A Bayesian framework for stable isotope mixing models

Erik B. Erhardt & Edward J. Bedrick

Environmental and Ecological Statistics

ISSN 1352-8505

Environ Ecol Stat
DOI 10.1007/s10651-012-0224-1

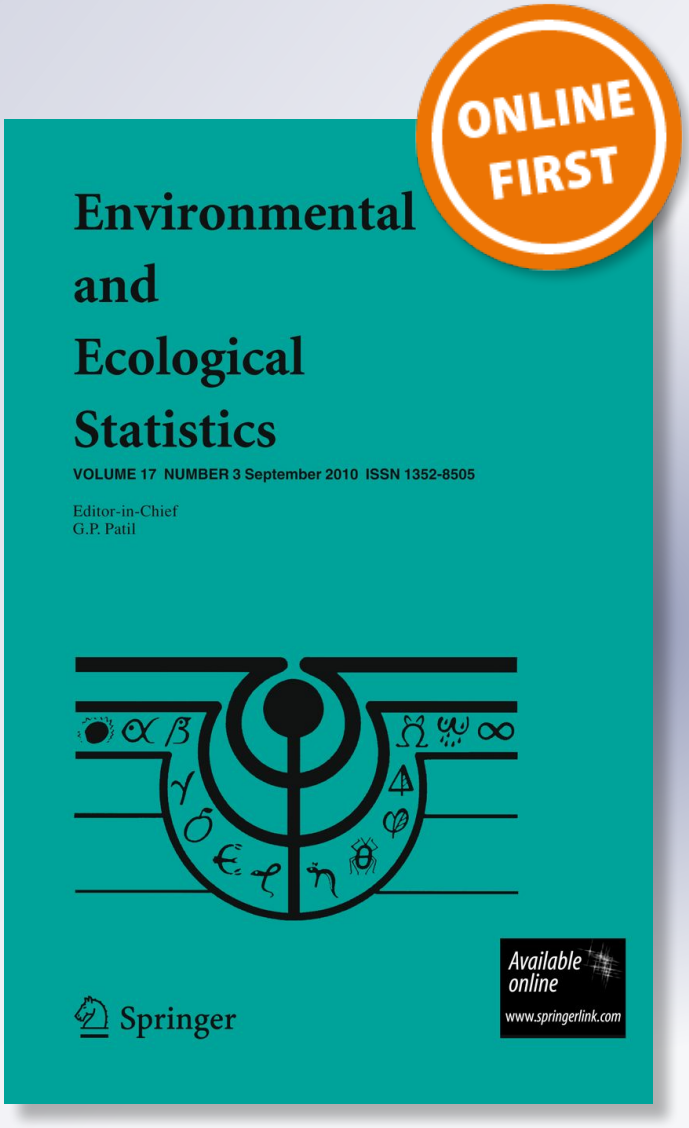

Springer

Environ Ecol Stat
DOI 10.1007/s10651-012-0224-1

A Bayesian framework for stable isotope mixing models

Erik B. Erhardt · Edward J. Bedrick

Received: 19 February 2011 / Revised: 28 September 2012

Abstract Stable isotope sourcing is used to estimate proportional contributions of sources to a mixture, such as in the analysis of animal diets and plant nutrient use. Statistical methods for inference on the diet proportions using stable isotopes have focused on the linear mixing model. Existing frequentist methods provide inferences when the diet proportion vector can be uniquely solved for in terms of the isotope ratios. Bayesian methods apply for arbitrary numbers of isotopes and diet sources but existing models are somewhat limited as they assume that trophic fractionation or discrimination is estimated without error or that isotope ratios are uncorrelated. We present a Bayesian model for the estimation of mean diet that accounts for uncertainty in source means and discrimination and allows correlated isotope ratios. This model is easily extended to allow the diet proportion vector to depend on covariates, such as time. Two data sets are used to illustrate the methodology. Code is available for selected analyses.

Keywords Animal ecology · Basic mixing model · MCMC · Resource utilization

Handling Editor: Steve Rathbun.

E. B. Erhardt (✉)
Department of Mathematics and Statistics, University of New Mexico,
Albuquerque, NM 87131, USA
e-mail: erike@stat.unm.edu

E. J. Bedrick
Division of Epidemiology and Biostatistics, Department of Internal Medicine,
University of New Mexico Health Sciences Center, Albuquerque, NM 87131, USA

Published online: 23 October 2012

1 Introduction

The goal of stable isotope sourcing is to estimate the proportional contributions of sources to a mixture. Stable isotope sourcing models are increasingly used to help understand animal diets and foodwebs, water sources in soils, plants, or water bodies, geological sources for soils or marine systems, decomposition and soil organic matter dynamics, tracing animal migration patterns, and forensics (Phillips 2001; Phillips and Gregg 2003; Martínez del Rio and Wolf 2005; Bickford et al. 2009; Hobson and Wassenaar 2008; Ehleringer et al. 2008). Animal ecology offers a rich complexity because of the preferential assimilation of elements from given sources into different tissues, so we focus our attention here. Isotopically, the consumer "is what it eats", so we aim to answer the question: each source contributes to *what proportion* of the consumer's mean diet?

Stable isotope analyses of a consumer animal's tissues (the **mixture**) and their potential prey and diet (the **sources**) is a powerful and well-studied means of quantifying relative contributions of isotopically distinct dietary components providing many benefits in comparison with traditional methods for quantifying diet, such as the analysis of stomach and fecal contents (Hobson and Wassenaar 2008, 1999). Inference on diet using stable isotopes focuses on the simplest mass-balance model, which we call the basic mixing model (BMM) (Phillips 2001). The basic mixing model states that the isotope ratio in the mixture or consumer is a convex combination of the mean isotope ratios in the source populations after correcting for discrimination. Discrimination is the mean difference of the isotope ratio in the source and how it appears in the consumer's tissues due to the assimilation process. The weights in the convex combination are the proportional contributions of the sources to the consumer's diet.

Estimating the vector of diet proportions in the BMM is a challenge because the mean isotope ratios in the source populations and the discrimination are typically unknown. Existing frequentist methods only apply when the BMM has a unique solution for the diet vector in terms of isotope ratios (Phillips and Gregg 2001). These methods can not be used in our examples. A Bayesian approach applies more generally and is especially useful in the typical setting where the number of diet sources is greater than the number of isotopes plus one, leading to a BMM that is underconstrained. Moore and Semmens (2008), Parnell and Jackson (2008), Semmens et al. (2009), Parnell et al. (2010) and Ward et al. (2010) consider Bayesian estimation in the BMM. Semmens et al. (2009) model subject-specific diets using a shrinkage estimator whereas other models focus on estimating mean diet. These models are somewhat restrictive as they assume the source isotope ratio or discrimination parameters are estimated without error and that the multivariate isotope ratio data has independent components. A further concern is their choice to specify their model with shared random effects.

We present a multivariate framework for Bayesian estimation of mean diet. Our model has three submodels, one for the consumer and one each for the estimation of source means and discrimination. The three submodels allow correlated isotope ratio data, and unlike existing models, are linked only through the defining relation for the means given by the BMM equation. To be consistent with our examples, we assume that discrimination is estimated from single-source diet experiments. Our

basic model assumes multivariate normal sampling models and conjugate priors, and for these choices a simple importance sampling algorithm is provided for posterior inferences on the diet proportion vector. Our framework provides a fairly transparent way to modify the sampling models, to consider alternative experimental methods for the estimation of discrimination, and to incorporate covariates into the analysis. We motivate the basic model with two examples and illustrate two generalizations. Section 3.8 provides a concrete discussion of differences between our model and existing Bayesian models.

Two extensions of the BMM are the concentration mixing model (CMM), which accounts for differences in the elemental concentrations in each source (Phillips and Koch 2002), and the extended mixing model (EMM), which accounts for differences in digestive efficiencies for different food types in the consumer (Martínez del Rio and Wolf 2005). The modeling challenges inherent in the BMM remain even if the additional parameters necessary in the model extensions are considered constant, as some other methods do (e.g., Parnell et al. 2010). Therefore, while these two extensions increase the realism and accuracy of the BMM, we focus on the BMM to put this fundamental statistical model on firmer statistical ground.

2 Background

Elements can exist in both stable and unstable (radioactive) forms. Each form of an element, or isotope, has the same number of protons but different numbers of neutrons. Elements of biological interest such as hydrogen, carbon, nitrogen, oxygen, and sulfur (H, C, N, O, and S) have two or more stable isotopes with the lightest of these present in much greater abundance than the others. For example, carbon and nitrogen each have one heavy stable isotope (^{13}C and ^{15}N) with a natural abundance of $\sim$1% or less and one light stable isotope (^{12}C and ^{14}N) that makes up the remainder. Carbon also has a radioactive isotope, ^{14}C. The isotope ratio, $\delta = 1000(R_{\text{sample}}/R_{\text{standard}} - 1)^{0}/_{00}$, is a normalized ratio of the number of rarer to the most common stable isotope in a sample (R_{sample}) relative to an international standard (R_{standard}) given in parts per thousand (Kendall and McDonnell 1998). A mass spectrometer is typically used to measure isotope ratios from tissue or blood samples after vaporization and ionization.

Carbon and nitrogen isotope ratios have found widespread use as biological tracers in studies of animal diets (Rundel et al. 1989). Carbon, δ^{13}C, varies among primary producers (plants and algae) with different photosynthetic pathways (C3 versus C4 photosynthesis), but changes little with trophic transfer, so is useful for identifying the source of dietary carbon (DeNiro and Epstein 1981; Inger and Bearhop 2008; Peterson and Fry 1987; Post 2002). Nitrogen, δ^{15}N, increases stepwise with trophic transfer so is used to estimate position in a foodweb (Minagawa and Wada 1984; Peterson and Fry 1987; Post 2002). Sulphur, δ^{34}S, varies among primary producers, but changes little with trophic transfer, so is used similarly to carbon, especially in marine systems (Currin et al. 1995; Peterson and Howarth 1987; Jones et al. 2010). Oxygen and hydrogen, $\delta^{18}O$ and δ^{2}H, vary across multiple spatial scales and environmental gradients (Bowen and Revenaugh 2003; Deines et al. 2009; Finlay et al. 2010; Solomon et al.

2011, 2009). For additional information, see Newsome et al. (2007) and Oulhote et al. (2010).

We formulate the BMM in terms of population means. In particular, the population mean isotope ratio among consumers β is assumed to be a convex combination of the mean isotope ratios (δ_s) in the source populations after correcting for discrimination (Δ_s) or trophic fractionation. Assuming there are I isotopes and S sources, the BMM is $\beta = \sum_{s=1}^{S} \pi_s \delta'_s$, where $\pi = (\pi_1, \ldots, \pi_S)^\top$ is the vector of proportional contributions of the S sources to the consumer's mean diet and $\delta'_s = \delta_s + \Delta_s$ for $s = 1, \ldots, S$. The source-specific discrimination terms, Δ_s, account for the consumer's ingestion, metabolization, and excretion of their diet (Minagawa and Wada 1984). Each vector of source and discrimination parameters has I elements, one for each isotope.

2.1 Motivating example: dunlin diet

To set ideas, we use data from Evans Ogden et al. (2005) to quantify the proportional use that $J = 174$ *Calidris alpina pacifica* (dunlin, a small migratory sea bird) made of farmland and marine resources on the Fraser River Delta, British Columbia, from January to April, 2000. The two sources ($S = 2$) in the dunlin diet represent protein from invertebrates that feed on plants with distinct photosynthetic pathways and isotope ratio ranges (Hobson 1999). The terrestrial source represents C3 plants, so-called because the first organic carbon compound made in photosynthesis contains three carbon atoms, while the marine source represents C4 plants.

Figure 1 provides a plot of carbon ($\delta^{13}C$) and nitrogen ($\delta^{15}N$) isotope ratios (i.e., $I = 2$) for the $J = 174$ dunlin and for samples of $K_2 = 20$ marine and $K_1 = 16$

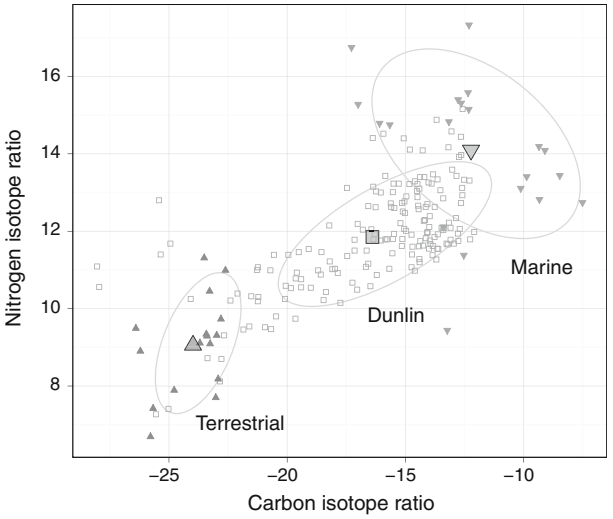

Fig. 1 Dunlin consumer and terrestrial and marine discrimination-corrected source δ^{13}C and δ^{15}N observations with means and bivariate normal 75 % probability ellipses for reference (The ellipses are summaries based on sampling model assumptions and are not implied by the linear mixing model.)

Table 1 Dunlin example: $J = 174$ observations of dunlin blood as a mixture of $S = 2$ sources using $I = 2$ isotopes of carbon and nitrogen

			Carbon		Nitrogen		
			Mean	SD	Mean	SD	Corr
Mixture		J					
Dunlin		174	−16.40	3.42	11.85	1.34	0.67
Sources		K_s					
Terrestrial	$(s = 1)$	16	−25.36	1.27	6.05	1.22	0.45
Marine	$(s = 2)$	20	−13.60	2.75	11.09	1.81	−0.35
Diet experiment		K					
Tissue (Dunlin on C3 diet)		4	−23.28	0.28	6.50	0.09	0.45
Diet (C3)		29	−24.66	0.40	3.50	0.40	0.34
Discrimination		(Tissue–Diet)	1.38		3.00		

Summaries are sample sizes, means, standard deviations, and correlations for the mixture, sources, and diet experiment discriminations (Evans Ogden et al. 2005)

terrestrial invertebrates from the sources. The data are summarized in Table 1. The plotted source data are the discrimination-corrected isotope ratio pairs $(\delta^{13}C, \delta^{15}N)_{sk} + \hat{\Delta}^\top$, where the discrimination $\hat{\Delta} = (\hat{\Delta}_C, \hat{\Delta}_N)^\top$ is assumed to be identical for the two sources and estimated as discussed in Sects. 3.2 and 4.1. The sample mean isotope ratios are superimposed on the plot. If the BMM held and the means were estimated without error then the dunlin mean would lie on a line segment joining the two discrimination-corrected source means, thus determining the population mean proportion vector $\pi = (\pi_1, \pi_2)^\top$. The goal is to estimate π accounting for uncertainty in the mean isotope ratios and discrimination.

3 Bayesian model

Our model has three submodels to estimate source, discrimination, and consumer parameters. Each submodel has a sampling model for the data given parameters and a prior distribution for the parameters. Our basic sampling models assume random samples $\mathcal{D}$ from multivariate $N(\mu, \Sigma)$ distributions. We use $g(\cdot)$ to identify a generic probability distribution.

3.1 Source model

The source mean isotope ratios $\delta_1, \ldots, \delta_S$ are estimated using samples from the S source populations. While sources may vary with time, for example, sources may come and go from the environment, we assume we know the possible sources and those sources are available over the study time. The model discussed here assumes source means are constant with respect to covariates, such as time, but the basic model can be extended to include such covariates. Let $\mathcal{D}_s = (d_{s1}, \ldots, d_{sK_s})$ be a random sample of size K_s from the sth source population, with $d_{sk} \sim N(\delta_s, \Sigma_s)$. Each d_{sk}

has I elements, one for each isotope. Assuming independence across samples and priors, the source model is $L_{\mathrm{S}} = \prod_{s=1}^{S} g(\mathcal{D}_s|\delta_s, \Sigma_s)g(\delta_s, \Sigma_s)$, where $g(\mathcal{D}_s|\delta_s, \Sigma_s)$ is a product of K_s multivariate normal densities and the prior $g(\delta_s, \Sigma_s)$ is to be specified. Let $\bar{\mathbf{d}}_s$ and $\hat{\Sigma}_s$ be the sample mean vector and covariance matrix for the sth source. All sample covariance matrices are maximum likelihood estimates (i.e., divisor is K_s rather than $K_s - 1$).

3.2 Diet experiment model for estimating discrimination

To be consistent with our examples, we assume that discrimination is estimated from E single-source diet experiments. In a single-source diet experiment consumers are fed for an extended period of time a controlled diet representative of a source. Discrimination is the mean difference between the isotope ratios in the diet source and in tissues of the consumer at diet equilibrium when there is no residual effect of previous diet on the isotope ratios (Caut et al. 2009). Single-source diet experiments are costly so sample sizes are typically small.

Discrimination is often assumed to be identical for similar sources so the number of diet experiments E is often less than S. From the eth diet experiment, we obtain a sample $\mathcal{D}_{\mathrm{D}e}$ of $K_{\mathrm{D}e}$ isotope ratio (vector) measurements for the proxy diet source and a sample $\mathcal{D}_{\mathrm{T}e}$ of $K_{\mathrm{T}e}$ from the tissue of consumers.

The population of diet and tissue isotope ratios have $N(\delta_{\mathrm{D}e}, \Sigma_{\mathrm{D}e})$ and $N(\delta_{\mathrm{T}e}, \Sigma_{\mathrm{T}e})$ distributions, respectively. Discrimination is $\Delta_e = \delta_{\mathrm{T}e} - \delta_{\mathrm{D}e}$. Let $\bar{\mathbf{d}}_{\mathrm{D}e}$ and $\bar{\mathbf{d}}_{\mathrm{T}e}$ be the sample means and let $\hat{\Sigma}_{\mathrm{D}e}$ and $\hat{\Sigma}_{\mathrm{T}e}$ be the sample covariance matrices for the diet and tissue samples from the eth diet experiment. The estimated discrimination from this experiment is $\hat{\Delta}_e = \bar{\mathbf{d}}_{\mathrm{T}e} - \bar{\mathbf{d}}_{\mathrm{D}e}$. Assuming independence within and among diet experiments, the diet model is

$$L_{\mathrm{D}} = \prod_{e=1}^{E} g(\mathcal{D}_{\mathrm{T}e}|\delta_{\mathrm{T}e}, \Sigma_{\mathrm{T}e})g(\delta_{\mathrm{T}e}, \Sigma_{\mathrm{T}e})g(\mathcal{D}_{\mathrm{D}e}|\delta_{\mathrm{D}e}, \Sigma_{\mathrm{D}e})g(\delta_{\mathrm{D}e}, \Sigma_{\mathrm{D}e}).$$

The combined source and diet model $L_{\mathrm{S}}L_{\mathrm{D}}$ depends on a set of covariance matrices Σ_* and a set of population means

$$\Theta = (\delta_1, \ldots, \delta_S, \delta_{\mathrm{D}1}, \ldots, \delta_{\mathrm{D}E}, \delta_{\mathrm{T}1}, \ldots, \delta_{\mathrm{T}E}).$$

3.3 Consumer model

We obtain a random sample $\mathcal{B} = (\mathbf{b}_1, \ldots, \mathbf{b}_J)$ of responses from consumers, with $\mathbf{b}_j|(\pi, \Sigma_b, \Theta) \sim N(\beta, \Sigma_b)$. The population mean response β depends on π and Θ through the BMM so the normal sampling model $g(\mathcal{B}|\pi, \Sigma_b, \Theta)$ is conditional on (π, Σ_b, Θ) but independent of Σ_*. The consumer model is $L_{\mathrm{C}} = g(\mathcal{B}|\pi, \Sigma_b, \Theta)g(\pi)g(\Sigma_b)$, where the priors on π and Σ_b are independent and independent of Θ. Let $\bar{\mathbf{b}}$ and $\hat{\Sigma}_b$ be the sample mean vector and covariance matrix.

3.4 Prior distributions

Each source and diet experiment sample $\mathcal{D}$ assumes a $N(\mu, \Sigma)$ model. Let $(K, \bar{\mathbf{d}}, \hat{\Sigma})$ be the sample size, mean, and covariance matrix based on $\mathcal{D}$. We follow common practice and use the conjugate prior $\mu|\Sigma \sim N(\mu_0, \Sigma/\nu_0)$ and $\Sigma \sim \text{IW}(\Sigma_0, m_0)$ (i.e., Inverse-Wishart) with density

$$g(\Sigma) \propto |\Sigma|^{-.5(m_0+I+1)} \exp\{-.5\text{trace}(\Sigma^{-1}\Sigma_0)\}.$$

For later reference note that $g(\mu, \Sigma|\mathcal{D}) = g(\mu|\mathcal{D})g(\Sigma|\mu, \mathcal{D})$, where $\mu|\mathcal{D} \sim$ Student-$t(df_P, \mu_P, \Sigma_P/df_P)$, with degrees of freedom $df_P = m_0 + 1 + K - I$, center $\mu_P = (K\bar{\mathbf{d}} + \nu_0\mu_0)/(K + \nu_0)$ and scale Σ_P/df_P where

$$(K + \nu_0)\Sigma_P = K\hat{\Sigma} + \Sigma_0 + \frac{K\nu_0}{K + \nu_0}(\bar{\mathbf{d}} - \mu_0)(\bar{\mathbf{d}} - \mu_0)^\top.$$

Furthermore, $\Sigma|(\mu, \mathcal{D}) \sim \text{IW}((K+\nu_0)(\Sigma_P + (\mu_P - \mu)(\mu_P - \mu)^\top), m_0 + K + 1)$. A limiting form of the conjugate prior is Jeffrey's prior $g(\mu, \Sigma) \propto |\Sigma|^{-.5(I+1)}$ for which $\mu|\mathcal{D} \sim t(K-I, \bar{\mathbf{d}}, \hat{\Sigma}/(K-I))$ and $\Sigma|(\mu, \mathcal{D}) \sim \text{IW}(K\hat{\Sigma} + K(\bar{\mathbf{d}}-\mu)(\bar{\mathbf{d}}-\mu)^\top, K)$. We will refer to $g(\mu, \Sigma|\mathcal{D})$ as an "intermediate posterior" or alternatively as an "updated prior".

For the consumer model, we assume $\pi \sim \text{Dirichlet}(\alpha)$, abbreviated $Dir(\alpha)$, where $\alpha = (\alpha_1, \ldots, \alpha_S)^\top$ is fixed and $\alpha_s = T\pi_{s0}$ where π_{s0} is a prior guess for the sth source proportion and T is the effective sample size for the prior. With the Dirichlet distribution location and precision (and covariance to some degree) of the diet proportion vector π can be specified. We also assume $\Sigma_b \sim \text{IW}(\Sigma_{b0}, n_b)$ or the limiting Jeffrey's prior $g(\Sigma_b) \propto |\Sigma_b|^{-.5I}$ which formally corresponds to $\Sigma_{b0} = 0$ and $n_b = -1$.

3.5 Identifiability and related issues

The linear system $\beta = \sum_{s=1}^{S} \pi_s \delta_s'$ that defines the BMM is underconstrained in the typical setting where $S > I + 1$, implying that π is not identifiable. A practical advantage of the Bayesian approach is that inferences about π do not require identifiability with the clear implication that the prior for π will impact certain features of the posterior distribution. A related issue is that the posterior distribution of π will assign mass to diet proportions for sources that are proposed but are not accessed. Thus the diet sources need to be carefully considered, especially when the BMM is underconstrained. Although we consider Dirichlet distributions, this issue applies regardless of the prior distribution for π.

3.6 Bayesian inference: general considerations and computing

Many of the parameters of the full posterior distribution $g(\text{parameters}|\text{data}) \propto L_S L_D L_C$ are of minor interest. A more focused analysis considers the parameters (π, Θ, Σ_b) that index the sampling distribution of the consumer isotope ratio

distribution. The source samples and diet experiments can then be viewed as primarily needed to generate prior information for Θ, which is required to estimate π, the feature of primary interest. In particular, each source and diet experiment sample $\mathcal{D}$ contributes $g(\mathcal{D}|\mu, \Sigma)g(\mu, \Sigma) = g(\mu, \Sigma|\mathcal{D})g(\mathcal{D})$ to the joint distribution for some (μ, Σ). As Σ does not appear in the consumer model, it can be integrated out of $g(\mu, \Sigma|\mathcal{D})$, giving $g(\mu|\mathcal{D})$ which can be used as an "updated prior" (see Sect. 3.4) along with the consumer model. That is,

$$g(\pi, \Theta, \Sigma_b|\text{data}) \propto g(\mathcal{B}|\pi, \Sigma_b, \Theta)g(\Sigma_b)g(\pi) \times \prod_{s=1}^{S} g(\delta_s|\mathcal{D}_s) \prod_{e=1}^{E} g(\delta_{Te}|\mathcal{D}_{Te})g(\delta_{De}|\mathcal{D}_{De}). \quad (1)$$

Moore and Semmens (2008) and Parnell et al. (2010) use a similar updating of priors in their models. A further reduction is obtained by integrating Σ_b out of the posterior for (π, Θ, Σ_b), leaving

$$g(\pi, \Theta|\text{data}) \propto h(\beta)g(\pi) \prod_{s=1}^{S} g(\delta_s|\mathcal{D}_s) \prod_{e=1}^{E} g(\delta_{\text{Te}}|\mathcal{D}_{\text{Te}})g(\delta_{\text{De}}|\mathcal{D}_{\text{De}}), \quad (2)$$

where $h(\beta)$ is a $t(df, \bar{\mathbf{b}}, (J\hat{\Sigma}_b + \Sigma_{b0})/Jdf)$ density evaluated at β and $df = n_b + J - I$.

The full posterior distribution can be simulated using Markov chain Monte Carlo (MCMC) in WinBUGS (Lunn et al. 2000) provided all priors are proper. However, the Markov chains have a tendency to mix poorly when the source means and discrimination δ_s and Δ_s are weakly informed, in part, because the consumer likelihood depends on Θ only through the discrimination-corrected source means $\delta'_s = \delta_s + \Delta_s$. The poor mixing is typically avoided by computing the posterior of (π, Θ, Σ_b) in (1) based on the "updated priors" for the sources and the diet experiments. This computation is easy to implement in WinBUGS and allows improper Jeffrey's priors in the source and diet submodels.

Alternatively, an importance sampling algorithm can be devised to simulate the posterior distribution of (π, Θ). The algorithm is easily programmed and only requires summary data (sample sizes, means, variances, and correlations), allowing for secondary analyses of published data. Assuming Jeffrey's priors, the steps in the algorithm are: (1) For sources $s = 1, \ldots, S$, generate $\delta^*_s \sim t(K_s - I, \bar{\mathbf{d}}_s, \hat{\Sigma}_s/(K_s - I))$. (2) For diet experiments $e = 1, \ldots, E$, generate $\delta^*_{\text{De}} \sim t(K_{\text{De}} - I, \bar{\mathbf{d}}_{\text{De}}, \hat{\Sigma}_{\text{De}}/(K_{\text{De}} - I))$ and $\delta^*_{\text{Te}} \sim t(K_{\text{Te}} - I, \bar{\mathbf{d}}_{\text{Te}}, \hat{\Sigma}_{\text{Te}}/(K_{\text{Te}} - I))$ and compute discrimination $\Delta^*_e = \delta^*_{\text{Te}} - \delta^*_{\text{De}}$. (3) For sources $s = 1, \ldots, S$, compute the discrimination-corrected source means $\delta^{'*}_s$ by adding δ^*_s and the discrimination Δ^*_e from the appropriate diet experiment. (4) Generate a source proportion vector from the prior distribution $\pi^* = (\pi^*_1, \ldots, \pi^*_S)^\top \sim Dir(\alpha)$. (5) Compute weight $w = h(b^*)$, the $t(J - I - 1, \bar{\mathbf{b}}, \hat{\Sigma}_b/(J - I - 1))$ density evaluated at $b^* = \sum_{s=1}^{S} \pi^*_s \delta^{'*}_s$. Steps 1–5 are repeated R times, giving a set $(\pi^*_r, \Theta^*_r, w_r)$ of simulated source proportion vectors, source and diet experiment mean vectors, and weights for $r = 1, \ldots, R$. Normalizing the weights $w^*_r = w_r / \sum_{l=1}^{R} w_l$ gives a sample (π^*_r, Θ^*_r) from the posterior of (π, Θ) with weights w^*_r for $r =$

$1, \ldots, R$. This algorithm is inefficient when most of the probability weight is assigned to a small fraction of the simulated samples. This can occur when the consumer sample size J is large and only a small fraction of the simulated consumer responses b^* are close to the sample mean $\bar{\mathbf{b}}$ relative to the scale of the reference t-distribution, leading to spikes in the weights $h(b^*)$. We have not encountered this problem in any of our analyses.

WinBUGS and MATLAB (MATLAB 2010) code to compute posteriors for the dunlin analysis in Sect. 4.1 is provided at

`StatAcumen.com/pub/2012_ErhardtBedrick_BFSIMM_code.zip`

Code for the other analyses is available from the authors upon request.

3.7 Model checking

Our basic model has multivariate normal sampling distributions with conjugate priors and a structural component given by the BMM. Each of the modeling assumptions should be assessed and changes made to the model as needed. A sensitivity analysis that addresses whether the priors have a significant effect on posterior inferences is an important part of this assessment. The prior on π deserves special scrutiny when the BMM is underconstrained. A key feature of the modular structure of our model is that distributional assumptions can be readily modified, as emphasized in Sect. 3.8 and in several examples.

Graphical and posterior predictive model checks are helpful for model checking. For simplicity, we consider an omnibus measure of lack of fit that is sensitive to heavy-tailed distributions. If $\mathcal{D} = (y_1, y_2, \ldots, y_n)$ is a sample from a $N(\mu, \Sigma)$ population, then a posterior predictive p-value (PPV) can be based on the discrepancy measure $Q(\mathcal{D}|\mu, \Sigma) = \sum_{i=1}^{n} c(\mathbf{y}_i|\mu, \Sigma)^2$, where $c(\mathbf{y}_i|\mu, \Sigma) = (\mathbf{y}_i - \mu)^\top \Sigma^{-1}(\mathbf{y}_i - \mu)$. The PPV is

$$\frac{1}{R}\sum_{r=1}^{R} 1_{\{Q(\mathcal{D}^r|\mu^r,\Sigma^r) > Q(\mathcal{D}_{obs}|\mu^r,\Sigma^r)\}},$$

where $\mathcal{D}^r$ is a sample of size n from a $N(\mu^r, \Sigma^r)$ distribution and (μ^r, Σ^r) is sampled from the posterior of (μ, Σ) for $r = 1, \ldots, R$. Gelman et al. (2000) discusses a variety of other predictive model checks.

The full posterior distribution can be used to obtain PPVs for consumer, source, and diet experiment model predictive distributions. Small p-values may result from misspecification of the BMM or the corresponding sampling model. Predictive p-values can also be based on the intermediate posteriors from the source and diet experiment samples; see Sect. 3.4. These predictive p-values provide checks on the source and diet experiment models without assuming the BMM holds. As the consumer model informs the estimation of the source and diet experiment model parameters through the BMM, another simple check of the BMM is to compare the posterior means for these parameters from the full posterior to the means from the intermediate posteriors.

3.8 Discussion of model and comparison with existing models

As implied by Phillips and Gregg (2001), we define the BMM as a linear relationship between the consumer population mean isotope ratios and the discrimination-corrected source population means. We assume multivariate normal sampling distributions and conjugate priors but any priors and any sampling distribution with finite means may be used in the consumer, source, and discrimination models (i.e., in L_{C}, L_{S}, and L_{D}). These models and the inferential strategies outlined above can also be modified to allow likelihood-based stratified, clustered, or spatial sampling of consumer and source populations and alternative experimental designs to estimate discrimination, such as using regression over a range of sources (Felicetti et al. 2003). For example, if the source model is changed to L_{S}^{*} then step 1 in the importance sampling algorithm has δ_s^{*} drawn from the corresponding distribution of $\delta_s|\mathcal{D}_s$ for $s = 1, \ldots, S$. Section 5 illustrates modifications of the basic model, including the modeling of diet as a function of covariates.

Our basic model is similar to the model of Parnell et al. (2010), which builds on earlier efforts by Moore and Semmens (2008). There are, however, two important differences between our models. First, we treat isotope ratio data as multivariate, allowing correlation among isotope ratios. Parnell et al. (2010) assume independence across isotopes. A more fundamental distinction concerns the specification of the model. The natural extension of Parnell et al. (2010) to our multivariate setting assumes the consumer isotope ratio responses satisfy $\mathbf{b}_j = \sum_{s=1}^{S} \pi_s (S_s + F_s) + e_j$, where the source and fractionation (or discrimination) effects are independent conditional on π with $S_s \sim N(\delta_s, \Sigma_s)$ and $F_s \sim N(\Delta_s, \Sigma_{\Delta_s})$, respectively, and $e_j \sim N(0, \Sigma_e)$. The source and fractionation random effects are shared by all consumers. The residual is subject specific. The likelihood function of Parnell et al. (2010) is based on the $N(\sum_{s=1}^{S} \pi_s(\delta_s + \Delta_s), \Sigma_e + \sum_{s=1}^{S} \pi_s^2(\Sigma_s + \Sigma_{\Delta_s}))$ marginal distribution of $\mathbf{b}_j$ given π, which is obtained by averaging out the shared random effects. Data-based updated priors on source and discrimination parameters are used to vary the parameters in the likelihood, with the resulting likelihoods averaged relative to the prior for π. The approach of Ward et al. (2010) is similar but their model does not include a subject-specific residual and assumes discrimination is known.

Our basic model has the same structure for the consumer population mean isotope ratios as the models of Parnell et al. (2010) and Ward et al. (2010) but our model is not defined in terms of shared random effects. We believe that shared random effect models are useful but do not reflect the complexity that is characteristic of many methods that are used for the estimation of discrimination, such as the diet experiments in our examples and regression over a range of sources. Bond and Diamond (2011) note that it is crucial to appropriately incorporate discrimination into the basic mixing model. A related concern is it is unclear why the model for consumer responses should be directly tied to the distribution of source and discrimination isotope ratios. Thus, for example, there is no obvious modification of the random effects models that would allow complex designs for sampling sources without impacting the distribution of the consumer responses. Our consumer model depends on sources and discrimination only through the means, allowing the distributional form for the consumer model to be unrelated to source and discrimination models, and each model to be modified

essentially independently of the others. As such, it is clear in our framework that the diet experiments are unnecessary when the population mean discriminations are known. In contrast, the random effects models include a stochastic component for fractionation in the distribution of consumer responses even when the corresponding population means are known.

4 Examples

4.1 Dunlin diet

Evans Ogden et al. (2004) assume discrimination is the same for both dunlin diet sources ($S = 2$) and estimate discrimination by conducting a diet experiment with foods of terrestrial C3 origin. The isotope ratio data from 29 diet samples and from 4 control dunlin used in the diet experiment are summarized in Table 1.

As $S = 2$, we can focus on π_1, the proportion of dunlin diet attributable to terrestrial sources. The proportion attributable to marine sources satisfies $\pi_2 = 1 - \pi_1$. Our importance sampling algorithm was used with our basic model, assuming a uniform prior on π_1 and Jeffrey's priors for the remaining parameters. The posterior mean and standard deviation of π_1 were $(0.385, 0.046)$ based on 50,000 importance samples. The posterior mean indicates that dunlin diet is approximately 40 % terrestrial and 60 % marine invertebrates. An equal tail 95 % posterior interval for π_1 is $(0.293, 0.463)$. Figure 2 plots a kernel-smoothed density estimate for π_1 based on 10000 bootstrap resamples from the importance distribution. The posterior distribution is unimodal and slightly skewed to the right.

For comparison, we used WinBUGS to compute the posterior for (π, Σ_b, Θ). The posterior mean and standard deviation of π_1 based on 50,000 samples after a 20,000 sample burn-in agreed with the importance sample summaries to 0.001. The standard error of the estimated posterior mean of π_1 was 0.002 with importance sampling and 0.001 using WinBUGS.

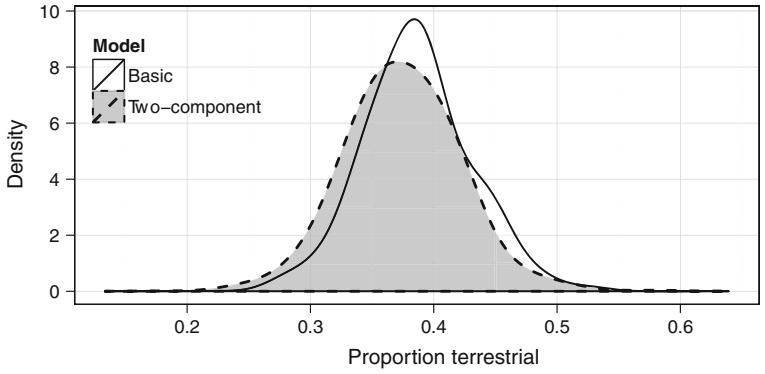

Fig. 2 Estimated posterior density of proportion of dunlin diet attributable to terrestrial sources using basic model prior (*solid line*, Sect. 4.1) and two-component mixture model (*dashed*, Sect. 5.2)

Boxplots and two-dimensional scatterplots of the consumer, source, and diet experiment samples show modest deviations from normality. Most notably, the dunlin carbon isotope ratio distribution is somewhat heavy tailed and negatively skewed. The PPV for the consumer model is 0.04. Predictive p-values based on the intermediate posteriors for the diet experiment and source models range from 0.09 for the tissue sample in the diet experiment to 0.62 for the terrestrial source. Section 5 considers alternative models for these data.

4.2 Mink diet

Table 2 and Fig. 3 summarize carbon and nitrogen stable isotope ratios ($I = 2$) from a sample of five live-captured coastal mink and seven sources ($S = 7$) of their diet, including tidal fish, blue mussels, crabs, shrimps, rodents, amphipods, and ducks (Ben-David et al. 1997). Table 2 includes summaries from two diet experiments to estimate discrimination. A beef diet was used to estimate the common discrimination for rodent and duck sources. A fish diet was used to estimate the common discrimination for the five other sources. The diet experiment data were given to us as summary statistics with no correlation information. We set the correlation to zero in our initial analyses. One very extreme outlier in the original mink sample was omitted from Table 2 and our analysis. Boxplots and two dimensional scatterplots of the consumer and source samples do not show any strong suggestion of nonnormality.

Table 2 Mink example, five observations of mink blood as a mixture of $S = 7$ sources using $I = 2$ isotopes of carbon and nitrogen

			Carbon		Nitrogen		
			Mean	SD	Mean	SD	Corr
Mixture		J					
Mink		5	−15.11	0.543	13.81	0.683	0.641
Sources		K_s					
Tidal fish	$(s = 1)$	14	−15.41	0.471	12.71	0.329	−0.744
Blue mussels	$(s = 2)$	11	−19.51	0.810	7.74	0.462	0.757
Crabs	$(s = 3)$	20	−16.27	0.968	9.20	0.664	−0.335
Shrimps	$(s = 4)$	6	−17.90	0.873	9.96	0.461	−0.105
Rodents	$(s = 5)$	15	−27.37	1.208	7.45	0.818	0.053
Amphipods	$(s = 6)$	25	−19.68	0.865	12.00	0.934	0.284
Ducks	$(s = 7)$	6	−23.38	2.374	11.29	2.182	0.510
Diet experiment		K					
Tissue (mink, beef diet)		7	−20.09	0.171	10.50	0.245	–
Diet (beef)		7	−24.22	0.728	6.19	0.142	–
Difference (T–D)			4.13	4.31			
Tissue (mink, fish diet)		10	−18.95	0.120	13.54	0.240	–
Diet (fish)		10	−20.52	0.609	11.73	0.384	–
Difference (T–D)			1.57		1.81		

Summaries are sample sizes, means, standard deviations, and correlations for mixture, source, and diet samples (Ben-David et al. 1997)

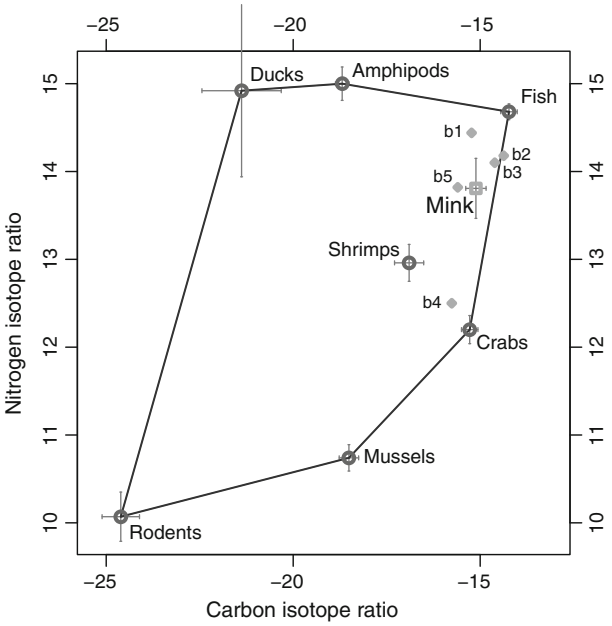

Fig. 3 Mink individuals and mean and discrimination-corrected source δ^{13}C and δ^{15}N means with marginal one standard *error bars*

We considered Jeffrey's priors and a $Dir(1, 1/6, 1/6, 1/6, 1/6, 1/6, 1/6)$ prior for π, with sources ordered as in Table 2. The Dirichlet prior is worth two observations, with a prior mean of $1/2$ for tidal fish and prior means of $1/12$ for each of the other six sources. The estimated posterior means and standard deviations for the proportion of diet attributable to the seven sources based on 50,000 importance samples are $(0.70, 0.05, 0.05, 0.08, 0.03, 0.06, 0.04)$ and $(0.16, 0.08, 0.09, 0.13, 0.05, 0.10, 0.06)$, respectively. With this prior, tidal fish and shrimp contribute roughly 70 and 10 % of an average mink's diet, respectively. The balance is made up of roughly equal contributions from the other five sources. The posterior standard deviations are large relative to the means, which is not surprising given the small mink sample. Figure 4 provides bivariate scatter plots of posterior samples for each pair of sources along with marginal histograms. The plots are based on 10,000 bootstrap resamples from the importance sampling distribution. Except for plots concerning tidal fish, the marginal and bivariate distributions are right skewed and concentrated near zero. Predictive p-values for the sources based on intermediate posteriors were between 0.34 and 0.69.

A sensitivity analysis was conducted to assess the impact of two model features on inferences. As the correlation between carbon and nitrogen isotope ratios was not reported with the diet experiments, we repeated the analysis assuming the sample correlations were either -0.50 or 0.50. This had no discernable effect on the posterior of π. We also recognized that the priors for π and Σ_b are critical because the BMM is underconstrained and the consumer sample has only five mink. We considered a

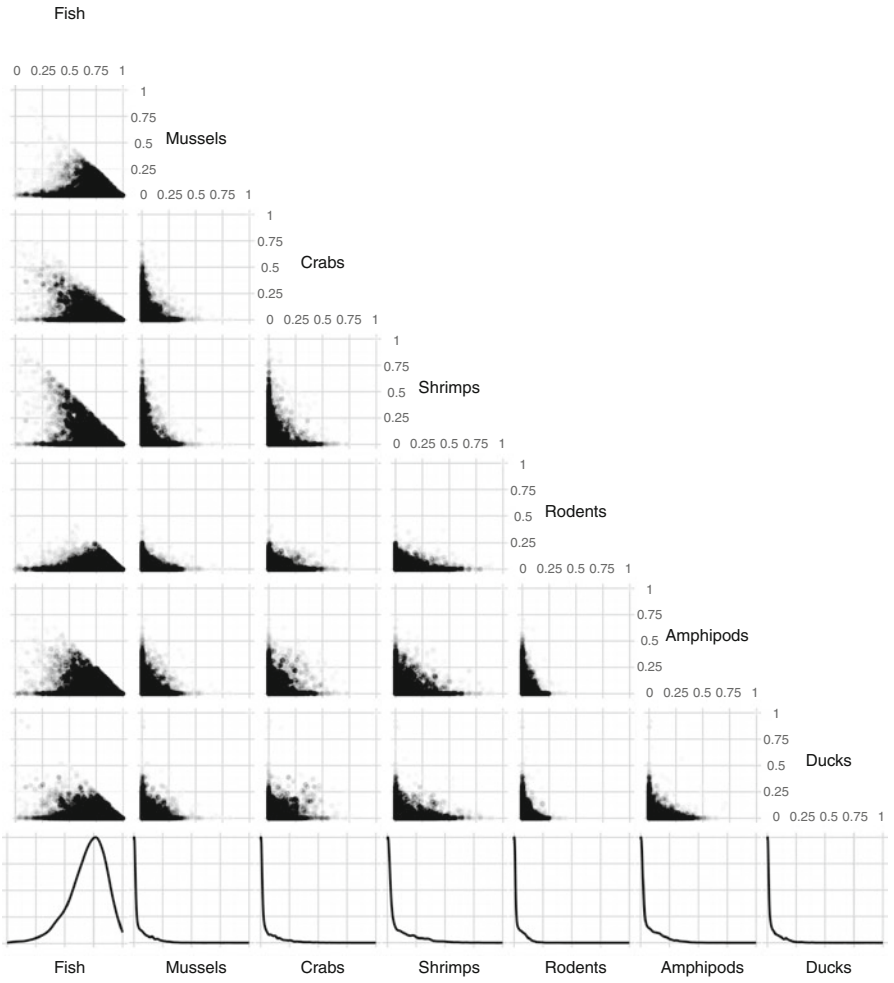

Fig. 4 Bivariate plots of posterior samples for mink diet proportion vector with posterior marginal density plots for individual sources along *bottom row*

$Dir(2/7, 2/7, 2/7, 2/7, 2/7, 2/7, 2/7)$ prior on π which is again worth two observations but assigns equal prior means of 1/7 to each diet source. With this prior, the posterior mean and standard deviation of π are (0.55, 0.06, 0.10, 0.12, 0.03, 0.08, 0.06) and (0.23, 0.09, 0.14, 0.17, 0.05, 0.11, 0.08), respectively. If, in addition, we use $\Sigma_b \sim \text{IW}(5\mathbf{I}_2, 5)$ instead of a Jeffrey's prior, the posterior mean and standard deviation of π are (0.62, 0.05, 0.08, 0.09, 0.03, 0.07, 0.06) and (0.17, 0.06, 0.11, 0.13, 0.04, 0.10, 0.09). Note that the prior mode for Σ_b is $5\mathbf{I}_2/(5+2+1) \approx 0.6\mathbf{I}_2$, which, except for the correlation, is fairly consistent with the mink consumer sample covariance matrix. Specifying $\Sigma_b \sim \text{IW}(2\mathbf{I}_2, 2)$ instead has little effect on these posterior summaries.

The main differences among the analyses rests in the contribution of tidal fish to diet, which is estimated to be between 55 and 70 %. Clearly, the priors on π on Σ_b

must be carefully considered. Small variations in the source and diet model priors had little effect on the posterior of π.

For comparison, we used WinBUGS to compute the posterior for (π, Σ_b, Θ) based on the last two sets of priors considered in the sensitivity analysis. Summaries for π based on 50,000 samples after a 20,000 sample burn-in agreed with our importance sample output to two decimal places. The Monte Carlo error in the estimated posterior mean of π was comparable for the two methods. Posterior predictive p-values for the consumer model exceeded 0.80.

5 Extending the basic model

5.1 Robustness of inferences using the basic model

The BMM is defined in terms of population means without reference to a particular distributional form. As our multivariate normal sampling distributions have functionally independent mean and covariance parameters, we expect that the estimated diet proportion vector given by the posterior mean of π might be somewhat robust to misspecification of the consumer, source, and diet experiment models. We used the dunlin data to examine this issue. Various sample sizes and diet proportion vectors were considered, treating the observed means and covariance matrices as population values except that the consumer population mean β satisfied the BMM given π and the population mean source and diet experiment isotope ratios. We estimated the posterior mean of π_1 in each of 100 samples assuming our basic model with Jeffreys' priors and a uniform prior on π_1. The "frequentist risk" measured by the root mean squared error (RMSE) and the frequentist bias in the estimate were obtained. The process was repeated assuming that the sampling distributions for the consumer and sources were linear transformations of independent exponential random variables with mean one. The transformations were chosen to match the mean and covariance structures used under normality. The RMSEs and biases were comparable to those obtained under normality. For example, the bias and RMSE when $\pi_1 = 0.20$ and 0.50 were $(-0.003, 0.039)$ and $(0.002, 0.035)$, respectively, for normal samples of the same size as the observed samples. The corresponding values were $(-0.012, 0.034)$ and $(0.002, 0.042)$ when normality was assumed but the data were exponential. Similar conclusions were obtained when the sampling distributions were mixtures of multivariate normal distributions.

This analysis suggests that our basic model provides some protection to misspecification of the sampling distributions when estimating π. A better approach to studying this issue is to directly compare posterior summaries from the basic model to alternative models, as illustrated below.

5.2 Mixture models for sources in dunlin diet

The source samples for the dunlin data are heterogeneous, reflecting subsamples of different species of invertebrates (Evans Ogden et al. 2004). Although the sample sizes

are relatively small, some suggestion of clustering is seen in the source samples plotted in Fig. 1. To assess the robustness of our initial inferences, we considered 2-component normal mixture models $p_s N(\mu_{s1}, \Sigma_{s1}) + (1-p_s) N(\mu_{s2}, \Sigma_{s2})$ as the sampling models for sources $s = 1$ and 2. The population source means are $\delta_s = p_s \mu_{s1} + (1 - p_s)\mu_{s2}$ under this model. We assumed conjugate priors $\mu_{sl}|\Sigma_{sl} \sim N(\mu_s, 100\Sigma_{sl})$, $\Sigma_{sl} \sim \text{IW}(5\mathbf{I}_2, 5)$, and $p_s \sim \text{Beta}(5, 5)$ for $s = 1, 2$. We centered the priors for the means at the observed source means: $\mu_s = \bar{\mathbf{d}}_s$ so that $\delta_s|\mathcal{D}_s$ is essentially centered at $\bar{\mathbf{d}}_s$, as with the Jeffrey's prior used in the original analysis. All other model specifications were unchanged. Updated priors for $\delta_s|\mathcal{D}_s$ were computed and used with the importance sampling algorithm to evaluate the posterior of (π_1, Θ). The primary effect of the mixture of normals model was to inflate the variability in $\delta_1|\mathcal{D}_1$ by a factor of 2.50 with no significant effect on $\delta_2|\mathcal{D}_2$. The posterior mean and standard deviation of π_1 are 0.374 and 0.047. A 95 % posterior interval for π_1 is (0.279, 0.474). Figure 2 provides a plot of a kernel-smoothed posterior density for π_1. The differences between the posterior summaries from the original model and the revised model are relatively minor.

As an alternative modification of the basic model, we computed the posterior of (π, Σ_b, Θ) under a heavier-tailed Student-$t(6, \beta, \Sigma_b)$ consumer sampling distribution. All other model specifications were unchanged. The PPV for the consumer model was 0.44, which suggests a better fit than was obtained assuming normality. However, the posterior mean and standard deviation of 0.353 and 0.047 are similar to that obtained from the basic model.

5.3 Temporal regression model for dunlin diet

Our previous analyses treated the dunlin diet as static but it is plausible to expect fluctuations in diet with changing weather. Evans Ogden et al. (2005) predict an increase in terrestrial habitat use by dunlin during periods of heavy rain and an increase in field feeding during lower temperatures and higher wind speeds. Thus, variation in mean dunlin isotope ratios might be expected over the 96 day sampling period for our data reflecting temporal changes in diet. Figure 5b plots the dunlin carbon and nitrogen isotope ratios by sampling day. The plots shows a quadratic time trend. More generally, typical diet may also depend on age and sex. However, the mean source isotope ratios should remain relatively constant over this sampling period but could depend on covariates.

To generalize the BMM, we focus on the temporal component of diet and assume that the mean diet proportion vector $\pi(t|\gamma)$ depends on time t and a parameter vector γ but that the discrimination-corrected source means are constant. To fit this model within our framework we only need to modify the consumer model, leaving the source and diet experiment models as originally specified. We assume that the consumer isotope ratios are independent given the sampling times with $\mathbf{b}_j|\gamma, \Sigma_b, \Theta, t_j \sim N(\beta(t_j|\gamma), \Sigma_b)$ for $e = 1, \ldots, J$, where $\beta(t|\gamma) = \sum_{s=1}^{S} \pi_s(t|\gamma)\delta_s'$. The consumer model is $L_C = g(\mathcal{B}|\gamma, \Sigma_b, \Theta, \mathbf{t})g(\gamma)g(\Sigma_b)$, where $g(\mathcal{B}|\gamma, \Sigma_b, \Theta, \mathbf{t})$ is a product of independent $N(\beta(t_j|\gamma), \Sigma_b)$ densities and $\mathbf{t} = (t_1, \ldots, t_J)^\top$ is the vector of sampling times. To complete the model we need to specify $\pi(t|\gamma)$. As $S = 2$ we focus on $\pi_1(t|\gamma)$

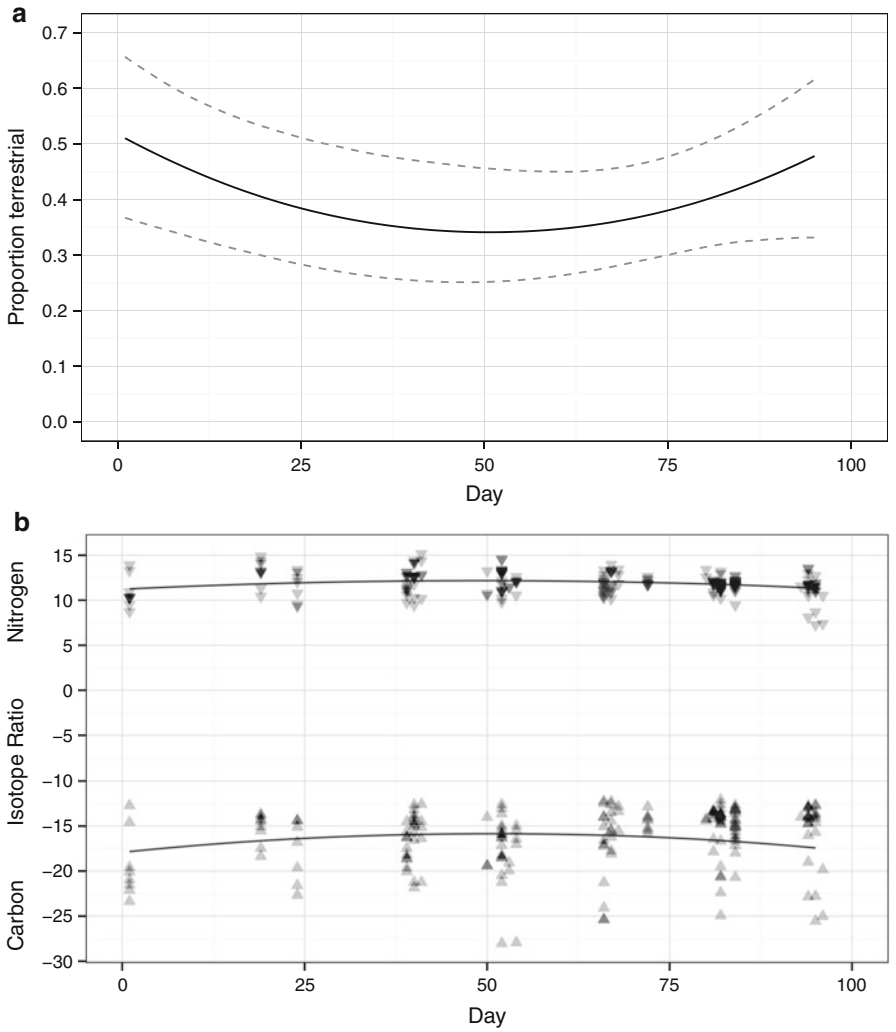

Fig. 5 **a** Estimated population diet proportion (*solid*) associated with the terrestrial source as a function of time, with pointwise 90 % posterior intervals (*dashed*) based on quadratic logit model. **b** Observed and estimated mean carbon and nitrogen isotope ratios based on quadratic fit to logit of population diet proportion associated with terrestrial source

and assume $\text{logit}(\pi_1(t|\gamma)) = \gamma_0 + .01\gamma_1(t-50) + .001\gamma_2(t-50)^2$. This choice makes empirical sense because the logistic function is approximately linear over the range $0.30 - 0.70$ in which we expect $\pi_1(t|\gamma)$ to fall and so a quadratic logistic model implies that $\beta(t|\gamma)$ is roughly a quadratic in time, which is consistent with the data.

We computed the posterior for $(\gamma, \Sigma_b, \Theta)$ in WinBUGS using the updated Jeffrey's priors for the consumer and diet experiment means. The prior for Σ_b is not that critical given the large consumer sample and we assume $\Sigma_b \sim \text{IW}(2\mathbf{I}_2, 2)$. Following principles outlined in Bedrick et al. (1996), we placed independent Beta(.8, 1.2)

priors on π_1 at t =25, 50 and 75 days and used the 1-to-1 mapping between γ and $(\pi_1(25|\gamma), \pi_1(50|\gamma), \pi_1(75|\gamma))$ to induce a prior for γ. Our prior guess for $\pi_1(t|\gamma)$ based on two prior observations is 0.40 at $t = 25$, 50 and 75 days, implying that there is no time effect on diet *a priori*. The induced prior for γ_0 is centered at approximately $\text{logit}(0.40) = -0.41$ whereas the priors for γ_1 and γ_2 are centered at approximately 0, consistent with no regression effect.

The regression vector γ has an estimated posterior mean of $(-0.67, -0.03, 0.29)$. The corresponding estimated posterior standard deviations are $(0.28, 0.51, 0.16)$. Summaries are based on 50,000 samples after a 20,000 sample burn-in. The Markov chains mixed reasonably well. Although the posterior distribution of γ_1 is concentrated near zero, there is a fairly strong suggestion that π_1 depends on time as $\Pr(\gamma_2 > 0 \mid \text{data}) = 0.96$. Figure 5a plots the estimated posterior mean of $\pi_1(t|\gamma)$ as a function of time along with pointwise 90 % posterior probability intervals for $\pi_1(t|\gamma)$. Figure 5b plots the estimated posterior mean of the dunlin mean isotope ratio $\beta(t|\gamma)$ as a function of time along with the observed carbon and nitrogen isotope ratio responses.

The bivariate normal dunlin consumer model is more consistent with the data than it was in the original analysis where time was ignored. However, the marginal distributions of the dunlin carbon isotope ratios are somewhat skewed and have outliers on certain sampling days. On days 52, 54, and 66, the PPVs indicate deficiencies with normality (all PPV = 0.01), but note that the minimum and average of the PPVs for the other 10 sampling days are 0.13 and 0.68, respectively. Although some fine tuning of the model is warranted, Fig. 5b shows that the estimated posterior mean isotope ratio $\beta(t|\gamma)$ mimics the observed trend reasonably well. This estimated trend also closely agrees with a non-parametric least squares fit, which suggests that minor modifications of the model might have minimal effect on the posterior mean of the diet proportion.

6 Concluding remarks

We presented a framework for Bayesian estimation of mean diet through the specification of three submodels, one for the consumer and one each for the estimation of source means and discrimination. The three submodels allow correlated isotope ratio data and, unlike existing models, are linked only through the defining relation for the means given by the BMM equations. This framework is flexible, and can be extended in many ways not illustrated in this paper. For example, the extended mixing model (EMM) modifies the BMM by recognizing that elemental concentration and digestive efficiencies of consumers for different food types can vary considerably (Martínez del Rio and Wolf 2005). The EMM has the same linear form as the BMM, but with additional parameters. Inference for the EMM is possible, in principle, with appropriate modification of our basic model.

Acknowledgments We are grateful to Keith Hobson for dunlin data and Merav Ben-David for mink data. We thank the editorial staff and a referee for numerous suggestions that improved the manuscript

References

Bedrick EJ, Christensen R, Johnson W (1996) A new perspective on priors for generalized linear models. J Am Stat Assoc 910(436):1450–1460

Ben-David M, Hanley TA, Klein DR, Schell DM (1997) Seasonal changes in diets of coastal and riverine mink: the role of spawning pacific salmon. Can J Zool 75:803–811

Bickford CP, Mcdowell NG, Erhardt EB, Powers HH, Hanson DT (2009) High frequency field measurements of diurnal carbon isotope discrimination and internal conductance in a semi-arid species, juniperus monosperma. Plant Cell Environ 320(7):796–810

Bond A, Diamond A (2011) Recent bayesian stable-isotope mixing models are highly sensitive to variation in discrimination factors. Ecol Appl 21:1017–1023

Bowen GJ, Revenaugh J (2003) the isotopic composition of modern meteoric precipitation. Water Resour Res 390(10):1299

Caut S, Angulo E, Courchamp F (2009) Variation in discrimination factors ($\Delta^{15}N$ and $\Delta^{13}C$): the effect of diet isotopic values and applications for diet reconstruction. J Appl Ecol 460(2):443–453

Currin CA, Newell SY, Paerl HW (1995) The role of standing dead spartina alterniflora and benthic microalgae in salt marsh food webs: considerations based on multiple stable isotope analysis. Marine Ecol Prog Ser Oldendorf 1210(1):99–116

Deines P, Wooller MJ, Grey J. (2009) Unravelling complexities in benthic food webs using a dual stable isotope (hydrogen and carbon) approach. Freshw Biol 540(11):2243–2251

DeNiro M, Epstein v. (1981) Influence of the diet on the distribution of nitrogen isotopes in animals. Geochim Cosmochim Acta 48:341–351

Ehleringer JR, Bowen GJ, Chesson LA, West AG, Podlesak DW, Cerling TE (2008) Hydrogen and oxygen isotope ratios in human hair are related to geography. Proc Natl Acad Sci 1050(8):2788

Evans Ogden LJ, Hobson KA, Lank DB (2004) Blood isotopic (δ^{13}C and δ^{15}N) turnover and diet-tissue fractionation factors in captive dunlin (Calidris alpina pacifica). The Auk 1210(1):170–177

Evans Ogden LJ, Hobson KA, Lank DB, Bittman S (2005) Stable isotope analysis reveals that agricultural habitat provides an important dietary component for nonbreeding Dunlin. Avian Conserv Ecol 10(3). http://www.ace-eco.org/vol1/iss1/art3/

Felicetti LA, Schwartz CC, Rye RO, Haroldson MA, Gunther KA, Phillips DL, Robbins CT (2003) Use of sulfur and nitrogen stable isotopes to determine the importance of whitebark pine nuts to yellowstone grizzly bears. Can J Zool 81:763–770

Finlay JC, Doucett RR, McNeely C (2010) Tracing energy flow in stream food webs using stable isotopes of hydrogen. Freshw Biol 550(5):941–951

Gelman A, Goegebeur Y, Tuerlinckx F, Van Mechelen I (2000) Diagnostic checks for discrete data regression models using posterior predictive simulations. J R Stat Soc Ser C (Appl Stat) 490(2):247–268

Hobson KA (1999) Tracing origins and migration of wildlife using stable isotopes: a review. Oecologia 1200(3):314–326

Hobson K, Wassenaar L (1999) Stable isotope ecology: an introduction Oecologia, Springer 120(3):312–313

Hobson KA, Wassenaar LI (eds) (2008) Tracking animal migration with stable isotopes. Academic Press, London

Inger R, Bearhop S (2008) Applications of stable isotope analyses to avian ecology. Ibis 1500(3):447–461

Jones AW, Dalton CM, Stowe ES, Post DM. (2010) Contribution of declining anadromous fishes to the reproductive investment of a common piscivorous seabird, the double-crested cormorant (phalacrocorax auritus). The Auk 1270(3):696–703

Kendall C, McDonnell JJ (eds) (1998) Isotope tracers in catchment hydrology. Elsevier, Amsterdam

Lunn DJ, Thomas A, Best N, Spiegelhalter D (2000) WinBUGS—a Bayesian modelling framework: concepts, structure, and extensibility. Stat Comput 100(4):325–337

Martínez del Rio C, Wolf BO Mass-balance models for animal isotopic ecology. In: Matthias Starck J, Tobias Wang , Tobias Wang (eds) (February 2005) Physiological and ecological adaptations to feeding in vertebrates, chapter 6. Science Publishers, Enfield pp 141–174

MATLAB (2010) Version 7.10.0 (R2010a) The MathWorks, Natick

Minagawa M, Wada E (1984) Stepwise enrichment of ^{15}N along food chains: further evidence and the relation between δ^{15}N and animal age. Geochim Cosmochim Acta 48:1135–1140

Moore JW, Semmens BX (2008) Incorporating uncertainty and prior information into stable isotope mixing models. Ecol Lett 11:470–480

Newsome SD, Martinez del Rio C, Bearhop S, Phillips DL (2007) A niche for isotopic ecology. Frontiers Ecol Environ 50(8):429–436
Oulhote Y, Bot BL, Deguen S, Glorennec P (2010) Using and interpreting isotope data for source identification. TrAC Trends Anal Chem
Parnell A, Jackson A (2008) Siar: stable isotope analysis in R 2008. R package version 3.3.
Parnell AC, Inger R, Bearhop S, Jackson AL (2010) Source partitioning using stable isotopes: coping with too much variation. PLoS ONE 50(3):e9672
Peterson BJ, Fry B (1987) Stable isotopes in ecosystem studies. Annu Rev Ecol Syst 18:293–320
Peterson BJ, Howarth RW (1987) Sulfur, carbon, and nitrogen isotopes used to trace organic matter flow in the salt-marsh estuaries of sapelo island, georgia. Limnol oceanogr 32(6):1195–1213
Phillips DL (2001) Mixing models in analyses of diet using multiple stable isotopes: a critique. Oecologia 127:166–170
Phillips DL, Gregg JW (2001) Uncertainty in source partitioning using stable isotopes. Oecologia 127: 171–179
Phillips DL, Gregg JW (2003) Source partitioning using stable isotopes: coping with too many sources. Oecologia 1360(2):261–269
Phillips DL, Koch PL (2002) Incorporating concentration dependence in stable isotope mixing models. Oecologia 130:114–125
Post DM (2002) Using stable isotopes to estimate trophic position: models, methods, and assumption. Ecology 83:703–718
Rundel PW, Ehleringer JR, Nagy KA (1989) Stable isotopes in ecological research. Springer, New York
Semmens BX, Ward EJ, Moore JW, Darimont CT (2009) Quantifying inter-and intra-population niche variability using hierarchical Bayesian stable isotope mixing models. PLoS One 40(7):e6187
Solomon CT, Cole JJ, Doucett RR, Pace ML, Preston ND, Smith LE, Weidel BC (2009) The influence of environmental water on the hydrogen stable isotope ratio in aquatic consumers. Oecologia 1610(2):313–324
Solomon CT, Carpenter SR, Clayton MK, Cole JJ, Coloso JJ, Pace ML, Vander Zanden MJ, Weidel BC (2011) Terrestrial, benthic, and pelagic resource use in lakes: results from a three-isotope Bayesian mixing model. Ecology 920(5):1115–1125
Ward EJ, Semmens BX, Schindler DE (2010) Including source uncertainty and prior information in the analysis of stable isotope mixing models. Environ Sci Technol 440(12):4645–4650. ISSN 0013-936X

Author Biographies

Erik B. Erhardt, PhD is an Assistant Professor of Statistics at the University of New Mexico Department of Mathematics and Statistics, where he also serves as Director of the Statistics Consulting Clinic. His research interests include Bayesian and Frequentist statistical methods for stable isotope sourcing and brain imaging. Erik is a Howard Hughes Medical Institute Interfaces Scholar collaborating in interdisciplinary research and offering consulting services in statistics.

Edward J. Bedrick, PhD is a Professor of Biostatistics at the University of New Mexico Health Sciences Center, where he serves as Director of the Biostatistics and Bioinformatics Shared Resource of the UNM Cancer Center. His research interests include model selection, general Bayesian methods and survival analysis. Ed is a Fellow of the American Statistical Association and former President of WNAR.