Methods for the integrated meta-analysis of mean and variation effects

Alistair M. Senior[1,2,3], Tim Dodgson[1,2], Malgorzata Lagisz[4,5,6], Yefeng Yang[4,5], and Shinichi Nakagawa[4,5,6]

1. Charles Perkins Centre, University of Sydney, Camperdown, NSW, 2006, Australia.
2. School of Life and Environmental Sciences, University of Sydney, Camperdown, NSW, 2006, Australia.
3. Sydney Precision Data Science Centre, University of Sydney, Camperdown, NSW, 2006, Australia.
4. Evolution and Ecology Research Centre, University of New South Wales, Sydney, NSW, 2052, Australia.
5. School of Biological Earth and Environmental Sciences, University of New South Wales, Sydney, NSW, 2052, Australia.
6. Department of Biological Sciences, University of Alberta, Edmonton, AB T6G 2E9, Canada.

Corresponding Author: A. M. Senior

Email: alistair.senior@sydney.edu.au

Tel: +64 (0) 286 270 703

Keywords: Coefficient of Variation, Diet Mixing, Effect Size, Heterogeneity, Ratio of Means, Research Synthesis

Open Research Statement: All code and data are available at https://github.com/AlistairMcNairSenior/IMAMV_Vignette.

Abstract

Meta-analyses in ecology and evolution typically focus on population means *via* effect sizes such as the log response ratio. Recently, there has been interest in quantifying effects on variability using the log variability ratio and the log coefficient of variation ratio. Until now, testing for the effects on group means and variabilities has necessitated two separate models. We present a workflow for one integrated meta-analysis of mean and variation effects, or 'IMAMV'. In a worked example, using data from the diet-mixing literature we show how the focal parameters from IMAMV match those from the equivalent two-model analysis. A common limitation to meta-analysis of variation, is unreported variance values in the primary literature. IMAMV can increase the power to detect effects on variation in meta-analytics datasets with missing variance values through 'borrowing of strength'. We show, for example, that in a dataset with 20% missing variance values, IMAMV increased the precision of the meta-analytic estimate on the variation effect by 10% compared to the conventional two-model approach. IMAMV can be implemented in commonly used software and requires no additional data beyond that used in the analysis of group means.

Introduction

Meta-analysis is widely used in many fields including ecology and evolutionary biology [1]. For meta-analysis, the user first quantifies the results from a set of comparable studies using a common effect-size metric. Analysing these effect sizes, one then estimates the overall sign and magnitude, as well as replicability, of effects in the literature [1,2]. The most widely used effect size in ecology is the log ratio of sample means (a.k.a., the log response ratio, lnRR) [3-5]. The second most popular is the standardised mean difference [3,4], which is also a mean-centric effect size. Hence, ecological meta-analysts have most often studied how phenomena affect group means, such as the difference between control and experimental treatment groups.

Over the past decade there has been increasing interest in looking beyond means to understand effects on inter-individual variability [6-9]. This interest has been supported by the development of effect sizes for meta-analysing variation effects [10,11]. Differences in variability between groups may be quantified using effect sizes such as the log variability ratio (lnVR) and log coefficient of variation ratio (lnCVR) [10,11]. Recent applications include assessing how inter-individual variation is affected by light pollution, sexual selection, and immune threats [12-14]. These methods originated in ecology and evolution [10], but have now become widespread e.g., in psychiatric medicine; [15,16-18].

Quantifying lnVR or lnCVR requires no additional data to that used for lnRR [10]. Therefore, any dataset that assesses mean effects *via* lnRR can also test for effects on variability.

In rare cases, studies have performed a two step-analysis, first analysing the mean, and then the variation e.g., [19,20,21] (Figure 1A). In other instances, the variation effects have been reported in subsequent re-analyses of data originally gathered to assess mean effects e.g., [22,23]. However, compared to lnRR, effect sizes for variation seem to be used rarely [4,23], implying that most variation effects go untested or at least unreported. Possible reasons for the underutilisation of the methods might include lack of awareness, the perceived effort of undertaking a second analysis, and/or under-reporting of variance values in the primary literature.

Here, we present an approach for the simultaneous meta-analysis of mean and variability effects using a single model; we refer to this approach as integrated meta-analysis of mean and variation effects (IMAMV; Figure 1B). We demonstrate how IMAMV offers the convenience of a single integrated model and also provides information on the correlation of mean- and variation-effects. Importantly, we show that for datasets with missing-variance values IMAMV boosts power to detect effects on the variation through 'borrowing of strength' (Figure 1B). IMAMV can be implemented in freely available and widely used R packages. For example, all models in the main text have been implemented in *metafor* [24], though packages such as *brms* or *MCMCglmm* [24-27] may also be used for more complex analyses.

As a case study, we apply IMAMV to the effects of diet-mixing on mean and among-animal variation in reproductive function. Diet-mixing studies compare the performance of groups of animals held on single- *vs* mixed-food diets. These data are well suited because they contain layers of non-independence typically seen in other

eco-evolutionary meta-analyses. The dataset contains data on the sample mean and sample variability for reproductive function in 282 groups of animals clustered into 69 experiments. *A priori*, we expect single-food diets to decrease means and increase among-animal variability [28].

Accompanying this paper is a vignette that describes the implementation of IMAMV in R (a pdf has been supplied for review "IMAMV_Vignette.pdf"), and all code and data are available at https://github.com/AlistairMcNairSenior/IMAMV_Vignette.

Key Effect Sizes and Estimators

A relatively unbiased estimator of the log population mean based on the sample mean, which we refer to as $\ln\bar{x}$, and its sampling variance ($v_{\ln\bar{x}}$) are [29]:

$$\ln\bar{x} = \log(\bar{x}) + \frac{1}{2}\left(\frac{(s/\bar{x})^2}{n}\right) = \log(\bar{x}) + \frac{1}{2}\left(\frac{\mathrm{CV}^2}{n}\right), \quad (1)$$

$$v_{\ln\bar{x}} = \frac{(s/\bar{x})^2}{n} + \frac{(s/\bar{x})^4}{2n^2} = \frac{\mathrm{CV}^2}{n} + \frac{\mathrm{CV}^4}{2n^2}, \quad (2)$$

where $\bar{x}$ is the sample mean, *s* is the sample standard deviation (SD), *n* is the sample size, and CV is the coefficient of variation (i.e., $s/\bar{x}$). The lnRR contrasts the sample means of two groups, and along with its sampling variance, can be calculated as [5]:

$$\mathrm{lnRR} = \ln\bar{x}_{\mathrm{E}} - \ln\bar{x}_{\mathrm{C}}, \quad (3)$$

$$v_{\ln RR} = v_{\ln \bar{x}_E} + v_{\ln \bar{x}_C}, \quad (4)$$

where $\ln \bar{x}_E$ and $\ln \bar{x}_C$ are the log sample means for the experimental and control groups and $v_{\ln \bar{x}_E}$ and $v_{\ln \bar{x}_C}$ are the sampling variances (note, this formulation assumes the samples are independent).

A relatively unbiased estimator of the log population SD, based on the sample SD and its sampling variance ($v_{\ln s}$) is [11,30]:

$$\ln s = \log(s) + \frac{1}{2(n-1)}, \quad (5)$$

$$v_{\ln s} = \frac{1}{2}\frac{n}{(n-1)^2}, \quad (6)$$

where all notation is as above. To quantify the difference in variation between two samples, such as experimental and control groups, one may use the lnVR [10]:

$$\ln VR = \ln s_E - \ln s_C, \quad (7)$$

$$v_{\ln VR} = v_{\ln s_E} + v_{\ln s_C}. \quad (8)$$

It may be desirable to meta-analyse studies where there is an association between the mean and the variance of the data. For example, there is a very strong association

between the log mean and log SD within the diet-mixing dataset (Figure 2A). Positive associations persist after correcting for inter-study unitary differences by centring each group on the within-study averages for log mean and log SD (Figure 2B). Hence, one may wish to understand how treatments affect variation, after correcting for any effects on the group mean. In such cases one may analyse the log CV see [31] for a discussion on the use of CV as a measure of variation. The log CV and its sampling variance can be estimated as [10]:

$$\ln\mathrm{CV} = \ln s - \ln\bar{x}, \tag{9}$$

$$v_{\ln\mathrm{CV}} = v_{\ln s} + v_{\ln\bar{x}}. \tag{10}$$

Up until now, the most common approach for the meta-analysis of variation has been to compute an effect size for difference in the log CV of two groups as lnCVR [11]:

$$\ln\mathrm{CVR} = \ln\mathrm{CV_E} - \ln\mathrm{CV_C} = \ln\mathrm{VR} - \ln\mathrm{RR}, \tag{11}$$

$$v_{\ln\mathrm{CVR}} = v_{\ln\mathrm{CV_E}} + v_{\ln\mathrm{CV_C}} = v_{\ln\mathrm{RR}} + v_{\ln\mathrm{VR}}. \tag{12}$$

We now demonstrate how these effect sizes for the mean and the variance can be analysed within a single statistical model, 'IMAMV'.

Contrast-Based IMAMV

The standard random-effects meta-analysis for a conventional 'contrast-based' analysis of effects on the mean using the lnRR can be written as:

$$\mathrm{lnRR}_j = \theta + a_j + m_j, \quad (13)$$

$$a_j \sim N(0, \sigma_a^2), \quad (14)$$

$$m_j \sim N\left(0, v_{\mathrm{lnRR}_j}\right), \quad (15)$$

where lnRR_j is the sampled effect size in the *j*th study (i.e., *j* = 1 ... *J* effect sizes) as estimated *via* eqn. 3, θ is the meta-analytic estimate of lnRR (i.e., μ_{lnRR}), a is the deviation of the true effect in the *j*th study from θ, and m_j is the deviation of lnRR_j from the true effect due to sampling. a_j is assumed to be normally distributed as per eqn. 14; σ_a is the estimated SD in effects among studies, and its square is often referred to as the heterogeneity, τ^2 (i.e., $\tau^2 = \sigma_a^2$). m_j is assumed to be normally distributed with mean 0 and SD $\sqrt{v_{\mathrm{lnRR}_j}}$ (eqn. 4). Where one wants to understand effects on variation, lnRR is substituted for lnCVR (i.e., eqns 11 and 12; Figure 1A).

We have applied eqn. 13 to the lnRR and lnCVR for the diet-mixing data using *metafor* (see Vignette). We calculated effects such that negative values indicate lower measures in the single-food group and *vice versa*. On average, single-food feeding leads to reductions in mean reproductive output (Two-Model Analysis, Table 1). The effect on

mean reproductive output is large, amounting to an 26% reduction on a single-food diet (i.e., $1 - e^{-0.30} = 0.26$). In contrast, single-food diets typically increase the CV by around 17% (i.e., $e^{0.16} = 1.17$). For both lnRR and lnCVR estimated heterogeneity is more than double the estimated effect (Table 1), suggesting considerable variance in the distribution of effects reported within the literature [32].

IMAMV is an alternative to performing separate analyses of the lnRR and lnCVR. We propose a bivariate meta-analysis that simultaneously estimates (1) the lnRR and (2) the paired differences between lnVR and lnRR. Importantly, this difference can be considered an estimate of lnCVR (see eqn. 11; Figure 1B).

In the current worked example each pair of samples has both an lnRR and lnVR, and an IMAMV version of eqn. 13 of these paired effect sizes can be written as:

$$y_{ij} = \alpha + a_j + \left(\beta_{\mathrm{lnVR}} + b_j\right) \times S_{ij} + m_{ij}, \quad (16)$$

$$\begin{bmatrix} a_k \\ b_k \end{bmatrix} \sim N\left(\begin{bmatrix} 0 \\ 0 \end{bmatrix}, \begin{bmatrix} \sigma_a^2 & \rho_{ab}\sigma_a\sigma_b \\ \rho_{ab}\sigma_a\sigma_b & \sigma_b^2 \end{bmatrix}\right), \quad (17)$$

$$m_{ij} \sim N\left(0, v_{y_{ij}}\right), \quad (18)$$

where y_{ij} is the ith effect type (i.e., i = 1 = lnRR, i = 2 = lnVR) from the jth study. Here α is the meta-analytic intercept which corresponds to the estimate for lnRR (i.e., $\alpha = \mu_{\mathrm{lnRR}}$), β_{lnVR} is the meta-analytic difference between lnVR and lnRR, and S_{ij} is a dummy

206 predictor coded as 0 if y_{ij} is a sample lnRR and 1 for lnVR. Because the term β_{lnVR} is an
207 estimate of lnVR – lnRR, it is also an estimate of lnCVR (eqn. 11). Note that, the analysis
208 explicitly pairs instances of lnRR and lnVR from the same samples, and the terms a_k and
209 b_k then give deviations of average lnRR and lnCVR for study *j* from α and β_{lnVR}. The *j*th
210 deviations are assumed bi-variate normally distributed as per eqn. 17: σ_a and σ_b give
211 the among-study SDs in lnRR and lnCVR. The term ρ_{ab} gives the correlation between
212 effects on the mean and variation effects at the between-study level. $v_{y_{ij}}$ is the
213 sampling variance for y_{ij}, estimated by eqn. 4 or 8 for instances of lnRR or lnVR,
214 respectively. We have applied eqn. 16 to the diet mixing data and the estimated point-
215 estimates for effect magnitude are identical to those from the two-model analysis
216 (Table 1). In addition, the correlation between lnRR and lnCVR was estimated to be -
217 0.58, suggesting that as single food diets generate more negative effects on the mean,
218 they simultaneously generate more variation.
219

220 Non-Independence and Random Effects in IMAMV

221 The models in eqns 13 and 16 assume that all effect sizes are independent, such that
222 each study/experiment only contains one control and one experimental group. In
223 ecology and evolution this assumption often is invalid because most datasets contain
224 non-independent effect sizes [3,4,33,34]. For example, in the diet-mixing dataset most
225 experiments yield more than one effect size. The most widely used solution to this issue
226 of non-independence is to include an additional term, to form a multi-level meta-
227 analytic model sensu [3]. In the case of IMAMV, we can formulate a multilevel extension
228 of eqn. 16 as:

$$y_{ijk} = \alpha + a_k + e_{jk} + \left(\beta_{\mathrm{lnVR}} + b_k + f_{jk}\right) \times S_{ijk} + m_{ijk}, \quad (19)$$

$$\begin{bmatrix} a_k \\ b_k \end{bmatrix} \sim N\left(\begin{bmatrix} 0 \\ 0 \end{bmatrix}, \begin{bmatrix} \sigma_a^2 & \rho_{ab}\sigma_a\sigma_b \\ \rho_{ab}\sigma_a\sigma_b & \sigma_b^2 \end{bmatrix}\right), \quad (20)$$

$$\begin{bmatrix} e_{jk} \\ f_{jk} \end{bmatrix} \sim N\left(\begin{bmatrix} 0 \\ 0 \end{bmatrix}, \begin{bmatrix} \sigma_e^2 & \rho_{ef}\sigma_e\sigma_f \\ \rho_{ef}\sigma_e\sigma_f & \sigma_f^2 \end{bmatrix}\right), \quad (21)$$

$$m_{ijk} \sim N\left(0, v_{y_{ijk}}\right), \quad (22)$$

where y_{ijk} is the *i*th effect type (i.e., *i* = 1 = lnRR, *i* = 2 = lnVR) from the *j*th pairwise contrast of treatment groups in the *k*th study; i.e., the dataset contains *k* = 1 ... *K* studies, and the *k*th study contains *j* = 1 ... *J* pairwise contrasts). a_k and b_k give deviations of average lnRR and lnCVR for study *k* from α and β_{lnVR}, while e_{jk} and f_{jk} give the deviations for the *j*th pairwise contrast in experiment *k*. σ_a and σ_b give the among-study SDs in lnRR and lnCVR respectively, while σ_e and σ_f give the within-study (i.e., between group) SDs. The terms ρ_{ab} and ρ_{ef} give the correlations between effects on the mean and variation at the between- and within-study levels, respectively. The total heterogeneity for lnRR can be estimated as $\tau^2_{ln\mathrm{RR}} = \sigma_a^2 + \sigma_e^2$, and that for lnCVR as $\tau^2_{ln\mathrm{CVR}} = \sigma_b^2 + \sigma_f^2$. Table 1 shows that the application of this model to the diet mixing dataset yields similar point estimates to the preceding analyses, though the CIs are wider, having accounted for non-independence.

Moderator Variables in IMAMV

Most ecological and evolutionary meta-analyses detect high levels of heterogeneity [32,35]. Moderators are meta-variables related to the individual effect sizes that may explain this heterogeneity. For example, in the diet mixing dataset, we have coded effect sizes by whether the focal species is terrestrial or marine dwelling. There are two common approaches to testing moderator variables: (1) stratification, and (2) meta-regression. IMAMV is compatible with both approaches. With stratification one simply subsets the data by the levels of the moderator and analyses each using separate instances of IMAMV.

Meta-regression involves fitting the moderator variable as a predictor in a model that estimates differences in the overall effect size between levels of the moderator. The IMAMV framework above already uses meta-regression, where the β_{lnVR} term in eqn. 19, estimates the lnCVR. Incorporating a moderator variable involves including an interaction between the moderator and the term estimating lnCVR. In the case of the IMAMV in eqns 19 through 22, a meta-regression including a two-level moderator coded as 0 and 1 can be formulated as:

$$y_{ijk} = \alpha + a_k + e_{jk} + \beta_{\mathrm{Mod}} \times R_{ijk} + \left(\beta_{\mathrm{lnVR}} + b_k + f_{jk}\right) \times S_{ijk} + \beta_{\mathrm{Mod\Delta lnVR}} \times V_{ijk} + m_{ijk}, \quad (23)$$

where R_{ijk} is a dummy predictor coded 0 if y_{ijk} is an effect size associated with the reference level of the moderator and 1 otherwise, β_{Mod} is the meta-analytic estimate for the difference in lnRR between levels of the moderator, V_{ijk} is a dummy predictor coded as 1 if y_{ijk} is both an estimate of the lnVR and level 1 of the moderator (and coded as 0 otherwise), $\beta_{\text{Mod}\Delta\text{lnVR}}$ is an estimate of the interaction term for the moderator and effect size type, and all other terms are as in eqn. 19. The term $\beta_{\text{Mod}\Delta\text{lnVR}}$ can be interpreted as difference in lnCVR between levels of the moderator.

Here, we have applied the meta-regression described in eqn. 23 to habitat differences in the diet mixing dataset. There are no statistically significant differences among habitats for lnCVR ($\beta_{\text{Mod}\Delta\text{lnVR}}$ = $\text{lnCVR}_{\text{Terrestriall}}$ - $\text{lnCVR}_{\text{Marine}}$ = 0.12, CI = -0.10 to 0.35). However, the reductive effect of single food diets on mean reproductive function is estimated to be stronger in terrestrial than marine habitats (β_{Mod} = $\text{lnRR}_{\text{Terrestriall}}$ - $\text{lnRR}_{\text{Marine}}$ = -0.23, CI = -0.47 to -0.002).

Borrowing of Strength and Missing Data

A limitation to meta-analysis of variation effects is missing data. Unfortunately, it is relatively common for some of the primary literature to not report the among-replicate SDs (or related metrics), which are needed to calculate lnVR or lnCVR (e.g., Figure 1A). IMAMV can boost power to detect effects on variability in datasets with missing SDs through 'borrowing of strength'. Borrowing of strength can occur in multivariate meta-analyses of the effects of the treatment on a pair of correlated outcomes (e.g., effect of an intervention on both blood-pressure and the risk of stroke) [36]. In such cases, the

correlated outcomes provide indirect information about the effects of the treatment on one another, potentially increasing the precision on the estimates in the analysis [36,37]. Borrowing of strength is expected to be particularly beneficial in cases where a random subset of studies does provide effect sizes one of the outcomes [36,37], making it a very valuable phenomenon for meta-analysis of variation where there are missing SDs.

When faced with a meta-analytic dataset with missing among-replicate SDs, we propose that the user may estimate the sampling variance of lnRR for studies with missing SDs using established methods e.g., [38], before applying IMAMV to the full lnRR and partial lnVR dataset. IMAMV is then expected to yield more precise estimate of the effects on the variation than would be obtained from a univariate analysis of the partial lnVR or lnCVR dataset.

To demonstrate this benefit of IMAMV, we have deleted the SD data from a random 20% of the entries within the diet-mixing dataset. For the complete cases, where SD was not missing, we estimated lnVR and lnCVR and the associated sampling variances as above. We estimated the lnRR and its sampling variance for every entry (i.e., including those with missing SD) following Nakagawa, Noble, Lagisz, Spake, Viechtbauer, Senior [38]. This resulted in a dataset with 331 instances of lnRR, but just 265 instances of lnCVR and lnVR. We then compared the results of the two-model analysis of lnRR and lnCVR with IMAMV (following eqn. 16). Both analyses provide comparable estimates of effect magnitude (Table 2). The estimates for lnRR also have the same standard error (SE; Table 2). However, the SE and CIs for the estimate of the lnCVR are narrower in IMAMV than in the two model analysis, despite the models containing the same effect sizes.

These differences in SE translate into a 10% gain in precision and an increase in efficiency for the IMAMV analysis (i.e., relative efficiency = 1.1 = $\frac{1/\mathrm{SE}^2_{\mathrm{IMAMV}}}{1/\mathrm{SE}^2_{2-\mathrm{model}}}$).

Additional Sources of Non-Independence

Beyond hierarchical data structures, eco-evolutionary datasets often contain other sources of non-independence. For example, by calculating multiple pair-wise effect sizes using the same control-group data contrast-based analyses can induce correlations among effect sizes; sometimes termed 'stochastic dependency' [33,39,40]. In the case of diet-mixing, consider a study that contains two single-food groups, A and B, and one mixed food group, C. In this case we have pairwise effect sizes contrasting AC and BC, duplicating the use of group C data. To a degree, contrast-based analyses can be corrected for stochastic dependency by including the estimated covariance among-correlated effect sizes [34,40]. Another solution to stochastic-dependency is to use an 'arm-based' model. Arm-based models circumvent the calculation of pairwise effect sizes prior to fitting the model. Rather, one fits the sample statistics from individual groups as outcomes (i.e., $\ln\bar{x}$ and $\ln s$ rather than lnRR and lnVR) and uses a meta-regression model to estimate the difference between treatment conditions. IMAMV is compatible with both different corrections for stochastic dependency, and the vignette gives a worked example of an arm-based IMAMV applied to the diet-mixing dataset (see Supplementary Materials and Vignette).

Another common problem in eco-evolutionary meta-analyses is phylogenetic non-independence [34,41], where we might expect more closely related taxa to display more

similar effect sizes. A solution, with which IMAMV is compatible, is to apply a phylogenetic meta-analysis. The vignette associated with this paper contains code for a phylogenetic IMAMV. The phylogenetic effects were accounted for by creating a phylogenetic covariance matrix for all species within each analysis rotl package [42], and including that matrix as a term in the model e.g., as in [27]. We have implemented the phylogenetic IMAMV using the package *metafor,* but note that there are limits to the complexity of the IMAMV that can be fitted in this package. More complex models could be implemented in *brms* or *MCMCglmm* [25] (see supplementary materials and Vignette).

Discussion

Here, we present a framework for the integrated meta-analysis of mean and variation effects (IMAMV). This approach allows the user to simultaneously meta-analyse effects on group means and variabilities, which previously necessitated two analyses. While the bivariate models presented here appear complex, the key terms from IMAMV can be interpreted equivalently to those coming from the univariate models of lnRR and lnCVR currently in use. What is more, the models themselves are very closely related to two analyses with which ecologists and evolutionary biologists may already be familiar. The first is the linear mixed-effects model (LMM) [43,44]; the most basic IMAMV in eqn. 16 has a structural similarity to a LMM containing a random-regression at a single-level. The second method with which IMAMV is similar is network meta-analysis (NMA). NMA is increasingly common in medical research [45], but still emerging in ecology and evolution c.f., [46]. Through NMA one estimates the effects of different factors/treatments on an outcome of interest, even if those factors/treatments are not directly compared in

underlying literature; e.g., one may estimate the effects of B *vs* C, from studies comparing A *vs* B and A *vs* C through their common control treatment A. The IMAMV model in eqn. 16 is structurally comparable to NMA model [47], but rather than fitting effect sizes from different treatments on the same outcomes, one fits mean and variance effect sizes from the same samples.

We have shown that IMAMV can is compatible with many of the tools that eco-evolutionary meta-analysts use to account for complex data structures (e.g., hierarchical and phylogenetic non-independence) [34]. Additionally, the usual frameworks for reporting meta-analyses are transferable. For example, tools for visualisation of contrasts-based effect sizes, such as forest plots and orchard plots [48,49], can be applied to lnRR and lnCVR. With regards to reporting heterogeneity statistics, the usual heterogeneity statistics, such as Q and $I^{2,}$ are estimable. However, we urge users to also think about direct derivatives of the estimated variation itself, such as the prediction intervals e.g., in orchard plots; [49], the 'coefficient of heterogeneity' and the closely related 'M statistic' [32]. These metrics convey a sense of the expected distribution of future effect sizes.

In addition to offering the convenience of a single analysis, IMAMV offers two explicit benefits over the two-model approach. The first is that the user gets an estimate of the correlation between effects on the mean and the variance. In some cases, this correlation itself may be of biological relevance. For example, correlations between effects on intra-genomic trait means and variances maybe indicative adaptation to a fluctuating environment through bet-hedging [50]. The second benefit of IMAMV (which

390 flows indirectly from the estimated correlation) is borrowing of strength. Through
391 borrowing of strength IMAMV can supply more precise estimates of effects on
392 variability, thus yielding increased power. As we demonstrate these benefits manifest
393 where there are missing SDs in the dataset. The problem of missing SDs is pervasive in
394 meta-analysis. One survey found around 70% of meta-analytic datasets in ecology and
395 evolution contain missing SDs, with the rate of missingness in a given dataset affecting
396 up to 30% of effect sizes [38,51]. Methods have been developed to deal with missing SD
397 data reviewed in [51]. However, these methods are designed to estimate the sampling
398 variance for mean-focussed effect sizes (e.g., lnRR and SMD) when SDs are missing and
399 are not considered appropriate when the SD is focus of the point estimate (e.g., lnVR or
400 lnCVR). As far as we can tell IMAMV is the first method proposed to help boost the
401 power of variation-focussed analyses in the presence of missing SDs. Note that to gain
402 this benefit of IMAMV the user must collect data on sample means from all studies,
403 including those that do not report SD data. Put another way, counterintuitively missing
404 SDs should not be considered an exclusion criteria for meta-analysis of the variation *via*
405 IMAMV.
406
407 Biologists work hard to gather and curate their meta-analytic datasets. However, many
408 of these datasets have not been used to their full potential by testing for effects on
409 variability. This is particularly surprising in the eco-evolutionary space, where biological
410 variability underpins core theoretical concepts e.g., niche breadth, and natural
411 selection [52]. We hope that IMAMV will allow eco-evolutionary users of meta-analysis to
412 conveniently test for the effects on both population means and within-population
413 variability.

Acknowledgements

AMS received funding from the Australian Research Council (FT230100240) as well as the University of Sydney.

Author Contributions

AMS and SN conceived the study. TD and AMS collected data. ML and YY contributed to the development of the code. AMS write the first draft of the paper. All authors contributed to the final draft and interpretation of results.

References

1. Gurevitch J, Koricheva J, Nakagawa S, Stewart G. Meta-analysis and the science of research synthesis. *Nature*. 2018;555(7695):175-182.
2. O'Dea RE, Lagisz M, Jennions MD, et al. Preferred reporting items for systematic reviews and meta-analyses in ecology and evolutionary biology: a PRISMA extension. *Biological Reviews*. 2021;96(5):1695-1722.
3. Nakagawa S, Santos ESA. Methodological issues and advances in biological meta-analysis. *Evolutionary Ecology*. 2012;26(5):1253-1274.
4. Nakagawa S, Yang Y, Macartney EL, Spake R, Lagisz M. Quantitative evidence synthesis: a practical guide on meta-analysis, meta-regression, and publication bias tests for environmental sciences. *Environmental Evidence*. 2023;12(1):8.
5. Hedges LV, Gurevitch J, Curtis PS. The meta-analysis of response ratios in experimental ecology. *Ecology*. 1999;80(4):1150-1156.
6. Nakagawa S, Schielzeth H. The mean strikes back: mean-variance relationships and heteroscedasticity. *Trends in Ecology & Evolution*. 2012;27(9):474-475.
7. Violle C, Enquist BJ, McGill BJ, et al. The return of the variance: intraspecific variability in community ecology. *Trends in Ecology & Evolution*. 2012;27(4):224-252.
8. Girard-Tercieux C, Vieilledent G, Clark A, et al. Beyond variance: simple random distributions are not a good proxy for intraspecific variability in systems with environmental structure. *Peer Community Journal*. 2024;4.
9. Des Roches S, Post DM, Turley NE, et al. The ecological importance of intraspecific variation. *Nature Ecology & Evolution*. 2018;2(1):57-64.
10. Nakagawa S, Poulin R, Mengersen K, et al. Meta-analysis of variation: ecological and evolutionary applications and beyond. *Methods in Ecology and Evolution*. 2015;6(2):143-152.
11. Senior AM, Viechtbauer W, Nakagawa S. Revisiting and expanding the meta-analysis of variation: The log coefficient of variation ratio. *Research Synthesis Methods*. 2020;11(4):553-567.
12. Yang Y, Liu Q, Pan C, et al. Species sensitivities to artificial light at night: A phylogenetically controlled multilevel meta-analysis on melatonin suppression. *Ecology Letters*. 2024;27(2):e14387.
13. Gómez-Llano M, Faria GS, García-Roa R, Noble DWA, Carazo P. Male harm suppresses female fitness, affecting the dynamics of adaptation and evolutionary rescue. *Evolution Letters*. 2024;8(1):149-160.
14. Foo YZ, Lagisz M, O'Dea RE, Nakagawa S. The influence of immune challenges on the mean and variance in reproductive investment: a meta-analysis of the terminal investment hypothesis. *BMC Biology*. 2023;21(1):107.
15. Lee M, Cernvall M, Borg J, et al. Cognitive Function and Variability in Antipsychotic Drug–Naive Patients With First-Episode Psychosis: A Systematic Review and Meta-Analysis. *JAMA Psychiatry*. 2024;81(5):468-476.
16. Brugger SP, Howes OD. Heterogeneity and Homogeneity of Regional Brain Structure in Schizophrenia: A Meta-analysis. *JAMA Psychiatry*. 2017;74(11):1104-1111.

17. Mizuno Y, McCutcheon RA, Brugger SP, Howes OD. Heterogeneity and efficacy of antipsychotic treatment for schizophrenia with or without treatment resistance: a meta-analysis. *Neuropsychopharmacology*. 2020;45(4):622-631.
18. McCutcheon RA, Pillinger T, Mizuno Y, et al. The efficacy and heterogeneity of antipsychotic response in schizophrenia: A meta-analysis. *Molecular Psychiatry*. 2021;26(4):1310-1320.
19. Moran NP, Sánchez-Tójar A, Schielzeth H, Reinhold K. Poor nutritional condition promotes high-risk behaviours: a systematic review and meta-analysis. *Biological Reviews*. 2021;96(1):269-288.
20. Trepel J, le Roux E, Abraham AJ, et al. Meta-analysis shows that wild large herbivores shape ecosystem properties and promote spatial heterogeneity. *Nature Ecology & Evolution*. 2024;8(4):705-716.
21. Hossain MAR, Olden JD. Global meta-analysis reveals diverse effects of microplastics on freshwater and marine fishes. *Fish and Fisheries*. 2022;23(6):1439-1454.
22. Senior AM, Nakagawa S, Raubenheimer D, Simpson SJ, Noble DWA. Dietary restriction increases variability in longevity. *Biol Letts*. 2017;13(3):20170057.
23. Sánchez-Tójar A, Moran NP, O'Dea RE, Reinhold K, Nakagawa S. Illustrating the importance of meta-analysing variances alongside means in ecology and evolution. *Journal of Evolutionary Biology*.
24. Viechtbauer W. Conducting meta-analyses in R with the metafor package. *J Stat Soft*. 2010;36:1-48.
25. Bürkner P-C. Advanced Bayesian multilevel modeling with the R package brms. *arXiv preprint arXiv:170511123*. 2017.
26. Hadfield JD. MCMC methods for multi-response generalized linear mixed models: The MCMCglmm R package. *J Stat Soft*. 2010;33(2):1-22.
27. Hadfield JD, Nakagawa S. General quantitative genetic methods for comparative biology: phylogenies, taxonomies and multi-trait models for continuous and categorical characters. *Journal of Evolutionary Biology*. 2010;23(3):494-508.
28. Senior AM, Nakagawa S, Lihoreau M, Simpson SJ, Raubenheimer D. An overlooked consequence of dietary mixing: a varied diet reduces inter-individual variance in fitness. *The American Naturalist*. 2015;186(5):649-659.
29. Lajeunesse MJ. Bias and correction for the log response ratio in ecological meta-analysis. *Ecology*. 2015;96(8):2056-2063.
30. Raudenbush SW, Bryk AS. Examining correlates of diversity. *Journal of Educational Statistics*. 1987;12:241-269.
31. Pélabon C, Hilde CH, Einum S, Gamelon M. On the use of the coefficient of variation to quantify and compare trait variation. *Evolution Letters*. 2020;4(3):180-188.
32. Yang Y, Noble DW, Spake R, Senior AM, Lagisz M, Nakagawa S. A pluralistic framework for measuring and stratifying heterogeneity in meta-analyses. 2023.
33. Nakagawa S, Noble DWA, Senior AM, Lagisz M. Meta-evaluation of meta-analysis: ten appraisal questions for biologists. *BMC Biology*. 2017;15(1):18.
34. Noble DWA, Lagisz M, O'Dea RE, Nakagawa S. Nonindependence and sensitivity analyses in ecological and evolutionary meta-analyses. *Molecular Ecology*. 2017;26(9):2410-2425.

35. Senior AM, Grueber CE, Kamiya T, et al. Heterogeneity in ecological and evolutionary meta-analyses: its magnitude and implications. *Ecology*. 2016;97(12):3293-3299.
36. Jackson D, White IR, Price M, Copas J, Riley RD. Borrowing of strength and study weights in multivariate and network meta-analysis. *Stat Methods Med Res*. 2017;26(6):2853-2868.
37. Riley RD, Abrams KR, Lambert PC, Sutton AJ, Thompson JR. An evaluation of bivariate random-effects meta-analysis for the joint synthesis of two correlated outcomes. *Stat Med*. 2007;26(1):78-97.
38. Nakagawa S, Noble DWA, Lagisz M, Spake R, Viechtbauer W, Senior AM. A robust and readily implementable method for the meta-analysis of response ratios with and without missing standard deviations. *Ecology Letters*. 2023;26(2):232-244.
39. Gleser LJ, Olkin I. Stochastically dependent effect sizes. In: Cooper H, Hedges LV, Valentine JC, eds. *The Handbook of Research Synthesis and Meta-Analysis*. New York: Russell Sage Foundation; 2009:357-376.
40. Lajeunesse MJ. On the meta-analysis of response ratios for studies with correlated and multi-group designs. *Ecology*. 2011;92(11):2049-2055.
41. Pottier P, Noble DWA, Seebacher F, et al. New horizons for comparative studies and meta-analyses. *Trends in Ecology & Evolution*. 2024;39(5):435-445.
42. Michonneau F, Brown JW, Winter DJ. rotl: an R package to interact with the Open Tree of Life data. *Methods in Ecology and Evolution*. 2016;7(12):1476-1481.
43. Robinson GK. That BLUP is a Good Thing: The Estimation of Random Effects. *Statistical Science*. 1991;6(1):15-32.
44. Gelman A, Hill J. *Data Analysis using regression and multilevel/hierarchichal models*. New York: Cambridge University Press; 2007.
45. Rouse B, Chaimani A, Li T. Network meta-analysis: an introduction for clinicians. *Intern Emerg Med*. 2017;12(1):103-111.
46. Xing S, Fayle TM. The rise of ecological network meta-analyses: Problems and prospects. *Global Ecology and Conservation*. 2021;30:e01805.
47. Lu G, Ades AE. Assessing Evidence Inconsistency in Mixed Treatment Comparisons. *Journal of the American Statistical Association*. 2006;101(474):447-459.
48. Nakagawa S, Lagisz M, O'Dea RE, et al. The orchard plot: Cultivating a forest plot for use in ecology, evolution, and beyond. *Research Synthesis Methods*. 2021;12:4-12.
49. Nakagawa S, Lagisz M, O'Dea RE, et al. orchaRd 2.0: An R package for visualising meta-analyses with orchard plots. *Methods in Ecology and Evolution*. 2023;14(8):2003-2010.
50. Bruijning M, Metcalf CJE, Jongejans E, Ayroles JF. The Evolution of Variance Control. *Trends Ecol Evol*. 2020;35(1):22-33.
51. Kambach S, Bruelheide H, Gerstner K, Gurevitch J, Beckmann M, Seppelt R. Consequences of multiple imputation of missing standard deviations and sample sizes in meta-analysis. *Ecology and Evolution*. 2020;10(20):11699-11712.
52. Hallgrimsson B, Brian KH. Variation and Variability: Central concepts in Biology. In: Hallgrimsson B, Hall BK, eds. *Variation: A Central Concept in Biology*. Academic Press; 2005:1-7.

560

Tables

Table 1. Estimated parameters from different analyses of the lnRR and lnCVR including contrast-based integrated meta-analyses of mean and variation effects (IMAMV). CI = 95% confidence interval, τ= heterogeneity.

Analysis	lnRR (CI)	τ_{lnRR}	lnCVR (CI)	τ_{lnCVR}
Two-Model Analysis	-0.30 (-0.37 to -0.22)	0.68	0.16 (0.09 to 0.22)	0.53
Contrast-Based IMAMV	-0.30 (-0.37 to -0.22)	0.67	0.16 (0.09 to 0.22)	0.52
Multilevel Contrast-Based IMAMV	-0.31 (-0.43 to -0.20)	0.69	0.17 (0.06 to 0.28)	0.54

Table 2. Estimated parameters from two-model analyses and contrast-based integrated meta-analyses of mean and variation effects (IMAMV) with 20% missing standard-deviation data. SE = standard error, CI = 95% confidence interval.

Analysis	lnRR (CI)	SE_{lnRR}	lnCVR (CI)	SE_{lnCVR}
Two-Model Analysis	-0.32 (-0.40 to -0.24)	0.040	0.18 (0.10 to 0.26)	0.040
Contrast-Based IMAMV	-0.32 (-0.40 to -0.24)	0.040	0.18 (0.11 to 0.26)	0.038

Figures

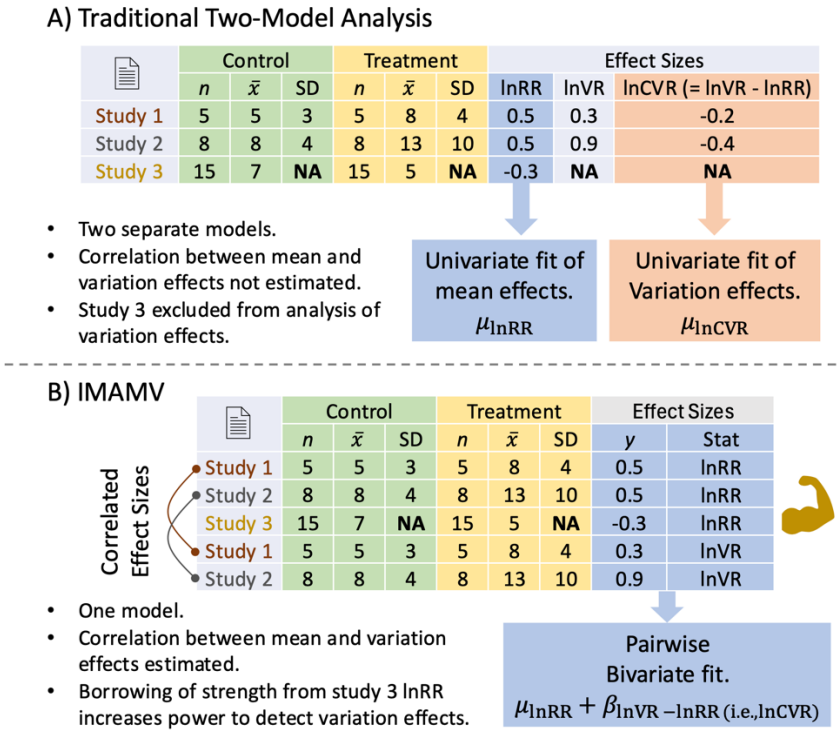

Figure 1. Hypothetical contrast of the current and proposed approach. A) the traditional method independently analyses effects on the mean and variation, overlooking any correlations among effect types. B) IMAMV uses bivariate meta-analysis to analyse mean and variation effects at the same time, thereby estimating any correlations and increasing strength for datasets with missing variance values.

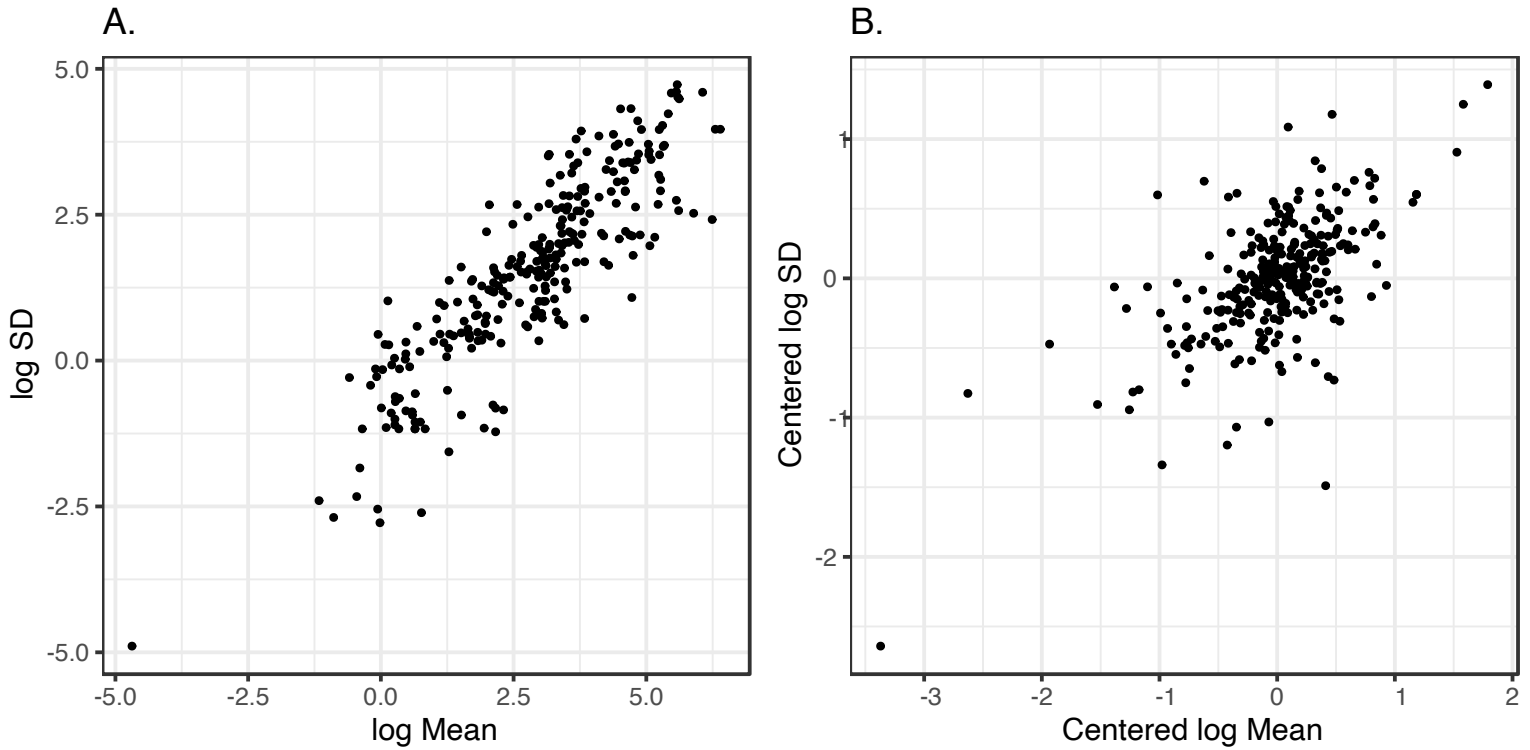

Figure 2. Association between log mean and log SD for reproductive data in the diet-mixing data set. In A) the data are in reported units. In B) the data have been mean-centred within each experiment.

Supplement Information for the Integrated Meta-Analysis of Mean and Variation Effects in Ecology and Evolution

Arm-Based IMAMV

As discussed in the main text, contrast-based analyses is that they can induce an additional layer of non-independence by calculating multiple pair-wise effect sizes using the same control-group data; sometimes termed 'stochastic dependency' (1-3). Arm-based models are free from this issue as one fits the sample statistics from individual groups as outcomes (i.e., $\ln\bar{x}$ and $\ln s$ rather than lnRR and lnVR) and uses a meta-regression model to estimate the difference between treatment conditions. An arm-based multi-level IMAMV can be written as:

$$y_{ijk} = \alpha + a_k + e_{jk} + \left(\beta_{\ln s} + b_k + f_{jk}\right) \times S_{ijk} + \left(\beta_{\mathrm{E}} + c_k + g_{jk}\right) \times T_{ijk} + \left(\beta_{\ln s\Delta\mathrm{E}} + d_k + h_{jk}\right) \times U_{ik} + m_{ik}, \quad \text{(S1)}$$

$$\begin{bmatrix} a_k \\ \vdots \\ d_k \end{bmatrix} \sim N\left(\begin{bmatrix} 0 \\ \vdots \\ 0 \end{bmatrix}, \begin{bmatrix} \sigma_a^2 & \cdots & \rho_{ad}\sigma_a\sigma_d \\ \vdots & \ddots & \vdots \\ \rho_{ad}\sigma_a\sigma_d & \cdots & \sigma_d^2 \end{bmatrix}\right), \quad \text{(S2)}$$

$$\begin{bmatrix} e_{jk} \\ \vdots \\ h_{jk} \end{bmatrix} \sim N\left(\begin{bmatrix} 0 \\ \vdots \\ 0 \end{bmatrix}, \begin{bmatrix} \sigma_e^2 & \cdots & \rho_{eh}\sigma_e\sigma_h \\ \vdots & \ddots & \vdots \\ \rho_{eh}\sigma_e\sigma_h & \cdots & \sigma_h^2 \end{bmatrix}\right), \quad \text{(S3)}$$

$$m_{ik} \sim N\left(0, v_{y_{ik}}\right), \quad \text{(S4)}$$

where, y_{ijk} is the ith sample statistic (i.e., i = 1 = $\ln\bar{x}$, i = 2 = ln$s$) of the jth sample in the kth experiment. Here, the α and $\beta_{\ln s}$ are interpretable as meta-analytic estimates of the log mean and log CV (i.e., $\beta_{\ln s}$ = $\mu_{\ln s - \ln\bar{x}} = \mu_{\mathrm{lnCV}}$; eqn. 9) under the control condition, and S_{ijk} is a dummy predictor coded as 0 if y_{ijk} is a sample $\ln\bar{x}$ and 1 if y_{ijk} is ln$s$. β_{E} is a meta-analytic estimate of difference in $\ln\bar{x}$ in the experimental and control conditions and can thus equivalent to lnRR (i.e., β_{E} = $\mu_{\ln\bar{x}_{\mathrm{E}} - \ln\bar{x}_{\mathrm{C}}} = \mu_{\mathrm{lnRR}}$; eqn. 3), and T_{ijk} is a dummy predictor coded as 0 if y_{ijk} is from a control condition and 1 for an experimental group. $\beta_{\ln s\Delta\mathrm{E}}$ is an interaction term giving the difference in $\beta_{\ln s}$ under the experimental and control conditions, and is therefore equivalent to lnCVR ($\beta_{\ln s\Delta\mathrm{E}} = \mathrm{lnCV_E} - \mathrm{lnCV_C} =$ lnCVR; eqn. 11), with U_{ik} as a third dummy predictor coded as 1 if y_{ijk} is an estimate of ln$s$ under the experimental condition, and 0 otherwise. a_k through d_k give deviations of the true values of each group from estimated parameters, and e_{jk} through h_{jk} give the within-sample deviations. Finally, m_{jk} gives the deviation of the sample from the group-specific effect, which has a SD estimated as $v_{y_{ik}}$(eqn. 2 if y_{ijk} is $\ln\bar{x}$ and eqn. 6 for ln$s$). The key variance components for meta-analytic interpretation are σ_c and σ_d, which give the among-experiment SD in lnRR and lnCVR, while the within-experiment SDs are σ_g and σ_h; $\tau^2_{ln\mathrm{RR}} = \sigma_c^2 + \sigma_g^2$ and $\tau^2_{ln\mathrm{CVR}} = \sigma_d^2 + \sigma_h^2$.

We have applied the arm-based IMAMV shown in eqn. S1 to the diet-mixing data using *metafor*. The point estimates and statistical significance of the overall effects of diet mixing on the log mean and log CV of traits are nearly identical using all methods (Table

S1). However, the arm-based estimates of heterogeneity are lower than those estimated by the contrast-based models in all cases but one. The most likely explanation for this difference is that the contrast-based model has treated a relatively large number of dependent effect sizes as independent, thereby inflating the variation among effects.

One can also include moderators in an arm-based IMAMV. We have not written the equation in full here, but the key terms of interest added to eqn. S1 would be: $\beta_{\mathrm{Mod}\Delta\mathrm{E}}$, which is an interaction interpretable as the difference in lnRR between levels of the moderator; and $\beta_{\mathrm{Mod}\Delta\mathrm{lns}\Delta\mathrm{E}}$, which is a three-way interaction that is interpretable as the difference in lnCVR between levels of the moderator.

Dual Formula Implementation of Integrated Meta-Analysis of Mean and Variation Effects (IMAMV)

In the main text we have implemented IMAMV models by simultaneously fitting statistics related to the log group means and SDs as a single response and using a moderator variable or fixed effect with appropriate error structure (e.g., a random-slope) to induce a bivariate model. The advantage of this approach is that these models can implemented in any software that can fit multi-level meta regression, such as the popular R package *metafor*.

An alternative offered by some software is to implement a bivariate model using a 'dual-formula' specification. One such example is the Bayesian R package *brms* (4), which offers dual-formula model specification.

A dual formula specification for a simple random-effects contrast-based IMAMV for the lnRR and lnCVR is:

$$\begin{pmatrix} \mathrm{lnRR}_i \\ \mathrm{lnCVR}_i \end{pmatrix} = \begin{pmatrix} \theta_{\mathrm{lnRR}} + a_i + p_i \\ \theta_{\mathrm{lnCVR}} + c_i + q_i \end{pmatrix}, \qquad \text{(S5)}$$

$$\begin{pmatrix} a_i \\ c_i \end{pmatrix} \sim N\left(\begin{pmatrix} 0 \\ 0 \end{pmatrix}, \begin{pmatrix} \sigma_a^2 & \sigma_{ac} \\ \sigma_{ac} & \sigma_c^2 \end{pmatrix} \right), \qquad \text{(S6)}$$

$$\begin{pmatrix} p_i \\ q_i \end{pmatrix} \sim N\left(\begin{pmatrix} 0 \\ 0 \end{pmatrix}, \begin{pmatrix} v_{\mathrm{lnRR}_i} \\ v_{\mathrm{lnVR}_i} \end{pmatrix} \right), \qquad \text{(S7)}$$

where lnRR_i and lnCVR_i are the effect sizes in the *i*th study as estimated by eqns 3 and 11 in the main text, θ_{lnRR} and θ_{lnCVR} are the meta-analytic overall effects as estimated by the model, a_i and c_i give the deviation of the population effect from θ_{lnRR} and θ_{lnCVR} in the *i*th study p_i and q_i give the deviations of the sampled effects from the population effects due to sampling in the *i*th study. Both a_i and c_i are assumed to be multi-variate normally distributed as shown in eqn. S2, where σ_a^2 and σ_c^2 give the among-study heterogeneity in lnRR and lnCVR, respectively, as estimated by the model, and σ_{ac} gives the covariance between effects at the level of the study. Also p_i and q_i are assumed to

be normally distributed as per eqn. S3, where the sampling variances are estimated via eqns 4 and 12 in the main text.

An arm-based IMAMV may also be implemented in a dual formula context as:

$$\begin{pmatrix} \ln\bar{x}_{ij} \\ \mathrm{lnCV}_{ij} \end{pmatrix} = \begin{pmatrix} \alpha_{\ln\bar{x}} + a_i + (\beta_{\ln\bar{x}} + b_i) \times T_{ij} + p_{ij} \\ \alpha_{\mathrm{lnCV}} + c_i + (\beta_{\mathrm{lnCV}} + d_i) \times T_{ij} + q_{ij} \end{pmatrix}, \quad \text{(S8)}$$

$$\begin{pmatrix} a_i \\ \vdots \\ d_i \end{pmatrix} \sim N\left(\begin{pmatrix} 0 \\ \vdots \\ 0 \end{pmatrix}, \begin{pmatrix} \sigma_a^2 & \cdots & \sigma_{ad} \\ \vdots & \ddots & \vdots \\ \sigma_{da} & \cdots & \sigma_d^2 \end{pmatrix} \right), \quad \text{(S9)}$$

$$\begin{pmatrix} p_{ij} \\ q_{ij} \end{pmatrix} \sim N\left(\begin{pmatrix} 0 \\ 0 \end{pmatrix}, \begin{pmatrix} v_{\ln\bar{x}_{ij}} \\ v_{\mathrm{lnCV}_{ij}} \end{pmatrix} \right), \quad \text{(S10)}$$

where $\ln\bar{x}_{ij}$ and lnCV_{ij} are the log mean and log CV from the jth group (i.e., *j* = 1, 2, where 1 = control and 2 = treatment) in the *i*th study as estimated by eqns 1 and 9 in the main text, $\alpha_{\ln\bar{x}}$ and α_{lnCV} are the meta-analytic overall estimates of ln*x* and lnCV in the control condition, and $\beta_{\ln\bar{x}}$ and β_{lnCV} are the effects of the treatment on ln*x* and lnCV. $\underline{T_{ij}}$ is a dummy predictor, coded as 0 where $\ln\bar{x}_{ij}$ and $\ln s_{ij}$ is form a control group and 1 otherwise. a_i and c_i give the deviation of the population effect from $\alpha_{\ln\bar{x}}$ and α_{lnCV} in the *i*th study, while b_i and d_i give the deviation of the population effect from $\beta_{\ln\bar{x}}$ and β_{lnCV} in the *i*th study. p_{ij} and q_{ij} give the deviations of the sampled statistics from the population statistic due to sampling in the *i*th study for the *j*th group. a_i through d_i are assumed to be multi-variate normally distributed as shown in eqn. S5, where σ_a^2 through σ_d^2 give the among-study heterogeneities, with σ_b^2 through σ_d^2 respectively, being interpretable as the estimated heterogeneity in the effect size. Also, p_{ij} and q_{ij} are assumed to be normally distributed as per eqn. S6, where the sampling variances are estimated via eqns 2 and 10 in the main text.

Table S1

Example of full reporting of all terms from an arm-based integrated meta-analysis of mean and variation effects (IMAMV). Results comes from the multi-level arm-based IMAMV in the diet-mixing dataset. For each term we report the interpretation and equivalent term in eqn. 24 in the main text. For fixed effects we also include 95% confidence intervals (CIs).

Fixed Effects			
Interpretation	Eqn. 24	Est.	CI
$\ln\bar{x}_{\text{Control}}$	α	2.92	2.52 to 3.32
$(\ln s - \ln\bar{x})_{\text{Control}} = \ln\text{CV}_{\text{Control}}$	β_{lns}	-1.40	-0.32 to -0.19
$\ln\bar{x}_{\text{Exp.}} - \ln\bar{x}_{\text{Control}} = \ln\text{RR}$	β_{E}	0.34	-0.47 to -0.20
$(\ln s - \ln\bar{x})_{\text{Exp.}} - (\ln s - \ln\bar{x})_{\text{Control}} = \ln\text{CVR}$	$\beta_{\text{lnsΔE}}$	0.19	0.07 to 0.30
Random Effects / Correlations			
Interpretation	Eqn. 24	Est.	
Among-study variance in α	σ_a^2	2.72	
Among-study variance in β_{lns}	σ_b^2	0.50	
Among-study variance in β_{E}	σ_c^2	0.03	
Among-study variance in $\beta_{\text{lnsΔE}}$	σ_d^2	0.04	
Among-study correlation between α and β_{lns}	ρ_{ab}	-0.13	
Among-study correlation between α and β_{E}	ρ_{ac}	0.83	
Among-study correlation between α and $\beta_{\text{lnsΔE}}$	ρ_{ad}	-0.57	
Among-study correlation between β_{lns} and β_{E}	ρ_{bc}	-0.66	
Among-study correlation between β_{lns} and $\beta_{\text{lnsΔE}}$	ρ_{bd}	-0.10	
Among-study correlation between β_{E} and $\beta_{\text{lnsΔE}}$	ρ_{cd}	-0.41	
Within-study variance in α	σ_e^2	0.14	
Within-study variance in β_{lns}	σ_f^2	0.08	
Within-study variance in β_{E}	σ_g^2	0.42	
Within-study variance in $\beta_{\text{lnsΔE}}$	σ_h^2	0.09	
Within-study correlation between α and β_{lns}	ρ_{ef}	-0.39	
Within-study correlation between α and β_{E}	ρ_{eg}	-0.26	
Within-study correlation between α and $\beta_{\text{lnsΔE}}$	ρ_{eh}	0.38	
Within-study correlation between β_{lns} and β_{E}	ρ_{fg}	-0.31	
Within-study correlation between β_{lns} and $\beta_{\text{lnsΔE}}$	ρ_{fh}	-0.11	
Within-study correlation between β_{E} and $\beta_{\text{lnsΔE}}$	ρ_{gh}	-0.73	

References

1. Gleser LJ, Olkin I. 2009. Stochastically dependent effect sizes. In *The Handbook of Research Synthesis and Meta-Analysis*, ed. H Cooper, LV Hedges, JC Valentine:357-76. New York: Russell Sage Foundation. Number of 357-76 pp.
2. Lajeunesse MJ. 2011. On the meta-analysis of response ratios for studies with correlated and multi-group designs. *Ecology* 92:2049-55
3. Nakagawa S, Noble DWA, Senior AM, Lagisz M. 2017. Meta-evaluation of meta-analysis: ten appraisal questions for biologists. *BMC Biology* 15:18
4. Bürkner P-C. 2017. Advanced Bayesian multilevel modeling with the R package brms. *arXiv preprint arXiv:1705.11123*

Integrated Meta-Analysis of Mean and Variation Effects (IMAMV) Vignette

XXXX

2025-06-24

Overview

This vignette is written to accompany the paper “Methods for the Integrated Meta-analysis of Mean and Variation Effects in Ecology and Evolution”.

Integrated meta-analysis of mean and variation effects (IMAMV) is a proposed analytical framework for simultaneously meta-analysing effects of treatments/groupings on sample means and sample variation. The method proposes to use a bivariate meta-analysis that analyses mean effects via the log ratio of sample means (lnRR), and variation via the log variance ratio (lnVR) and log coefficient of variation ratio (lnCVR). The method exploits the fact that lnVR - lnRR = lnCVR, and that this difference is estimable via meta-regression.

This vignette works through three different approaches to the meta-analysis of variation: (1) the standard two-model approach, which uses separate meta-analyses of the lnRR and lnCVR, (2) contrast-based IMAMV, and (3) arm-based IMAMV. Versions of each analysis that account for non-independence and test for moderator variables are explored, as are phylogenetic models. Mathematical descriptions of the models are made available in the accompanying paper.

As a working example, we study a meta-analytic dataset of experiments on diet-mixing. The models here are implemented in `metafor` and `brms`, both of which are often used for meta-analysis in the eco-evo fields. However, they can be implemented in any package that allows for meta-regression with random-effects, such as `MCMCglmm`.

Dataset

Diet mixing experiments test for the effects of single- vs mixed-food diets on animal performance. We have compiled a database of diet-mixing studies looking at reproductive data from 282 groups of animals in 69 experiments. The dataset was compiled by searching and updating the dataset analysed in “An Overlooked Consequence of Dietary Mixing: A Varied Diet Reduces Interindividual Variance in Fitness. Am. Nat. 2015. 186, 649-659. DOI: 10.1086/683182”. The data contains several layers of non-independence, including that most experiments contain more than two treatment groups that can be justifiably contrasted with one another (i.e., more than two effect sizes per experiment).

For the sake of this vignette, effects sizes have already been calculated and the data processed in different ways for the different variants of the analyses presented. Effect sizes were calculated using the formulas in the main text of the accompanying paper. There are four different data formats: contrast- vs arm-based, and each in long- vs wide-format. The processed dataframes have been stored in a list in the Rdata file `data_list.Rda`. In any given analysis a user of IMAMV would only need to format their data in one of these ways, depending on the analysis of choice.

Finally, the `data_list` object also contains a phylogenetic covariance matrix, used in the models toward the end of this vignette. The full contents of the list is shown here.

```
load("data_list.Rda")
names(data_list)
```

```
## [1] "arm_wide"      "arm_long"      "contrast_wide" "contrast_long"
## [5] "phylo"
```

Analysis 1: The Two-Model Approach

The conventional approach to meta-analysis of variation has most often used contrast-based effect sizes. The effects on the mean being estimated via the lnRR, then the effects on the variation using a separate model applied to the lnCVR (or in some instances the lnVR).

For the diet-mixing data, we have pre-processed data in format amenable to these analyses. It can be accessed in the data list via `data_list$contrast_wide`.

```
head(data_list$contrast_wide)
```

```
##   Article.ID Author              Journal Year               Consumer.Sp
## 1       art1 Alajmi Aquaculture Research 2015 Parvocalanus_crassirostris
## 2       art1 Alajmi Aquaculture Research 2015 Parvocalanus_crassirostris
## 3       art1 Alajmi Aquaculture Research 2015 Parvocalanus_crassirostris
## 4       art1 Alajmi Aquaculture Research 2015 Parvocalanus_crassirostris
## 5       art1 Alajmi Aquaculture Research 2015 Parvocalanus_crassirostris
## 6       art1 Alajmi Aquaculture Research 2015 Parvocalanus_crassirostris
##   Habitat Experiment.ID Data.ID       lnRR      v_lnRR       lnVR    v_lnVR
## 1  Marine            c1    dat1  0.9387648 0.039645511 -0.7375989 0.1946240
## 2  Marine            c1    dat2  0.3849543 0.047678185 -0.3022809 0.1946240
## 3  Marine            c1    dat3  0.7386473 0.039141741 -1.1895841 0.1946240
## 4  Marine            c1    dat4 -0.3069166 0.003327347 -0.5465437 0.1583058
## 5  Marine            c1    dat5 -0.8607272 0.011360020 -0.1112256 0.1583058
## 6  Marine            c1    dat6 -0.5070341 0.002823576 -0.9985288 0.1583058
##        lnCVR   v_lnCVR
## 1 -1.6763638 0.2342695
## 2 -0.6872352 0.2423021
## 3 -1.9282314 0.2337657
## 4 -0.2396271 0.1616331
## 5  0.7495015 0.1696658
## 6 -0.4914947 0.1611294
```

```
dim(data_list$contrast_wide)
```

```
## [1] 331  14
```

The most relevant columns for now are:

- `lnRR` gives the effect of a single-food diet relative to a mixed-food diet on the sample mean.
- `lnCVR` gives the effect of a single-food diet relative to a mixed-food diet on the sample CV.
- `v_lnRR` gives the sampling variances for the lnRR.
- `v_lnCVR` gives the sampling variances for the lnCVR.
- `Data.ID` is a unit-level variable with 1:n effect sizes.

Other columns give meta-variables related to the article and species associated with each effect size.

A pair of random-effects meta-analyses using the package `metafor` can be specified as follows.

```
library(metafor)
```

```
rma.mv(yi = lnRR, V = v_lnRR, random = list(~1|Data.ID), data = data_list$contrast_wi
de)
```

```
##
## Multivariate Meta-Analysis Model (k = 331; method: REML)
##
## Variance Components:
##
##             estim    sqrt  nlvls  fixed   factor
## sigma^2    0.4564  0.6755    331     no  Data.ID
##
## Test for Heterogeneity:
## Q(df = 330) = 23693.7200, p-val < .0001
##
## Model Results:
##
## estimate      se     zval    pval    ci.lb    ci.ub
##  -0.2961  0.0381  -7.7721  <.0001  -0.3707  -0.2214  ***
##
## ---
## Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
```

```
rma.mv(yi = lnCVR, V = v_lnCVR, random = list(~1|Data.ID), data = data_list$contrast_
wide)
```

```
##
## Multivariate Meta-Analysis Model (k = 331; method: REML)
##
## Variance Components:
##
##             estim    sqrt  nlvls  fixed   factor
## sigma^2    0.2798  0.5290    331     no  Data.ID
##
## Test for Heterogeneity:
## Q(df = 330) = 1561.2967, p-val < .0001
##
## Model Results:
##
## estimate      se    zval    pval   ci.lb   ci.ub
##   0.1590  0.0332  4.7826  <.0001  0.0938  0.2241  ***
##
## ---
## Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
```

The model estimates give a negative lnRR (`estimate = -0.2961`), which is statistically significant (i.e., the CI does not span 0). This estimate suggests that single food diets decrease mean reproductive function to around 74% of that on mixed-diets (i.e., `exp(-0.2961) = 0.7437`). However, there is a positive lnCVR (`estimate = 0.1590`), suggesting single food diets increase the CV by around 17% (i.e., `exp(0.1590) = 1.1723`). For lnRR the heterogeneity is `0.6755` , while that for lnCVR is `0.5290` , and in both cases this can be considered substantial (i.e., `sqrt(sigma^2) > abs(estimate)`)

Analysis 2: Contrast-Based IMAMV

IMAMV is an alternative to implementing two separate analyses. Rather a bivariate model is fitted to both the lnRR and lnVR data. The model uses a meta-regression (similar to a 'random-regression') to estimate paired differences between lnVR and lnRR from the same samples, thus yielding an estimate of lnCVR.

The data for a contrast based IMAMV are effectively the same as those for the two-model analysis above, but in long format. This means that effect sizes of different types (i.e., lnRR and lnVR) are mixed in the same column, with type identified by a dummy variable.

In the case of the diet-mixing dataset, we have formatted these data and they are available in the data list as follows.

```
head(data_list$contrast_long)
```

```
##           yi          vi stat Experiment.ID Data.ID               Consumer.Sp
## 1  0.9387648 0.039645511 lnRR            c1    dat1 Parvocalanus_crassirostris
## 2  0.3849543 0.047678185 lnRR            c1    dat2 Parvocalanus_crassirostris
## 3  0.7386473 0.039141741 lnRR            c1    dat3 Parvocalanus_crassirostris
## 4 -0.3069166 0.003327347 lnRR            c1    dat4 Parvocalanus_crassirostris
## 5 -0.8607272 0.011360020 lnRR            c1    dat5 Parvocalanus_crassirostris
## 6 -0.5070341 0.002823576 lnRR            c1    dat6 Parvocalanus_crassirostris
##   Habitat
## 1  Marine
## 2  Marine
## 3  Marine
## 4  Marine
## 5  Marine
## 6  Marine
```

```
dim(data_list$contrast_long)
```

```
## [1] 662   7
```

The key columns for analysis are:

- `yi` contains a mix of lnRR and lnVR data.
- `vi` is the associated sampling variances.
- `stat` is a dummy variable indicating whether the row contains an instance of lnRR or lnVR.
- `Data.ID` indicates those lnRR and lnVR that are calculated from the same pair of samples.

We can implement a contrast-based IMAMV in `metafor` as follows.

```
rma.mv(yi=yi, V=vi, mods=~stat, random=list(~stat|Data.ID), struct="GEN", data = data
_list$contrast_long)
```

```
##
## Multivariate Meta-Analysis Model (k = 662; method: REML)
##
## Variance Components:
##
## outer factor: Data.ID (nlvls = 331)
## inner term:   ~stat   (nlvls = 2)
##
##            estim    sqrt  fixed  rho:  intr     stVR
## intrcpt   0.4504  0.6711     no           -  -0.5805
## statlnVR  0.2751  0.5245     no          no        -
##
## Test for Residual Heterogeneity:
## QE(df = 660) = 26176.4495, p-val < .0001
##
## Test of Moderators (coefficient 2):
## QM(df = 1) = 22.8889, p-val < .0001
##
## Model Results:
##
##           estimate      se     zval    pval    ci.lb    ci.ub
## intrcpt    -0.2940  0.0378  -7.7757  <.0001  -0.3681  -0.2199  ***
## statlnVR    0.1574  0.0329   4.7842  <.0001   0.0929   0.2219  ***
##
## ---
## Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
```

The estimates for the effect of diet mixing on the mean, as quantified by lnRR, are given as the `intrcpt`. The effect on variation, as quantified by lnCVR, are given as `statlnVR`. The overall effects are almost identical to that estimated by the two-model analysis; here we have lnRR = `-0.2940` and lnCVR = `0.1574`, and again the CIs exclude 0. The total heterogeneity estimates are also nearly identical to those from a two-model analysis. For lnRR this is `0.6711` and for lnCVR `0.5245`.

We have fitted lnVR, so how is the model output is interpretable as lnCVR?

The model here has fitted the `stat` dummy variable in a meta-regression, estimating the difference in magnitude between effect sizes that are coded as lnVR and those coded as lnRR, which is another way of estimating lnCVR; lnVR - lnRR = lnCVR. Importantly, when we specified the model, we ensured that the lnVR-lnRR differences were estimated at the level of the individual sample pairs. The argument `struct = "GEN"` specifies a model that is similar to 'random-regression' mixed effects model, estimating the slopes for lnVR-lnRR at the level of individual sample pairs (i.e., via `~stat|Data.ID`).

The contrast-based IMAMV has estimated similar to terms to those from the two-model analysis, but also yields additional estimates of the correlation between lnRR and lnCVR; `stVR = -0.5805`. One can interpret this as more negative estimates of lnRR are associated with more positive effects of lnCVR. Put another way, the bigger the reductive effect of the diet on mean reproductive function, the more variation it generates.

Non-Indepdendence and Multi-Level IMAMV

This basic random-effects model has assumed independence of effect sizes. However, it is common in eco-evolutionary meta-analyses for the same experiment to generate multiple effects sizes, meaning those effects from the same experiment are non-independent. Multi-level models can be used to account for this non-independence by clustering effect sizes based on some level of grouping.

In the diet-mixing dataset, the column `Experiment.ID` codes each effect size by its experiment of origin. This can be added to our contrast-based IMAMV under the `random` argument to create a multi-level model as follows.

```
rma.mv(yi=yi, V=vi, mods=~stat, random=list(~stat|Experiment.ID, ~stat|Data.ID), stru
ct="GEN", data = data_list$contrast_long)
```

```
##
## Multivariate Meta-Analysis Model (k = 662; method: REML)
##
## Variance Components:
##
## outer factor: Experiment.ID (nlvls = 69)
## inner term:   ~stat         (nlvls = 2)
##
##            estim    sqrt  fixed  rho:  intr     stVR
## intrcpt   0.1152  0.3395     no           -  -0.3896
## statlnVR  0.1402  0.3745     no          no        -
##
## outer factor: Data.ID (nlvls = 331)
## inner term:   ~stat   (nlvls = 2)
##
##            estim    sqrt  fixed  phi:  intr     stVR
## intrcpt   0.3570  0.5975     no           -  -0.6992
## statlnVR  0.1491  0.3862     no          no        -
##
## Test for Residual Heterogeneity:
## QE(df = 660) = 26176.4495, p-val < .0001
##
## Test of Moderators (coefficient 2):
## QM(df = 1) = 8.7982, p-val = 0.0030
##
## Model Results:
##
##           estimate      se     zval    pval    ci.lb    ci.ub
## intrcpt    -0.3145  0.0596  -5.2734  <.0001  -0.4314  -0.1976  ***
## statlnVR    0.1706  0.0575   2.9662  0.0030   0.0579   0.2833   **
##
## ---
## Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
```

The sign and statistical significance of the estimates in the multi-level models match those in the random-effects model, though the magnitude of the effects are slightly larger (i.e., deviate more from 0) and precision slightly lower (i.e., the CIs are wider). Here the heterogeneity has been partitioned between the among- and within-experiment levels (i.e., `Experiment.ID` vs `Data.ID`). The total estimated heterogeneity can be calculated as the square root of the sum of the `sigma^2` estimates from the two levels. For lnRR this is `sqrt(0.1152 + 0.3570) = 0.6872` and for lnCVR it is `sqrt(0.1402 + 0.1491) = 0.5379`, largely matching the estimates from the random-effects model above.

Meta-Regression in Contrast-Based IMAMV

Moderator variables, which might explain heterogeneity in effects, can be included in IMAMV. In the diet-mixing dataset, for example, we have coded effect sizes by whether the focal species is terrestrial vs marine dwelling in the column `Habitat`.

In `metafor` we can include the moderator by specifying it in interaction with the `stat` variable in the `mods` argument.

```
rma.mv(yi=yi, V=vi, mods=~stat + Habitat + stat:Habitat, random=list(~stat|Experimen
t.ID, ~stat|Data.ID), struct="GEN", data = data_list$contrast_long)
```

```
##
## Multivariate Meta-Analysis Model (k = 662; method: REML)
##
## Variance Components:
##
## outer factor: Experiment.ID (nlvls = 69)
## inner term:   ~stat         (nlvls = 2)
##
##            estim    sqrt  fixed  rho:  intr     stVR
## intrcpt   0.1170  0.3421     no           -  -0.3591
## statlnVR  0.1383  0.3720     no          no        -
##
## outer factor: Data.ID (nlvls = 331)
## inner term:   ~stat   (nlvls = 2)
##
##            estim    sqrt  fixed  phi:  intr     stVR
## intrcpt   0.3532  0.5943     no           -  -0.6984
## statlnVR  0.1490  0.3860     no          no        -
##
## Test for Residual Heterogeneity:
## QE(df = 658) = 26171.7780, p-val < .0001
##
## Test of Moderators (coefficients 2:4):
## QM(df = 3) = 12.7709, p-val = 0.0052
##
## Model Results:
##
##                              estimate      se     zval    pval    ci.lb
## intrcpt                       -0.1910  0.0867  -2.2014  0.0277  -0.3610
## statlnVR                       0.1055  0.0857   1.2302  0.2186  -0.0626
## HabitatTerrestrial            -0.2361  0.1197  -1.9732  0.0485  -0.4707
## statlnVR:HabitatTerrestrial    0.1248  0.1153   1.0825  0.2790  -0.1011
##                                ci.ub
## intrcpt                      -0.0209  *
## statlnVR                      0.2735
## HabitatTerrestrial           -0.0016  *
## statlnVR:HabitatTerrestrial   0.3507
##
## ---
## Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
```

In the meta-regression the marine category has been fitted as the reference group, and so the different estimates can be interpreted as follows.

- `intrcpt = -0.1910` is the lnRR in marine animals.
- `statlnVR = 0.1055` is the lnCVR in marine animals.

- `HabitatTerrestrial = -0.2361` gives the difference in lnRR between terrestrial and marine organisms. The estimate is negative and statistically significant, suggesting that single-food diets reduce the mean reproductive function in terrestrial more than in marine habitats.
- `statlnVR:HabitatTerrestrial = 0.1248` is the interaction between the `stat` term and the `Habitat` moderator. Although it is an interaction it can quite straightforwardly be interpreted as the difference in lnCVR between habitats. While the estimate is positive, suggesting a larger effect of the diet on the CV in the terrestrial habitat, it is non-significant.

Analysis 3: Arm-Based IMAMV

It is common for contrast-based analyses to contain an additional layer of non-independence that arises when two effect sizes are based on the same control sample, sometimes termed 'stochastic dependency'. In the diet-mixing dataset, for instance, we calculated all pair-wise effect sizes within the same experiment. Therefore, two effect sizes that compare different single-food treatments to the same mixed-food treatment are correlated.

Stochastic dependency arises because contrast-based models calculate effect-sizes prior to model fitting. Arm-based models circumvent this non-independence by fitting the sample statistic from each group as the outcome, and using the model to estimate differences between treatment groups.

For an arm-based analysis the individual sample log means and log SDs or log CVs for each group are fitted, and a dummy variable specifies, which treatment the group was exposed to. We have formatted the diet-mixing data this way in the data list here.

```
head(data_list$arm_wide)
```

```
##   Article.ID Author              Journal Year                Consumer.Sp
## 1       art1 Alajmi Aquaculture Research 2015 Parvocalanus_crassirostris
## 2       art1 Alajmi Aquaculture Research 2015 Parvocalanus_crassirostris
## 3       art1 Alajmi Aquaculture Research 2015 Parvocalanus_crassirostris
## 4       art1 Alajmi Aquaculture Research 2015 Parvocalanus_crassirostris
## 5       art1 Alajmi Aquaculture Research 2015 Parvocalanus_crassirostris
## 6       art1 Alajmi Aquaculture Research 2015 Parvocalanus_crassirostris
##   Habitat Experiment.ID      mean n        sd  treat      lnX        v_lnX
## 1  Marine            c1  89.03226 5  7.096774 single 4.489634 0.0012715488
## 2  Marine            c1  50.96774 5 10.967742 single 3.935824 0.0093042225
## 3  Marine            c1  72.90323 5  4.516129 single 4.289517 0.0007677783
## 4  Marine            c1  34.19355 5 14.838710    mix 3.550869 0.0383739623
## 5  Marine            c1 120.96774 5 12.258065    mix 4.796551 0.0020557977
## 6  Marine            c1  19.35484 5  6.451613    mix 2.974054 0.0224691358
##       lnSD  v_lnSD       lnCV    v_lnCV Data.ID
## 1 2.084640 0.15625 -2.4049938 0.1575215    dat1
## 2 2.519958 0.15625 -1.4158652 0.1655542    dat2
## 3 1.632655 0.15625 -2.6568614 0.1570178    dat3
## 4 2.822239 0.15625 -0.7286300 0.1946240    dat4
## 5 2.631184 0.15625 -2.1653667 0.1583058    dat5
## 6 1.989330 0.15625 -0.9847234 0.1787191    dat6
```

```
dim(data_list$arm_wide)
```

```
## [1] 282  18
```

Each row contains the sample statistics from an individual group of animals, with the variable `treat` specifying whether the group was exposed to a single- or mixed-food diet.

A very basic arm-based analysis of the mean effects using the `lnX` and `v_lnX` data can specified in `metafor` as here.

```
rma.mv(yi = lnX, V = v_lnX, mods=~treat, random = list(~treat|Experiment.ID, ~treat|D
ata.ID), struct="GEN",  data = data_list$arm_wide)
```

```
##
## Multivariate Meta-Analysis Model (k = 282; method: REML)
##
## Variance Components:
##
## outer factor: Experiment.ID (nlvls = 69)
## inner term:   ~treat        (nlvls = 2)
##
##              estim    sqrt  fixed  rho:  intr    trts
## intrcpt      2.7342  1.6535     no           -  1.0000
## treatsingle  0.0183  0.1353     no          no       -
##
## outer factor: Data.ID (nlvls = 282)
## inner term:   ~treat  (nlvls = 2)
##
##              estim    sqrt  fixed  phi:  intr     trts
## intrcpt      0.1387  0.3725     no           -  -0.1968
## treatsingle  0.3868  0.6219     no          no        -
##
## Test for Residual Heterogeneity:
## QE(df = 280) = 612141.6429, p-val < .0001
##
## Test of Moderators (coefficient 2):
## QM(df = 1) = 23.1014, p-val < .0001
##
## Model Results:
##
##              estimate      se     zval    pval    ci.lb    ci.ub
## intrcpt        2.9053  0.2027  14.3297  <.0001   2.5079   3.3027  ***
## treatsingle   -0.3279  0.0682  -4.8064  <.0001  -0.4616  -0.1942  ***
##
## ---
## Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
```

The estimate for `intrcpt` is the average log mean of control groups, and is somewhat meaningless as we have averaged over many different species and measures of reproductive function. However, the estimate for `treatsingle` is the difference in log sample means between single- and mixed-food groups, and thus interpretable as the lnRR. Note samples from the same experiment etc. were paired in the analysis, thus retaining the principle of concurrent control. The value `-0.3228` is similar to those coming from the contrast-based models above, though this analysis can be considered to have better accounted for non-independence as it is free of stochastic dependency. The heterogeneity can be estimated by adding up the `treatsingle` rows in the `Experiment.ID` and `Data.ID` parts of the output; `sqrt(0.0183 + 0.3868) = 0.6365`

For an arm-based IMAMV we need combine the above approach with bivariate model. For this analysis the arm-based data are best transformed in to long-format with the log sample means (`lnX`) and log sample SDs (`lnSD`) in the same column and a dummy variable specifying which is which.

For the diet mixing data, these are available pre-formatted here.

```
head(data_list$arm_long)
```

```
##   Article.ID Author              Journal Year               Consumer.Sp
## 1       art1 Alajmi Aquaculture Research 2015 Parvocalanus_crassirostris
## 2       art1 Alajmi Aquaculture Research 2015 Parvocalanus_crassirostris
## 3       art1 Alajmi Aquaculture Research 2015 Parvocalanus_crassirostris
## 4       art1 Alajmi Aquaculture Research 2015 Parvocalanus_crassirostris
## 5       art1 Alajmi Aquaculture Research 2015 Parvocalanus_crassirostris
## 6       art1 Alajmi Aquaculture Research 2015 Parvocalanus_crassirostris
##   Habitat Experiment.ID  treat Data.ID       yi      vi stat
## 1  Marine            c1 single    dat1 2.084640 0.15625 lnSD
## 2  Marine            c1 single    dat2 2.519958 0.15625 lnSD
## 3  Marine            c1 single    dat3 1.632655 0.15625 lnSD
## 4  Marine            c1    mix    dat4 2.822239 0.15625 lnSD
## 5  Marine            c1    mix    dat5 2.631184 0.15625 lnSD
## 6  Marine            c1    mix    dat6 1.989330 0.15625 lnSD
```

```
dim(data_list$arm_long)
```

```
## [1] 564  12
```

The columns `yi` and `vi` give the log sample statistics and their sampling variances, `stat` is a dummy variable identifying whether the sample statistic is an instance of the log mean or the log SD, and `treat` specifies whether the sample is from a group on a single- or mixed-food diet. Here `Data.ID` pairs instances of the log mean and log SD that are from the same sample.

A multi-level arm-based IMAMV can be fit as follows.

```
rma.mv(yi = yi, V = vi, mods=~treat*stat, random = list(~treat*stat|Experiment.ID, ~t
reat*stat|Data.ID), struct="GEN",  data = data_list$arm_long)
```

```
##
## Multivariate Meta-Analysis Model (k = 564; method: REML)
##
## Variance Components:
##
## outer factor: Experiment.ID (nlvls = 69)
## inner term:   ~treat * stat (nlvls = 4)
##
##                        estim    sqrt  fixed  rho:  intr    trts     stSD
## intrcpt               2.7207  1.6494     no           -  0.8294  -0.1341
## treatsingle           0.0288  0.1698     no          no       -  -0.6633
## statlnSD              0.5015  0.7082     no          no      no        -
## treatsingle:statlnSD  0.0438  0.2092     no          no      no       no
##                          t:SD
## intrcpt               -0.5710
## treatsingle           -0.4058
## statlnSD              -0.1024
## treatsingle:statlnSD       -
##
## outer factor: Data.ID       (nlvls = 282)
## inner term:   ~treat * stat (nlvls = 4)
##
##                        estim    sqrt  fixed  phi:  intr     trts     stSD
## intrcpt               0.1389  0.3727     no           -  -0.2607  -0.3852
## treatsingle           0.4198  0.6479     no          no        -  -0.3122
## statlnSD              0.0776  0.2786     no          no       no        -
## treatsingle:statlnSD  0.0913  0.3021     no          no       no       no
##                          t:SD
## intrcpt               0.3795
## treatsingle           -0.7280
## statlnSD              -0.1100
## treatsingle:statlnSD       -
##
## Test for Residual Heterogeneity:
## QE(df = 560) = 629739.6516, p-val < .0001
##
## Test of Moderators (coefficients 2:4):
## QM(df = 3) = 306.6336, p-val < .0001
##
## Model Results:
##
##                      estimate      se      zval    pval    ci.lb    ci.ub
## intrcpt                2.9207  0.2022   14.4430  <.0001   2.5244   3.3171  ***
## treatsingle           -0.3358  0.0691   -4.8575  <.0001  -0.4713  -0.2003  ***
## statlnSD              -1.4072  0.0929  -15.1505  <.0001  -1.5892  -1.2252  ***
## treatsingle:statlnSD   0.1870  0.0581    3.2190  0.0013   0.0731   0.3008   **
##
## ---
## Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
```

This model is arguably more free of (or corrected for) non-independence than any of the preceding analyses.

The output looks quite complex, but actually has a relatively easy interpretation that maps on to the contrast-based models above.

- `intrcpt = 2.9207` is the log sample mean for control, in this instance mixed-food, groups. The estimate is somewhat meaningless.
- `treatsingle = −0.3358` is the lnRR, here the difference in log means between single- and mixed-food groups. The sign and magnitude of effect is consistent with all other analyses of the lnRR above, and remains statistically significant.
- `statlnSD = −1.4072` is the log CV for the control group (actually log SD - log mean, but that is equal to the log CV), and again the value is a bit meaningless.
- `treatsingle:statlnSD = 0.1870` is the estimate of the lnCVR, which again matches the contrast-based analyses above in terms of sign, magnitude and statistical significance.

The relevant heterogeneities for the lnRR and lnCVR are available in the rows labelled `treatsingle` and `treatsingle:statlnSD` in the `Experiment.ID` and `Data.ID` parts of the analysis. Correlations between effect sizes at the among- and within-study levels have also been estimated, and can found by reading across the row `treatsingle` to the column `t:SD`. At both levels the correlations are negative (among-experiment = `−0.4058`, within-experiment = `−0.7280`), matching the estimates from the contrast-based analyses.

Meta-Regression in Arm-Based IMAMV

Again we can include/test moderator variables in an arm-based IMAMV. Here is an arm-based model that tests the moderating effect of Habitat again.

```
rma.mv(yi = yi, V = vi, mods=~treat*stat*Habitat, random = list(~treat*stat|Experimen
t.ID, ~treat*stat|Data.ID), struct="GEN",  data = data_list$arm_long)
```

```
##
## Multivariate Meta-Analysis Model (k = 564; method: REML)
##
## Variance Components:
##
## outer factor: Experiment.ID (nlvls = 69)
## inner term:   ~treat * stat (nlvls = 4)
##
##                        estim    sqrt  fixed  rho:  intr    trts     stSD
## intrcpt               2.7234  1.6503     no           -  0.8385  -0.1125
## treatsingle           0.0236  0.1537     no          no       -  -0.6358
## statlnSD              0.4880  0.6986     no          no      no        -
## treatsingle:statlnSD  0.0441  0.2099     no          no      no       no
##                          t:SD
## intrcpt               -0.5511
## treatsingle           -0.3476
## statlnSD              -0.1458
## treatsingle:statlnSD       -
##
## outer factor: Data.ID       (nlvls = 282)
## inner term:   ~treat * stat (nlvls = 4)
##
##                        estim    sqrt  fixed  phi:  intr     trts     stSD
## intrcpt               0.1406  0.3750     no           -  -0.0060  -0.3903
## treatsingle           0.2919  0.5403     no          no        -  -0.3167
## statlnSD              0.0778  0.2789     no          no       no        -
## treatsingle:statlnSD  0.0298  0.1726     no          no       no       no
##                          t:SD
## intrcpt               -0.4041
## treatsingle           -0.8601
## statlnSD               0.4447
## treatsingle:statlnSD       -
##
## Test for Residual Heterogeneity:
## QE(df = 556) = 589589.8407, p-val < .0001
##
## Test of Moderators (coefficients 2:8):
## QM(df = 7) = 313.0065, p-val < .0001
##
## Model Results:
##
##                                      estimate      se      zval    pval
## intrcpt                                3.1690  0.3224    9.8304  <.0001
## treatsingle                           -0.2257  0.0985   -2.2898  0.0220
## statlnSD                              -1.5932  0.1454  -10.9538  <.0001
## HabitatTerrestrial                    -0.4043  0.4142   -0.9761  0.3290
## treatsingle:statlnSD                   0.1380  0.0860    1.6057  0.1084
## treatsingle:HabitatTerrestrial        -0.2007  0.1374   -1.4610  0.1440
## statlnSD:HabitatTerrestrial            0.3043  0.1876    1.6221  0.1048
## treatsingle:statlnSD:HabitatTerrestrial  0.0908  0.1169    0.7767  0.4373
##                                        ci.lb    ci.ub
## intrcpt                               2.5372   3.8009  ***
## treatsingle                          -0.4188  -0.0325    *
## statlnSD                             -1.8782  -1.3081  ***
## HabitatTerrestrial                   -1.2161   0.4075
```

```
## treatsingle:statlnSD                       -0.0305   0.3065
## treatsingle:HabitatTerrestrial             -0.4701   0.0686
## statlnSD:HabitatTerrestrial                -0.0634   0.6721
## treatsingle:statlnSD:HabitatTerrestrial    -0.1383   0.3198
##
## ---
## Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
```

The long list of interaction terms in the model looks daunting, but again can be directly mapped on to a contrast-based model in terms of effect size. The most important terms to look at are:

- `treatsingle = -0.2257` , which gives the lnRR for marine species.
- `treatsingle:statlnSD = 0.1380` , which gives the lnCVR in the marine species.
- `treatsingle:HabitatTerrestrial = -0.2007` is the difference in lnRR between terrestrial and marine species. This estimate is similar in magnitude to the contrast-based meta-regression, but is non-significant.
- `treatsingle:statlnSD:HabitatTerrestrial = 0.0908` (the three-way-interactive term) is the difference in lnCVR between the different habitats.

The less interesting terms can be interpreted as follows. `intrcpt` and `statlnSD` is the average log mean and log CV of the control groups for all marine species, while `HabitatTerrestrial` and `statlnSD:HabitatTerrestrial` are differences between terrestrial and marine taxa for those log means and log CVs.

Phylogenetic Models

Multi-species meta-analyses such as are explored here contain phylogenetic non-independence as some of the species are more closely related, while others are more distantly related. It is common for ecologists and evolutionary biologists to correct for this non-independence using a phylogenetic model. IMAMV is compatible with such models. We now demonstrate two options for fitting phylogentic IMAMV.

First we must load and solve a matrix which gives the relatedness among the different species in the diet-mixing dataset, which we created using the R package `rotl` . The matrix is stored at the end of the data list.

```
phyloM<-data_list$phylo
A<-solve(as(phyloM, "dgCMatrix"))
```

The matrix contains a row/column for each of the species in the dataset listed under the column `Consumer.Sp` . A phylogenetic correction can then applied to the arm-based IMAMV in `metafor` as follows using the `R` argument.

```
rma.mv(yi = yi, V = vi, mods=~treat*stat, random = list(~treat*stat|Consumer.Sp, ~tre
at*stat|Data.ID), struct="GEN",  data = data_list$arm_long, R = list(Consumer.Sp =
A))
```

```
##
## Multivariate Meta-Analysis Model (k = 564; method: REML)
##
## Variance Components:
##
## outer factor: Consumer.Sp   (nlvls = 51)
## inner term:   ~treat * stat (nlvls = 4)
##
##                        estim    sqrt  fixed  rho:  intr    trts     stSD
## intrcpt               2.2229  1.4909     no           -  0.9429  -0.2335
## treatsingle           0.0278  0.1666     no          no       -  -0.5441
## statlnSD              0.5236  0.7236     no          no      no        -
## treatsingle:statlnSD  0.0206  0.1435     no          no      no       no
##                          t:SD
## intrcpt               -0.9760
## treatsingle           -0.8478
## statlnSD               0.0163
## treatsingle:statlnSD        -
##
## outer factor: Data.ID       (nlvls = 282)
## inner term:   ~treat * stat (nlvls = 4)
##
##                        estim    sqrt  fixed  phi:  intr     trts     stSD
## intrcpt               0.7893  0.8884     no           -  -0.0302   0.0159
## treatsingle           0.3247  0.5698     no          no        -  -0.2615
## statlnSD              0.1518  0.3896     no          no       no        -
## treatsingle:statlnSD  0.1534  0.3917     no          no       no       no
##                          t:SD
## intrcpt               -0.3784
## treatsingle           -0.2281
## statlnSD              -0.2183
## treatsingle:statlnSD        -
##
## Test for Residual Heterogeneity:
## QE(df = 560) = 629739.6516, p-val < .0001
##
## Test of Moderators (coefficients 2:4):
## QM(df = 3) = 166.7641, p-val < .0001
##
## Model Results:
##
##                      estimate      se      zval    pval    ci.lb    ci.ub
## intrcpt                3.0229  0.2275   13.2904  <.0001   2.5771   3.4687  ***
## treatsingle           -0.3305  0.1220   -2.7092  0.0067  -0.5696  -0.0914   **
## statlnSD              -1.3713  0.1121  -12.2311  <.0001  -1.5910  -1.1515  ***
## treatsingle:statlnSD   0.1829  0.0659    2.7740  0.0055   0.0537   0.3121   **
##
## ---
## Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
```

The key estimates from for the main effects (`treatsingle = -0.3305` for lnRR and `treatsingle:statlnSD = 0.1829` for lnCVR) match those from the non-phylogenetic model.

The user may have noted that here we swapped the experiment-level random effect for the species. This is a limit of the current implementation of `rma.mv`, which only permits two random effects with the inner terms (i.e., `inner|outer`) used to pair data from the same samples.

We now demonstrate how to implement an IMAMV model with more than two random effects using `brms`. See `?brms` to get started with this package.

```
library(brms)
```

`brms` allows 'dual-formula' specification, which some users may find more straight forward. Here we are implementing an arm-based model and thus we specify an equation for the log mean and for the log SD. The data are best suited to this function in wide format (i.e., with the mean and variance-related statistics in separate columns).

The formula for the model is specified here. Important parts to note are that (1) the `treat` is fitted as a fixed effect, and random slope at each of the multiple levels, (2) that the phylogenetic correlation matrix is associated with the `Consumer.Sp` level of the model, (3) we have specified `lnCV` as the variation-related metric and (4) that the formula is duplicated for both the log mean and log CV

```
arm_form<-bf(lnX |se(sqrt(v_lnX)) ~ treat + (1+treat|a|gr(Consumer.Sp, cov=A)) + (1+t
reat|b|Experiment.ID) + (1+treat|c|Data.ID)) +
            bf(lnCV |se(sqrt(v_lnCV)) ~ treat + (1+treat|a|gr(Consumer.Sp, cov=A)) +
(1+treat|b|Experiment.ID) + (1+treat|c|Data.ID)) +
            set_rescor(FALSE)
```

The model fits using an MCMC alogirthm. The user must specify how long the model runs for (`iter` and `warmup` arguments), how frequently samples are taken (`thin` argument), and how many replicate chains are run (`chains` argument). Here we have specified a model that will run relatively quickly for the sake of the vignette. MCMC algorithms must be checked in a series of diagnostic tests for behavior and convergence (e.g., see `?gelman.diag`). A robust analysis that passes all such checks would likely need a longer run and with chains run in, at least, triplicate.

Finally we run the model, noting I have set the seed to increase the reproducibility of results.

```
set.seed(123)
phylo_IMAMV<-brm(formula=arm_form, data=data_list$arm_wide, data2=list(A=A), family=g
aussian, cores=1, chains=1, iter=3000, warmup=2000, thin=1)
```

```
summary(phylo_IMAMV)
```

```
## Warning: Parts of the model have not converged (some Rhats are > 1.05). Be
## careful when analysing the results! We recommend running more iterations and/or
## setting stronger priors.
```

```
## Warning: There were 38 divergent transitions after warmup. Increasing
## adapt_delta above 0.8 may help. See
## http://mc-stan.org/misc/warnings.html#divergent-transitions-after-warmup
```

```
##  Family: MV(gaussian, gaussian)
##   Links: mu = identity; sigma = identity
##          mu = identity; sigma = identity
## Formula: lnX | se(sqrt(v_lnX)) ~ treat + (1 + treat | a | gr(Consumer.Sp, cov =
A)) + (1 + treat | b | Experiment.ID) + (1 + treat | c | Data.ID)
##          lnCV | se(sqrt(v_lnCV)) ~ treat + (1 + treat | a | gr(Consumer.Sp, cov =
A)) + (1 + treat | b | Experiment.ID) + (1 + treat | c | Data.ID)
##    Data: data_list$arm_wide (Number of observations: 282)
##   Draws: 1 chains, each with iter = 3000; warmup = 2000; thin = 1;
##          total post-warmup draws = 1000
##
## Multilevel Hyperparameters:
## ~Consumer.Sp (Number of levels: 51)
##                                     Estimate Est.Error l-95% CI u-95% CI Rhat
## sd(lnX_Intercept)                       0.18      0.10     0.01     0.38 1.00
## sd(lnX_treatsingle)                     0.03      0.02     0.00     0.08 1.00
## sd(lnCV_Intercept)                      0.08      0.05     0.01     0.18 1.00
## sd(lnCV_treatsingle)                    0.06      0.03     0.01     0.13 1.00
## cor(lnX_Intercept,lnX_treatsingle)      0.05      0.44    -0.77     0.83 1.00
## cor(lnX_Intercept,lnCV_Intercept)      -0.05      0.41    -0.76     0.77 1.00
## cor(lnX_treatsingle,lnCV_Intercept)    -0.01      0.46    -0.82     0.81 1.00
## cor(lnX_Intercept,lnCV_treatsingle)    -0.24      0.43    -0.91     0.71 1.00
## cor(lnX_treatsingle,lnCV_treatsingle)  -0.23      0.43    -0.89     0.72 1.00
## cor(lnCV_Intercept,lnCV_treatsingle)   -0.07      0.44    -0.82     0.77 1.00
##                                     Bulk_ESS Tail_ESS
## sd(lnX_Intercept)                         59      171
## sd(lnX_treatsingle)                      154      369
## sd(lnCV_Intercept)                       105      178
## sd(lnCV_treatsingle)                      78      308
## cor(lnX_Intercept,lnX_treatsingle)       398      476
## cor(lnX_Intercept,lnCV_Intercept)        242      255
## cor(lnX_treatsingle,lnCV_Intercept)      185      182
## cor(lnX_Intercept,lnCV_treatsingle)      156      248
## cor(lnX_treatsingle,lnCV_treatsingle)    193      328
## cor(lnCV_Intercept,lnCV_treatsingle)     198      481
##
## ~Data.ID (Number of levels: 282)
##                                     Estimate Est.Error l-95% CI u-95% CI Rhat
## sd(lnX_Intercept)                       0.38      0.04     0.31     0.46 1.01
## sd(lnX_treatsingle)                     0.50      0.14     0.25     0.74 1.01
## sd(lnCV_Intercept)                      0.30      0.05     0.21     0.40 1.04
## sd(lnCV_treatsingle)                    0.26      0.13     0.03     0.55 1.13
## cor(lnX_Intercept,lnX_treatsingle)      0.19      0.37    -0.40     0.91 1.03
## cor(lnX_Intercept,lnCV_Intercept)      -0.57      0.13    -0.80    -0.30 1.01
## cor(lnX_treatsingle,lnCV_Intercept)    -0.47      0.33    -0.94     0.41 1.04
## cor(lnX_Intercept,lnCV_treatsingle)    -0.35      0.37    -0.88     0.48 1.02
## cor(lnX_treatsingle,lnCV_treatsingle)  -0.49      0.34    -0.95     0.30 1.00
## cor(lnCV_Intercept,lnCV_treatsingle)    0.12      0.39    -0.61     0.83 1.05
##                                     Bulk_ESS Tail_ESS
## sd(lnX_Intercept)                        216      370
## sd(lnX_treatsingle)                       15      170
## sd(lnCV_Intercept)                        55       76
## sd(lnCV_treatsingle)                      15       36
## cor(lnX_Intercept,lnX_treatsingle)        14       62
## cor(lnX_Intercept,lnCV_Intercept)        126      255
```

```
## cor(lnX_treatsingle,lnCV_Intercept)         24       41
## cor(lnX_Intercept,lnCV_treatsingle)         44      112
## cor(lnX_treatsingle,lnCV_treatsingle)       53      127
## cor(lnCV_Intercept,lnCV_treatsingle)        53      117
##
## ~Experiment.ID (Number of levels: 69)
##                                     Estimate Est.Error l-95% CI u-95% CI Rhat
## sd(lnX_Intercept)                       1.61      0.19     1.23     2.01 1.00
## sd(lnX_treatsingle)                     0.12      0.08     0.01     0.30 1.02
## sd(lnCV_Intercept)                      0.68      0.09     0.50     0.84 1.01
## sd(lnCV_treatsingle)                    0.10      0.07     0.01     0.26 1.00
## cor(lnX_Intercept,lnX_treatsingle)      0.23      0.37    -0.56     0.83 1.00
## cor(lnX_Intercept,lnCV_Intercept)      -0.17      0.15    -0.44     0.14 1.01
## cor(lnX_treatsingle,lnCV_Intercept)    -0.30      0.42    -0.89     0.64 1.03
## cor(lnX_Intercept,lnCV_treatsingle)    -0.17      0.40    -0.83     0.66 1.01
## cor(lnX_treatsingle,lnCV_treatsingle)  -0.10      0.43    -0.84     0.75 1.00
## cor(lnCV_Intercept,lnCV_treatsingle)   -0.15      0.43    -0.84     0.76 1.00
##                                     Bulk_ESS Tail_ESS
## sd(lnX_Intercept)                        141      339
## sd(lnX_treatsingle)                       73      327
## sd(lnCV_Intercept)                       216      286
## sd(lnCV_treatsingle)                     165      302
## cor(lnX_Intercept,lnX_treatsingle)       527      544
## cor(lnX_Intercept,lnCV_Intercept)        179      357
## cor(lnX_treatsingle,lnCV_Intercept)       18       61
## cor(lnX_Intercept,lnCV_treatsingle)      436      623
## cor(lnX_treatsingle,lnCV_treatsingle)    355      533
## cor(lnCV_Intercept,lnCV_treatsingle)     349      508
##
## Regression Coefficients:
##                  Estimate Est.Error l-95% CI u-95% CI Rhat Bulk_ESS Tail_ESS
## lnX_Intercept        2.93      0.21     2.54     3.35 1.00      123      200
## lnCV_Intercept      -1.39      0.10    -1.59    -1.19 1.00      269      517
## lnX_treatsingle     -0.34      0.07    -0.48    -0.20 1.02      189      378
## lnCV_treatsingle     0.21      0.07     0.08     0.35 1.00      311      522
##
## Further Distributional Parameters:
##            Estimate Est.Error l-95% CI u-95% CI Rhat Bulk_ESS Tail_ESS
## sigma_lnX      0.00      0.00     0.00     0.00   NA       NA       NA
## sigma_lnCV     0.00      0.00     0.00     0.00   NA       NA       NA
##
## Draws were sampled using sampling(NUTS). For each parameter, Bulk_ESS
## and Tail_ESS are effective sample size measures, and Rhat is the potential
## scale reduction factor on split chains (at convergence, Rhat = 1).
```

The first thing to note is the warnings at the top of the output. This model would clearly need a good deal of tinkering with the MCMC specifications before we would consider the results robust. The fixes to these issues will be data set specific and do not affect dual formula specification for IMAMV. Hence, we do not present a full exploration of how to adjust the MCMC part of the model. Some solutions to explore might include adjusting `max_treedepth` via the `control` argument, running the model for longer (`iter` , `warmup` and `thin` arguments), and specifying a different prior (`prior` argumnet). There are many good tutorials on using the `brms` package that cover model specification in detail; see here https://ayumi-495.github.io/multinomial-GLMM-tutorial/ (https://ayumi-495.github.io/multinomial-GLMM-tutorial/).

Putting these concerns aside for the sake of the vignette, we can interpret the results as follows. The main overall estimates for the effects of diet mixing can be found under the `Regresson Coefficients` part of the output. The term `lnX_treatsingle = -0.35` is the lnRR, while `lnCV_treatsingle = 0.21` is the lnCVR. Again these estimates are similar to those seen in the phylogenetic model from `metafor`. The heterogeneity estimates can be found under the `Multilevel Hyperparameters` part of the output. The rows labelled `sd(lnX_treatsingle)` and `sd(lnCV_treatsingle)` give the SD among lnRR and lnCVR at different levels of the analysis. Note the estimates for `Consumer.Sp` are considerably smaller than those at other levels of the analysis, indicating that phylogenetic effects are likely to be weak.

Data Visualization

Contrast-based IMAMV can use the already wide array of tools for data-visualization (e.g., forest plots and orchard plots). These must just be applied to the lnRR and lnCVR (or lnVR if that is the analysis of choice).

There are no widely used formats for the visualization of arm-based meta-analyses. Here we propose a new tool, based on estimation plots and orchard plots. We have written a function called `gg_mestimation`, which requires the packages `ggplot2` and `ggbeeswarm` to work.

The function can be loaded from the header file `mestimation_functions.R`.

```
source("mestimation_function.R")
```

The key arguments to pass to the function are as follows:

- `data`: a dataframe containing, at least, the sample statistics to be plotted, the group of each statistic (e.g., control or experimental) and the sample size of each.
- `group`: the column name in `data` giving the group for each statistic being plotted.
- `stat`: the column name in `data` giving the sample statistics being plotted.
- `n`: the column name in `data` giving the sample sizes for each statistic being plotted.
- `control_mu`: the estimate of the meta-analytic mean in the control group.
- `mu`: the overall effect estimated by meta-analysis.
- `ci_l` and `ci_u`: the confidence intervals on `mu`.
- `tau`: the total heterogeneity associated with `mu`.

Here for example, we plot the analysis in the first arm-based IMAMV from above. This model was specified as:

```
rma.mv(yi = yi, V = vi, mods=~treat*stat, random = list(~treat*stat|Experiment.ID, ~t
reat*stat|Data.ID), struct="GEN",  data = data_list$arm_long)
```

```
##
## Multivariate Meta-Analysis Model (k = 564; method: REML)
##
## Variance Components:
##
## outer factor: Experiment.ID (nlvls = 69)
## inner term:   ~treat * stat (nlvls = 4)
##
##                        estim    sqrt  fixed  rho:  intr    trts     stSD
## intrcpt               2.7207  1.6494     no           -  0.8294  -0.1341
## treatsingle           0.0288  0.1698     no          no       -  -0.6633
## statlnSD              0.5015  0.7082     no          no      no        -
## treatsingle:statlnSD  0.0438  0.2092     no          no      no       no
##                          t:SD
## intrcpt               -0.5710
## treatsingle           -0.4058
## statlnSD              -0.1024
## treatsingle:statlnSD       -
##
## outer factor: Data.ID       (nlvls = 282)
## inner term:   ~treat * stat (nlvls = 4)
##
##                        estim    sqrt  fixed  phi:  intr     trts     stSD
## intrcpt               0.1389  0.3727     no           -  -0.2607  -0.3852
## treatsingle           0.4198  0.6479     no          no        -  -0.3122
## statlnSD              0.0776  0.2786     no          no       no        -
## treatsingle:statlnSD  0.0913  0.3021     no          no       no       no
##                          t:SD
## intrcpt                0.3795
## treatsingle           -0.7280
## statlnSD              -0.1100
## treatsingle:statlnSD       -
##
## Test for Residual Heterogeneity:
## QE(df = 560) = 629739.6516, p-val < .0001
##
## Test of Moderators (coefficients 2:4):
## QM(df = 3) = 306.6336, p-val < .0001
##
## Model Results:
##
##                      estimate      se      zval    pval    ci.lb    ci.ub
## intrcpt                2.9207  0.2022   14.4430  <.0001   2.5244   3.3171  ***
## treatsingle           -0.3358  0.0691   -4.8575  <.0001  -0.4713  -0.2003  ***
## statlnSD              -1.4072  0.0929  -15.1505  <.0001  -1.5892  -1.2252  ***
## treatsingle:statlnSD   0.1870  0.0581    3.2190  0.0013   0.0731   0.3008   **
##
## ---
## Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
```

Data in wide-format are well suited for use in this function, and we start by plotting the effects on the log mean (data in `lnX`). The estimate for the mean in the control group for `control_mu` is given under the analysis `intrcpt = 2.9207`. The effect of diet mixing on the mean (i.e., the lnRR) for the `mu` argument is

given under `treatsingle = −0.3358`, along with its CIs (`ci.lb = −0.4713`, `ci.ub = −0.2003`). The total heterogeneity for `tau`, estimated as the sum of the `treatsingle` terms at the two levels of the analysis, is `sqrt(0.4198 + 0.0288) = 0.6698`.

```
gg_mestimation(data = data_list$arm_wide, group = "treat", stat = "lnX", n = "n", con
trol_mu = 2.9207, mu = −0.3358, ci_l = −0.4713, ci_u = −0.2003, tau = 0.6698)
```

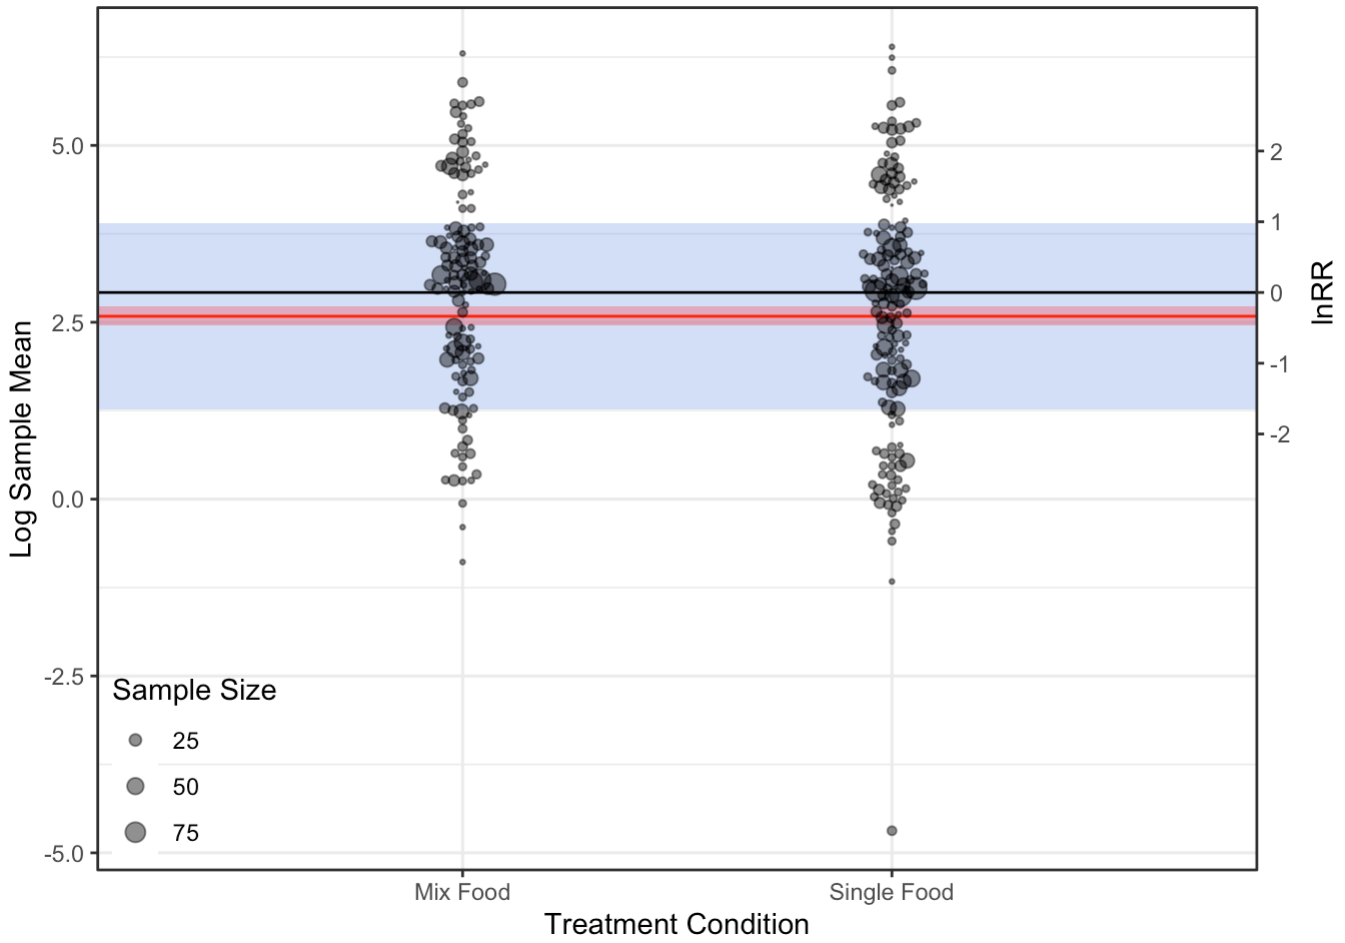

We see the log mean (y-axis) as a function of the dietary treatment (x-axis). Individual sample statistics are scaled by their sample size. The right-hand axis gives the projection in to the effect size (lnRR) space. In the effect size space, the line of no effect is the black horizontal line, while the estimated effect is the red horizontal line. The red shading gives the 95% CI on the effect on the mean, which excludes 0 in this case. The blue shaded zone visualizes the 95% prediction interval for future effects, which is based on the estimated heterogeneity (`tau`).

We can then reapply this to the lnCVR, using the estimates for `statlnSD` in the model as the `control_mu`, and `treatsingle:statlnSD` in the model as `mu`. The heterogeneity is `sqrt(0.0913 + 0.0438) = 0.3676`.

```
gg_mestimation(data = data_list$arm_wide, group = "treat", stat = "lnCV", n = "n", co
ntrol_mu = −1.4072, mu = 0.1870, ci_l = 0.0731, ci_u = 0.3008, tau = 0.3676)
```

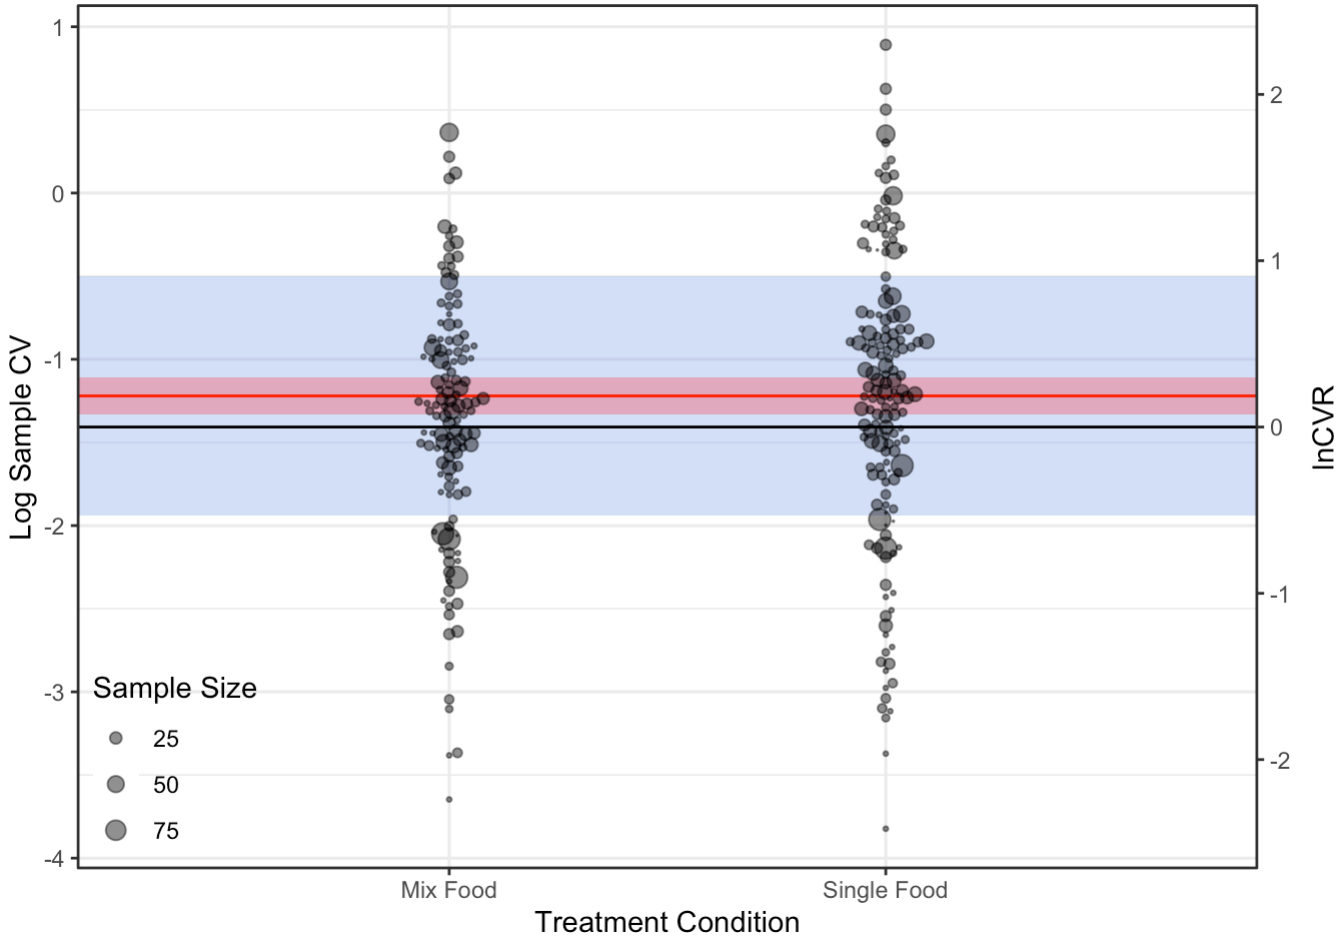

Session Info

```
library(sessioninfo)
session_info()
```

```
## ─ Session info ─────────────────────────────────────────────────────────────
##  setting  value
##  version  R version 4.5.1 (2025-06-13)
##  os       macOS Sequoia 15.5
##  system   aarch64, darwin20
##  ui       X11
##  language (EN)
##  collate  en_US.UTF-8
##  ctype    en_US.UTF-8
##  tz       Australia/Sydney
##  date     2025-07-31
##  pandoc   3.4 @ /Applications/RStudio.app/Contents/Resources/app/quarto/bin/tools/
aarch64/ (via rmarkdown)
##  quarto   1.6.42 @ /Applications/RStudio.app/Contents/Resources/app/quarto/bin/qua
rto
##
## ─ Packages ─────────────────────────────────────────────────────────────────
##  package        * version    date (UTC) lib source
##  abind            1.4-8      2024-09-12 [1] CRAN (R 4.5.0)
##  backports        1.5.0      2024-05-23 [1] CRAN (R 4.5.0)
##  bayesplot        1.13.0     2025-06-18 [1] CRAN (R 4.5.0)
##  beeswarm         0.4.0      2021-06-01 [1] CRAN (R 4.5.0)
##  bridgesampling   1.1-2      2021-04-16 [1] CRAN (R 4.5.0)
##  brms           * 2.22.0     2024-09-23 [1] CRAN (R 4.5.0)
##  Brobdingnag      1.2-9      2022-10-19 [1] CRAN (R 4.5.0)
##  bslib            0.9.0      2025-01-30 [1] CRAN (R 4.5.0)
##  cachem           1.1.0      2024-05-16 [1] CRAN (R 4.5.0)
##  callr            3.7.6      2024-03-25 [1] CRAN (R 4.5.0)
##  checkmate        2.3.2      2024-07-29 [1] CRAN (R 4.5.0)
##  cli              3.6.5      2025-04-23 [1] CRAN (R 4.5.0)
##  coda             0.19-4.1   2024-01-31 [1] CRAN (R 4.5.0)
##  codetools        0.2-20     2024-03-31 [1] CRAN (R 4.5.1)
##  digest           0.6.37     2024-08-19 [1] CRAN (R 4.5.0)
##  distributional   0.5.0      2024-09-17 [1] CRAN (R 4.5.0)
##  dplyr            1.1.4      2023-11-17 [1] CRAN (R 4.5.0)
##  evaluate         1.0.4      2025-06-18 [1] CRAN (R 4.5.0)
##  farver           2.1.2      2024-05-13 [1] CRAN (R 4.5.0)
##  fastmap          1.2.0      2024-05-15 [1] CRAN (R 4.5.0)
##  generics         0.1.4      2025-05-09 [1] CRAN (R 4.5.0)
##  ggbeeswarm     * 0.7.2      2023-04-29 [1] CRAN (R 4.5.0)
##  ggplot2        * 3.5.2      2025-04-09 [1] CRAN (R 4.5.0)
##  glue             1.8.0      2024-09-30 [1] CRAN (R 4.5.0)
##  gridExtra        2.3        2017-09-09 [1] CRAN (R 4.5.0)
##  gtable           0.3.6      2024-10-25 [1] CRAN (R 4.5.0)
##  htmltools        0.5.8.1    2024-04-04 [1] CRAN (R 4.5.0)
##  inline           0.3.21     2025-01-09 [1] CRAN (R 4.5.0)
##  jquerylib        0.1.4      2021-04-26 [1] CRAN (R 4.5.0)
##  jsonlite         2.0.0      2025-03-27 [1] CRAN (R 4.5.0)
##  knitr            1.50       2025-03-16 [1] CRAN (R 4.5.0)
##  labeling         0.4.3      2023-08-29 [1] CRAN (R 4.5.0)
##  lattice          0.22-7     2025-04-02 [1] CRAN (R 4.5.1)
##  lifecycle        1.0.4      2023-11-07 [1] CRAN (R 4.5.0)
##  loo              2.8.0      2024-07-03 [1] CRAN (R 4.5.0)
##  magrittr         2.0.3      2022-03-30 [1] CRAN (R 4.5.0)
##  mathjaxr         1.8-0      2025-04-30 [1] CRAN (R 4.5.0)
```

```
##  Matrix         * 1.7-3      2025-03-11 [1] CRAN (R 4.5.1)
##  matrixStats      1.5.0      2025-01-07 [1] CRAN (R 4.5.0)
##  metadat        * 1.4-0      2025-02-04 [1] CRAN (R 4.5.0)
##  metafor        * 4.8-0      2025-01-28 [1] CRAN (R 4.5.0)
##  mvtnorm          1.3-3      2025-01-10 [1] CRAN (R 4.5.0)
##  nlme             3.1-168    2025-03-31 [1] CRAN (R 4.5.1)
##  numDeriv       * 2016.8-1.1 2019-06-06 [1] CRAN (R 4.5.0)
##  pillar           1.10.2     2025-04-05 [1] CRAN (R 4.5.0)
##  pkgbuild         1.4.8      2025-05-26 [1] CRAN (R 4.5.0)
##  pkgconfig        2.0.3      2019-09-22 [1] CRAN (R 4.5.0)
##  plyr             1.8.9      2023-10-02 [1] CRAN (R 4.5.0)
##  posterior        1.6.1      2025-02-27 [1] CRAN (R 4.5.0)
##  processx         3.8.6      2025-02-21 [1] CRAN (R 4.5.0)
##  ps               1.9.1      2025-04-12 [1] CRAN (R 4.5.0)
##  QuickJSR         1.8.0      2025-06-09 [1] CRAN (R 4.5.0)
##  R6               2.6.1      2025-02-15 [1] CRAN (R 4.5.0)
##  RColorBrewer     1.1-3      2022-04-03 [1] CRAN (R 4.5.0)
##  Rcpp           * 1.0.14     2025-01-12 [1] CRAN (R 4.5.0)
##  RcppParallel     5.1.10     2025-01-24 [1] CRAN (R 4.5.0)
##  reshape2         1.4.4      2020-04-09 [1] CRAN (R 4.5.0)
##  rlang            1.1.6      2025-04-11 [1] CRAN (R 4.5.0)
##  rmarkdown        2.29       2024-11-04 [1] CRAN (R 4.5.0)
##  rstan            2.32.7     2025-03-10 [1] CRAN (R 4.5.0)
##  rstantools       2.4.0      2024-01-31 [1] CRAN (R 4.5.0)
##  sass             0.4.10     2025-04-11 [1] CRAN (R 4.5.0)
##  scales           1.4.0      2025-04-24 [1] CRAN (R 4.5.0)
##  sessioninfo    * 1.2.3      2025-02-05 [1] CRAN (R 4.5.0)
##  StanHeaders      2.32.10    2024-07-15 [1] CRAN (R 4.5.1)
##  stringi          1.8.7      2025-03-27 [1] CRAN (R 4.5.0)
##  stringr          1.5.1      2023-11-14 [1] CRAN (R 4.5.0)
##  tensorA          0.36.2.1   2023-12-13 [1] CRAN (R 4.5.0)
##  tibble           3.3.0      2025-06-08 [1] CRAN (R 4.5.0)
##  tidyselect       1.2.1      2024-03-11 [1] CRAN (R 4.5.0)
##  vctrs            0.6.5      2023-12-01 [1] CRAN (R 4.5.0)
##  vipor            0.4.7      2023-12-18 [1] CRAN (R 4.5.0)
##  withr            3.0.2      2024-10-28 [1] CRAN (R 4.5.0)
##  xfun             0.52       2025-04-02 [1] CRAN (R 4.5.0)
##  yaml             2.3.10     2024-07-26 [1] CRAN (R 4.5.0)
##
##  [1] /Library/Frameworks/R.framework/Versions/4.5-arm64/Resources/library
##  * ── Packages attached to the search path.
##
## ──────────────────────────────────────────────────────────────────────────────
```