The Stata Journal (2011)
11, Number 2, pp. 255–270

Multivariate random-effects meta-regression: Updates to mvmeta

Ian R. White
MRC Biostatistics Unit
Cambridge, UK
ian.white@mrc-bsu.cam.ac.uk

Abstract. An extension of `mvmeta`, my program for multivariate random-effects meta-analysis, is described. The extension handles meta-regression. Estimation methods available are restricted maximum likelihood, maximum likelihood, method of moments, and fixed effects. The program also allows a wider range of models (Riley's overall correlation model and structured between-studies covariance); better estimation (using Mata for speed and correctly allowing for missing data); and new postestimation facilities (I-squared, standard errors and confidence intervals for between-studies standard deviations and correlations, and identification of the best intervention). The program is illustrated using a multiple-treatments meta-analysis.

Keywords: st0156_1, mvmeta, meta-analysis, meta-regression, I-squared

1 Introduction

Stata software for meta-analysis is well advanced and has been described in a recent collection of articles (Sterne 2009). Most software is designed for univariate meta-analysis, in which each study contributes an estimate of a single quantity; but there has been recent interest in multivariate meta-analysis, in which some studies contribute estimates of more than one quantity: for example, intervention effects on different outcomes or differences in one outcome among three or more groups (van Houwelingen, Arends, and Stijnen 2003; Jackson, Riley, and White Forthcoming). I have previously described a Stata routine, `mvmeta` (White 2009), that fits the multivariate random-effects meta-analysis model using restricted maximum likelihood (REML), maximum likelihood (ML), or the method of moments (MM).

This article presents various extensions to `mvmeta`. Covariates are allowed so that a multivariate meta-regression is performed. For the case in which within-study correlations are unknown, Riley's overall correlation model can be fit (Riley, Thompson, and Abrams 2008). The between-studies covariance matrix may now be structured. The I-squared statistic, which measures the impact of heterogeneity on a meta-analysis (Higgins and Thompson 2002), has been extended to the multivariate case and implemented. Confidence intervals are available both for variance components and for I-squared. Finally, in the case of comparisons of multiple interventions, the probability that each is the best intervention can be estimated.

The model considered is

$$\begin{aligned} y_i &\sim N(\mu_i, S_i) \\ \mu_i &\sim N(\beta X_i, \Sigma) \end{aligned}$$

where y_i is a vector of estimates from the ith study, S_i is their variance–covariance matrix, μ_i is the study-specific mean vector, and X_i is a matrix of study-specific covariates. In this model, the data are y_i, S_i, and X_i, and we aim to estimate the regression coefficients β and the between-studies variance–covariance matrix Σ.

I describe the new `mvmeta` options in section 2 and give technical details in section 3. The command is illustrated in a multiple-treatments meta-analysis in section 4, and limitations and possible extensions are discussed in section 5.

2 mvmeta: Multivariate random-effects meta-regression

2.1 Syntax

`mvmeta` *b V xvars* [*if*] [*in*] [, *old_options new_options*]

where *old_options* are the options for `mvmeta` described in White (2009) and *new_options* are described below.

2.2 New model and estimation options

`wscorr(riley)` can be used when within-study correlations are unknown. It uses the alternative model of Riley, Thompson, and Abrams (2008) to estimate an overall correlation; see section 3.5. Riley (2009) discusses other ways to handle unknown within-study correlations.

`bscovariance(`*string*`)` specifies the between-studies covariance structure; see section 2.4.

`equations(`*yvar1*`:`*xvars1* [`,` *yvar2*`:`*xvars2* [`, ...`]]`)` allows different outcomes to have different regression models. For example, for two-dimensional *b*, `mvmeta b V x` is the same as `mvmeta b V, eq(b1:x,b2:x)`.

`noconstant` suppresses the constant in meta-regression.

`longparm` estimates the results as one regression model for each outcome. Without covariates, this is usually less convenient than the default (in which all outcomes form a single regression model) but is required if the `pbest()` option will be used. With covariates, `longparm` is the default and cannot be changed.

Other new estimation options, which are described in the help file, are `noposdef`, `psdcrit(#)`, *maximize_options*, `augment`, and `augquiet`.

2.3 New output options

For regression parameters

`dof(`*expression*`)` specifies the degrees of freedom for t tests and confidence intervals on the regression parameters. The expression may include `n`, the number of observations. The default is to use a normal distribution.

`pbest(min|max` [*if*] [*in*]`,` [`reps(`*#1*`) zero gen(`*string*`) seed(`*#2*`) format(%`*fmt*`) id(`*varlist*`)`]`)` requests estimation of the probability that each linear predictor is the best—that is, the maximum or minimum, depending on the first argument of `pbest()`. The probability is estimated under a Bayesian model with flat priors, assuming that the posterior distribution of the parameter estimates is approximated by a normal distribution with mean and variance equal to the frequentist estimates and variance–covariance matrix. Rankings are constructed by drawing the coefficients *#1* times (default is 100) from their approximate posterior density. For each draw, the linear predictor is evaluated for each study, and the largest linear predictor is noted. The `zero` option specifies that 0 be considered another linear predictor; its use is illustrated in the example in section 4. `gen()` specifies that the probabilities be saved in variables with the prefix *string*. `seed()` specifies the random-number seed, `format()` specifies the output format, and `id()` specifies identifiers for the output. Although the default behaviour is to rank linear predictors, the `predict` option ranks the true effects in a future study with the same covariates, thus allowing for heterogeneity as well as parameter uncertainty, as in the calculation of prediction intervals (Higgins, Thompson, and Spiegelhalter 2009). For models without covariates, `pbest()` is only available if `longparm` was specified when the model was fit.

For between-studies variance parameters

`i2` reports the between-studies variance τ^2_j and the I-squared statistic (Higgins and Thompson 2002) for each outcome, together with confidence intervals. See section 3.6 for details.

`i2fmt(%`*fmt*`)` specifies an output format for the I-squared statistics.

`nc`chi2 uses the option of `heterogi` in computing confidence intervals for τ^2_j and I-squared. It is only relevant after MM estimation. See section 3.6 for details.

`ciscale(sd|logsd|logh)` determines the scale on which confidence intervals for τ^2_j and I-squared are computed. The default is `ciscale(sd)`. See section 3.6 for details.

`test`sigma is only allowed after ML or REML estimation. It performs a likelihood-ratio test or restricted likelihood-ratio test of $\Sigma = 0$. The latter is valid because the models compared have the same fixed part (Verbeke and Molenberghs 2000).

2.4 Covariance structures

`bscovariance(unstructured)` estimates an unstructured Σ and is the default. Starting values for Σ may be specified explicitly by `start(`*matrix_expression*`)`. `start(mm)` (the default) specifies that the starting value be computed by the `mm` method. `start(0)` uses a starting value of 0.001 times the default, because a starting value of 0 leads to nonconvergence (White 2009). The starting value for β is derived from Σ using (1) below.

`bscovariance(proportional` *matexp*`)` models $\Sigma = \tau^2\Sigma_0$, where τ is an unknown parameter and Σ_0 is a known matrix expression (for example, a matrix name or `I(2)`). `start(#)` then specifies the starting value for the scalar τ.

`bscovariance(equals` *matexp*`)` forces $\Sigma = \Sigma_0$, where Σ_0 is a known matrix expression (for example, a matrix name or `I(2)`).

`bscovariance(correlation` *matexp*`)` models $\Sigma = \mathbf{D} \times matexp \times \mathbf{D}$, where *matexp* is a known matrix expression containing the between-study correlations and $\mathbf{D}$ is an unknown diagonal matrix containing the between-studies standard deviations. `start(`*rowvector*`)` specifies the starting values for the diagonal of $\mathbf{D}$.

2.5 Other changes in version 2

The `showchol` option has been renamed `showall`, the `corr()` option has been renamed `wscorr()`, and the `bscorr` and `bscov` options have been renamed `print(bscorr)` and `print(bscov)`, respectively. `mvmeta` typed without specifying b and V redisplays the latest estimation results, and output options (including `showall`, `eform`, `nouncertainv`, `print()`, `level()`, `dof`, `i2`, and `pbest()`) may be used.

3 Details

3.1 Notation

The data for `mvmeta` for the rth outcome from the ith study ($i = 1, \ldots, n$) are the point estimate y_{ir} (a scalar) and the covariates x_{ir} (a $q_r \times 1$ vector). For standard meta-analysis, $x_{ir} = (1)$, a vector of ones. The mean of y_{ir} is assumed to be $\beta_r x_{ir}$, where β_r is a $1 \times q_r$ (row)vector. Thus we have the marginal models

$$\begin{aligned} y_{ir} &\sim N(\mu_{ir}, s_{ir}^2) \\ \mu_{ir} &\sim N(\beta_r x_{ir}, \tau_r^2) \end{aligned}$$

In matrix notation, we write the (row)vector outcome $y_i = (y_{i1}, y_{i2}, \ldots, y_{ip})$. We also know the within-study variance–covariance matrix S_i (a $p \times p$ matrix). We assume the following joint model:

$$
\begin{aligned}
y_i &\sim N(\mu_i, S_i) \\
\mu_i &\sim N(\beta X_i, \Sigma) \\
S_i &= \begin{pmatrix} s_{i1}^2 & \rho_{i12}s_{i1}s_{i2} & \cdots & \rho_{i1p}s_{i1}s_{ip} \\ \rho_{i12}s_{i1}s_{i2} & s_{i2}^2 & \cdots & \rho_{i2p}s_{i2}s_{ip} \\ \vdots & \vdots & \ddots & \vdots \\ \rho_{i1p}s_{i1}s_{ip} & \rho_{i2p}s_{i2}s_{ip} & \cdots & s_{ip}^2 \end{pmatrix} \\
\beta &= (\beta_1, \beta_2, \ldots, \beta_p) \\
X_i &= \begin{pmatrix} x_{i1} & 0 & \cdots & 0 \\ 0 & x_{i2} & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & x_{ip} \end{pmatrix} \\
\Sigma &= \begin{pmatrix} \tau_1^2 & \kappa_{12}\tau_1\tau_2 & \cdots & \kappa_{1p}\tau_1\tau_p \\ \kappa_{12}\tau_1\tau_2 & \tau_2^2 & \cdots & \kappa_{2p}\tau_2\tau_p \\ \vdots & \vdots & \ddots & \vdots \\ \kappa_{1p}\tau_1\tau_p & \kappa_{2p}\tau_2\tau_p & \cdots & \tau_p^2 \end{pmatrix}
\end{aligned}
$$

where y_i and μ_i are $1 \times p$, S_i and Σ are $p \times p$, β is $1 \times q_+$, and X_i is $q_+ \times p$. Σ may be constrained as explained in section 2.4. Our aim is to estimate β and Σ.

3.2 Estimating β, knowing Σ

We set $W_i = (\Sigma + S_i)^{-1}$. Then

$$\hat{\beta} = \left(\sum_i y_i W_i X_i'\right)\left(\sum_i X_i W_i X_i'\right)^{-1} \tag{1}$$

3.3 Estimating Σ: likelihood-based methods

We still use the notation $W_i = (\Sigma + S_i)^{-1}$, noting that this now depends on the unknown Σ. The log likelihood and restricted log likelihood, respectively, are

$$-2L = \sum_i \{\log|\Sigma + S_i| + (y_i - X_i\beta)W_i(y_i - X_i\beta)' + p_i \log 2\pi\} \tag{2}$$

$$-2RL = -2L + \log\left|\sum_i X_i W_i X_i'\right| - q_+ \log 2\pi \tag{3}$$

where p_i is the number of observed outcomes in y_i. Where a study reports only a subset of outcomes, y_i and S_i are of reduced dimension, so Σ in (2) and (3) is replaced by its corresponding submatrix. This makes unnecessary the augmentation procedures in the

previous version (White 2009), but they can still be implemented using the `augment` option.

The (restricted) log likelihood is maximized by a Newton–Raphson algorithm using Stata's `ml` procedure. The code has been speeded up by computing the log likelihood using Mata. For unstructured Σ, the basic model parameters are taken as the elements of a Cholesky decomposition of Σ, ensuring that Σ is nonnegative definite (Riley et al. 2007). For the model $\Sigma = \tau^2\Sigma_0$, the basic parameter is τ.

3.4 Estimating Σ: method of moments

Jackson, White, and Thompson (2010) define a matrix generalization of the univariate Q statistic of DerSimonian and Laird (1986). With unstructured Σ, this satisfies $E(Q_{rs}) = A_{rs} + B_{rs}\Sigma_{rs}$ for $r, s = 1, \ldots, p$, where A and B are matrices that can be computed from the observed data. Estimation of Σ is therefore straightforward and fast. The MM has not yet been developed for meta-regression with structured Σ or for the overall correlation model described in section 3.5 below.

3.5 Unknown within-study correlations

When within-study correlations are unknown, various options are available, including sensitivity analysis over alternative values (Riley 2009). Alternatively, Riley, Thompson, and Abrams (2008) proposed an "overall correlation model" that does not involve the within-study correlations. Let $\text{var}(y_i) = V_i$; the standard model of section 3.1 has $V_i = S_i + \Sigma$. The alternative model has the same diagonal elements, $V_{irr} = S_{irr} + \Sigma_{rr}$, but off-diagonal elements $V_{irs} = \rho^O_{rs}\sqrt{V_{irr}V_{iss}}$ for $r \neq s$. Here ρ^O_{rs} represents an overall correlation between outcomes r and s.

3.6 I-squared

I-squared measures the impact of heterogeneity on the meta-analysis (Higgins and Thompson 2002). In univariate meta-analysis, I-squared is computed as the ratio of a "between" variance (the appropriate element of Σ) to the sum of the "between" variance and a "within" variance given by (9) of Higgins and Thompson (2002). To generalize this, I propose computing I-squared separately for each outcome and handling covariates by defining I-squared for the rth outcome as

$$I^2_r = \frac{\tau^2_r}{A_{rr}/B_{rr} + \tau^2_r} \tag{4}$$

where A_{rr}/B_{rr} is a "typical" squared standard error, and A_{rr} and B_{rr} are as defined in section 3.4. If there are no covariates, then (4) corresponds exactly to the definition of Higgins and Thompson (2002). If, further, τ^2 is estimated by MM, then (4) gives the standard quantity $I^2_r = \max[0, \{Q_{rr} - (n_r - 1)/Q_{rr}\}]$ (for example, as output by `metan`), where n_r is the number of studies reporting outcome y_r. However, (4) applies

equally well if τ^2 is estimated by REML or ML. This definition of I-squared does not account for "borrowing strength" between outcomes.

When estimation uses the MM, confidence intervals for I_r^2 are computed on the scale of $\log(H_r)$, where $H_r^2 = (1 - I_r^2)^{-1} = Q_{rr}/(n_r - 1)$, as suggested by Higgins and Thompson (2002; they also called them "uncertainty intervals") and implemented in Stata by `heterogi`. The noncentral chi-squared option of `heterogi` is available through the `ncchi2` option. Confidence intervals for τ_r^2 are derived from (4). Exact methods (Biggerstaff and Jackson 2008) are computationally intensive and have not been implemented in Stata.

When estimation uses REML or ML, confidence intervals are first estimated for τ_r^2 using the estimated standard errors. The confidence interval may be computed on the scale of τ (the default), $\log(\tau)$, or $\log(H_r)$. Confidence intervals for I_r^2 are then derived using (4). With unstructured Σ, confidence intervals for the between-studies correlations κ_{rs} are also available; they are computed on the scale of $\log\{(1 + \kappa_{rs})/(1 - \kappa_{rs})\}$. All standard errors are computed using Stata's `nlcom` command. When one or more basic variance parameters is estimated as zero, the corresponding zero term is dropped from the expression for τ_r^2 to avoid causing `nlcom` to fail. This fix can be checked by changing the order of the variables (using `mvmeta`'s `vars()` option), which often avoids causing `nlcom` to fail. In my experience, the two methods always give the same confidence intervals.

4 Example

4.1 Data

We use data from a multiple-treatments meta-analysis comparing four interventions to promote smoking cessation. These data have been previously presented and analyzed by Lu and Ades (2004). The interventions are coded A, B, C, and D, and the data for each trial arm are summarized as the number of individuals and the number who quit smoking. The original dataset is

```
. use smoking_raw, clear
(Smoking data from Lu & Ades (2006))
. list, noo clean
```

study	design	dA	nA	dB	nB	dC	nC	dD	nD
1	ACD	9	140	.	.	23	140	10	138
2	BCD	.	.	11	78	12	85	29	170
3	AB	79	702	77	694	.	.	.	.
4	AB	18	671	21	535	.	.	.	.
5	AB	8	116	19	146	.	.	.	.
6	AC	75	731	.	.	363	714	.	.
7	AC	2	106	.	.	9	205	.	.
8	AC	58	549	.	.	237	1561	.	.
9	AC	0	33	.	.	9	48	.	.
10	AC	3	100	.	.	31	98	.	.
11	AC	1	31	.	.	26	95	.	.
12	AC	6	39	.	.	17	77	.	.
13	AC	95	1107	.	.	134	1031	.	.
14	AC	15	187	.	.	35	504	.	.
15	AC	78	584	.	.	73	675	.	.
16	AC	69	1177	.	.	54	888	.	.
17	AC	64	642	.	.	107	761	.	.
18	AC	5	62	.	.	8	90	.	.
19	AC	20	234	.	.	34	237	.	.
20	AD	0	20	.	.	.	.	9	20
21	BC	.	.	20	49	16	43	.	.
22	BD	.	.	7	66	.	.	32	127
23	CD	.	.	.	.	12	76	20	74
24	CD	.	.	.	.	9	55	3	26

The first stage of analysis constructs a dataset of estimated intervention effects and their variance–covariance matrices. We choose A as the reference category. Trials without an arm A (trials 2 and 21–24) are augmented with an arm A with 0.01 individuals and 0.001 successes. Trials containing zero cells (trials 9 and 20) have 1 individual with 0.5 successes added to each arm. This leads to the following augmented dataset:

```
. use smoking_aug
(Smoking data from Lu & Ades (2006))
. list, noo clean
```

study	design	dA	nA	dB	nB	dC	nC	dD	nD
1	ACD	9	140	.	.	23	140	10	138
2	BCD	.001	.01	11	78	12	85	29	170
3	AB	79	702	77	694	.	.	.	.
4	AB	18	671	21	535	.	.	.	.
5	AB	8	116	19	146	.	.	.	.
6	AC	75	731	.	.	363	714	.	.
7	AC	2	106	.	.	9	205	.	.
8	AC	58	549	.	.	237	1561	.	.
9	AC	.5	34	.	.	9.5	49	.	.
10	AC	3	100	.	.	31	98	.	.
11	AC	1	31	.	.	26	95	.	.
12	AC	6	39	.	.	17	77	.	.
13	AC	95	1107	.	.	134	1031	.	.
14	AC	15	187	.	.	35	504	.	.
15	AC	78	584	.	.	73	675	.	.
16	AC	69	1177	.	.	54	888	.	.
17	AC	64	642	.	.	107	761	.	.
18	AC	5	62	.	.	8	90	.	.
19	AC	20	234	.	.	34	237	.	.
20	AD	.5	21	.	.	.	.	9.5	21
21	BC	.001	.01	20	49	16	43	.	.
22	BD	.001	.01	7	66	.	.	32	127
23	CD	.001	.01	.	.	12	76	20	74
24	CD	.001	.01	.	.	9	55	3	26

We now compute the log odds-ratios for arms `B`, `C`, and `D` relative to arm `A`, as well as the variance–covariance matrix of these three estimates. We could use `mvmeta_make` (White 2009), but it is easy to run the loop

```
foreach trt in A B C D {
      if "`trt´"=="A" continue
      gen y`trt´ = log(d`trt´/(n`trt´-d`trt´)) - log(dA/(nA-dA))
      gen S`trt´`trt´ = 1/d`trt´ + 1/(n`trt´-d`trt´) + 1/dA + 1/(nA-dA)
      foreach trt2 in A B C D {
          if "`trt2´"=="A" continue
          if "`trt2´">"`trt´" gen S`trt´`trt2´ = 1/dA + 1/(nA-dA) ///
              if !mi(d`trt´) & !mi(d`trt2´)
      }
}
format y* S* %6.2g
```

which yields the following data:

```
. list study design y* S*, noo clean
```

study	design	yB	yC	yD	SBB	SBC	SBD	SCC	SCD	SDD
1	ACD	.	1.1	.13	.	.	.	.17	.12	.23
2	BCD	.39	.39	.62	1111	1111	1111	1111	1111	1111
3	AB	-.016	.	.	.029	.	.	.	.	.
4	AB	.39	.	.	.11	.	.	.	.	.
5	AB	.7	.	.	.19	.	.	.	.	.
6	AC	.	2.2	.	.	.	.	.02	.	.
7	AC	.	.87	.	.	.	.	.63	.	.
8	AC	.	.42	.	.	.	.	.024	.	.
9	AC	.	2.8	.	.	.	.	2.2	.	.
10	AC	.	2.7	.	.	.	.	.39	.	.
11	AC	.	2.4	.	.	.	.	1.1	.	.
12	AC	.	.44	.	.	.	.	.27	.	.
13	AC	.	.46	.	.	.	.	.02	.	.
14	AC	.	-.16	.	.	.	.	.1	.	.
15	AC	.	-.24	.	.	.	.	.03	.	.
16	AC	.	.039	.	.	.	.	.035	.	.
17	AC	.	.39	.	.	.	.	.028	.	.
18	AC	.	.11	.	.	.	.	.35	.	.
19	AC	.	.58	.	.	.	.	.089	.	.
20	AD	.	.	3.5	.	.	.	.	.	2.2
21	BC	1.8	1.7	.	1111	1111	.	1111	.	.
22	BD	.066	.	1.1	1111	.	1111	.	.	1111
23	CD	.	.52	1.2	.	.	.	1111	1111	1111
24	CD	.	.57	.16	.	.	.	1111	1111	1111

The first stage of analysis is now complete. In the second stage of analysis, we use `mvmeta` to model the intervention effects across studies, using consistency and inconsistency models.

4.2 Consistency model

A consistency model (Lu and Ades 2004) allows the intervention effects to be heterogeneous between studies but assumes that there are no systematic differences between designs. It is easy to fit:

```
. mvmeta y S
Note: using method reml
Regressing yB on
Regressing yC on
Regressing yD on
Note: using variables yB yC yD
Note: 24 observations on 3 variables
Variance-covariance matrix: unstructured
```

(*output omitted*)

```
Multivariate meta-analysis
Variance-covariance matrix = unstructured
Method = reml                                  Number of dimensions    =     3
Restricted log likelihood = -53.826928         Number of observations  =    24
```

	Coef.	Std. Err.	z	P>|z|	[95% Conf.	Interval]
Overall_mean						
yB	.3326048	.3048747	1.09	0.275	-.2649385	.9301482
yC	.6810167	.218959	3.11	0.002	.2518649	1.110168
yD	.8357459	.3664475	2.28	0.023	.117522	1.55397

```
Estimated between-studies SDs and correlation matrix:
           SD         yB         yC         yD
yB  .31410047          1          .          .
yC   .7497773  .9362371          1          .
yD  .72247338  .85588029  .61958804          1
```

We see that, under the consistency assumption, interventions C and D are significantly superior to A, and D appears to be the best. We could perform significance tests between B, C, and D using `lincom`. The heterogeneity (between-studies variation) is larger for C versus A and D versus A than for B versus A.

It is often of interest to find the best intervention. We can do this using

```
. mvmeta y S, longparm pbest(max in 1, zero reps(1000) seed(478))
```

(*output omitted*)

```
Estimated probabilities (%) of being the maximum
(allows for parameter uncertainty):
```

	zero	yB	yC	yD
1.	0.0	3.1	31.9	65.0

Let the intervention effects be μ_B, μ_C, and μ_D, all representing contrasts from the reference intervention A. Positive values indicate better interventions in this dataset, so if μ_B, μ_C, and μ_D are all negative, then A is best; otherwise, the intervention with the largest μ is best. Thus we want to find the largest member of the set $\{0, \mu_B, \mu_C, \mu_D\}$, which is coded using `max` to find the largest and `zero` to include 0 in the set. We specify `in 1` to output results for the first study only; because there are no covariates, the results for all studies are the same.

In the output, the columns headed `zero`, `yB`, `yC`, and `yD` each indicate the posterior probability that intervention A, B, C, or D is the best, respectively. The best intervention is probably D and is very likely to be either C or D.

4.3 Estimating I^2

We can estimate the contribution of between-studies heterogeneity to the meta-analyses:

```
. mvmeta, i2
  (output omitted)
Approximate confidence intervals for between-studies SDs and I^2:
```

Variable	SD	[95% Conf.	Interval]	I^2	[95% Conf.	Interval]
yB	.31410107	0	.90877776	31	0	79
yC	.74977167	.39813052	1.1014128	88	68	94
yD	.72246304	0	1.8312487	8	0	35

```
Note: I^2 computed from estimated between-studies and typical within-studies
> variances
Note: CI computed on SD scale
Note: one or more CIs for I^2 were computed by dropping zero terms
Between-study correlations:
```

Variables	Correl.	[95% Conf.	Interval]
yB & yC	.93624729	-.99999763	1
yB & yD	.85586496	-.99994596	.99999967
yC & yD	.61958754	-.83474606	.99011097

```
Note: CI computed on log((1+corr)/(1-corr)) scale
```

The main contribution of between-studies heterogeneity appears to arise from the A–C contrast. Note that the between-studies correlations are very poorly estimated. In fact, the unstructured $\mathbf{\Sigma}$ matrix is barely identified in this problem. We next consider a structured $\mathbf{\Sigma}$ matrix.

4.4 Structured Σ

In sparser problems, it may be useful to assume that the heterogeneity variance is the same for each intervention contrast. This can be done by forcing Σ to be proportional to the matrix `P` defined below (Salanti et al. 2008):

```
. mat P = I(3) + J(3,3,1)
. mat l P
symmetric P[3,3]
    c1  c2  c3
r1   2
r2   1   2
r3   1   1   2
```

```
. mvmeta y S, bscov(prop P)
Note: using method reml
Regressing yB on
Regressing yC on
Regressing yD on
Note: using variables yB yC yD
Note: 24 observations on 3 variables
Variance-covariance matrix: proportional to P
```

(*output omitted*)

```
Multivariate meta-analysis
Variance-covariance matrix = proportional P
Method = reml                                   Number of dimensions    =     3
Restricted log likelihood = -54.946189          Number of observations  =    24
```

	Coef.	Std. Err.	z	P>|z|	[95% Conf.	Interval]
Overall_mean						
yB	.3984951	.3310639	1.20	0.229	-.2503782	1.047368
yC	.7023595	.1990896	3.53	0.000	.312151	1.092568
yD	.8658847	.3762281	2.30	0.021	.1284912	1.603278

```
Estimated between-studies SDs and correlation matrix:
          SD        yB        yC        yD
yB  .6744175         1         .         .
yC  .6744175        .5         1         .
yD  .6744175        .5        .5         1
```

4.5 Inconsistency model

An inconsistency model allows intervention effects to differ between designs (to a greater extent than can be explained by the heterogeneity). It therefore requires a multivariate meta-regression, with particular dummy variables for design as covariates. There are many ways to parameterize this model: we choose the two-arm designs involving A as basic contrasts, and we introduce one extra effect for each two-arm design that does not include A and two extra effects for each three-arm design.

```
. tab design, gen(des)
```

design	Freq.	Percent	Cum.
ACD	1	4.17	4.17
BCD	1	4.17	8.33
AB	3	12.50	20.83
AC	14	58.33	79.17
AD	1	4.17	83.33
BC	1	4.17	87.50
BD	1	4.17	91.67
CD	2	8.33	100.00
Total	24	100.00	

```
. mvmeta y S, bscov(prop P) eq(yC: des1 des2 des6, yD: des1 des2 des7 des8)
Note: using method reml
Regressing yB on
Regressing yC on des1 des2 des6
Regressing yD on des1 des2 des7 des8
Note: using variables yB yC yD
Note: 24 observations on 3 variables
Variance-covariance matrix: proportional to P
```

(*output omitted*)

```
Multivariate meta-analysis
Variance-covariance matrix = proportional P
Method = reml                                   Number of dimensions    =     3
Restricted log likelihood = -45.783933          Number of observations  =    24
```

	Coef.	Std. Err.	z	P>|z|	[95% Conf.	Interval]
yB						
_cons	.3303086	.4673829	0.71	0.480	-.5857451	1.246362
yC						
des1	.3468573	.882037	0.39	0.694	-1.381903	2.075618
des2	-.3728619	1.013567	-0.37	0.713	-2.359417	1.613693
des6	-.5253268	1.004197	-0.52	0.601	-2.493516	1.442862
_cons	.7044357	.2347562	3.00	0.003	.2443219	1.164549
yD						
des1	-3.393989	1.889914	-1.80	0.073	-7.098153	.3101744
des2	-2.966854	1.926324	-1.54	0.124	-6.742379	.8086707
des7	-2.148826	1.940325	-1.11	0.268	-5.951792	1.654141
des8	-2.576181	1.80985	-1.42	0.155	-6.123422	.9710605
_cons	3.522517	1.67126	2.11	0.035	.2469077	6.798126

```
Estimated between-studies SDs and correlation matrix:
          SD        yB        yC        yD
yB  .7430402         1         .         .
yC  .7430402        .5         1         .
yD  .7430402        .5        .5         1
```

We can now test for inconsistency by jointly testing the seven inconsistency parameters:

```
. test ([yC]: des1 des2 des6) ([yD]: des1 des2 des7 des8)
 ( 1)  [yC]des1 = 0
 ( 2)  [yC]des2 = 0
 ( 3)  [yC]des6 = 0
 ( 4)  [yD]des1 = 0
 ( 5)  [yD]des2 = 0
 ( 6)  [yD]des7 = 0
 ( 7)  [yD]des8 = 0
           chi2(  7) =    5.11
         Prob > chi2 =    0.6464
```

There is no evidence of inconsistency here. It is not valid to test consistency by comparing restricted likelihoods between models, because the models' fixed parts differ—but we could instead reestimate the models by maximum likelihood and perform a likelihood-ratio test.

5 Difficulties and limitations

`mvmeta` implements a two-stage meta-analysis procedure. This is common practice, but it does involve a quadratic approximation to the log likelihood, which may perform poorly with sparse data. One-stage procedures are possible with individual participant data (Smith, Williamson, and Marson 2005). They are implemented for Stata by `metandi` for diagnostic test data (Harbord and Whiting 2009), but they are not implemented more generally.

The MM is a fast alternative to REML, but further research is required to extend it to new situations, including structured Σ and Riley's overall correlation model.

6 Acknowledgments

I was supported by MRC grant U1052.00.006. I would like to thank Stephen Kaptoge, Dan Jackson, Richard Riley, Julian Higgins, Jessica Barrett, and Antonio Gasparrini for their help and encouragement.

7 References

Biggerstaff, B. J., and D. Jackson. 2008. The exact distribution of Cochran's heterogeneity statistic in one-way random effects meta-analysis. *Statistics in Medicine* 27: 6093–6110.

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

Harbord, R. M., and P. Whiting. 2009. metandi: Meta-analysis of diagnostic accuracy using hierarchical logistic regression. *Stata Journal* 9: 211–229.

Higgins, J. P. T., and S. G. Thompson. 2002. Quantifying heterogeneity in a meta-analysis. *Statistics in Medicine* 21: 1539–1558.

Higgins, J. P. T., S. G. Thompson, and D. J. Spiegelhalter. 2009. A re-evaluation of random-effects meta-analysis. *Journal of the Royal Statistical Society, Series A* 172: 137–159.

Jackson, D., R. Riley, and I. R. White. Forthcoming. Multivariate meta-analysis: Potential and promise. *Statistics in Medicine*.

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

Lu, G., and A. E. Ades. 2004. Combination of direct and indirect evidence in mixed treatment comparisons. *Statistics in Medicine* 23: 3105–3124.

Riley, R. D. 2009. Multivariate meta-analysis: The effect of ignoring within-study correlation. *Journal of the Royal Statistical Society, Series A* 172: 789–811.

Riley, R. D., K. R. Abrams, A. J. Sutton, P. C. Lambert, and J. R. Thompson. 2007. Bivariate random-effects meta-analysis and the estimation of between-study correlation. *BMC Medical Research Methodology* 7: 3.

Riley, R. D., J. R. Thompson, and K. R. Abrams. 2008. An alternative model for bivariate random-effects meta-analysis when the within-study correlations are unknown. *Biostatistics* 9: 172–186.

Salanti, G., J. P. Higgins, A. Ades, and J. P. Ioannidis. 2008. Evaluation of networks of randomized trials. *Statistical Methods in Medical Research* 17: 279–301.

Smith, C. T., P. R. Williamson, and A. G. Marson. 2005. Investigating heterogeneity in an individual patient data meta-analysis of time to event outcomes. *Statistics in Medicine* 24: 1307–1319.

Sterne, J. A. C., ed. 2009. *Meta-Analysis in Stata: An Updated Collection from the Stata Journal*. College Station, TX: Stata Press.

van Houwelingen, H. C., L. R. Arends, and T. Stijnen. 2003. Advanced methods in meta-analysis: multivariate approach and meta-regression. *Statistics in Medicine* 21: 589–624.

Verbeke, G., and G. Molenberghs. 2000. *Linear Mixed Models for Longitudinal Data*. New York: Springer.

White, I. R. 2009. Multivariate random-effects meta-analysis. *Stata Journal* 9: 40–56.

About the author

Ian White is a program leader at the MRC Biostatistics Unit in Cambridge, UK. His research interests focus on handling missing data, noncompliance and measurement error in the analysis of clinical trials, observational studies, and meta-analysis. He was originally motivated to write software for multivariate meta-analysis by his work with the Emerging Risk Factors Collaboration (http://ceu.phpc.cam.ac.uk/research/erfc).