T \beta_{\omega}, \sigma_{\omega}), i = 1, \dots, n, \end{aligned} \tag{7a}$$

with mean (E) and variance (Var), respectively,

$$E(Y|\mathbf{x}, \hat{\sigma}^2; \zeta) = \int y dF(y|\mathbf{x}, \hat{\sigma}^2; \zeta) \quad \text{Var}(Y|\mathbf{x}, \hat{\sigma}^2; \zeta) = \int \{y - E(Y|\mathbf{x}, \hat{\sigma}^2; \zeta)\}^2 dF(y|\mathbf{x}, \hat{\sigma}^2; \zeta).$$

Using standard terminology for discrete mixture models (e.g., McLachlan and Peel, 2000), $G_{\mathbf{x}}$ is the (discrete) mixing distribution, which depends on covariates \mathbf{x} ; the component indices are given by $j = 0, \pm 1, \pm 2, \dots$; the component (kernel) probability densities are given by normal densities, $(n(y_i|\mu_{0j} + \mathbf{x}_i^T \beta, \phi \hat{\sigma}_i^2))_{j=-\infty}^{\infty}$, with component parameters $(\mu_{0j})_{j=-\infty}^{\infty}$, and with mixing weights $(\omega_j(\mathbf{x}_i^T \beta_{\omega}, \sigma_{\omega}))_{j=-\infty}^{\infty}$ that sum to 1 at every $\mathbf{x} \in \mathcal{X}$. Specifically, in the model, the mixture weights $\omega_j(\mathbf{x}_i^T \beta_{\omega}, \sigma_{\omega})$ are defined by a difference between two Normal(0,1) CDFs:

$$\omega_j(\mathbf{x}_i^T \beta_{\omega}, \sigma_{\omega}) = \Phi(\{j - \mathbf{x}_i^T \beta_{\omega}\} / \sigma_{\omega}) - \Phi(\{j - 1 - \mathbf{x}_i^T \beta_{\omega}\} / \sigma_{\omega}).$$

Thus, the mixture weights can be viewed as the categorical probabilities of a cumulative-probits regression model (e.g., McCullagh, 1980), for infinitely many ordered categories $j = 0, \pm 1, \pm 2, \dots$. Thus, it is easy to see that the mixture weights $\omega_j(\mathbf{x}_i^T \beta_{\omega}, \sigma_{\omega})$ sum to 1 at every $\mathbf{x} \in \mathcal{X}$. Finally, our mixture model (7a) provides a flexible (infinite) mixture of normal densities, conditional on any covariates \mathbf{x} of interest. The development of our model is, in part, motivated by the well known fact that any smooth probability density can be approximated arbitrarily well by a suitable discrete mixture of normal densities (e.g., Lo, 1984).

As shown in equation (7a), for a set of data $\mathcal{D}_n = \{(y_i, \mathbf{x}_i, \hat{\sigma}_i^2)\}_{i=1}^n$, the model assumes that each effect-size y_i is distributed by according to a probability density $f(y_i|\mathbf{x}_i, \hat{\sigma}_i^2; \zeta)$ that is constructed by a mixture of an infinite number of normal densities $n(y_i|\mu_{0j} + \mathbf{x}_i^T \beta, \phi \hat{\sigma}_i^2)$, having corresponding means $\mu_{0j} + \mathbf{x}_i^T \beta$ and mixture weights

$\omega_j(\mathbf{x}_i^T \boldsymbol{\beta}_\omega, \sigma_\omega)$ (for $j=0, \pm 1, \pm 2, \dots$). Therefore, given covariates \mathbf{x} and model parameters ζ , the model is flexible enough to allow the shape of the effect-size distribution (density) $f(y|\mathbf{x}, \hat{\sigma}^2; \zeta)$ to take on virtually any form; this density can be unimodal symmetric, or skewed, or more multimodal. Moreover, the model allows the entire shape and location of the effect-size distribution (density) $f(y|\mathbf{x}, \hat{\sigma}^2; \zeta)$ to change flexibly with the covariates \mathbf{x} . The model has these flexibilities, because it models the effect-size density $f(y|\mathbf{x}, \hat{\sigma}^2; \zeta)$ by infinitely many random intercept parameters $\mu_0 = (\mu_{0j})_{j=-\infty}^{\infty}$, corresponding to infinitely many covariate-dependent mixture weights $\{\omega_j(\mathbf{x}_i^T \boldsymbol{\beta}_\omega, \sigma_\omega) : j = 0, \pm 1, \pm 2, \dots\}$. Moreover, the density $f(y|\mathbf{x}, \hat{\sigma}^2; \zeta)$ of our model, and thus its mean and variance, can change as a function of \mathbf{x} , even when all slope parameters in $\boldsymbol{\beta}$ are zero, because then the corresponding slope parameters in $\boldsymbol{\beta}_\omega$ can be non-zero, for the mixture weights. Also, if $\hat{\sigma}_i^2 \neq \hat{\sigma}_k^2$, then $\text{Var}(Y|\mathbf{x}, \hat{\sigma}_i^2; \zeta) \neq \text{Var}(Y|\mathbf{x}, \hat{\sigma}_k^2; \zeta)$, even though the variance parameter ϕ is the same over all components ($j=0, \pm 1, \pm 2, \dots$), an assumption that is not uncommon in practice.

In the Bayesian nonparametric model, the parameter σ_ω controls the level of multimodality of $f(y|\mathbf{x}, \hat{\sigma}^2; \zeta)$. To explain, assume for the moment that $\mathbf{x}_i^T \boldsymbol{\beta} = 0$ and $f(y|\mathbf{x}) = f(y|\mathbf{x}, \hat{\sigma}^2 = 1; \zeta)$, for simplicity, and with no loss of generality. On the one hand, a small value of σ_ω indicates that $f(y|\mathbf{x})$ is unimodal, that is, modeled as a unimodal normal density $n(y_i|\mu_j, \hat{\sigma}_j^2)$ for a j satisfying $j-1 < \mathbf{x}^T \boldsymbol{\beta}_\omega < j$, with mixture weight $\omega_j(\mathbf{x}^T \boldsymbol{\beta}_\omega, \sigma_\omega)$ near 1. This is because the function $\Phi(\mathbf{x}^T \boldsymbol{\beta}_\omega / \sigma_\omega)$ is approximately 0 for $\mathbf{x}^T \boldsymbol{\beta}_\omega < 0$, while it is approximately 1 for $\mathbf{x}^T \boldsymbol{\beta}_\omega > 0$. As σ_ω approaches infinity, the mixture weights become more spread out, and then $f(y|\mathbf{x})$ becomes multimodal, with each mixture weight $\omega_j(\mathbf{x}^T \boldsymbol{\beta}_\omega, \sigma_\omega)$ above zero and much less than 1. These ideas are illustrated in Figure 2, which plots the mixture weights and the corresponding density of our model, $f(y|\mathbf{x})$, for a range of σ_ω , given $\mathbf{x}^T \boldsymbol{\beta}_\omega = .7$, given samples of (μ_j, σ_j^2) from a normal-gamma distribution. As shown, the conditional density $f(y|\mathbf{x})$ is unimodal when σ_ω is small, and $f(y|\mathbf{x})$ becomes more multimodal as σ_ω increases. The level of multimodality in the data is indicated by the posterior distribution of σ_ω , under the model.

The Bayesian nonparametric meta-analytic model (7a) is completed by the specification of a joint proper prior density $\pi(\zeta)$ for the infinitely many model parameters $\zeta = (\boldsymbol{\beta}, \gamma, (\mu_{0j})_{j=-\infty}^{\infty}, \phi, \boldsymbol{\beta}_\omega, \sigma_\omega)$, according to the joint prior distributions:

$$\mu_{0j} | \sigma_0^2 \sim n(0, \sigma_0^2), \quad j = 0, \pm 1, \pm 2, \dots \tag{8a}$$

$$\boldsymbol{\beta}_0 \sim n(0, v \rightarrow \infty) \tag{8b}$$

$$(\beta_k, \gamma_k) | \phi \sim n(\beta | 0, \phi v_1^{\gamma_k} v_0^{1-\gamma_k}) \cdot .5^{\gamma_k} (1 - .5)^{1-\gamma_k}, \quad k = 1, \dots, p \tag{8c}$$

$$\phi^{-1} \sim \text{ga}(a_\phi/2, a_\phi/2) \tag{8d}$$

$$\sigma_0^2 \sim \text{un}(0, b_0) \tag{8e}$$

$$(\boldsymbol{\beta}_\omega, \sigma_\omega) \sim n_{p+1}(\boldsymbol{\beta}_\omega | 0, \sigma_\omega^2 10^5 \mathbf{I}_{p+1}) \text{ga}(\sigma_\omega^{-2} | 1, 1). \tag{8f}$$

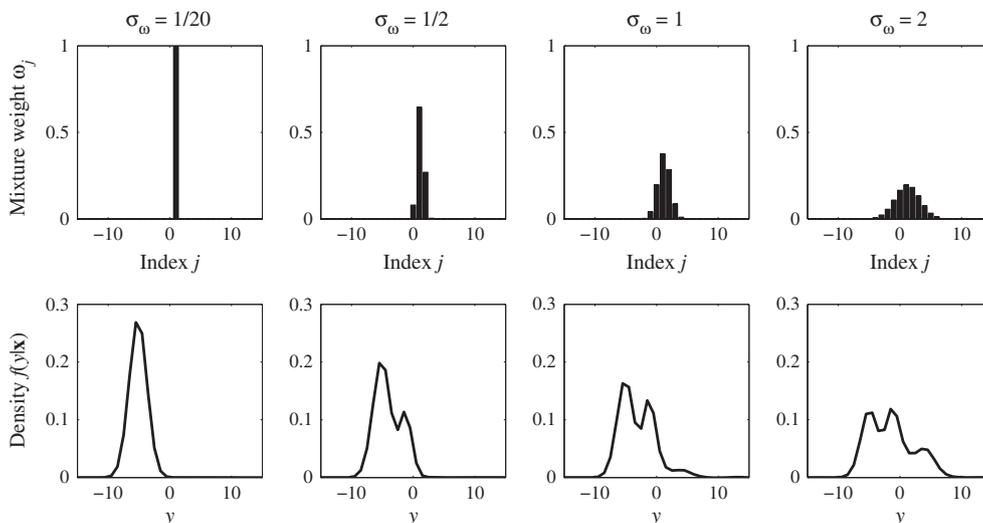


Figure 2. Effect-size density $f(y|\mathbf{x})=f(y|\mathbf{x}; \hat{\sigma}^2=1; \zeta)$ for $\sigma_\omega = 1/20, 1/2, 1, 2$, given $\mathbf{x}^T \boldsymbol{\beta} = .7$ and given sampled values of (μ_j, σ_j^2) .

As shown, a diffuse prior is assigned to the overall mean effect-size parameter β_0 . Also, as shown in (8c), we adapt the default spike-and-slab priors, to enable automatic covariate (predictor) selection in the posterior distribution of our model (George and McCulloch, 1997). The gamma prior for the inverse dispersion parameter ϕ^{-1} has mean $E(\phi^{-1}) = 1$ and variance $\text{Var}(\phi^{-1}) = \frac{2}{a_\phi}$, with a_ϕ indicating the degree of 'belief' in this prior (Nam et al., 2003). Furthermore, we specify uniform prior density $\text{un}(\sigma_0|0, b_0)$ for the variance σ_0^2 of the random intercepts $(\mu_{0j})_{j=-\infty}^{\infty}$, for a reasonably large value b_0 , as consistent with previous recommendations (Gelman, 2006, Section 7.1). Alternatively, one may specify a half-t prior density for σ_0 . Most of the prior distributions in (8) represent default and rather diffuse choices of prior, which can be used in general meta-analytic applications where prior information is typically limited. Of course, if for a given meta-analytic data set there is more prior (e.g., scientific) information available about one or more of the model parameters, then the prior distributions can be modified accordingly.

It is instructive to relate the parameters of the general normal meta-analytic model of equation (1) that are assigned a general prior density of equation (4) (Sections 2.2 and 2.4), with the parameters of the Bayesian nonparametric meta-analytic model. Across both models, the linear regression coefficients β , including the overall effect-size mean β_0 , have the same interpretation of how the mean effect-size depends on covariates. In fact, as previously shown (Karabatsos and Walker, 2012), the estimation of the regression coefficients β in the Bayesian nonparametric meta-analytic model follows from the standard theory in Bayesian normal linear models (fixed-effects or random-effects). Also, the parameters $\gamma = (\gamma_1, \dots, \gamma_p)^T$ have the same interpretation as (random) indicators of which covariates are included as significant predictors of the models, and the parameter σ_0^2 has the same interpretation as the variance of the level-2 random intercepts. A key difference is that the Bayesian nonparametric model is a discrete mixture model, which specifies a covariate-dependent infinite-mixture distribution G_x for the random intercept parameter μ_0 , as opposed to a hierarchical model or a random effects model. In contrast, a normal random-effects model is a hierarchical model, which specifies a normal mixture distribution G for the random intercept parameter, such that the mixture distribution is not covariate-dependent. Moreover, for the Bayesian nonparametric model, the discrete mixture distribution G_x induces (random) clusterings among the n study reports y_i (via the posterior distribution of the model), in terms of the random intercept parameter μ_0 . This clustering feature of the Bayesian nonparametric model enables the model to account for correlations among the study reports y_i ($i = 1, \dots, n$), lessening the need to specify a non-diagonal sampling covariance matrix $\hat{\Sigma}_n$ to account for correlated effect sizes.

Following Bayes' theorem, the data likelihood $L(\mathcal{D}_n) = \prod_{i=1}^n f(y_i|\mathbf{x}_i, \hat{\sigma}_i^2; \zeta)$ updates the prior density $\pi(\zeta)$, to a posterior density $\pi(\zeta|\mathcal{D}_n)$, given by equation (2). Then the posterior predictive density is given by equation (3), which gives an estimator of the true effect-size density in the study population, given data \mathcal{D}_n and covariates \mathbf{x} of interest. Also, recall that for the task of covariate selection, the k th covariate can be viewed as a 'significant predictor,' when the posterior inclusion probability of the covariate, $\text{Pr}[\gamma_k = 1|\mathcal{D}_n]$, is at least .5 (Barbieri and Berger, 2004). Finally, the level of multimodality in the density $f(y|\mathbf{x}, \hat{\sigma}^2, \zeta)$ is indicated by the posterior distribution of σ_ω .

Karabatsos and Walker (2012) describe the MCMC methods can be used to perform inference of the posterior density $\pi(\zeta|\mathcal{D}_n)$, of the posterior predictive density $f_n(y|\mathbf{x}, \hat{\sigma}^2)$ of the model, and to perform inference of any functional of these densities.

4. Bayesian predictive model assessment methods

Model selection is the practice of comparing different models that are fitted to a common sample data set, and then identifying the single model that best describes or predicts the underlying population distribution of the sample data. In meta-analytic practice, it is often of interest to perform model selection (e.g., Sutton, 2000, Section 11.7.3). For example, model selection is used in meta-analysis to choose between the fixed-effects and random-effects model (Borenstein et al., 2010), or to select important predictors of the effect-size in a regression setting (Higgins and Thompson, 2004).

After M meta-analytic models are fit to a data set, \mathcal{D}_n , the predictive performance of each Bayesian model $m \in \{1, \dots, M\}$ can be assessed by the mean-square posterior predictive-error criterion (Laud and Ibrahim, 1995; Gelfand and Ghosh, 1998)

$$D(m) = \sum_{i=1}^n \{y_i - E_n(Y_i|\mathbf{x}_i, \hat{\sigma}_i^2, m)\}^2 + \sum_{i=1}^n \text{Var}_n(Y_i|\mathbf{x}_i, \hat{\sigma}_i^2, m) \tag{9}$$

$$= \sum_{i=1}^n \int (y - y_i)^2 f_n(y|\mathbf{x}_i, \hat{\sigma}_i^2, m) dy = \sum_{i=1}^n D_i(m). \tag{10}$$

The criterion (9) is a standard criterion that is often used for the assessment and comparison of Bayesian models (e.g., Gelfand and Banerjee, 2010). Among the M Bayesian meta-analytic models that are compared, the model with the smallest value of $D(m)$ is identified as the one that best describes the underlying population

distribution of the given sample data set \mathcal{D}_n . The first term of (9) measures data goodness-of-fit, and the second term is a penalty that is large for models, which either over-fit or under-fit the data, as in other classical model selection criteria. Taking the square root, $\sqrt{D(m)}$, makes the criterion interpretable on the original scale of the effect size (y). Similarly, the individual square-root quantities $\sqrt{D_i(m)}$ (for $i=1, \dots, n$) can provide a detailed assessment about a model's predictive performance. A large value of $\sqrt{D_i(m)}$ would indicate that the observed effect-size y_i is an outlier under the model.

5. Illustration

In this section, we illustrate all the methods that we presented in Sections 2–4, through the meta-analysis of a large real data set involving 24 covariates. In this data analysis, we use the $D(m)$ predictive mean-square error criterion to compare the predictive accuracy of the Bayesian nonparametric meta-analysis model, and the various Bayesian normal fixed-effects and normal random-effects models. We also compare some of the parameter estimates between these Bayesian models, as well as the parameter estimates of normal fixed-effects and random-effects models estimated either under full maximum likelihood or restricted maximum likelihood. For each data analysis, each covariate was previously z-standardized to have mean 0 and variance 1, by taking $x_{ki} = (x'_{ki} - \hat{\mu}'_k) / \hat{\sigma}'_{X(k)}$ for $i=1, \dots, n$, given the mean and standard deviation $(\hat{\mu}'_k, \hat{\sigma}'_{X(k)})$ of the original ovariate data $(x'_{k1}, \dots, x'_{kn})$. Then the estimate of the intercept parameter (β_0) is interpretable as the mean study effect-size, and the β coefficients are all interpretable on a common scale.

In total, we consider 16 Bayesian meta-analytic models for the data set, including the Bayesian nonparametric model, along with various normal fixed-effects models and 2-level or 3-level normal random-effects models, which differ as to whether or not they have covariates and whether or not they have spike-and-slab priors for covariate selection. For all of these models, we specified the same prior densities for parameters that the models shared in common, in order to place the Bayesian model comparisons on a rather equal-footing. These priors were generally consistent with the recommendations of the previous literature (Sections 2.4, 3). Specifically, for all Bayesian models, as follows, we assigned the normal prior density $\pi(\beta_0) = n(\beta_0 | 0, v \rightarrow \infty)$ (with $v = 10^5$); we assigned diffuse normal priors $\pi(\beta_k) = n(\beta_k | 0, v \rightarrow \infty)$ (with $v = 10^5$), $k = 1, \dots, p = 24$, to the slope coefficients of all models with covariates and without spike-and-slab priors; we assigned hyper-prior variances $v_0 = .001$ and $v_1 = 10$, and Bernoulli prior parameters $Pr(\gamma_k = 1) = \frac{1}{2}$ ($k = 1, \dots, p = 24$), for all models with covariates and spike-and-slab priors (in terms of equation (6)); we assigned the uniform prior density $\pi(\sigma_0) = un(\sigma_0 | 0, 100)$ with large scale (100) for the variance parameter of the Bayesian nonparametric model and all normal 2-level random effects models; we specified the rather non-informative prior for (σ_0^2, ψ) (equation (5), Section 2.4), for all Bayesian 2-level normal random-effects models that allow for correlated random intercepts (Stevens and Taylor, 2009); and we specified the uniform prior density $\pi(\sigma_{00}) = un(\sigma_{00} | 0, 100)$ with large scale (100) for the level-3 variance parameter σ_{00}^2 , for all 3-level normal random-effects models. For the Bayesian nonparametric model in particular, we specified rather diffuse (high-variance) priors $\phi^{-1} \sim ga(\sigma_\omega^{-2} | .5/2, .5/2)$, $\beta_\omega | \sigma_\omega^2 \sim n_{p+1}(\beta_\omega | 0, \sigma_\omega^2 10^5 \mathbf{I}_{p+1})$, and $\sigma_\omega^{-2} \sim ga(\sigma_\omega^{-2} | 1, 1)$ (equations (8d) and (8f), Section 3).

For the application of each Bayesian model, the posterior distribution of the parameters was estimated on the basis of 200,000 MCMC samples, after trace plots indicated that MCMC samples of key model parameters and of the $D(m)$ criterion stabilized and mixed well, and after 95% Monte Carlo confidence intervals of these quantities attained half-widths that were small for practical purposes, that is, that typically ranged between .01 and .05, and not exceeding .1. Again, these procedures accord with previous recommendations for checking the quality of MCMC-based posterior estimates, based on a single MCMC run (Geyer, 1992, 2011, Chapter 1). All Bayesian meta-analytic models were estimated using menu-driven software that the first author developed (based on work on an earlier paper, Karabatsos and Walker, 2012), and is available on his web page. Finally, each of the twelve 2-level and 3-level random-effects models, assuming uncorrelated random intercepts, was also estimated by full maximum likelihood and restricted maximum likelihood, using the *nml* package (Pinheiro *et al.*, 2010) of the R statistical software (R Development Core Team, 2012).

5.1. Behavioral genetics data

Antisocial behavior, which includes aggression, willingness to violate rules and laws, defiance of adult authority, and violation of social norms (Walker *et al.*, 2004), is the most frequent reason that children are referred for mental health services in schools (Adelman and Taylor, 2010). Yet, it is the most intractable of all behavior and mental health problems, is challenging to treat, and must be addressed across the lifespan (Moffitt, 1993, 2005).

To advance understanding and treatment, many behavioral genetic studies have investigated the heritability of antisocial behavior, by correlating ratings of antisocial behavior among monozygotic (MZ) identical twin pairs, and among dizygotic (DZ) fraternal twin pairs. In each study, the ratings were carried out either by the mother, father, teacher, self, or an observer. Then the heritability, defined as the proportion of phenotypic variance explained by genetic factors, is estimated by twice the difference between the MZ correlation and DZ correlation, for twins of

the same sex. Specifically, for a given gender, suppose that n_{MZ} monozygotic identical twin pairs yield a correlation $\hat{\rho}_{MZ}$ on an antisocial behavior trait, such as conduct disorder, aggression, delinquency, and externalizing behavior. Also, suppose that n_{DZ} dizygotic fraternal twin pairs yield a correlation $\hat{\rho}_{DZ}$ on the same trait. Then the heritability of the antisocial behavior trait is estimated by (Falconer and Mackay, 1996)

$$\hat{h}^2 = 2(\hat{\rho}_{MZ} - \hat{\rho}_{DZ}) \tag{11}$$

This effect-size statistic (11) has sampling variance:

$$\hat{\sigma}^2 = 4 \left[\left\{ (1 - \hat{\rho}_{MZ}^2)^2 / n_{MZ} \right\} + \left\{ (1 - \hat{\rho}_{DZ}^2)^2 / n_{DZ} \right\} \right]$$

We identified 29 independent studies that provided the information necessary to estimate antisocial behavior heritability, for the $n=71$ independent samples (study reports) of MZ-DZ twin comparisons (Talbot *et al.*, 2012). These studies were published during years 1966–2009, and their full references are listed in the Appendix. There were two to three heritability estimates per study on average, and each study provided between 1 to 10 estimates. The left panel of Figure 3 presents the heritability estimates of antisocial behavior (i.e., the effect-size observations), stratified by gender, by rater type, and by the studies, which were numerically identified by publication year order (Appendix). The one slightly negative estimate (−.06) may have resulted from sampling error (Gill and Jensen, 1968), as suggested by its relatively large variance.

Here, it is of interest to perform a meta-analysis of the studies, to learn about the overall heritability (effect-size) distribution for the underlying study population, as well as to learn how heritability changes with key study-level covariates. Again, analyses will be performed using the various fixed-effects and random-effects models under maximum-likelihood estimation, and using our Bayesian nonparametric model. A total of 24 covariates were identified. They include publication year, the square root of the heritability variance (denoted SE(ES)) to provide an investigation of publication bias; indicators (0–1) of female status (49%) versus male; ten indicators of antisocial behavior ratings carried out by mother (mean = .53), father (.06), teacher (.24), self (.15), independent observer (.04), and ratings carried out on conduct disorder (.03), aggression (.40), delinquency (.10), and externalizing (.48) antisocial behavior; an indicator of whether a weighted average of heritability measures was taken within study over different groups of raters who rated the same twins (28% of cases, for which the 10 indicator covariates are scored as group proportions), mean age of the study subjects in months (overall mean = 119.8, s.d. = 49.5); and indicators of hi-majority ($\geq 60\%$) white twins in study (94%), zygosity obtained by questionnaire (80%), or through DNA samples (68%), study inclusion of low socioeconomic (SES) status level subjects (20%) and mid-to-high SES subjects (90%), missing SES information (10%), representative sample (85%), longitudinal sample (85%), and location of the study in terms of latitude and longitude.

Given the structure of the data, with the $n=71$ heritability (effect-size) estimate reports nested within the 29 studies, any one of at least 15 normal fixed-effects models or normal random-effects models can be considered for the purposes of meta-analysis. Specifically, they include the following: fixed-effects models; 2-level random-effects models, with level-2 random intercepts μ_0 assumed to be independent via the specification of a multivariate normal $n_n(\mu_0|0, \sigma_0^2 I_n)$ distribution, with either a structure that has each of the 71 independent samples of MZ-DZ twin comparisons defining its own group, or a grouping structure has each of the 29 studies

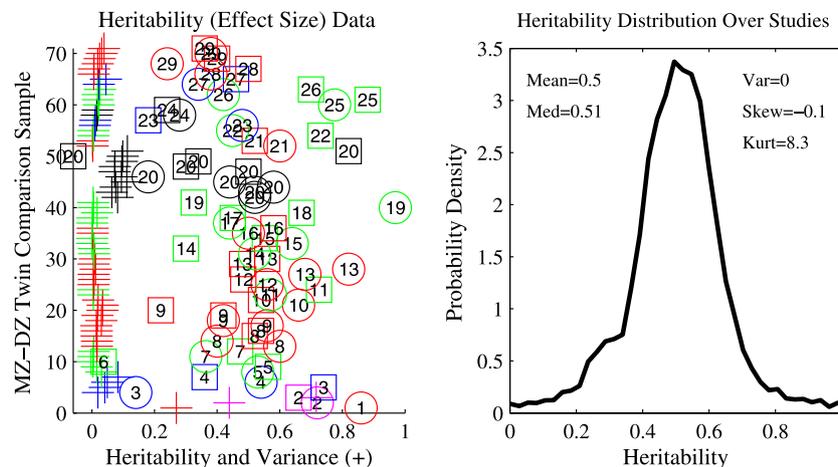


Figure 3. Left panel: heritability estimate, and its variance (+), for each of the 71 independent samples of MZ and DZ twin comparisons, provided by 29 studies. A circle refers to women, square refers to men. Also, red refers to mother rater, black to teacher rater, blue to self-rater, magenta to observer rater, and green to mixed raters. The study identification number is located in the given square or circle, and this numbering is according to the order of publication date. Right panel: Bayesian estimate of the heritability distribution, over the universe of studies. Med: median; Var: variance; Skew: skewness; Kurt: kurtosis.

defining its own group; 2-level random-effects models that allow for dependent level-2 random intercepts via the specification of a multivariate normal $n_n(\boldsymbol{\mu}_0|0, \sigma_0^2\mathbf{I}_n + \psi\mathbf{M}_n)$ distribution (Stevens and Taylor, 2009) with the binary (0–1) matrix \mathbf{M}_n indicating which pairs of the 71 study reports belong in the same study; 3-level random-effects models, with the 71 heritability (effect-size) estimate reports (level 2) nested within the 29 studies (level 2), and with random intercepts modeled, respectively, by the multivariate normal densities $n_n(\boldsymbol{\mu}_0|0, \sigma_0^2\mathbf{I}_n)$ and $n_n(\boldsymbol{\mu}_{00}|0, \sigma_{00}^2\mathbf{I}_n)$; and with each model having either no covariates, or having all 24 covariates with no spike-and-slab priors for automatic covariate selection, or having all 24 covariates along with spike-and-slab priors for automatic covariate selection via the posterior distribution. Finally, we also analyze the data set using the Bayesian meta-analytic model, which includes all 24 covariates, and which assigns spike-and-slab priors for automatic covariate selection via the posterior distribution. Alternatively, we could have also specified study indicators (0,1) as covariates in the model, in order to explicitly account for the multiple, dependent effect sizes within studies. But we excluded these predictors in order to keep the illustration of our model as simple as possible.

For the behavioral genetics data, Table 2 compares the estimate of the mean-squared predictive-error criterion, $D(m)$, between the 15 Bayesian normal fixed-effects models and normal random-effects models, and the Bayesian nonparametric meta-analytic model. Among all the models compared, the Bayesian nonparametric model attained the smallest value of the predictive mean-square error criterion $D(m)$, by a relatively large margin. Hence, among all the models compared, the Bayesian nonparametric model provides the best description of the study-population effect-size distribution that underlies the (sample) behavioral genetics data set. The meta-analytic models that attained the second best value of the criterion had about eight times the mean-squared predictive error, compared to our Bayesian nonparametric model. Also, for the Bayesian nonparametric model, the 5-number summary of the estimates of the predictive residuals $\sqrt{D_i(m)}$ was (.0, .0, .1, .1, .3), over the 71 heritability effect-size observations y_i . So none of the observations appeared to be outliers under the model.

The mean heritability (effect-size) estimate ($\hat{\beta}_0$) was quite similar among all the 16 Bayesian models, ranging from .49 to .51; while the posterior mean estimates ($\hat{\sigma}_0^2, \hat{\sigma}_{00}^2, \hat{\psi}$) of the random intercept variance parameters ranged between .00 to .02. Similar estimates were obtained from normal fixed-effects and random-effects models that were estimated either under maximum likelihood, or by restricted maximum likelihood. For all the Bayesian models assigned spike-and-slab priors, the posterior inclusion probabilities $Pr[\gamma = 1 | \mathcal{D}_{71}]$ did not exceed .05 for all 24 covariates, well below the significance threshold of .5. For the Bayesian nonparametric model, the marginal posterior mean (standard deviation) estimate of the dispersion parameter ϕ was .09 (.03).

As mentioned, the Bayesian nonparametric model is an infinite mixture model that is able to account for all possible shapes and locations of effect-size distributions, including all normal distributions. Meanwhile, in terms of the predictive mean-square error criterion $D(m)$, the Bayesian nonparametric model far-outperformed all other normal fixed-effects and normal random-effects models, which assume more strictly assume normal effect-size densities. These facts together suggest that the data set violates the assumptions of effect-size normality. For the Bayesian nonparametric model, the right panel of Figure 4 presents the posterior predictive estimate of the overall heritability (effect-size) distribution, for the underlying population of studies. This estimate is conditional on the covariates $\mathbf{x} = \mathbf{x}_0 = (1, 0, \dots, 0)^T$, and also, it is conditioned on the minimum effect-size variance σ_i^2 of .0001 over all 71 heritability reports, so that this distribution reflects information from a large-sample study. According to this estimate of the overall heritability (effect-size) distribution, there is some evidence of skewness (–.1), with the overall mean (.50) and median (.51) heritability (effect-size) being slightly different. Moreover, the figure suggests that there are two modes in this distribution, one at about the mean of .50, and the other at about .35, meaning there are about two ‘significant’ heritabilities (effect sizes) in the population, not only one.

Table 2. For the behavioral genetics data, a comparison of the predictive mean square error criterion, between the Bayesian nonparametric model, and various normal fixed-effects (FE) models, and various normal 2-level (2L), dependent 2-level (D2L), and 3-level random effects (RE) models, each of which either has no covariates (0), or 24 covariates (x), or 24 covariates with stochastic search variable selection (ss); also, each 2L or D2L model groups the random intercept parameters by the 29 studies (by Study), or groups them by the 71 independent samples of twins (by MZ-DZ).

Model	$D(m)$	Model	$D(m)$
BNP-ss	0.6	D2L-x	5.5
D2L-0	4.8	3L-x	5.5
2L-0, by MZ-DZ	4.8	2L-0, by Study	5.8
3L-0	4.8	FE-0	5.9
D2L-ss	5.4	2L-ss, by Study	6.0
2L-ss, by MZ-DZ	5.4	FE-ss	6.0
3L-ss	5.4	2L-x, by Study	6.0
2L-x, by MZ-DZ	5.5	FE-x	6.0

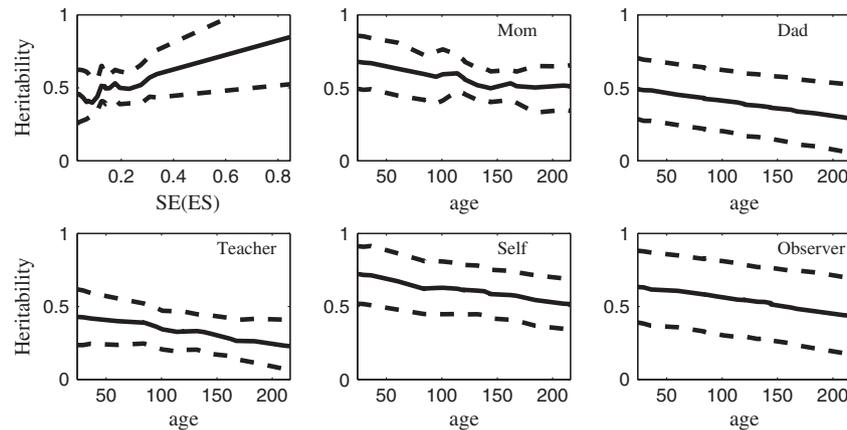


Figure 4. For each of the six panels, the estimated posterior predictive median (solid line) and interquartile range (dashed lines) of antisocial behavior heritability (effect-size), given values of a covariate, while controlling for all other covariates by fixing their values to zero. Mom: mother rater; Dad: father rater; Teacher: teacher rater; Self: self-rater; Observer: observer rater.

Upon a close inspection of the numerical estimates of the predictive density, the modes are at heritability values of .38 and .51. So both modes can provide information that contributes to the accumulation of evidence about the overall heritability (effect-size) for the substantive researchers of behavioral genetics. Assuming that this effect-size distribution is normal would lead to the mistaken conclusion that there is only one homogeneous group of heterabilities (i.e., one mode) in the effect-size distribution.

The first panel of Figure 4 shows the median (50%ile) and interquartile range (i.e., 25%ile and 75%ile) of the posterior predictive estimate of the heritability (effect-size) distribution, by SE(ES) (and by corresponding effect-size variance $\{SE(ES)\}^2 = \hat{\sigma}^2$). As shown, the median effect-size has a slight nonlinear relationship with SE(ES), but the lack of strong relationship of the median effect-size with SE(ES) further confirms a lack of evidence of publication bias in the data.

Finally, another important issue in this area of behavioral genetics deals with the issue of informant discrepancy; that is, the issue of whether the heritability (effect-size) estimates are the same across raters, or whether they depend on rater type, and further the heritability estimates may also depend on rater age. Recall that each of the 71 heritability (effect-size) estimates were obtained from ratings of antisocial behavior that were made by one of five rater types, namely, the mother, the father, the teacher, the self, or an independent observer. To address this research issue in detail, the remaining five panels of Figure 4 present the posterior predictive estimates of the effect-size distributions, conditional on covariates \mathbf{x} that indicate rater type and rater age, while conditioning on the effect-size variance estimate $\hat{\sigma}_i^2 = .0001$ that reflects a large-sample study. As shown in these panels, the heritability-age correlation seems to be slightly negative for all rater types. Moreover, the heritability distributions seemed to be similar among mother, father, self, and independent observer raters, while the teacher raters have noticeably different heritability distributions.

In conclusion, the Bayesian nonparametric model, which attained superior predictive utility of the data according to the $D(m)$ criterion, was compared to 15 other fixed-effects and normal random effects models. We interpret results of this formal model comparison as implying that the effect-size distribution is not symmetric and unimodal, conditionally on covariates. This is because, conditionally on covariates, the Bayesian nonparametric model can account for all possible effect-size distributions. In contrast, a normal fixed-effects or normal random-effects model assumes that the effect-size distribution is symmetric and unimodal. Indeed, the posterior predictive density estimates under the Bayesian nonparametric model, shown in the previous two figures, show that the effect-size distributions are multimodal and lack symmetry. Moreover, a close inspection of the numerical density estimates of Figure 4 reveals that there are two modes (clusters) of antisocial heritability, not just one, as concluded in previous meta-analyses, which had assumed unimodal normal meta-analytic models (Miles and Carey, 1997; Bergen *et al.*, 2007; Burt, 2009a, 2009b; Rhee and Waldman, 2002). This multimodal statistical conclusion implies that there are (at least) two distinct groups in the population of individuals with antisocial behavior. This has the practical implication of informing that these individuals can be better treated by approaching them as a heterogeneous population, rather than a homogeneous one. Finally, we find that the results of the model are not impacted by publication bias, given the non-significant relationship between SE(ES) and heritability effect size. This result has the practical implication of allowing us to interpret the estimates of the model with more confidence.

6. Discussion

In this paper, we have proposed a Bayesian nonparametric model for meta-analysis, and demonstrated its suitability for meta-analytic data sets that give rise to asymmetric and more multimodal effect-size population

distributions. As mentioned, the traditional normal fixed- and random-effects models, while frequently used for meta-analysis, are not fully satisfactory because they make empirically falsifiable assumptions about the data. They include the assumption that the effect sizes are normally distributed (conditionally on model parameters). As we have shown, empirical violations of such an assumption can negatively affect the accuracy of prediction of meta-analytic data.

In contrast to the traditional models, our proposed Bayesian nonparametric model flexibly accounts for all distributions of the effect sizes, including all normal distributions. For the real data set that was analyzed in this paper, this flexibility enabled the Bayesian model to provide a better description of the underlying effect-size distribution of the underlying study population. At the same time, the model provides a richer description of meta-analytic data, by allowing the data analyst to infer the whole distribution of effect sizes, over studies, and to infer how the whole distribution changes as a function of key study-level covariates. Thus, the model goes beyond the mean as the measure of an overall effect-size. Furthermore, for the given meta-analytic data set at hand, the Bayesian nonparametric model automatically identifies important study-level predictors of the mean effect-size.

Appendix A: Behavioral-Genetic Studies

The following numerical list provides citations for all 29 behavioral-genetic studies, which were subject to a meta-analysis. The list is given in the order of publication year. Each item of the list presents the identification number assigned to the given study. The identification numbers of the studies are also shown in Figure 3.

1. Scarr, S. 1966. Genetic factors in activity motivation. *Child Development* **37**: 663–673.
2. Owen, D.R., & Sines, J.O. 1970. Heritability of personality in children. *Behavior Genetics* **1**: 235–248.
3. Loehlin, J.C., & Nichols, R.C. 1976. *Heredity, environment, and personality*. Austin: University of Texas Press.
4. Rowe, D.C. 1983. Biometrical genetic models of self-reported delinquent behavior: A twin study. *Behavior Genetics* **13**: 473–489.
5. Graham, P., & Stevenson, J. 1985. A twin study of genetic influences on behavioral deviance. *Journal of the American Academy of Child Psychiatry* **24**: 33–41.
6. Lytton, H., Watts, D., & Dunn, B.E. 1988. Stability of genetic determination from age 2 to age 9: A longitudinal twin study. *Social Biology* **35**: 62–73.
7. Eaves, L.J., Silberg, J.L., Meyer, J.M., Maes, H.H., Simonoff, E., Pickles, A., Rutter, M., Neale, M.C., Reynolds, C.A., Erikson, M.T., Heath, A.C., Loeber, R., Truett, K.R., & Hewitt, J.K. 1997. Genetics and developmental psychopathology: 2. The main effects of genes and environment on behavioral problems in the Virginia twin study of adolescent behavioral development. *Journal of Child Psychology and Psychiatry* **38**: 965–980.
8. Gjone, H., & Stevenson, J. 1997. The association between internalizing and externalizing behavior in childhood and early adolescence: Genetic or environmental common influences? *Journal of Abnormal Child Psychology* **25**: 277–286.
9. Eley, T.C., Lichtenstein, P., & Stevenson, J. 1999. Sex differences in the etiology of aggressive and nonaggressive antisocial behavior: Results from two twin studies. *Child Development* **70**: 155–168.
10. Hudziak, J.J., Rudiger, L.P., Neale, M.C., Heath, A.C., & Todd, R.D. 2000. A twin study of inattentive, aggressive, and anxious/depressed behaviors. *Journal of the American Academy of Child and Adolescent Psychiatry* **39**: 469–476.
11. Arseneault, L., Moffitt, T.E., Caspi, A., Taylor, A., Rijdsdijk, F.V., & Jaffee, S.R. et al. 2003. Strong genetic effects on cross-situational antisocial behavior among 5-year-old children according to mothers, teachers, examiner-observers, and twins' self reports. *Journal of Child Psychology and Psychiatry*, **44**: 832–848.
12. Eley, T.C., Lichtenstein, P., & Moffitt, T.E. 2003. A longitudinal behavioral genetic analysis of the etiology of aggressive and nonaggressive antisocial behavior. *Development and Psychopathology*, **15**: 383–402.
13. Van Beijsterveldt, C.E.M., Bartels, M., Hudziak, J.J., & Boomsma, D.I. 2003. Causes of stability of aggression from early childhood to adolescence: A longitudinal genetic analysis in Dutch twins. *Behavior Genetics* **33**: 591–605.
14. Vierikko, E., Pulkkinen, L., Kaprio, J., Viken, R., & Rose, R.J. 2003. Sex differences in genetic and environmental effects on aggression. *Aggressive Behavior* **29**, 55–68.
15. Derks, E.M., Hudziak, J.J., Van Beijsterveldt, C.E.M., Dolan, C.V., & Boomsma, D.I. 2004. A study of genetic and environmental influences on maternal and paternal CBCL syndrome scores in a large sample of 3-year-old Dutch Twins. *Behavior Genetics* **34**: 571–583.
16. Gregory, A.M., Eley, T.C., & Plomin, R. 2004. Exploring the association between anxiety and conduct problems in a large sample of twins aged 2–4. *Journal of Abnormal Child Psychology* **32**: 111–122.
17. Van Beijsterveldt, C.E.M., Verhulst, F.C., Molenaar, P.C.M., & Boomsma, D.I. 2004. The genetic basis of problem behavior in 5-year-old Dutch twin pairs. *Behavior Genetics* **34**: 229–242.
18. Blonigen, D.M., Hicks, B.M., Krueger, R.F., Patrick, C.J., & Iacono, W.G. 2005. Psychopathic personality traits: Heritability and genetic overlap with internalizing and externalizing psychopathology. *Psychological Medicine* **35**, 637–648.

19. Burt, S.A., McGue, M., Krueger, R.F., & Iacono, W.G. 2005. Sources of covariation among the child-externalizing disorders: Informant effects and the shared environment. *Psychological Medicine* **35**, 1133–1144.
20. Haberstick, B.C., Schmitz, S., Young, S.E., & Hewitt, J.K. 2005. Contributions of genes and environments to stability and change in externalizing and internalizing problems during elementary and middle school. *Behavior Genetics* **35**: 381–396.
21. Ligthart, L., Bartels, M., Hoekstra, R.A., Hudziak, J.J., & Boomsma, D.I. 2005. Genetic contributions to subtypes of aggression. *Twin Research and Human Genetics* **8**: 483–491.
22. Saudino, K.J., Ronald, A., & Plomin, R. 2005. The etiology of behavior problems in 7-year-old twins: Substantial genetic influence and negligible shared environmental influence for parent ratings and ratings by same and different teachers. *Journal of Abnormal Child Psychology*, **33**: 113–130.
23. Tuvblad, C., Grann, M., & Lichtenstein, P. 2006. Heritability for adolescent antisocial behavior differs with socioeconomic status: gene-environment interaction. *Journal of Child Psychology and Psychiatry* **47**: 734–743.
24. Vierikko, E., Pulkkinen, L., Kaprio, J., & Rose, R.J. 2006. Genetic and environmental sources of continuity and change in teacher-rated aggression during early adolescence. *Aggressive Behavior* **32** : 308–320.
25. Van Hulle, C.A., Lemery-Chalfant, K., & Goldsmith, H.H. 2007. Genetic and environmental influences on socio-emotional behavior in toddlers: an initial twin study of the infant-toddler social and emotional assessment. *Journal of Child Psychology and Psychiatry* **48**: 1014–1024.
26. Ball, H.A., Arseneault, L., Taylor, A., Maughan, B., Caspi, A., & Moffitt, T.E. 2008. Genetic and environmental influences on victims, bullies, and bully-victims in childhood. *Journal of Child Psychology and Psychiatry*, **49**: 104–112.
27. Button, T.M. M., Lau, J.Y.F., Maughan, B., & Eley, T.C. 2008. Parental punitive discipline, negative life events and gene-environment interplay in the development of externalizing behavior. *Psychological Medicine* **38**: 29–39.
28. Saudino, K.J., Carter, A.S., Purper-Ouakil, D., & Gorwood, P. 2008. The etiology of behavioral problems and competencies in very young twins. *Journal of Abnormal Psychology* **117**: 48–62.
29. Tuvblad, C., Raine, A., Zheng, M., & Baker, L.A. 2009. Genetic and environmental stability differs in reactive and proactive aggression. *Aggressive Behavior* **35**: 437–452.

Acknowledgements

This research is supported by grant SES-1156372, from the NSF program in Methodology, Measurement, and Statistics. The author thanks the Editor, and the anonymous Associate Editor and two reviewers, for many detailed suggestions that have helped improve the presentation of this manuscript.

References

- Adelman H, Taylor L. 2010. *Mental Health in Schools: Engaging Learners, Preventing Problems, Improving Schools*. Corwin Press, Thousand Oaks, CA.
- Aitchison J. 1975. Goodness of prediction fit. *Biometrika* **62**: 547–554.
- Andrews D, Mallows C. 1974. Scale mixtures of normal distributions. *Journal of the Royal Statistical Society: Series B: Methodological* **36**: 99–102.
- Barbieri M, Berger J. 2004. Optimal predictive model selection. *Annals of Statistics* **32**: 870–897.
- Bazzett T. 2008. *An introduction to behavior genetics*. Sinauer Associates: Sunderland, MA.
- Bergen S, Gardner C, Kendler K. 2007. Age-related changes in heritability of behavioral phenotypes over adolescence and young adulthood: A meta-analysis. *Twin Research and Human Genetics* **10**: 423–433.
- Berkey C, Hoaglin D, Mosteller F, Colditz G. 1995. A random-effects regression model for meta-analysis. *Statistics in Medicine* **14**: 395–411.
- Borenstein M. 2009. Effect sizes for continuous data. In *The Handbook of Research Synthesis and Meta-Analysis (Second Edition)*, Cooper H, Hedges L, Valentine J (eds.). Russell Sage Foundation, New York; 221–235.
- Borenstein M, Hedges L, Higgins J, Rothstein H. 2010. A basic introduction to fixed-effect and random-effects models for meta-analysis. *Research Synthesis Methods* **1**: 97–111.
- Branscum AJ, Hanson TE. 2008. Bayesian nonparametric meta-analysis using Polya tree mixture models. *Biometrics* **64**: 825–833.
- Brooks S, Gelman A, Jones G, Meng X. 2011. *Handbook of Markov Chain Monte Carlo*. Chapman and Hall/CRC, Boca Raton, FL.
- Burr D, Doss H. 2005. A Bayesian semiparametric model for random-effects meta-analysis. *Journal of the American Statistical Association* **100**: 242–251.
- Burt S. 2009a. Are there meaningful etiological differences within antisocial behavior? Results of a meta-analysis. *Clinical Psychology Review* **29**: 163–178.
- Burt S. 2009b. Rethinking environmental contributions to child and adolescent psychopathology: a meta-analysis of shared environmental influences. *Psychological Bulletin* **135**: 608–637.

- Chib S, Greenberg E. 1995. Understanding the Metropolis-Hastings algorithm. *American Statistician* **49**: 327–335.
- Cooper H, Hedges L. 2009. Research synthesis as a scientific process. In *The Handbook of Research Synthesis and Meta-Analysis (Second Edition)*, Cooper H, Hedges L, Valentine J (eds.). Russell Sage Foundation, New York; 3–18.
- Cooper H, Hedges L, Valentine J. 2009. *The Handbook of Research Synthesis and Meta-Analysis (Second Edition)*. Russell Sage Foundation, New York.
- Denison D, Holmes C, Mallick B, Smith A. 2002. *Bayesian Methods for Nonlinear Classification and Regression*. John Wiley and Sons, New York.
- DuMouchel W, Normand S-L. 2000. Computer-modeling and graphical strategies for meta-analysis. In *Meta-analysis in medicine and health policy*, Stangl D, Berry D (eds.). Marcel Dekker, New York; 127–178.
- Egger M, Smith GD, Schneider M, Minder C. 1997. Bias in meta-analysis detected by a simple, graphical test. *British Medical Journal* **315**: 629–634.
- Falconer D, Mackay T. 1996. *Introduction to quantitative genetics*. Longman, London.
- Fleiss J, Berlin J. 2009. Effect size for dichotomous data. In *The Handbook of Research Synthesis and Meta-Analysis (Second Edition)*, Cooper H, Hedges L, Valentine J (eds.). Russell Sage Foundation, New York; 237–253.
- Gelfand A, Banerjee S. 2010. Multivariate spatial process models. In *Handbook of Spatial Statistics*, Gelfand A, Diggle P, Guttorp P, Fuentes M (eds.). Chapman and Hall/CRC, Boca Raton; 495–515.
- Gelfand A, Ghosh J. 1998. Model choice: a minimum posterior predictive loss approach. *Biometrika* **85**: 1–11.
- Gelfand A, Smith A. 1990. Sampling-based approaches to calculating marginal densities. *Journal of the American Statistical Association* **85**: 398–409.
- Gelman A. 2006. Prior distributions for variance parameters in hierarchical models. *Bayesian Analysis* **3**: 515–533.
- Gelman A, Carlin A, Stern H, Rubin D. 2004. *Bayesian Data Analysis (Second Edition)*. Chapman and Hall, Boca Raton, Florida.
- George E, McCulloch R. 1997. Approaches for Bayesian variable selection. *Statistica Sinica* **7**: 339–373.
- Geyer C. 1992. Practical Markov chain Monte Carlo (with discussion). *Statistical Science* **7**: 473–511.
- Geyer C. 2011. Introduction to MCMC. In *Handbook of Markov Chain Monte Carlo*, Brooks S, Gelman A, Jones G, Meng X (eds.). CRC, Boca Raton, FL; 3–48.
- Gilks W, Wang C, Yvonnet B, Coursaget P. 1993. Random-effects models for longitudinal data using Gibbs sampling. *Biometrics* **49**: 441–453.
- Gill J, Jensen E. 1968. Probability of obtaining negative estimates of heritability. *Biometrics* **24**: 517–526.
- Glass G. 1976. Primary, secondary, and meta-analysis of research. *Educational Researcher* **5**: 3–8.
- Gleser L, Olkin I. 2009. Stochastically dependent effect sizes. In *The Handbook of Research Synthesis and Meta-Analysis (Second Edition)*, Cooper H, Hedges L, Valentine J (eds.). Russell Sage Foundation, New York; 357–376.
- Harville D. 1977. Maximum likelihood approaches to variance component estimation and to related problems. *Journal of the American Statistical Association* **72**: 320–338.
- Hedges L. 1981. Distribution theory for Glass's estimator of effect size and related estimators. *Journal of Educational and Behavioral Statistics* **6**: 107–128.
- Hedges L, Vevea J. 1998. Fixed and random effects models in meta analysis. *Psychological Methods* **3**: 486–504.
- Higgins J, Thompson S. 2004. Controlling the risk of spurious findings from meta-regression. *Statistics in Medicine* **23**: 1663–1682.
- Higgins J, Thompson S, Spiegelhalter D. 2009. A re-evaluation of random-effects meta-analysis. *Journal of the Royal Statistical Society, Series A* **172**: 137–159.
- Karabatsos G, Walker S. 2012. Adaptive-modal Bayesian nonparametric regression. *Electronic Journal of Statistics* **6**: 2038–2068.
- Konstantopoulos S. 2007. Introduction to meta-analysis. In *Best Practices in Quantitative Methods*, Osbourne J (ed.). Sage, Thousand Oaks, CA; 177–194.
- Konstantopoulos S. 2011. Fixed effects and variance components estimation in three-level meta-analysis. *Research Synthesis Methods* **2**: 61–76.
- Laud P, Ibrahim J. 1995. Predictive model selection. *Journal of the Royal Statistical Society, Series B* **57**: 247–262.
- Lo A. 1984. On a class of Bayesian nonparametric estimates. *Annals of Statistics* **12**: 351–357.
- Louis T, Zelterman D. 1994. Bayesian approaches to research synthesis. In *The Handbook of Research Synthesis*, Cooper H, Hedges L (eds.). Russell Sage Foundation, New York; 411–422.
- McCullagh P. 1980. Regression models for ordinal data (with discussion). *Journal of the Royal Statistical Society, Series B* **42**: 109–142.
- McLachlan G, Peel D. 2000. *Finite Mixture Models*. John Wiley and Sons, New York.
- Miles D, Carey G. 1997. Genetic and environmental architecture of human aggression. *Journal of Personality and Social Psychology* **72**: 207–217.
- Moffitt T. 1993. Adolescence-limited and life-course persistent antisocial behavior: a developmental taxonomy. *Psychological Review* **100**: 674–701.
- Moffitt T. 2005. The new look of behavioral genetics in developmental psychopathology: gene-environment interplay in antisocial behaviors. *Psychological Bulletin* **131**: 533–554.
- Müller P, Quintana F. 2004. Nonparametric Bayesian data analysis. *Statistical Science* **19**: 95–110.
- Nam I-S, Mengersen K, Garthwaite P. 2003. Multivariate meta-analysis. *Statistics in Medicine* **22**: 2309–2333.

- Pinheiro J, Bates D, DebRoy S, Sarkar D, R Development Core Team. 2010. nlme: linear and nonlinear mixed effects models. *R package version 3*: 1–97.
- R Development Core Team. 2012. R: A Language and Environment for Statistical Computing. R Foundation for Statistical Computing, Vienna, Austria.
- Raudenbush S, Bryk A. 2002. Hierarchical Linear Models: Applications and Data Analysis Methods (Second Edition). Sage, Thousand Oaks, CA.
- Rhee S, Waldman I. 2002. Genetic and environmental influences on antisocial behavior: a meta-analysis of twin and adoption studies. *Psychological bulletin* **128**: 490–529.
- Stevens J, Taylor A. 2009. Hierarchical dependence in meta-analysis. *Journal of Educational and Behavioral Statistics* **34**: 46–73.
- Sutton A. 2000. Methods for Meta-analysis in Medical Research. John Wiley, Chichester.
- Talbott B, Karabatsos G, Zurheide J. 2012. Sensitivity of raters and the assessment of heritability of antisocial behavior: a meta-analysis of twin, adoption, and sibling studies. *Technical report*, University of Illinois-Chicago.
- Thompson S, Higgins J. 2002. How should meta-regression analyses be undertaken and interpreted? *Statistics in Medicine* **21**: 1559–1573.
- Thompson S, Sharp S. 1999. Explaining heterogeneity in meta-analysis: a comparison of methods. *Statistics in Medicine* **18**: 2693–2708.
- Walker H, Ramsey E, Gresham F. 2004. Antisocial Behavior in Schools: Evidence-based Practices. Wadsworth Publishing, Belmont, CA.
- West M. 1984. Outlier models and prior distributions in Bayesian linear regression. *Journal of the Royal Statistical Society: Series B: Methodological* **46**: 431–439.