The role of secondary outcomes in multivariate meta-analysis

John B. Copas*
University of Warwick, UK
Dan Jackson
MRC Biostastics Unit, University of Cambridge, UK
Ian R. White
MRC Clinical Trials Unit at UCL, University College London, UK
Richard D. Riley
University of Keele, UK

Summary

Univariate meta-analysis concerns a single outcome of interest measured across a number of independent studies. However, many research studies will have also measured secondary outcomes. Multivariate meta-analysis allows us to take these secondary outcomes into account, and can also include studies where the primary outcome is missing. We define the efficiency (E) as the variance of the overall estimate from a multivariate meta-analysis relative to the variance of the overall estimate from a univariate meta-analysis. The extra information gained from a multivariate meta-analysis of n studies is then similar to the extra information gained if a univariate meta-analysis of the primary effect had a further n(1-E)/E studies. The variance contribution of a study's secondary outcomes (its borrowing of strength) can be thought of as a contrast between the variance matrix of the outcomes in that study and the set of variance matrices of all the studies in the meta-analysis. In the bivariate case this is given a simple graphical interpretation as the *borrowing of strength plot*. We discuss how these findings can also be used in the context of random effects meta-analysis. Our discussion is motivated by a published meta-analysis of ten anti-hypertension clinical trials.

Keywords: Multivariate meta-analysis; Borrowing of strength; Multiple outcomes.

*E-mail: jbc@stats.warwick.ac.uk

1 Introduction

Univariate meta-analysis is well established as a statistical tool for research synthesis, when a single outcome of primary interest is measured across several independent studies. Many research studies, however, report data on multiple outcomes, with the primary outcome supported by measures of one or more secondary outcomes. Multivariate meta-analysis offers the potential for more accurate estimation by also taking the data on these secondary outcomes into account. Another advantage of the multivariate approach is the potential for increasing the number of eligible research studies, since we can also include studies where the primary outcome is missing and data are only reported on some of the secondary outcomes.

A key question in the expanding literature on multivariate meta-analysis is the comparison between multivariate and univariate approaches — how much borrowing of strength do the secondary outcomes contribute to the estimation of the primary treatment effect? The empirical examples discussed by Sohn (2000), Simel and Bossuyt (2009) and Trikalinos *et al.* (2014) mostly show rather little difference between the results of multivariate and univariate meta-analysis, even though in some of these examples the outcomes are quite highly correlated. This has led some to question whether the multivariate approach is of any real practical value. Other examples, however, suggest that taking the secondary outcomes into account can make a useful contribution (Fibrinogen Studies Collaboration, 2009; Riley, 2009; Kirkham *et al.*, 2012). Why do these differences arise? What is it about the statistical properties of the studies in a meta-analysis that determine the contribution of the secondary outcomes?

By comparing the multivariate estimate of the primary treatment effect with the corresponding univariate estimate taking only the primary outcomes into account, Jackson *et al.* (2017) derive an expression for borrowing of strength, measuring the additional contribution which each study's secondary outcome estimates make to the variance of the summary primary treatment effect over and above the contribution of the study's primary outcome estimate. The corresponding expression for the total contribution of individual studies gives a measure of study weights, analogous to the familiar use of study weights in univariate meta-analysis. The aim of this paper is to re-examine Jackson's formulae, to explore some of their consequences and extensions, and to offer a more transparent understanding of how borrowing of strength depends on individual study characteristics. We generalize a number of points which earlier papers have made using examples and simulation studies. Data from a published meta-analysis of 10 clinical trials on the treatment of hypertension is taken as a motivating example.

Section 2 gives our basic set-up, showing that the borrowing of strength given by the secondary outcomes of the ith study can be written as an explicit function of two variance matrices, the within-study variance matrix V_i and the harmonic average $\bar{V}$of all the V_is in the meta-analysis. With an appropriate re-defining of V_i (Section 2.2) this also covers cases where one or more of the outcomes in the ith study is missing. Properties of the borrowing of strength function are most easily seen in the bivariate case, where the *borrowing of strength plot* is a useful way of interpreting the relative contributions of the two outcomes. The bivariate case is discussed in Section 3 and illustrated using data from the hypertension example. The bivariate case is generalized to the multivariate case in Section 4, leading to a general formulation of the necessary and sufficient conditions for a study to give borrowing of strength in multivariate meta-analysis.

Section 5 follows Jackson *et al.* (2017) by showing that, at least as descriptive measures, borrowing of strength in multivariate fixed effects models applies equally well to random effects models, thus allowing for between-studies heterogeneity in a way analogous to the DerSimonian-Laird method in univariate meta-analysis (DerSimonian and Laird, 1986). A simulation study based on the hypertension example shows the importance of distinguishing between borrowing of strength as a descriptive measure (describing the data to hand) and as an inferential measure (describing an underlying population model), a distinction which does not arise in the same way for fixed effects models.

The final Section 6 gives a brief discussion of some of the important assumptions being made in this paper.

2 The variance contribution of individual studies

2.1 Basic set-up

We consider a multivariate meta-analysis of n independent studies, each of which measures a $p \times 1$ vector y of treatment effect estimates corresponding to the p different outcomes. The standard multivariate fixed effects model is

$$y_i \sim N(\beta, V_i)\ ;\ i = 1, 2, \cdots, n. \tag{1}$$

The $p \times p$ variance matrix V_i in (1) is specific to each study, but the unknown mean parameter β is assumed to be the same for all studies (the fixed effects assumption). To start with, we assume that all p outcomes are measured in all n studies in the meta-analysis.

Treating each V_i as known (the usual assumption), the score function for the unknown parameter β (derivative of the log likelihood) is

$$\sum V_i^{-1}(y_i - \beta)\ , \tag{2}$$

and so the maximum likelihood estimate (MLE) of vector β is

$$\hat{\beta} = \Omega \sum V_i^{-1} y_i\ , \tag{3}$$

where Ω is the variance matrix of $\hat{\beta}$ given by

$$\Omega = \mathrm{Var}(\hat{\beta}) = (\sum V_i^{-1})^{-1}\ .$$

This can be rewritten as

$$\Omega = \frac{1}{n}\bar{V}\ ,$$

where

$$\bar{V} = (n^{-1} \sum V_i^{-1})^{-1}\ , \tag{4}$$

the harmonic average of the V_is. Whether we use the actual within-study variances V_i, or crudely approximate them all by $\bar{V}$, we end up with the same variance matrix of $\hat{\beta}$.

Even if all p components of y are observed, we focus interest on estimating the treatment effect for just one of these outcomes which, without loss of generality, we take to be the first. So from now on we will describe, for each study, y_{i1} as the scalar treatment effect estimate for the *primary* outcome and the remaining components of y_i as the $(p-1) \times 1$ vector of estimates for the *secondary* outcomes. In some cases the primary outcome may be clearly identified from the context. For example, the bivariate ($p = 2$) example in Fibrinogen Studies Collaboration (2009) concerned study estimates y_{i1} of a treatment effect adjusted for differences across a defined set of covariates, but also included estimates y_{i2} which are partially adjusted for just a subset of these covariates. The fully adjusted results are of primary interest, but the advantage of including the secondary outcomes is that we can also take account of studies which do not measure the full set of confounding covariates. In other cases, such as the bivariate example studied in Section 3.2, we may be interested in all the outcomes, in which case we can arbitrarily re-label the outcomes as appropriate. The essential assumption is that we are interested in the separate (marginal) inferences to be made for one or more of the outcomes rather than in the correlations between the meta-analysis estimates across different outcomes. So we assume from now on that our primary interest is in $\hat{\beta}_1$, with variance

$$\mathrm{Var}(\hat{\beta}_1) = \Omega_{11} = \left[(\sum V_i^{-1})^{-1}\right]_{11} = \frac{1}{n}\bar{V}_{11} \ ,$$

where $\bar{V}_{11}$ is the $(1,1)$ element of $\bar{V}$ in (4).

A natural comparison for the multivariate estimate of β_1 is univariate meta-analysis, which looks only at the values of y_{1i} and ignores the data on the secondary outcomes. The relevant univariate model would then be

$$y_{i1} \sim N(\beta_1, \sigma_i^2) \ ,$$

where

$$\sigma_i^2 = \mathrm{l}^T V_i \mathrm{l} \ ,$$

and l is the unit vector $\mathrm{l} = (1, 0, \cdots, 0)^T$. The univariate estimate is

$$\tilde{\beta}_1 = \frac{\sum \sigma_i^{-2} y_{i1}}{\sum \sigma_i^{-2}}$$

with variance

$$\mathrm{Var}(\tilde{\beta}_1) = \frac{1}{\sum \sigma_i^{-2}} \ .$$

Under model (1), both $\hat{\beta}_1$ and $\tilde{\beta}_1$ are unbiased and normally distributed estimates of β_1, and so to compare their statistical properties all we need to know is the efficiency, E, defined by

$$E = \frac{\mathrm{Var}(\hat{\beta}_1)}{\mathrm{Var}(\tilde{\beta}_1)} = \Omega_{11} \sum \sigma_i^{-2} = \frac{1}{n}\bar{V}_{11} \sum (\mathrm{l}^T V_i \mathrm{l})^{-1} \ . \tag{5}$$

Necessarily, $E \leq 1$ as the MLE $\hat{\beta}_1$ is fully efficient. The smaller is E, the greater is the relative contribution of the secondary outcomes, suggesting $1-E$ as a measure of the role of the secondary outcomes in the multivariate estimate of the primary treatment effect. This combines the information in the secondary outcomes of all the studies in the meta-analysis and so $1 - E$ can thought of as a measure of *total* borrowing of strength, equivalent to BoS_r^{RV} in the notation of Jackson *et al.*, 2017 (section 2.2). However, the simpler notation $1 - E$ emphasises its dependence on a basic statistical concept which may open up useful interpretations taken from other areas of statistics, a possibility not immediately obvious from the earlier notation.

A simple example here is the familiar interpretation of efficiency in terms of sample size: for an inefficient estimate (efficiency E) to match the accuracy that a fully efficient estimate (efficiency 1) can achieve with a sample size of n, the sample size would have to be increased from n to n/E. Similarly, in meta-analysis, the extra information which the secondary outcomes of n studies gives to the estimation of β_1 can be thought of as like the extra information we would get in univariate meta-analysis if we could measure the primary outcomes of a further $n(1-E)/E$ studies. For example, if there are 9 studies ($n = 9$) and $E = 0.9$, the advantage of using multivariate instead of univariate meta-analysis is like finding the data for one more study. This simple idea will be used several times in the analysis of the hypertension example in Section 3.2 below.

The last expression in (5) is the ratio of the $(1, 1)$ element of the harmonic mean of the V_is to the harmonic mean of the $(1, 1)$ elements of the V_is . These are the same thing if the V_is are all the same, in which case $E = 1$. If the V_is are different then $E \leq 1$, which suggests another interpretation of $(1 - E)$ as a measure of the variation of the matrices V_i about their harmonic average $\bar{V}$. This is analogous to the usual interpretation of the coefficient of variation (ratio of standard deviation to the mean) as a simple relative measure of the variation of a univariate sample about its arithmetic mean.

These calculations are comparing the relative contributions which the primary and secondary outcomes make to the estimation of β_1 using all n studies in the meta-analysis. To break this down into the contributions of individual studies, define, for any study with inverse variance matrix V^{-1},

$$T(V^{-1}) = \frac{\left[\Omega V^{-1}\Omega\right]_{11}}{\Omega_{11}} = \frac{1}{n}\frac{\left[\bar{V} V^{-1}\bar{V}\right]_{11}}{\bar{V}_{11}} \ . \tag{6}$$

We write the argument of (6) as V^{-1} rather than V to reflect the fact that all of the formulae for multivariate meta-analysis presented earlier involve the study variances V_i only through their inverses V_i^{-1}. As we shall see in Section 2.2, this also simplifies the notation in cases where there are missing data. Clearly, (6) is a function of two arguments, V^{-1} and $\bar{V}$, and so (6) has further simplified the notation by suppressing the second argument. We can do this because we are mainly interested in the contributions of individual studies within the context of a *given* observed meta analysis, in which case we can treat $\bar{V}$ as if it was fixed.

We use the function $T(V^{-1})$ to investigate the role of individual studies in three different ways, analogous to the definitions of influence in regression analysis:

- *Direct interpretation.* From (3),

$$\text{Var}(\hat{\beta}_1) = \Omega_{11} = \sum [\Omega V_i^{-1}\Omega]_{11} = \Omega_{11} \sum T(V_i^{-1}) \ .$$

Hence, $T(V_i^{-1})$ is the proportional contribution of the ith study to the variance of $\hat{\beta}_1$, proportional in the sense that

$$\sum T(V_i^{-1}) = 1 \ .$$

In univariate meta-analysis, $V_i^{-1} = \sigma_i^{-2}$ and (6) gives $T(V_i^{-1}) = \sigma_i^{-2}/\sum_i \sigma_i^{-2}$ which is just the weight of the ith study in the weighted average $\hat{\beta}_1$. When $p \geq 2$, $T(V_i^{-1})$ can still be interpreted as the weight of the ith study in multivariate meta-analysis, agreeing with the weight w_{ir} derived from an orthogonal decomposition of the score function in Jackson *et al.* (2017, section 3). However, the function $T(V^{-1})$ is not restricted to the V's which happen to be represented in the meta-analysis:

- *Add-one-in interpretation.* If n is large and the variance matrix V is of the same order of magnitude as $\bar{V}$, then under reasonable conditions on the matrices involved,

$$(I + n^{-1}\bar{V}V^{-1})^{-1} = I - n^{-1}\bar{V}V^{-1} + O(n^{-2}) \ . \tag{7}$$

  Post multiplying each side of (7) by $n^{-1}\bar{V}$ and using (4), we get the approximation

$$[(\sum_{j=1}^{n} V_j^{-1} + V^{-1})^{-1}]_{11} = \Omega_{11}(1 - T(V^{-1})) + O(n^{-3}) \ . \tag{8}$$

  The left-hand side of equation (8) is the updated variance of $\hat{\beta}_1$ if we add a new study with inverse variance V^{-1} to the meta-analysis. So, for large n, $T(V^{-1})$ is the proportional *decrease* in $\mathrm{Var}(\hat{\beta}_1)$.

- *Leave-one-out interpretation.* Replacing V by $-V_i$ in (8) similarly shows that $T(V_i^{-1})$ is the proportional *increase* in $\mathrm{Var}(\hat{\beta}_1)$ if the ith study is removed from the meta-analysis.

The first of these properties is exact, but the second and third are only asymptotic (large n) approximations. This reflects differences in the background studies being assumed for the add-one-in and hold-one-out calculations, i.e. differences in the second argument $\bar{V}$) in (6). For example, if a study we are thinking of adding in happens to have the same variance V as an existing study which we are thinking of leaving out, then the common value of $T(V^{-1})$ suggests that the two effects would be the same. But one is defining this study's contribution in terms of the difference between having $n+1$ studies and n studies, but the other is comparing $n-1$ with n studies. If n is large there is no material difference between the two. Essentially, the add-one-in and hold-one-out approximations are ignoring the effect that adding or subtracting studies has on the value of $\bar{V}$. These distinctions are analogous to the different definitions of residuals and influence in other areas of statistics. See Section 3 below for a clearer illustration of some of these points in the simpler context of bivariate meta-analysis ($p = 2$).

The definition of E in (5) arises from comparing $\mathrm{Var}(\hat{\beta}_1)$ with the value of this variance if only the primary outcomes had been measured across the whole of the meta-analysis. Similarly, for investigating the role of individual studies, we can ask what happens to

$\mathrm{Var}(\hat{\beta}_1)$ if we add in an extra study with variance matrix V, but only take account of its primary outcome estimate $y_1 \sim N(\beta_1, \sigma^2)$ with $\sigma^2 = \mathrm{l}^T V \mathrm{l}$. This will add $\sigma^{-2}(y_1 - \beta_1)$ to the score function (2) for the scalar β_1, but will add nothing to the score function for the secondary outcomes. Hence the contribution to the vector score function for the estimation of the complete vector β is

$$V_*^{-1}(y - \beta) ,$$

where the matrix V_*^{-1} is defined as

$$V_*^{-1} = \sigma^{-2}\mathrm{l}\mathrm{l}^T = (\mathrm{l}^T V \mathrm{l})^{-1}\mathrm{l}\mathrm{l}^T , \tag{9}$$

the $p \times p$ matrix with σ^{-2} in the $(1,1)$ position and zero everywhere else. The relative decrease in $\mathrm{Var}(\hat{\beta}_1)$ is therefore (approximately)

$$T(V_*^{-1}) = T\{(\mathrm{l}^T V \mathrm{l})^{-1}\mathrm{l}\mathrm{l}^T\} . \tag{10}$$

We define the *borrowing of strength*, $B(V^{-1})$, of a study with variance matrix V to be the difference between (6) and (10),

$$B(V^{-1}) = T(V^{-1}) - T\{(\mathrm{l}^T V \mathrm{l})^{-1}\mathrm{l}\mathrm{l}^T\} . \tag{11}$$

This measures the contribution that the secondary outcomes of this particular study makes to $\mathrm{Var}(\hat{\beta}_1)$ over and above the contribution made by its primary outcome. If $B(V^{-1})$ is zero then nothing is gained by observing the secondary outcomes. The notation $T(V^{-1})$ refers to the (T)otal contribution of a study; the notation $B(V^{-1})$ refers to the (B)orrowing of strength, how much of this proportional increase in precision is contributed by the secondary outcomes.

Although the formula for $T(V^{-1})$ is only an asymptotic approximation for the variance effect of adding a new study, as noted above we get exact results when adding over the existing studies. We can similarly add the univariate contributions (10) over the existing studies to give Ω_{11}^{-1} times

$$\left[\Omega(\sum \sigma_i^{-2}\mathrm{l}\mathrm{l}^T)\Omega\right]_{11} = \left[(\Omega\mathrm{l})(\Omega\mathrm{l})^T\right]_{11} \sum \sigma_i^{-2} = (\Omega_{11})^2 \sum \sigma_i^{-2} .$$

It follows that

$$\sum T\{(\mathrm{l}^T V_i \mathrm{l})^{-1}\mathrm{l}\mathrm{l}^T\} = \Omega_{11} \sum \sigma_i^{-2} = E ,$$

and so

$$\sum B(V_i^{-1}) = 1 - E . \tag{12}$$

This confirms that the efficiency of univariate meta-analysis can be interpreted as the total of the proportional variance contributions of all the primary outcomes, and that the sum of the borrowing of strengths of these studies is the proportion of the total variance which is attributable to the secondary outcomes. For studies within the meta-analysis, $B(V_i^{-1})$ is equivalent to BoS_{ir}^{SD} in Jackson's notation (Jackson *et al.*, 2017, section 2.4), and the additivity property (12) is implied by equations (11) and (12) of that section.

Both the functions $T(V^{-1})$ and $B(V^{-1})$ are linear functions in the sense that, for any positive scalar constant k,

$$T(kV^{-1}) = kT(V^{-1}) \tag{13}$$

and

$$B(kV^{-1}) = kB(V^{-1}) \quad . \tag{14}$$

Now multiplying the matrix V^{-1} by k is like increasing the study sample size by the factor k whilst keeping the relative magnitudes of the elements of V^{-1} the same. We can think of these relative magnitudes as determined by the design of the study — characteristics of the population from which we are sampling. The actual magnitudes of the elements of V^{-1} are then determined by the sample size. Property (13) confirms that if we add a new study onto the meta analysis and double its sample size, then the decrease in variance will double. Property (14) shows that if a study gives no borrowing of strength so that $B(V^{-1}) = 0$, then $B(kV^{-1}) = 0$ for all k. So whether or not a study offers any borrowing of strength depends only on the study's design and not on its sample size.

Riley (2009) noted that if the V_is are all the same then there is no borrowing of strength, and so the secondary outcomes are then irrelevant as far as estimating β_1 is concerned. This follows immediately from the above formulation, since (6) would then give

$$T(\bar{V}^{-1}) = \frac{1}{n} = T((1\!\!1^T \bar{V} 1\!\!1)^{-1} 1\!\!1 1\!\!1^T) \ , \tag{15}$$

and hence $B(\bar{V}^{-1}) = 0$. And so if all the V_is are the same, $V_i = \bar{V}$ and so $B(V_i^{-1}) = 0$ for all i, hence $E = 1$. This also follows from a simple argument of sufficiency: if $V_i = \bar{V}$ for all i then the score function (2) is exactly equivalent to that of a single study with $V = \bar{V}$ and $y = \sum y_i$. But for any single study the estimate of β_j is simply the jth treatment effect estimate y_j. As Riley (2009) implies, and already found here, borrowing of strength can only arise if there are differences between the V_is. More generally, for there to be any borrowing of strength, these differences must not be simply a matter of different sample sizes, but substantive differences in the background and research methods used in each study. This generalizes the special case of two groups of bivariate studies with proportional V_i's discussed in Jackson *et al.* (2017, section 2.2.1).

When the V_is differ and $E < 1$, as will usually be the case in practice, (15) still holds for any study with $V = k\bar{V}$ for some scalar k, and so such a study will also give no borrowing of strength. We could describe such a study as one with 'average design'. This suggests that it will tend to be the studies which are most atypical in terms of design which contribute most borrowing of strength. Studies whose designs are fairly typical of the meta-analysis as a whole are likely to give little or no borrowing of strength, regardless of their sample sizes. Exactly what this means will be investigated further in sections 3 and 4.

2.2 Missing outcomes

The univariate effect in (10) is for a study in which only the primary outcome is observed. More generally, suppose that only q of the p outcomes are observed, outcomes y_j with $j = j_1, j_2, \cdots, j_q$, with the remaining $(p - q)$ outcomes assumed to be missing at random. We can think of this as selecting a q-dimensional sub-vector from the $p \times 1$ vector y, which we can write as $J^T y$ where J is the $p \times q$ incidence matrix

$$J_{jk} = \begin{cases} 1 & \text{if } j = j_k \\ 0 & \text{otherwise} \end{cases} \quad \text{for } j = 1, 2, \cdots, p \ ; \ k = 1, 2, \cdots, q \ .$$

Matrix J is simply the matrix of zeroes and ones which picks out the required components — the first column has 1 in row j_1, the second has 1 in row j_2, and so on, with all other elements set to 0. Such a study's contribution to the score function for the corresponding sub-vector of β is then

$$(J^T V J)^{-1} J^T (y - \beta) \ . \tag{16}$$

There is no contribution to the score function for the missing outcomes, and so this study's contribution to the score function for the complete vector β is (16) padded out with zeros for each of the unobserved outcomes, namely

$$V_*^{-1}(y - \beta) \ ,$$

where now

$$V_*^{-1} = J(J^T V J)^{-1} J^T \ . \tag{17}$$

Thus to fit the multivariate meta-analysis model when one or more of the studies have missing outcome estimates, we simply use the complete data method as before but with the V_i's for the incomplete studies replaced by the appropriate matrix (17).

If all outcomes are measured, then $q = p$, J is the $p \times p$ identity matrix, and $V_*^{-1} = V^{-1}$ as expected. If only the primary outcome is measured, then $J = 1$ and V_*^{-1} is the previous case (9). Of particular interest is when only the secondary outcomes are measured, since in this case we have a study which cannot be included in a univariate analysis of the primary outcome, but can be included in a multivariate analysis which can then allow information about the unobserved primary outcome to be imputed from the observed values of the secondary outcomes. In this case, J is the $p \times (p-1)$ matrix consisting of the $(p-1) \times (p-1)$ identity matrix supplemented with a row of zeros along the top.

Some care is needed in interpreting the notation V_*^{-1}. By replacing V^{-1} with V_*^{-1} for studies with missing outcomes, the usual formulae for maximum likelihood estimation set out earlier in this section continue to apply even if some, or even all, of the studies in the meta-analysis have one or more missing outcomes. But, despite the notation, V_*^{-1} cannot be interpreted as a matrix inverse (it is singular), or as the known value of V^{-1} for an incomplete study. In reality, all the elements of V^{-1} are unknown parameters, but with complete data we follow the usual convention of assuming these are known because they can be consistently estimated from the within-study data. However, with incomplete outcomes, only the sub-matrix $J^T V J$ of the full matrix V is estimable, and so we have an estimate of V_*^{-1} but not of V^{-1}. The rows and columns of zeros in V_*^{-1} imply that various unidentifiable correlation parameters within V^{-1} are being artificially set to zero. An equation such as (10) means that the contribution of an incomplete study to the variance of $\hat{\beta}$ is *as if* $V^{-1} = V_*^{-1}$. It does not mean that $V^{-1} = V_*^{-1}$ in the usual sense of a mathematical equality.

This discussion gives a formal justification for the more informal data augmentation view taken by Riley (2009) and Jackson *et al.* (2011), who refer to missing outcomes as equivalent to setting their variances to ∞ and their correlations to zero.

3 Borrowing of strength in bivariate meta-analysis

3.1 The borrowing of strength plot

In bivariate meta-analysis, with only one secondary outcome, we can obtain reasonably simple explicit expressions for all of the quantities discussed in the last section. In particular, the finding that borrowing of strength depends on differences between the V_is can be given a constructive interpretation in terms of residuals in a regression model.

In the bivariate case, suppose that the variance matrix V_i of $y_i = (y_{i1}, y_{i2})^T$ is

$$V_i = \begin{pmatrix} \sigma_i^2 & \rho_i\sigma_i\nu_i \\ \rho_i\sigma_i\nu_i & \nu_i^2 \end{pmatrix} .$$

So (σ_i, ν_i) are the standard errors of (y_{i1}, y_{i2}), ρ_i is the correlation between them, and the inverse of V_i is

$$V_i^{-1} = \frac{1}{1-\rho_i^2} \begin{pmatrix} \sigma_i^{-2} & -\rho_i\sigma_i^{-1}\nu_i^{-1} \\ -\rho_i\sigma_i^{-1}\nu_i^{-1} & \nu_i^{-2} \end{pmatrix} . \tag{18}$$

Adding (18) over the n studies, and taking the inverse, gives the harmonic mean

$$\bar{V} = n\Omega = \frac{n}{s_{11}s_{22} - s_{12}^2} \begin{pmatrix} s_{22} & s_{12} \\ s_{12} & s_{11} \end{pmatrix} \tag{19}$$

where (s_{11}, s_{22}, s_{12}) are weighted between-studies sums of squares and products of the outcome accuracies $(\sigma_i^{-1}, \nu_i^{-1})$,

$$s_{11} = \sum \frac{\sigma_i^{-2}}{1-\rho_i^2} \tag{20}$$

$$s_{22} = \sum \frac{\nu_i^{-2}}{1-\rho_i^2} \tag{21}$$

$$s_{12} = \sum \frac{\rho_i}{1-\rho_i^2}\sigma_i^{-1}\nu_i^{-1} , \tag{22}$$

with weights depending on different functions of the within-study correlations ρ_i. As expected, each of these quantities retains the feature of a harmonic average.

Apart from these differences in the weights, (20) - (22) are like the second-order absolute sample moments of the n pairs $(\sigma_i^{-1}, \nu_i^{-1})$, suggesting a through-the-origin linear regression model in which we can examine the extent to which a study's primary accuracy σ^{-1} can be predicted from its secondary accuracy ν^{-1}. Allowing for the different weights, consider predicting u_i from v_i, where

$$u_i = \frac{\rho_i}{(1-\rho_i^2)^{\frac{1}{2}}}\sigma_i^{-1} \quad , \quad v_i = \frac{1}{(1-\rho_i^2)^{\frac{1}{2}}}\nu_i^{-1} . \tag{23}$$

If we plot the n observed values of u_i against the corresponding values of v_i, the least squares slope through the origin is

$$\frac{\sum u_i v_i}{\sum v_i^2} = \frac{s_{12}}{s_{22}} ,$$

and so the least squares prediction line is

$$\hat{u} = \frac{s_{12}}{s_{22}} v = \frac{s_{12}}{(1-\rho^2)^{\frac{1}{2}} s_{22}} \nu^{-1} \ . \tag{24}$$

Requiring the regression line to go though the origin is a natural requirement, since if we know that a study has a very small sample size then we know in advance that both u and v will be close to zero. The definitions of u and v in (23) have assumed complete data, but studies with missing data can also be included as in Section 2.2. If only the primary outcome estimate in the ith study is observed, then we take both ν_i^{-1} and ρ_i to be zero, and so $u_i = v_i = 0$. If only the secondary outcome estimate is observed, we take take σ_i^{-1} and ρ_i to be zero, leading to $u_i = 0$ and $v_i = \nu_i^{-1}$.

The plot of the n values of u_i against their predicted values $\hat{u}_i$ turns out to be closely related to the borrowing of strength function $B(V_i^{-1})$ defined in section 2. Using (18) and (19), and evaluating the required matrix terms explicitly, we get

$$T(V_i^{-1}) = \frac{[\Omega V_i^{-1} \Omega]_{11}}{\Omega_{11}} = \frac{s_{22}^2 \sigma_i^{-2} - 2 s_{12} s_{22} \rho_i \sigma_i^{-1} \nu_i^{-1} + s_{12}^2 \nu_i^{-2}}{s_{22}(s_{11}s_{22} - s_{12}^2)(1-\rho_i^2)} \ .$$

Rewriting ν_i and σ_i in terms of u_i and v_i, and completing the square, gives

$$T(V_i^{-1}) = \Omega_{11} \left\{ \frac{1-\rho_i^2}{\rho_i^2} u_i^2 + \left(u_i - \frac{s_{12}}{s_{22}} v_i \right)^2 \right\} \ .$$

The first term in the outer brackets is just σ_i^{-2}, proportional to the univariate variance contribution of the primary outcome, and so the borrowing of strength is just the second term

$$B(V_i^{-1}) = \Omega_{11} \left(u_i - \frac{s_{12}}{s_{22}} v_i \right)^2 = \Omega_{11} (u_i - \hat{u}_i)^2 \ . \tag{25}$$

Thus $B(V_i^{-1})$ is proportional to the squared residual of the point $(\hat{u}_i, u_i)$ from the diagonal prediction line $u = \hat{u}$. For any other study with inverse variance V^{-1}, $B(V^{-1})$ is similarly proportional to the squared residual of its point $(\hat{u}, u)$ from the line, and so indicates the (approximate) decrease in $\mathrm{Var}(\hat{\beta}_1)$ which we would get if we were to add this study into the meta-analysis. The proportionality factor is

$$\Omega_{11} = \mathrm{Var}(\hat{\beta}_1) = \frac{s_{22}}{s_{11}s_{22} - s_{12}^2} \ .$$

If the ith study has missing data, $(\hat{u}_i, u_i)$ is either (0,0) when the secondary outcome estimate is missing, or $((s_{12}/s_{22})\nu_i^{-1}), 0)$ when the primary outcome estimate is missing. In the first case, the point is always on the line and so, as expected, there can be no borrowing of strength. In the second case, the point is down on the horizontal axis and so will generally have a non-zero residual and so, again as expected, will contribute at least some borrowing of strength.

The $(\hat{u}_i, u_i)$ plot is easier to interpret if we first scale u_i and $\hat{u}_i$ by the factor $\Omega_{11}^{\frac{1}{2}}$, giving

$$w_i = \Omega_{11}^{\frac{1}{2}} u_i = \left(\frac{s_{22}}{s_{11}s_{22} - s_{12}^2} \right)^{\frac{1}{2}} \frac{\rho_i \sigma_i^{-1}}{(1-\rho_i^2)^{\frac{1}{2}}}$$

and

$$\hat{w}_i = \Omega_{11}^{\frac{1}{2}} \hat{u}_i = \left(\frac{s_{12}^2}{s_{22}(s_{11}s_{22} - s_{12}^2)} \right)^{\frac{1}{2}} \frac{\nu_i^{-1}}{(1-\rho_i^2)^{\frac{1}{2}}} \ .$$

We call the scatter plot of w_i against $\hat{w}_i$ the *Borrowing of Strength Plot*. Now the ith squared residual from the diagonal line, $(w_i - \hat{w}_i)^2$, is *equal* to $B(V_i^{-1})$. The combined variance contributions of the secondary outcomes in the meta-analysis is indicated by the scatter of the points about the diagonal regression line. If the points all lie on the line then $B(V_i^{-1}) = 0$ for all i and so $E = 1$. More generally, we can show from the earlier formulae that

$$1 - E = \sum (w_i - \hat{w}_i)^2 \ , \tag{26}$$

and so $1 - E$ is equal to the residual sum of squares of the points in the borrowing of strength plot.

To aid interpretation of the borrowing of strength plot, equation (26) means that, for efficiency E, the root mean squared distance of the points from the diagonal line $w = \hat{w}$ is

$$\bar{d} = \left(\frac{1-E}{n} \right)^{\frac{1}{2}} \ .$$

For example, to achieve 90% efficiency, the root mean squared distance is $\bar{d} = (10n)^{-\frac{1}{2}}$. This is indicated on the borrowing of strength plot by the two parallel lines

$$w = \hat{w} \pm \left(\frac{1}{10n} \right)^{\frac{1}{2}} \ . \tag{27}$$

These lines give a visual benchmark for interpreting residuals in terms of efficiency. If the points are predominantly inside, or predominantly outside, these lines, then the efficiency of univariate meta-analysis is likely to be greater than, or less than, 0.9. As noted previously in section 2.1, an efficiency of 90% indicates that the information gained from the secondary outcomes in multivariate meta-analysis is like the extra information which would be available in univariate meta-analysis if we had an additional $n/9$ studies.

Equation (25) also gives us the necessary and sufficient condition for a study to give no borrowing of strength. If $(\hat{u}_i, u_i)$ lies on the line, then $u_i = (s_{12}/s_{22})v_i$ and so

$$\frac{\rho_i \sigma_i \nu_i}{\sigma_i^2} = \frac{s_{12}}{s_{22}} \ . \tag{28}$$

The left hand side of (28) is the ratio of the covariance element in V_i (V_{i12}) to its primary diagonal element V_{i11}, whilst the right hand side is the ratio of the corresponding elements of Ω, or of $\bar{V}$. For no borrowing of strength these are equal, and so

$$\frac{V_{i12}}{V_{i11}} = \frac{\bar{V}_{12}}{\bar{V}_{11}} \ . \tag{29}$$

Previously we noted that a study with $V_i = k\bar{V}$ for some scalar constant k gives no borrowing of strength. This is a sufficient but not necessary condition — all we need is that the top row (or left hand column) of V_i is proportional to the top row (or left hand column) of

$\bar{V}$. In particular, there is no requirement on the secondary variance ν_i^2 *per se.* We show in Section 4 that this generalizes to any number of secondary outcomes.

The borrowing of strength plot also illustrates two other aspects of borrowing of strength which were discussed in Section 2. Firstly, for studies in the meta-analysis, $B(V_i^{-1})$ is the proportional contribution of the ith secondary outcome estimate to $\mathrm{Var}(\hat{\beta}_1)$ (the direct interpretation), but, for a study outside the meta-analysis, $B(V^{-1})$ is only the approximate (large n) contribution which the secondary outcome estimate would make if this study were added into the meta- analysis (the add-one-in interpretation). We see the nature of this approximation in the borrowing of strength plot. The line $w = \hat{w}$ is the least squares line of best fit (through the origin) for the n points $(\hat{w}_i, w_i)$. But if we add in the new study, the value of the scale factor $\Omega_{11}^{\frac{1}{2}}$ will change, affecting the coordinates for all the studies. So the residual of the new point from the line fitted by least squares to the enhanced data will not be the same as the residual from the line calculated from the original n studies alone. If n is large then adding one more study will only have a small effect on the fitted line, and so these two residuals will be similar.

Secondly, we have noted the linear property of the function $B(V^{-1})$ in (14). If we multiply V^{-1} by k then both w and $\hat{w}$ are scaled by the factor $\sqrt{k}$ and so the squared residual from the diagonal line is scaled by the original factor k, which means that $B(kV^{-1}) = kB(V^{-1})$ as required. If V^{-1} gives no borrowing of strength then the point will simply move up or down the diagonal line according to the value of k.

3.2 Example

Figure 1 illustrates data from 10 clinical trials designed to test the effectiveness of hypertension treatments in reducing the risk of subsequent diagnoses of cardiovascular disease (CVD) and stroke. This meta-analysis was originally published by Wang *et al.* (2005) and discussed further in Riley *et al.* (2015) and Jackson *et al.* (2017). Each randomized controlled trial was well balanced between active treatment and placebo, but varied widely in size, from under 200 patients in trial number 3 to almost 7000 patients in trial number 5 (trial numbers consistent with previous tables, for example Table 1 of Riley *et al.*, 2015). Figure 1 shows individual trial data for two outcomes, the estimated log hazard ratio (log HR) for CVD (y_1), and the estimated log HR for stroke (y_2). Values of y_1 (crosses) and y_2 (circles) are plotted against the within-study correlations ρ_i, with the corresponding pairs of within-study confidence intervals for β_1 and β_2 shown as the solid and dashed line segments respectively. The small numbers to the left of the confidence intervals identify the study numbers 1-10. The vertical coordinates of some of the data in Figure 1 have been slightly adjusted to aid clarity of the plot. Separate homogeneity tests of the values of y_1 and y_2 are both well-consistent with fixed effects models, leading to univariate combined confidence intervals of $(-0.374, -0.115)$ for CVD log HR, and $(-0.531, -0.235)$ for stroke log HR. It is not at all obvious from Figure 1 whether a bivariate approach, taking both outcomes into account, will lead to more accurate estimates and if so by how much.

If the log hazard ratio for CVD is taken as the primary outcome (y_1), the formulae in Section 2 give the respective univariate and multivariate estimates of β_1 and their variances to be

$$\tilde{\beta}_1 = -0.244 \ , \ \mathrm{Var}(\tilde{\beta}_1) = 0.00434 \ , \ \hat{\beta}_1 = -0.244 \ , \ \mathrm{Var}(\hat{\beta}_1) = 0.00427 \ . \tag{30}$$

The estimates are virtually identical. The ratio of the variances is the efficiency $E = 0.984$, showing that in this example the stroke data give very little extra information for the assessment of CVD risk reduction. The last two rows of Table 1 give the total variance contribution $T(V_i^{-1})$ and the borrowing of strength $B(V_i^{-1})$ for each of the ten studies, confirming that none of these studies gives any worthwhile contribution from the secondary outcome. We can check directly that the borrowing of strength figures add up to $1 - E$. Figure 2 shows the corresponding borrowing of strength plot. Again we can check the theory by showing that the least squares slope of these points is 1, and that the residual sum of squares is $1 - E = 0.016$. The two dotted lines are the 90% efficiency bars (27). All the points are well within these limits, confirming the high efficiency of univariate meta-analysis and the minimal contribution of the secondary outcomes in this case.

If the primary interest is to estimate the log HR for stroke instead of CVD, then we use exactly the same formulae but with the notation reversed appropriately, retaining the same values of ρ_i and s_{12} but interchanging y_{i1} with y_{i2}, σ_i with ν_i, and s_{11} with s_{22}. In terms of the original notation we are now estimating β_2, giving

$$\tilde{\beta}_2 = -0.383 \ , \ \mathrm{Var}(\tilde{\beta}_2) = 0.00569 \ , \ \hat{\beta}_2 = -0.381 \ , \ \mathrm{Var}(\hat{\beta}_2) = 0.00505 \ ,$$

with the new efficiency $E = 0.888$. The two estimates are again very similar, but the multivariate method is now noticeably more accurate. The borrowing of strength plot for estimating β_2 is shown in Figure 3, which now shows a much greater dispersion about the regression line than in Figure 2 (the mean squared spread of the residuals is now close to the dotted 90% efficiency lines). Figure 4 illustrates the proportional contributions which the studies make to $\mathrm{Var}(\hat{\beta}_2)$. This is a line plot, the upper (solid) line highlighting the values of $T(V_i^{-1})$ (total contributions), the lower (broken) line highlighting the corresponding values of $T(V_i^{-1}) - B(V_i^{-1})$ (univariate contributions). The distance between the two lines matches the squared residuals in Figure 3. The largest borrowing of strength comes from the ninth study, where the secondary outcome accounts for almost a third of the total variance contribution of that study. This study accounts for about a half of the total borrowing of strength of all the studies, although its sample size is by no means the largest (although it does have the largest correlation). The efficiency of 89% shows that the variance of the multivariate estimate of β_2 is about 10% lower than the variance of the univariate estimate, which is roughly what we might expect if we were able to increase the size of a univariate meta-analysis from 10 to 11 studies. In this sense, the value of including data on the 10 secondary outcomes can be likened to the value of having the primary outcome estimate of one additional study.

Comparing these two efficiencies shows that there is no symmetry in borrowing of strength: the values of y_1 make a modest contribution to the accuracy of $\hat{\beta}_2$ but the values of y_2 make almost no contribution to the accuracy of $\hat{\beta}_1$. More generally, we can show that if $E = 1$ (no borrowing of strength) when estimating β_1, then E will be strictly less than one (positive borrowing of strength) for estimating β_2 except in the special case of all the studies having the same correlation (as in section 3.3).

There is no missing data in these trials. To illustrate the impact that missing outcomes might have had on this analysis, and to demonstrate the use of multivariate meta-analysis when there are missing data, imagine that we wish to estimate the CVD risk β_1 when both outcomes are available in trials $1 - 5$ but only the stroke outcome is measured in the remaining trials $6-10$. Then we get

$$\tilde{\beta}_1 = -0.175\ ,\ \mathrm{Var}(\tilde{\beta}_1) = 0.0117\ ,\ \hat{\beta}_1 = -0.196\ ,\ \mathrm{Var}(\hat{\beta}_1) = 0.00956\ .$$

Inevitably, the variance of $\hat{\beta}_1$ is now considerably larger than the complete data case in (30). The efficiency of $E = 0.815$ now reflects the difference between univariate meta-analysis using only the first 5 trials, and multivariate meta-analysis using the information in all 10 trials. This value of E is roughly 5/6, which is the variance improvement we might expect to get if we were able to use univariate meta-analysis with the number of trials increased from 5 to 6. In this sense, the value of including the 5 trials with missing primary outcomes can be likened to the value of having one further trial with complete data.

Figure 5 is the borrowing of strength plot for this missing data example. The points $(\hat{w}_i, w_i)$ for studies 1 to 5 are the same as in Figure 2 except for a re-scaling of the axes, but the five points for the missing studies are all moved vertically down to the horizontal axis. This completely alters the size of the residuals and hence the borrowing of strength figures for all of the trials. Figure 2 showed that, for estimating β_1 with complete data, none of the 10 secondary values y_2 makes any useful contribution over and above the contribution of the corresponding observed values of y_1. So we might expect that with these missing data all of the borrowing of strength would come from trials $5-10$ since in these trials y_1 is no longer available. But this is not so, as shown in the variance contributions plot in Figure 6 (using the same format as Figure 4). Now we get

$$\sum_1^5 B(V_i^{-1}) = 0.144\ ,\ \sum_6^{10} B(V_{*i}^{-1}) = 0.041\ , \tag{31}$$

where V_{*i}^{-1} is the proxy matrix (9) for the ith trial, the 2×2 matrix with ν_i^{-2} as the lower diagonal element and zero's elsewhere. The sum of these two numbers in (31) is $0.185 = 1 - E$ as expected, but the missing studies only contribute 22% of the total borrowing of strength. This illustrates one of the main points in Section 2.1, that the borrowing of strength given by a particular study depends on how typical that study is of the meta analysis as a whole, and only indirectly on the statistical characteristics of the study itself. Changing the later studies leaves studies 1-5 exactly the same, but can drastically alter their borrowing of strength. We can also see a difference if we look at the estimation of β_2 with the same pattern of missing data. We are again leaving trials $1-5$ as before, but now trials $5-10$ measure only the primary outcome. Now the borrowing of strengths $B(V_i^{-1})$ for the first five trials add up to about 1%, less than the sum over the same trails in the complete data case of about 5% (Figure 4).

3.3 The special case of equal within-study correlations

A statistical understanding of the plotting coordinates $(\hat{w}_i, w_i)$ in the borrowing of strength plot is complicated by the fact that the weighted sums of squares and products in (20)-(22) use different weights, also reflected in the different factors appearing in u_i and v_i in (23). However, if the ρ_i's are constant, $\rho_i = \rho_0$ say, these differences in the weights can be absorbed into an overall scale factor, leading to a more transparent version of many of the formulae in Section 3.1. This special case is also of interest in its own right since, as will be discussed in Section 5, fitting the bivariate model with constant correlations can provide a useful sensitivity analysis in cases where the within-study correlations are not provided by the study reports (Jackson *et al.*, 2011).

Let $(s^*_{11}, s^*_{22}, s^*_{12})$ be the ordinary (un-weighted) sums of squares and products of the n accuracy pairs $(\sigma_i^{-1}, \nu_i^{-1})$. Then if we imagine a scatter plot of σ_i^{-1} against ν_i^{-1}, the least squares line of best fit through the origin has slope s^*_{12}/s^*_{22}. Thus, for any given value of ν^{-1}, the least squares prediction of σ^{-1} is

$$\widehat{\sigma^{-1}} = \frac{s^*_{12}}{s^*_{22}}\nu^{-1} \ .$$

Relating this to the earlier notation gives

$$(u_i, \hat{u}_i) = \frac{\rho_0}{(1-\rho_0^2)^{\frac{1}{2}}}(\sigma_i^{-1}, \widehat{\sigma_i^{-1}}) \ ,$$

and so the ith residual in the borrowing of strength plot is

$$(w_i - \hat{w}_i) = \left\{\frac{\rho_0^2 s^*_{22}}{s^*_{11}s^*_{22} - \rho_0^2 s^{*2}_{12}}\right\}^{\frac{1}{2}} (\sigma_i^{-1} - \widehat{\sigma_i^{-1}}) \ .$$

So, with this slightly different scale factor, we can think of the borrowing of strength plot as little more than a linear regression of the within-study accuracies of the primary outcomes plotted against the corresponding accuracies of the secondary outcomes. The fact that borrowing of strength is given by the least squares residuals again confirms that borrowing of strength is all a matter of how the variances of individual studies fit in with the overall pattern of variances in the meta-analysis as a whole.

3.4 Borrowing of strength as a within-study ratio

We have measured borrowing of strength in terms of $B(V^{-1})$, the variance contribution of a study's secondary outcome relative to the overall variance Ω_{11}. We could instead consider the ratio $R(V^{-1}) = B(V^{-1})/T(V^{-1})$, the contribution of the study's secondary outcome as a proportion of that study's total contribution to $\mathrm{Var}(\hat{\beta}_1)$. This removes the effect of any scale factor in V, so that for a given meta-analysis $R(V^{-1})$ is a function of just two quantities, ρ, the correlation between the outcomes, and z, the ratio of the standard errors

$$z = \frac{\sigma}{\nu} \ .$$

The earlier formulae now give

$$R(V^{-1}) = \frac{\rho^2(u - s_{22}^{-1}s_{12}v)^2}{\rho^2(u - s_{22}^{-1}s_{12}v)^2 + (1-\rho^2)u^2} = \frac{(\rho - s_{22}^{-1}s_{12}z)^2}{(\rho - s_{22}^{-1}s_{12}z)^2 + 1 - \rho^2} \ . \qquad (32)$$

A contour plot of (32) against ρ and z gives a complete picture of how, within a given meta analysis (i.e. for a given value of the slope parameter s_{12}/s_{22}) a study's borrowing of strength, defined in this way, depends on individual study characteristics.

Figure 7 shows a contour plot of $R(V^{-1})$ for $s_{12}/s_{22} = 0.660$, the value of the slope parameter found in the example in Section 3.2. Values of ρ are shown on the vertical axis, values of z are shown using a log scale on the horizontal axis. The contour values are labelled along the bottom and up the left hand side of the plot. The dashed line is the zero contour when $\rho = 0.660z$: at these values there is no borrowing of strength. The contour plot shows that R is large when either z is large (y_1 less accurate than y_2), or when z is small (y_1 more accurate than y_2) *and* ρ is large (outcomes highly correlated). The smaller plotting symbols $1-10$ on Figure 7 show the values of (z, ρ) for the 10 studies in the example. Most of the points are fairly close to the zero contour: for only 3 of these studies is $R(V^{-1}) > 0.1$, suggesting that E is close to 1, as found earlier. The plotting symbol X (in bold) indicates the point (z, ρ) for a study with $V = \bar{V}$ defined in (4). This point corresponds to the harmonic mean of the ten points labeled $1-10$, and, as expected, lies on the zero contour (no borrowing of strength). The interpretation of a study's borrowing of strength as a contrast between V and $\bar{V}$ can be seen on the graph as the distance between the study's (z, ρ) point and the harmonic mean point X, measured in the direction orthogonal to the contours in that region.

If one of the outcomes is missing, $R(V^{-1})$ becomes $R(V_*^{-1})$ and so the point (z, ρ) lies on the horizontal axis, to the extreme left if y_2 is missing and to the extreme right if y_1 is missing, giving $R(V_*^{-1}) = 0$ and 1 respectively, as expected.

4 Borrowing of strength in multivariate meta-analysis

4.1 Decomposing the variance contribution of an individual study

This section looks at the generalization of section 3 to the multivariate case with $p > 2$. Now the outcome estimates of the ith study are $y_i = (y_{i1}, y_{i2})$ with y_{i2}, the secondary outcome estimates, a $(p-1) \times 1$ vector. When $p = 2$, all the formulae in this section reduce to the corresponding expressions already seen in section 3.

In the multivariate case, we write $V_i = \text{Var}(y_i)$ as the partitioned matrix

$$V_i = \begin{pmatrix} \sigma_i^2 & \sigma_i \rho_i^T \Lambda_i \\ \sigma_i \Lambda_i \rho_i & \Lambda_i P_i \Lambda_i \end{pmatrix} , \tag{33}$$

where σ_i^2 is the variance of y_{i1} as before, ρ_i is the $(p-1) \times 1$ vector of correlation coefficients between y_{i1} and y_{i2}, P_i is the $(p-1) \times (p-1)$ correlation matrix of y_{i2}, and Λ_i is the $(p-1) \times (p-1)$ diagonal matrix of the standard deviations of the components of y_{i2}.

To simplify the algebra for calculating matrix inverses, define

$$a_i = \rho_i^T P_i^{-1} \rho_i \ , \ b_i = P_i^{-1} \rho_i \ , \ C_i = P_i^{-1} + \frac{b_i b_i^T}{1 - a_i} \ . \tag{34}$$

Then, using a standard formula for the inverse of a partitioned matrix,

$$V_i^{-1} = \frac{1}{1 - a_i} \begin{pmatrix} \sigma_i^{-2} & -\sigma_i^{-1} b_i^T \Lambda_i^{-1} \\ -\sigma_i^{-1} \Lambda_i^{-1} b_i & (1 - a_i) \Lambda_i^{-1} C_i \Lambda_i^{-1} \end{pmatrix} \ . \tag{35}$$

Some of this notation can be interpreted in terms of a multiple regression of the primary on the secondary outcome estimates within the ith study. The vector in the off-diagonal partition of (35) is proportional to the vector of regression coefficients, and $\sigma_i^2(1-a_i)$ is the residual mean square. Thus a_i can be interpreted as the multiple correlation (R^2) of this regression: $a_i = 0$ means that the primary and secondary outcome estimates are independent, $a_i = 1$ means that they are exactly linearly related.

Adding (35) over the n studies gives

$$\Omega^{-1} = \sum V_i^{-1} = n(\bar{V})^{-1} = \begin{pmatrix} s_{11} & -s_{12}^T \\ -s_{12} & S_{22} \end{pmatrix} , \tag{36}$$

where

$$\begin{aligned} s_{11} &= \sum \frac{\sigma_i^{-2}}{1-a_i} \\ S_{22} &= \sum \Lambda_i^{-1} C_i \Lambda_i^{-1} \\ s_{12} &= \sum \frac{\sigma_i^{-1}\Lambda_i^{-1} b_i}{1-a_i} . \end{aligned}$$

Thus the inverse of (36) is

$$\Omega = \frac{\bar{V}}{n} = \frac{1}{s_{11} - s_{12}^T S_{22}^{-1} s_{12}} \begin{pmatrix} 1 & s_{12}^T S_{22}^{-1} \\ S_{22}^{-1} s_{12} & (s_{11} - s_{12}^T S_{22}^{-1} s_{12}) S_{22}^{-1} + S_{22}^{-1} s_{12} s_{12}^T S_{22}^{-1} \end{pmatrix} \tag{37}$$

and so $\mathrm{Var}(\hat{\beta}_1)$ is

$$\Omega_{11} = \frac{1}{s_{11} - s_{12}^T S_{22}^{-1} s_{12}} . \tag{38}$$

As before, the components of Ω^{-1} in (36) are weighted sums of squares and products of the precisions of the components of y_i: σ_i^{-1} for the primary outcome and the diagonal elements of Λ_i^{-1} for the secondary outcomes. The scalar s_{11} is the same as in the bivariate case, S_{22} is the $(p-1) \times (p-1)$ matrix of weighted sums of squares and products for the secondary outcome precisions, and s_{12} is the corresponding $(p-1) \times 1$ vector of weighted sums of cross products between the primary and secondary precisions. When $p = 2$ these formulae reduce to the corresponding quantities in Section 3, with the matrix S_{22} becoming the scalar s_{22}. In the bivariate case, the weights involved in these sums are also the same, since when $p = 2$ the quantities defined in (34) reduce to the scalars

$$P_i = 1 \ , \ a_i = \rho_i^2 \ , \ b_i = \rho_i \ , \ C_i = \frac{1}{1-\rho_i^2} \ , \tag{39}$$

where ρ_i is now just the ordinary scalar correlation between the two outcome estimates in the ith study.

From (35), (37) and (38), the total variance contribution of the ith study is

$$\begin{aligned} T(V_i^{-1}) &= \frac{1}{s_{11} - s_{12}^T S_{22}^{-1} s_{12}} (1 \ \ s_{12}^T S_{22}^{-1}) V_i^{-1} (1 \ \ s_{12}^T S_{22}^{-1})^T \\ &= \Omega_{11} \left(\frac{\sigma_i^{-2}}{1-a_i} - 2\sigma_i^{-1} \frac{f_i^T b_i}{1-a_i} + f_i^T C_i f_i \right) , \end{aligned}$$

where f_i is the $(p-1)\times 1$ vector

$$f_i = \Lambda_i^{-1} S_{22}^{-1} s_{12} \ . \tag{40}$$

As before,

$$T(\sigma_i^{-2} 11^T) = \Omega_{11}\sigma_i^{-2} \ ,$$

and so

$$B(V_i^{-1}) = T(V_i^{-1}) - T(\sigma_i^{-2}11^T) = \Omega_{11}\left(\sigma_i^{-2}\frac{a_i}{1-a_i} - 2\sigma_i^{-1}\frac{f_i^T b_i}{1-a_i} + f_i^T C_i f_i\right) \ .$$

This is a quadratic function of σ_i^{-1}, the accuracy of the primary outcome. Completing the square gives

$$B(V_i^{-1}) = \Omega_{11}\left[\left\{\left(\frac{a_i}{1-a_i}\right)^{\frac{1}{2}}\sigma_i^{-1} - \frac{f_i^T b_i}{\{a_i(1-a_i)\}^{\frac{1}{2}}}\right\}^2 + f_i^T C_i f_i - \frac{(f_i^T b_i)^2}{a_i(1-a_i)}\right] \ . \tag{41}$$

For a given meta-analysis s_{12} and S_{22} are fixed, and so the vector f_i in (40) is just a linear function of the diagonal elements of Λ_i^{-1}, the accuracies of the secondary outcome estimates in the ith study.

For a simpler notation for (41), we extend the u_i and $\hat{u}_i$ notation in the bivariate case to

$$u_i = (\frac{a_i}{1-a_i})^{\frac{1}{2}}\sigma_i^{-1} \quad , \quad \hat{u}_i = \frac{f_i^T b_i}{\{a_i(1-a_i)\}^{\frac{1}{2}}} = s_{12}^T S_{22}^{-1}\frac{\Lambda_i^{-1} b_i}{\{a_i(1-a_i)\}^{\frac{1}{2}}} \ . \tag{42}$$

For a given meta-analysis (fixed values of s_{11}, s_{12} and S_{22}), u_i is proportional to the accuracy of the study's primary outcome and v_i is a scalar linear function of the accuracies of the secondary outcomes. If we define

$$g_i = \frac{(f_i^T b_i)^2}{a_i f^T P_i^{-1} f_i} \ ,$$

then, from (34),

$$f_i^T C_i f_i = f_i^T P_i^{-1} f_i + \frac{(f_i^T b_i)^2}{1-a_i} = \hat{u}_i^2\left(\frac{1-a_i}{g_i} + a_i\right) \ ,$$

from which we get

$$f_i^T C_i f_i - \frac{(f_i^T b_i)^2}{a_i(1-a_i)} = f_i^T C_i f_i - \hat{u}_i^2 = \frac{(1-a_i)(1-g_i)}{g_i}\hat{u}_i^2 \ .$$

Thus (41) is

$$B(V_i^{-1}) = \Omega_{11}\left\{(u_i - \hat{u}_i)^2 + \frac{(1-a_i)(1-g_i)}{g_i}\hat{u}_i^2\right\} \ , \tag{43}$$

and so

$$T(V_i^{-1}) = \Omega_{11}\left\{\frac{1-a_i}{a_i}u_i^2 + (u_i - \hat{u}_i)^2 + \frac{(1-a_i)(1-g_i)}{g_i}\hat{u}_i^2\right\} \ . \tag{44}$$

Up to the scale factor $\Omega_{11} = \mathrm{Var}(\hat{\beta}_1)$, (44) decomposes the total variance contribution of the ith study into three non-negative parts, analogous to the main effects and interaction of the two factors u_i (proportional to the accuracy of the primary outcome) and $\hat{u}_i$ (a linear function of the accuracies of the secondary outcomes). The three effects are

- the term in u_i^2. This is
$$\frac{1-a_i}{a_i} u_i^2 = \sigma_i^{-2} \ ,$$
the direct contribution of the primary outcome of the ith study as in univariate meta-analysis;

- the term in $(u_i - \hat{u}_i)^2$ as in bivariate meta-analysis, measuring the difference between the actual accuracy of the primary outcome and, in some sense, what might be expected from the pattern of u_is and $\hat{u}_i$s observed in the meta-analysis as a whole;

- the term in $\hat{u}_i^2$, the additional effect of the accuracies of the study's secondary outcomes. This is zero if $g_i = 1$.

The borrowing of strength for the ith study is proportional to the sum of the second and third terms of (44). The presence of the third term shows that there is a qualitative difference in borrowing of strength properties between the multivariate and bivariate cases. When $p = 2$, the quantities f_i, b_i and P_i are all scalars as in (39), and hence

$$g_i = \frac{(f_i^T b_i)^2}{a_i f_i^T P_i^{-1} f_i} = \frac{f_i^2 b_i^2}{a_i f_i^2 P_i^{-1}} = 1 \ .$$

Thus, when $p = 2$, $g_i = 1$ for all i and so the third term in (44) is zero.

To see the equivalence of the $(u_i - \hat{u}_i)^2$ term when $p = 2$, the quantity $\hat{u}_i$ in (44), when expressed in the notation of Section 3.1, becomes

$$\hat{u}_i = \frac{s_{12}}{(1-\rho_i^2)^{\frac{1}{2}} s_{22}} \nu_i^{-1} \ ,$$

which is just the same as (24). Hence, in the bivariate case, the residual $(u_i - \hat{u}_i)$ using the definition in (42) is exactly the same as the residual $(u_i - \hat{u}_i)$ defined earlier in (25). The third term in (44) is still zero in the multivariate case if $g_i = 1$, in which case the motivation of $\hat{u}_i$ as a least squares prediction of u_i continues to hold in the sense that $\sum u_i \hat{u}_i / \sum \hat{u}_i^2 = 1$.

4.2 Necessary and sufficient condition for no borrowing of strength

From $B(V_i^{-1})$ in (43), for there to be no borrowing of strength in the ith study we must have two conditions: $g_i = 1$ and $u_i = \hat{u}_i$. We can exclude the trivial case $a_i = 1$ which would mean there is an exact linear relationship between the ith study's primary and secondary outcome estimates.

For the first condition, $g_i = 1$ if $f_i = k\rho_i$ for any arbitrary scalar factor k. This is also a necessary condition for $g_i = 1$, as can be verified directly by using a Lagrange multiplier calculation to find the maximum value of g_i for different values of f_i. From (40) and (33), this means that

$$V_i = \begin{pmatrix} \sigma_i^2 & k^{-1}\sigma_i s_{12}^T S_{22}^{-1} \\ k^{-1}\sigma_i S_{22}^{-1} s_{12} & \Lambda_i P_i \Lambda_i \end{pmatrix} .$$

Comparing this with (37), the equivalent condition is that the covariance vector part of V_i in (33) is a scalar multiple of the corresponding covariance vector part of Ω (or of $\bar{V}$).

For the second condition, if $f_i = k\rho_i$ then

$$\hat{u}_i = \frac{f_i^T b_i}{\{a_i(1-a_i)\}^{\frac{1}{2}}} = \frac{k\rho_i^T P_i^{-1}\rho_i}{\{a_i(1-a_i)\}^{\frac{1}{2}}} = k\left(\frac{a_i}{1-a_i}\right)^{\frac{1}{2}} .$$

and so

$$(u_i - \hat{u}_i)^2 = \left(\frac{a_i}{1-a^i}\right)(\sigma_i^{-1} - k)^2 .$$

So if $u_i = \hat{u}_i$ then $k = \sigma_i^{-1}$. This extends the proportionality between the covariance vector parts of V_i and Ω required for $g_i = 1$ to also include the $(1,1)$ term. So the necessary and sufficient condition for no borrowing of strength is that the first row (or first column) of V_i must be a scalar multiple of the corresponding row or column of the harmonic mean matrix $\bar{V}$. Thus the necessary and sufficient condition (29) in the bivariate case generalizes directly to the multivariate case, where V_{i12} and $\bar{V}_{12}$ are now the covariance vector components of V_i and $\bar{V}$ respectively.

Note that condition (29) gives no constraint on the size of the scalar multiple k involved, and hence no constraint on the sample size of the trial. Both small and large trials may end up giving no borrowing of strength, including studies with large correlations between the primary and secondary outcomes. Note also that condition (29) imposes no constraint on the $(2,2)$ partition of V in (33), i.e. on the distribution of the estimates for the secondary outcomes *per se*.

We commented in Section 4.1 that the some of the components of V_i^{-1} in (35) can be interpreted in terms of a within-study multiple regression of y_{i1} on y_{i2}. We can similarly interpret (29) in terms of the regression the other way round, predicting the vector of secondary outcome estimates y_{2i} from the primary outcome estimate y_{1i}. The vector of regression coefficients for the ith study would then be V_{i12}/V_{i11}, which is just the left side of (29). Hence the necessary and sufficient condition for the ith study to give no borrowing of strength is that the within-study vector of regression coefficients for predicting the secondary from the primary estimates is the same as the corresponding regression vector for a study with the harmonic mean variance matrix $\bar{V}$.

5 Random effects models

The results in this paper depend on some important assumptions, most obviously the assumption of a fixed effects model, that all studies are modelled by (1). This strong assumption, that all the studies are estimating the same treatment effect β, is widely discussed in the univariate meta-analysis literature. Jackson *et al.* (2017) follows a number of other papers on multivariate meta-analysis by also including random effects models.

These papers generalize the usual two-stage approach to univariate random effects meta-analysis by first estimating a between-studies variance matrix Ψ by $\hat{\Psi}$ and then using the fixed effects model (1) with each V_i replaced by $V_i + \hat{\Psi}$. Jackson *et al.* (2010) shows how the familiar univariate DerSimonian-Laird (D-L) estimate (Der-Simonian and Laird, 1986) can be extended to the multivariate case, using the univariate D-L estimates for each outcome taken individually, and analogous method-of-moments estimates for each covariance component. Other methods of estimating Ψ have been discussed in several recent papers (Chen *et al.*, 2012; Jackson *et al.*, 2013; Ma and Mazumdar, 2011).

The borrowing-of-strength quantities E and $B_i = B(V_i^{-1})$ discussed earlier are descriptive measures of how much the multivariate estimation of β_1 has been influenced by the data on the secondary outcomes. In random effects models, the corresponding estimates $\hat{E}$ and $\hat{B}_i = B(\hat{V}_{REi}^{-1})$ calculated from the fitted marginal variance matrices $\hat{V}_{REi} = V_i + \hat{\Psi}$ are similarly descriptive measures of the role of the secondary outcomes within the fitted model. The definition of E in (5) is only a valid measure of efficiency if the variances of the two estimates being compared are based on a consistent model, which means that the diagonal element of $\hat{\Psi}$ for the primary outcome must be the same as the univariate random effects variance estimate we would get if we fitted a univariate random effects meta-analysis model to the data on the primary outcome alone. Only under this condition do we retain the same interpretation of $\hat{E}$ and $\hat{B}_i$ as discussed earlier for fixed effects models. In practice, D-L estimates are almost always used in univariate random effects meta-analysis, suggesting that $\hat{\Psi}$ should be estimated using a method-of-moments estimate which retains the univariate D-L estimates as its diagonal elements. A slight modification to the truncation step in Jackson *et al.* (2010) is needed to ensure that this is always the case, which for bivariate meta-analysis (as in the example below) simply amounts to truncating the estimated random effects correlation to its nearest value in the interval $[-1, 1]$. We can then retain the same interpretation of $\hat{E}$ and $\hat{B}_i$ as a direct comparison of the fitted variance of $\hat{\beta}_1$ using all of the data with the fitted variance we would get from a univariate meta-analysis using only the data on the primary outcomes. In this sense, the theory and interpretation of borrowing of strength statistics for fixed effects models applies in exactly the same way to random effects models, as implied by the discussion in Jackson *et al.* (2017, section 4).

As the variance matrices V_i are assumed known, the descriptive measures E and B_i can also be given an inferential interpretation as estimates of the borrowing of strength parameters of the true underlying model (1). However, applying this to random effects models raises different issues, since now the marginal variance matrices $\hat{V}_{REi}$ depend on $\hat{\Psi}$ which can exhibit substantial sampling uncertainty if n is small (Guolo and Varin, 2017). Arguably, $\hat{\Psi}$ has a greater influence on $\hat{E}$ and on $\hat{B}_i$ than it has on the more usual problem of estimating β, since $\hat{\beta}$ retains its unbiasedness property conditionally on all possible values of $\hat{\Psi}$. However, the example below suggests that $\hat{E}$ and $\hat{B}_i$ can still provide useful estimates of the borrowing of strength properties of the true underlying random effects model.

As a simple illustration in the bivariate case, suppose that the treatment effect estimates in the example of Section 3.2 were in fact generated from the bivariate random effects model

$$y_i \sim N(\beta, V_i + \Psi)\ ;\ i = 1, 2, \cdots, 10, \tag{45}$$

with $\Psi = \alpha \bar{V}$ and $\alpha \geq 0$, where the V_is are as in Table 1 and $\bar{V}$ is their harmonic mean as in (4). Then by increasing α from 0 (the fixed effects model) we get increasing between-study heterogeneity. A small value of α means that the fixed effects model slightly underestimates the variability of the y_is, and the assumed form of Ψ means that the pattern of variances remains reasonably similar to those observed in the data. We can then simulate vectors y_i from (45) and compare the borrowing of strength statistics calculated from the actual marginal variance matrices $V_{REi} = V_i + \alpha\bar{V}$ with the corresponding statistics calculated from the estimated marginal variances $\hat{V}_{REi} = V_i + \hat{\Psi}$. For the reason discussed above, we calculate $\hat{\Psi}$ using the slightly modified version of the method of Jackson *et al.* (2010) which was mentioned earlier.

Table 2 describes the results of a small simulation study based on 1000 replications for each of 5 values of α, ranging from $\alpha = 0$ (fixed effects) to $\alpha = 2$ (quite substantial heterogeneity). We assume that the primary interest is the value of β_2, the log hazard ratio for the risk of stroke. The first row of Table 2 shows the actual efficiencies E based on V_{REi}. As expected, the entry 0.888 for $\alpha = 0$ is just the fixed effects efficiency already quoted in Section 3.2. Adding the same variance matrix to each V_i has the effect of reducing the relative differences between them, which explains why the values of E tend to increase as α increases. The estimated efficiencies $\hat{E}$ based on $\hat{V}_{REi}$ vary randomly between simulations, but their sample medians across the 1000 simulations, shown in the second row of the table, also follow a similar pattern. We summarize the simulation results using medians rather than means because of skewness caused by the truncation of D-L estimates. Section 2.1 has shown that the actual study-specific borrowing of strength components B_i always add up to $1 - E$ and so when, in the random effects model, $\hat{E}$ is different from E we cannot expect the corresponding estimated and true borrowing of strength components to be exactly comparable. However, from a practical point of view, what we would hope to see is that the studies which show the greatest (or least) borrowing of strength under the estimated model are the same, or substantially the same, as the studies which give the greatest (or least) borrowing of strength under the true model. For each simulation, the extent to which this is so is measured by the rank correlation $rc(\hat{B}_i, B_i)$. The third row of Table 2 shows the sample medians of these rank correlations. These are satisfactorily high ($\geq 90\%$) for the smaller values of α but, as expected, tend to deteriorate slightly as the heterogeneity increases.

6 Discussion

In most statistical problems, taking into account data on relevant covariates or confounders leads to more accurate estimates and predictions, especially if the secondary variables are closely correlated with the main variable of interest. However, this is not necessarily the case in meta-analysis — multivariate meta-analysis can give little or no improvement over univariate methods even if the secondary outcomes are closely correlated with the primary outcome. By writing the borrowing of strength measure BoS_{i1}^{SD} proposed by Jackson *et al.* (2017) as the explicit function $B(V^{-1})$ in (11), and then evaluating some of this function's mathematical properties, we have tried to shed light on how and why individual study characteristics may or may not lead to a useful role for secondary outcomes in multivariate meta-analysis.

The paper has made a number of important assumptions. The fixed effects model (1) and its application to random effects models has been discussed in Section 5. We have also assumed that, by replacing a within-study inverse variance V^{-1} by V_*^{-1} in (17), the fixed effects formulae continue to apply when one or more of the outcomes is missing. This is only valid under the missing-at-random assumption, that the chance of an outcome being unreported can be modeled as an independent chance mechanism conditional on the outcome estimates which actually are observed. Acknowledging this assumption can be crucially important in meta-analysis, where outcome reporting bias, for example when several outcomes are measured but only those showing a statistically significant effect are reported, is a common problem (Kirkham *et al.*, 2010), although the simulations in Kirkham *et al.* (2012) suggest that in some circumstances multivariate methods can be more robust than univariate methods to departures from this assumption. Subjective assessments of the risk of outcome reporting bias (Kirkham *et al.*, 2010) can lead to useful univariate bias corrections (Copas *et al.*, 2013), and an extension to multivariate models may also be possible.

The paper has also assumed that the V_is (or the V_{*i}^{-1}s) are known, so that borrowing of strength measures can be evaluated explicitly. Riley (2009) emphasizes the importance of taking the within-study correlations into account, and discusses the problem when, in practice, authors of research papers may only report estimates and standard errors for the outcomes taken one at a time. In that case only the diagonal elements of V_i are provided directly in study reports. If we can obtain full data (individual patient data) for such studies, then consistent estimates of the within-study correlations can be calculated, but in practice this may be difficult or impossible. Various approaches to dealing with this issue have been suggested, such as sensitivity analyses that explore a variety of different within-study correlations (Jackson *et al.*, 2011). The special case of equal within-study correlations (Section 3.3) can be a useful starting point. Wei and Higgins (2013) examine ways in which it may be possible to estimate these correlations retrospectively from other information that might be available. A partial approach is to note our finding that most borrowing of strength comes from studies whose designs are most atypical of the studies as a whole, and so by comparing the research methods used in the studies it may be possible to at least roughly identify which studies might be worth following up. Concentrating on trying to obtain further data for just some of these studies, and using the missing data formula (17) for other studies, may give at least some indication of whether including secondary outcomes in a multivariate model offers the potential to improve the estimate of the primary treatment effect.

In practice the matrix V_i is calculated from the data in the ith study, and so the assumption that the V_i's are known is ignoring the sampling error in these variance estimates. Table 1 shows that the example in Section 3.2 is based on large sample sizes, but with smaller samples the resulting inferences can underestimate uncertainty and be biased in cases where the estimated variances are correlated with the values of the ys. In univariate meta-analysis this bias is particularly noticeable in the Egger test for funnel plot symmetry (Egger *et al.*, 1997), as demonstrated in several simulation studies. Copas and Lozada-Can (2009) give a general method for calculating bias corrections for such test statistics. Berkey (1995) suggests a simpler way of eliminating bias, by smoothing the variance estimates across the studies. Assuming that study variances are inversely proportional to study sample size, and estimating the proportionality factor from the studies as a whole, essentially eliminates the correlation between the outcomes and their variances. Harbord *et al.* (2006) suggests a similar idea. However, for estimating efficiency and borrowing of strength as discussed in this paper, such considerations of bias are not directly relevant as E and $B(V_i^{-1})$ depend only on the V_is and not on the actual values of the y_is. If each estimated V_i is consistent then so will be the estimates of the derived borrowing of strength quantities. It is important to avoid any smoothing of the V_is so that they properly reflect the characteristics of each individual study.

Acknowlegements

DJ, IRW and RDR are supported by funding from Medical Research Council grant MR/J013595/1. DJ and IRW are also supported by funding from Medical Research Council grants U105260558 and (for IRW) UU12023/21. We acknowledge helpful comments from the referees of this paper.

References

Berkey, C. S., Hoaglin, D. C., Mosteller, F., and Colditz, G. A. (1995). A random-effects regression model for meta-analysis. *Statistics in Medicine*, **14**, 385-411.

Chen, H., Manning, A. K. and Dupuis, J. (2012). A method of moments estimate for random effects multivariate meta analysis. *Biometrics*, **78**, 1278-1284.

Copas, J. B., Dwan, K., Kirkham, J. J. and Williamson, P. R. (2013). A model-based correction for outcome reporting bias in meta-analysis. *Biostatistics*, **15**, 370-383.

Copas, J. B. and Lozada-Can, C. (2009). The radial plot in meta-analysis: approximations and applications. *Applied Statistics*, **58**, 329-344.

DerSimonian, R and Laird, N (1986). Meta analysis in clinical trials. *Controlled Clinical Trials*. **7**, 177-188.

Egger, M., Smith, G. D., Schneider, M. and Minder, C. (1997). Bias in meta analysis detected by a simple graphical test. *Br. Med. J.*, **315**, 629-634.

Fibrinogen Studies Collaboration (2009). Systematically missing confounders in individual participant data meta-analysis of observational cohort studies. *Statistics in Medicine*, **28**, 1218-1237.

Guolo, A. and Varin, C. (2017) Random effects meta-analysis: the number of studies does matter. *Statistical Methods in Medical Research*, **26**, 1500-1518.

Harbord, R. M., Egger, M. and Sterne, J. A. C. (2006) A modified test for small study effects in meta analysis of controlled trials with binary endpoints. *Statistics in Medicine*, **25**, 3443-3457.

Jackson, D., White, I. R. and Thompson, S. G. (2010). Extending DerSimonian and Laird's methodology to perform multivariate random effects meta-analysis. *Statistics in Medicine*, **29**, 1282-1297.

Jackson, D., Riley, R, and White, I. R. (2011). Multivariate meta-analysis potential and promise. *Statistics in Medicine*, **30**, 2481-2498.

Jackson, D., White, I. R. and Riley, R. D. (2013). A matrix-based method of moments for fitting the multivariate random effects model for meta-analysis and meta-regression. *Biometrical Journal*, **55**, 231-245.

Jackson, D., White, I. R., Price, M., Copas, J. B. and Riley, R. D. (2017). Borrowing of strength and study weights in multivariate and network meta-analysis. *Statistical Methods in Medical Research*, in press.

Kirkham, J. J., Dwan, K., Dodd, S.,Altmann, D. G., Smyth, R., Jacoby, A., Gamble, C. and Williamson, P. R. (2010). The impact of outcome reporting bias in randomized controlled trials on a cohort of systematic reviews. *British Med. J.*, **340**, c365.

Kirkham, J. J., Riley, R. D. and Williamson, P. R. (2012). A multivariate meta-analysis approach for reducing the impact of outcome reporting bias in systematic reviews. *Statistics in Medicine*, **31**, 2179-2195.

Ma Y. and Mazumdar, M. (2011). Multivariate meta-analysis: a robust approach based on the theory of U-statistics. *Statistics in Medicine*, **30**, 2911-2929.

Riley, R. D. (2009). Multivariate meta-analysis: the effect of ignoring within-study correlation. *J. Roy. Statist. Soc., A*, **172**, 789-811.

Riley, R.D., Price, M.J., Jackson, D., Wardle, M., Gueyffier, F., Wang, J., Staessen, J.A. and White, I.R. (2015). Multivariate meta-analysis using individual patient data. *Research Synthesis Methods*, **6**, 157-174.

Simel, D. L. and Bossuy, P. M. M. (2009). Differences between univariate and bivariate models for summarizing diagnostic accuracy may not be large. *J. Clinical Epidemiology*, **62**, 1292-1300.

Sohn, S. Y. (2000). Multivariate meta-analysis with potentially correlated marketing study results. *Naval Research Logistics*, **47**, 500-510.

Trikalinos, T. A., Hoaglin, D. C. and Schmid, C. H. (2014). An empirical assessment of univariate and multivariate meta-analysis for categorical outcomes. *Statistics in Medicine*, **33**, 1441-1459.

Wang, J. G., Staessen, J. A., Franklin, S. S., Fagard, R. and Gueyffier, F. (2005). Systolic and diastolic blood pressure lowering as determinants of cardiovascular outcome. *Hypertension*, **45**, 907-913.

Wei, Y. and Higgins, J. P. (2013). Estimating within-study covariances in multivariate meta-analysis with multiple outcomes. *Statistics in Medicine*, **32**, 1191-1205.

Study	1	2	3	4	5	6	7	8	9	10	total
Sample size	1530	349	172	4798	6991	2651	4736	268	2391	4695	28581
σ	0.41	0.36	0.45	0.17	0.17	0.14	0.14	1.08	0.30	0.17	
ν	1.08	0.41	0.59	0.26	0.33	0.17	0.14	0.91	0.20	0.17	
ρ	0.16	0.64	0.10	0.52	0.42	0.62	0.69	0.35	0.78	0.62	
$100T(V^{-1})\%$	2.5	3.3	2.5	14.4	14.3	21.6	21.4	0.4	5.3	14.3	100.0
$100B(V^{-1})\%$	0.02	0.02	0.34	0.15	0.10	0.23	0.04	0.08	0.53	0.04	1.55

Table 1: Sample sizes, values of (σ, ν, ρ), and percentage values of $T(V_i^{-1})$ and $B(V_i^{-1})$ for estimating β_1 in the example.

α	0	0.5	1	1.5	2
True efficiency E_{RE}	0.888	0.932	0.948	0.957	0.961
Median $\hat{E}_{RE}$	0.903	0.928	0.945	0.947	0.954
Median rc$(BoS_i, \widehat{BoS}_i)$	0.927	0.903	0.891	0.867	0.842

Table 2: Simulation of a random effects variant of the example for estimating β_2, comparing the estimated efficiency and BoS for the fitted random effects model with their corresponding true values. Increasing values of α indicate increasing heterogeneity.

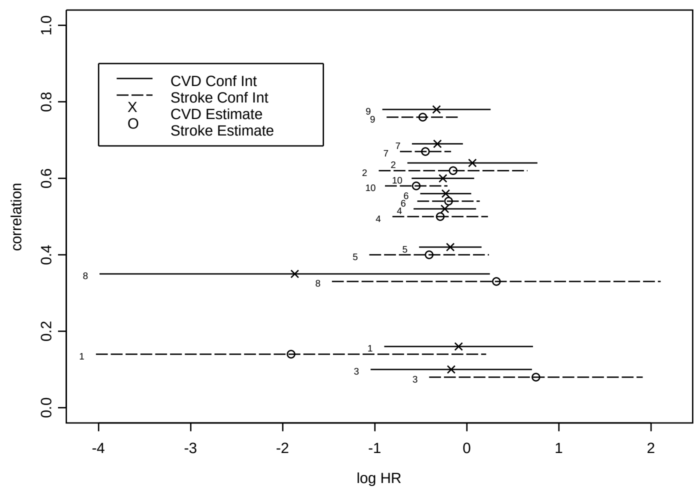

Figure 1: Graph illustrating the raw data for the example. The plotted points and horizontal line segments show the within-study estimates and 95% confidence intervals for the hazard ratios for CVD and for stroke within each of the ten trials, plotted against within-study correlation.

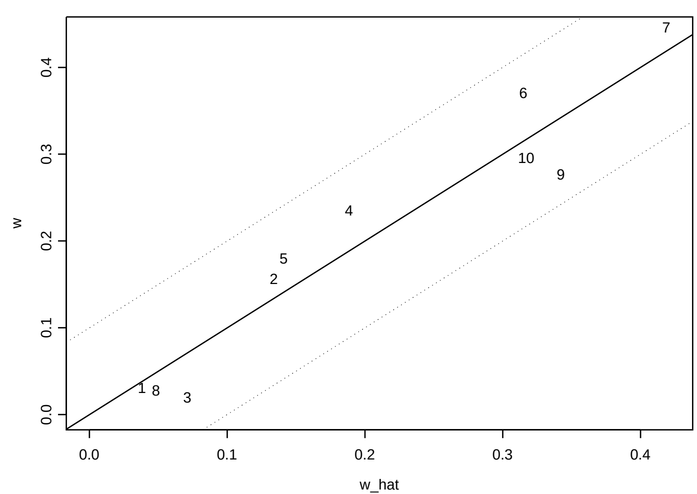

Figure 2: Borrowing of strength plot for estimating β_1 ($E = 0.984$)

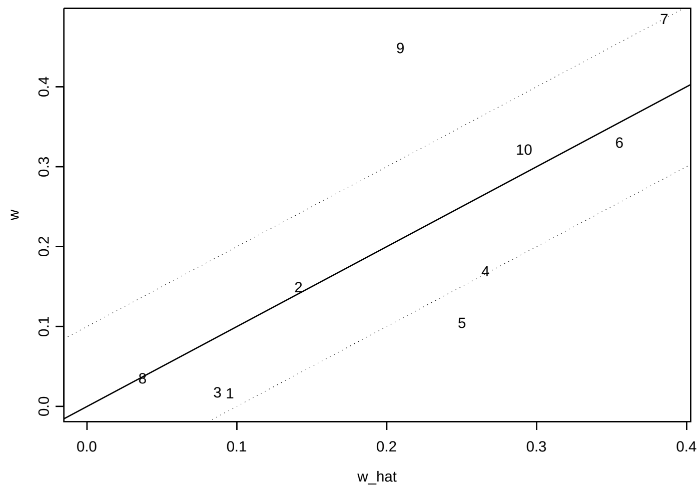

Figure 3: Borrowing of strength plot for estimating β_2 ($E = 0.888$)

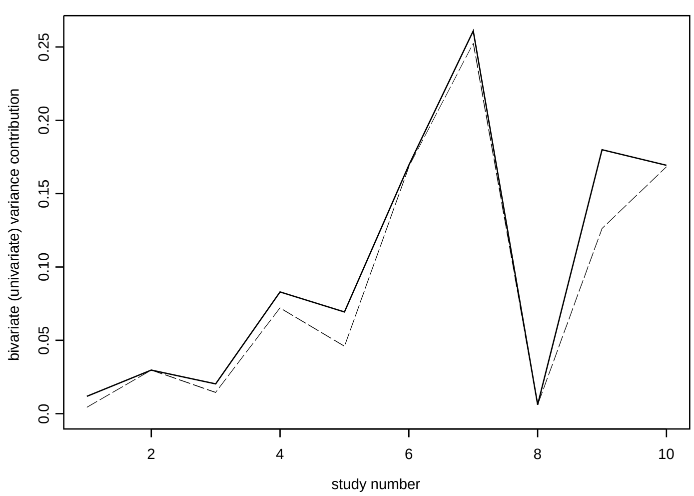

Figure 4: Study values of T (multivariate variance contribution, solid line) and $T - B$ (univariate variance contribution, dashed line) for estimating β_2 ($E = 0.888$)

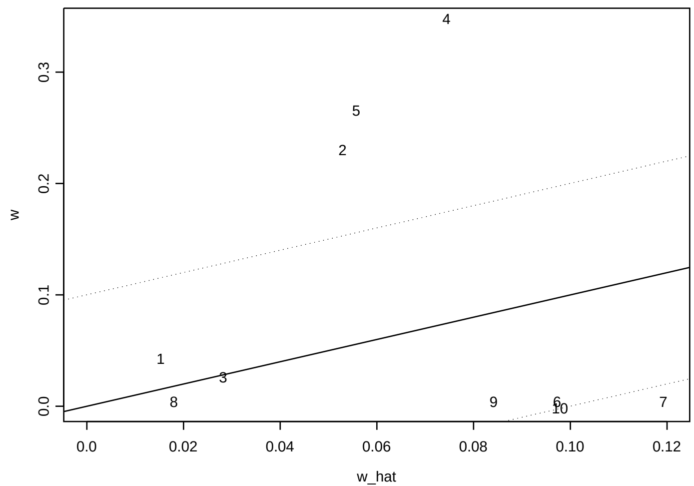

Figure 5: Borrowing of strength plot for estimating β_1 with missing outcomes ($E = 0.815$)

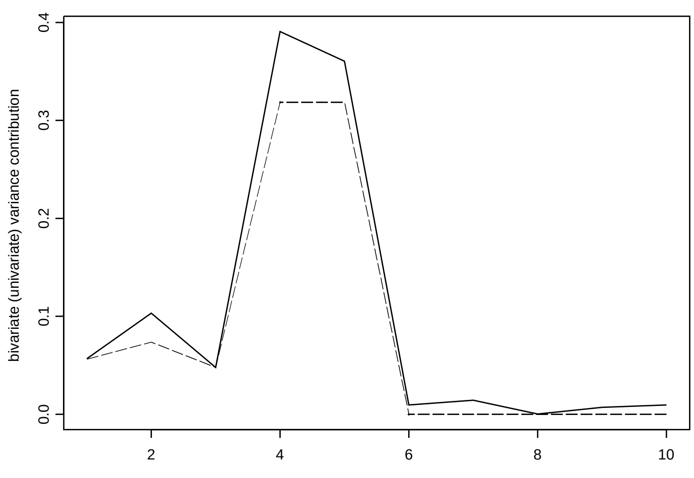

Figure 6: Study values of T (multivariate variance contribution, solid line) and $T - B$ (univariate variance contribution, dashed line) for estimating β_1 with missing outcomes ($E = 0.815$)

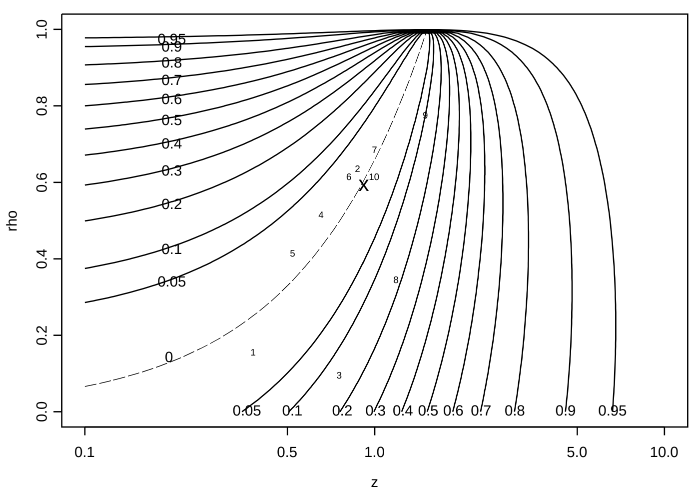

Figure 7: Contour plot of $R(V^{-1})$ for estimating β_1 ($E = 0.984$). The points labelled 1-10 show the values of (z, ρ) for the studies in the example. The point X indicates the corresponding point for matrix $\bar{V}$