Meta-Analysis with R

Book · November 2015

DOI: 10.1007/978-3-319-21416-0

CITATIONS
76

READS
9,075

3 authors, including:

Guido Schwarzer
Faculty of Medicine and Medical Center - University of Freiburg
184 PUBLICATIONS **8,481** CITATIONS
SEE PROFILE

Gerta Rücker
University Medical Center Freiburg
229 PUBLICATIONS **8,331** CITATIONS
SEE PROFILE

Guido Schwarzer • James R. Carpenter •
Gerta Rücker

Meta-Analysis with R

Springer

Guido Schwarzer
Institute for Medical Biometry
and Statistics
Medical Center – University of Freiburg
Freiburg, Germany

James R. Carpenter
MRC Clinical Trials Unit, London
and London School of Hygiene
and Tropical Medicine
London, United Kingdom

Gerta Rücker
Institute for Medical Biometry and Statistics
Medical Center – University of Freiburg
Freiburg, Germany

ISSN 2197-5736 ISSN 2197-5744 (electronic)
Use R!
ISBN 978-3-319-21415-3 ISBN 978-3-319-21416-0 (eBook)
DOI 10.1007/978-3-319-21416-0

Library of Congress Control Number: 2015949262

Springer Cham Heidelberg New York Dordrecht London

Printed on acid-free paper

Springer International Publishing AG Switzerland is part of Springer Science+Business Media (www.springer.com)

Preface

Meta-analysis plays a key role in evidence synthesis in many research disciplines, not least the social sciences, medicine and economics. The aim of this book is to equip those involved in such work (who are often not trained statisticians) to use R for meta-analysis, and thus promote both the use of R and the latest statistical methods in this area.

The attractions of R in this context (besides its free availability from http://www.r-project.org/) are its fast yet powerful and flexible graphics and its well-established algorithmic base.

The book assumes no prior knowledge of R, and takes readers through every step of the way from installing R, loading data from other packages, performing and interpreting the analyses. Parts I and II cover the essentials, while Part III considers more advanced topics, which remain the subject of active research.

Throughout, the ideas are illustrated with examples, and all the codes necessary to repeat these examples (including creating all the plots in the book) are either in the text itself or the web-appendix http://meta-analysis-with-r.org/. In selecting the code to include in the main text, we have assumed readers are relatively new to R. More experienced users can easily skip over familiar material.

Freiburg, Germany — Guido Schwarzer
London, UK — James R. Carpenter
Freiburg, Germany — Gerta Rücker
December, 2014

Contents

Chapter 2
Fixed Effect and Random Effects Meta-Analysis

In this chapter we describe the two main methods of meta-analysis, fixed effect model and random effects model, and how to perform the analysis in R. For both models the inverse variance method is introduced for estimation. The pros and cons of these methods in various contexts have been debated at length in the literature [9, 28, 29, 41], without any conclusive resolution. Here, we briefly describe each model, and how it is estimated in the R package **meta** [33, 34].[1]

An estimated treatment effect and its variance from each study are sufficient to apply the inverse variance method. Therefore, this method is sometimes called the generic inverse variance method. For the random effects model, various methods to estimate the between-study variance, the Hartung–Knapp adjustment and prediction intervals are briefly described.

We also show how to use R to generate forest plots. Along the way, we will show how the tabular and graphical summaries usually included in Cochrane reviews can be generated in R. We give examples using both base R and functions provided by our R package **meta**. The various methods of meta-analysis are best illustrated using base R; furthermore some basic R knowledge is gained from working with fundamental R functions. The R code using functions from the R package **meta** shows how routine manipulations and calculations can be automated. In practice a meta-analyst would like to do the analyses using the more sophisticated functions in the R package **meta**. Accordingly, readers not interested in the mathematical details could run over the examples using base R functions.

We will use a continuous outcome to introduce both fixed effect and random effects model. Accordingly, we start by describing the two most common effect measures for continuous outcomes, mean difference and standardised mean

[1] If you did not already install R package **meta** do so using R command `install.packages("meta")`.

G. Schwarzer et al., *Meta-Analysis with R*, Use R!,
DOI 10.1007/978-3-319-21416-0_2

difference. In Sect. 2.6, the generic inverse variance method is applied in meta-analyses with survival outcome, cross-over trials and adjusted estimates from regression models.

2.1 Effect Measures for Continuous Outcomes

Meta-analysis typically focuses on comparing two interventions, which we refer to as *experimental* and *control*. When the response is continuous (i.e. quantitative) typically the mean, standard deviation and sample size are reported for each group. Let $\hat{\mu}_{ek}, s^2_{ek}, n_{ek}$ and $\hat{\mu}_{ck}, s^2_{ck}, n_{ck}$ denote the observed mean, standard deviation and sample size for study k, $k = 1, \ldots, K$ (see Table 2.1).

We consider two different types of effect measures for continuous outcomes: mean difference and standardised mean difference. The mean difference is typically used when all studies report the outcome on the same scale. On the other hand, the standardised mean difference can be used when studies measure the outcome on different scales, e.g. different depression scales like the Hamilton Depression Rating Scale or the Hospital Anxiety and Depression Scale.

2.1.1 Mean Difference

For study k, the estimated mean difference is

$$\hat{\mu}_k = \hat{\mu}_{ek} - \hat{\mu}_{ck}, \tag{2.1}$$

Table 2.1 Variable names in R datasets for meta-analyses of continuous responses

Variable name	Notation	Description
`author`		First author of study
`year`		Year study published (if available)
`Ne`	n_e	Number of patients in the experimental (i.e. active) treatment arm
`Me`	$\hat{\mu}_e$	Mean response in the experimental treatment arm
`Se`	s_e	Standard deviation of the response in the experimental treatment arm
`Nc`	n_c	Number of patients in the control (often equivalent to placebo) arm
`Mc`	$\hat{\mu}_c$	Mean response in the control arm
`Sc`	s_c	Standard deviation of the response in the control arm

with variance estimate[2]

$$\widehat{\text{Var}}\,(\hat{\mu}_k) = \frac{s_{ek}^2}{n_{ek}} + \frac{s_{ck}^2}{n_{ck}}\,. \tag{2.2}$$

An approximate two-sided $(1-\alpha)$ confidence interval for the mean difference is given by

$$(\hat{\mu}_{ek} - \hat{\mu}_{ck}) \;\pm\; z_{1-\frac{\alpha}{2}} \sqrt{\frac{s_{ek}^2}{n_{ek}} + \frac{s_{ck}^2}{n_{ck}}} \tag{2.3}$$

with $z_{1-\frac{\alpha}{2}}$ denoting the $1-\frac{\alpha}{2}$ quantile of the standard normal distribution. For the usual 95 % confidence interval, $z_{1-\frac{0.05}{2}} = z_{0.975} = 1.96$, i.e. the 97.5 % point of the standard normal distribution.

Example 2.1 We return to the meta-analysis by Spooner et al. [37] comparing Nedocromil sodium with placebo for preventing exercise-induced bronchoconstriction which we already used in Chap. 1. Outcome of interest is the maximum fall in the forced expiratory volume in 1 second (FEV_1) over the course of follow-up, expressed as a percentage. Accordingly, all studies report the same outcome and the use of the mean difference is warranted.

The raw data consist of eight variables with headings in Table 2.1. Code to read in the data, together with the data, are shown in Fig. 1.2. From the data we see that the meta-analysis contains 17 studies, with sample sizes ranging between 16 (`Shaw 1985`; `DeBenedictis 1995`) and 48 (`Novembre 1994f`).

For each study (labelled by first author and date) mean values, standard deviations and sample sizes are given in Fig. 1.2. Thus for study 1 (`Boner 1988`) the estimated mean difference is $13.54 - 20.77 = -7.23$ and for study 2 (`Boner 1989`) it is $15.70 - 22.70 = -7.00$ (see Fig. 1.4). Accordingly, the maximum fall in FEV_1 is on average about 7 % in `Boner 1988` and `Boner 1989`. For study 1 (`Boner 1988`) the 95 % confidence interval (2.3) is

$$(13.54 - 20.77) \;\pm\; 1.96 \sqrt{\frac{13.85^2}{13} + \frac{21.46^2}{13}} \quad \text{giving} \quad (-21.11, 6.65).$$

We can use base R to calculate mean difference and 95 % confidence interval for the `Boner 1988` trial (assuming that the file `dataset01.csv` is in the current working directory; see Sect. 1.2 for details):

```
> # 1. Read in the data
> data1 <- read.csv("dataset01.csv", as.is=TRUE)
> # 2. Calculate mean difference and its standard error for
```

[2]Note we could use a pooled estimate of the sample variance, but this assumes that the response variance is the same in the two groups which will not be true in general.

We get the same result by using the `metacont` function with argument `sm="SMD"` (Standardised Mean Difference):

```
> print(metacont(Ne, Me, Se, Nc, Mc, Sc, sm="SMD",
+               data=data2, subset=1), digits=2)
  SMD        95%-CI     z  p-value
 -0.6 [-1.33; 0.13] -1.61   0.1083

Details:
- Inverse variance method
```

Once the standardised mean difference and its variance have been calculated using the formulae (2.4) and (2.5), the calculations for both fixed effect and random effects meta-analyses follow exactly as described in the next section. □

2.2 Fixed Effect Model

The fixed effect model assumes that the estimated effects from the component studies in a meta-analysis come from a single homogeneous population. In order to calculate an overall estimate, we therefore average the estimates from each study, allowing for the fact that some estimates are more precise than others (having come from larger studies).

More formally, let $k = 1, \ldots, K$ index study, $\hat{\theta}_k$ denote the intervention effect estimate from study k, and θ denote the intervention effect in the population, which we wish to estimate. Denote by $\hat{\sigma}_k^2$ the sample estimate of $\text{Var}\,(\hat{\theta}_k)$.

The fixed effect model is

$$\hat{\theta}_k = \theta + \sigma_k \epsilon_k, \quad \epsilon_k \overset{\text{i.i.d.}}{\sim} N(0,1). \tag{2.7}$$

We now consider the fixed effect estimate of θ, denoted by $\hat{\theta}_F$. Given estimates $(\hat{\theta}_k, \hat{\sigma}_k)$, $k = 1, \ldots, K$, the maximum-likelihood estimate under model (2.7) is

$$\hat{\theta}_F = \frac{\sum_{k=1}^{K} \hat{\theta}_k / \hat{\sigma}_k^2}{\sum_{k=1}^{K} 1/\hat{\sigma}_k^2} = \frac{\sum_{k=1}^{K} w_k \hat{\theta}_k}{\sum_{k=1}^{K} w_k}. \tag{2.8}$$

Accordingly, $\hat{\theta}_F$ is a weighted average of the individual effect estimates $\hat{\theta}_k$ with weights $w_k = 1/\hat{\sigma}_k^2$. Therefore, this method is called the *inverse variance method*.

The variance of $\hat{\theta}_F$ is estimated by

$$\widehat{\text{Var}}\,(\hat{\theta}_F) = \frac{1}{\sum_{k=1}^{K} w_k}. \tag{2.9}$$

A $(1-\alpha)$ confidence interval for $\hat{\theta}_F$ can be calculated by

$$\hat{\theta}_F \pm z_{1-\frac{\alpha}{2}} \text{S.E.}(\hat{\theta}_F) \tag{2.10}$$

with standard error $\text{S.E.}(\hat{\theta}_F) = \sqrt{\widehat{\text{Var}}(\hat{\theta}_F)}$ and $z_{1-\frac{\alpha}{2}}$ denoting the $1-\frac{\alpha}{2}$ quantile of the standard normal distribution. A corresponding test for an overall treatment effect can be constructed using $\hat{\theta}_F \Big/ \text{S.E.}(\hat{\theta}_F)$ as test statistic.

Example 2.3 The fixed effect estimate $\hat{\theta}_F$ and its 95 % confidence interval for the bronchoconstriction meta-analysis are given in Fig. 1.4; here we show how $\hat{\theta}_F$ can be calculated using R. Recall Eqs. (2.1) and (2.2) which give the mean difference $\hat{\mu}_k$ and its variance estimate $\widehat{\text{Var}}(\hat{\mu}_k)$. The fixed effect estimate $\hat{\theta}_F$ and its variance can be calculated using the following quantities:

$$\hat{\theta}_k = \hat{\mu}_k$$
$$\hat{\sigma}_k^2 = \widehat{\text{Var}}(\hat{\mu}_k).$$

The fixed effect estimate and its variance can be calculated using base R code:

```
> # 1. Calculate mean difference, variance and weights
> MD <- with(data1, Me - Mc)
> varMD <- with(data1, Se^2/Ne + Sc^2/Nc)
> weight <- 1/varMD
> # 2. Calculate the inverse variance estimator
> round(weighted.mean(MD, weight), 4)
[1] -15.514
> # 3. Calculate the variance
> round(1/sum(weight), 4)
[1] 1.4126
```

Note, the standard `weighted.mean` function is used to calculate $\hat{\theta}_F$.

The meta-analysis can be conducted much easier using the `metacont` function which yields identical results:

```
> mc1 <- metacont(Ne, Me, Se, Nc, Mc, Sc,
+                 data=data1,
+                 studlab=paste(author, year))
> round(c(mc1$TE.fixed, mc1$seTE.fixed^2), 4)
[1] -15.5140   1.4126
```

We select `mc1$TE.fixed`, i.e. the Treatment Estimate in the fixed effect model, and its standard error `mc1$seTE.fixed` from the meta-analysis object `mc1`. We can use the command `str(mc1)` to print the whole structure of the meta-analysis object `mc1` and look at the help page of the `metacont` function which describes the individual elements of `mc1`.

A complete printout for the meta-analysis is given in Fig. 2.2. The first thing the output gives is a table whose rows are the component studies in the meta-analysis.

```
                          MD          95%-CI %W(fixed) %W(random)
Boner 1988             -7.2 [-21.1;   6.7]      2.82       3.08
Boner 1989             -7.0 [-16.2;   2.2]      6.38       6.58
Chudry 1987           -18.4 [-28.8;  -8.0]      5.01       5.29
Comis 1993            -16.8 [-27.8;  -5.8]      4.50       4.78
DeBenedictis 1994a    -13.0 [-22.8;  -3.2]      5.68       5.93
DeBenedictis 1994b    -16.6 [-35.8;   2.6]      1.47       1.64
DeBenedictis 1995     -13.9 [-27.6;  -0.2]      2.87       3.13
Debelic 1986          -18.2 [-30.7;  -5.8]      3.52       3.80
Henriksen 1988        -29.7 [-41.6; -17.8]      3.83       4.11
Konig 1987            -14.2 [-25.0;  -3.4]      4.65       4.93
Morton 1992           -22.5 [-33.5; -11.5]      4.48       4.76
Novembre 1994f        -13.0 [-19.5;  -6.6]     12.98      12.15
Novembre 1994s        -15.1 [-23.8;  -6.4]      7.14       7.28
Oseid 1995            -14.8 [-23.7;  -5.9]      6.82       6.99
Roberts 1985          -20.0 [-36.9;  -3.1]      1.90       2.10
Shaw 1985             -24.2 [-33.2; -15.1]      6.67       6.85
Todaro 1993           -13.4 [-18.7;  -8.1]     19.29      16.58

Number of studies combined: k=17

                           MD          95%-CI      z  p-value
Fixed effect model    -15.5 [-17.8; -13.2] -13.1 < 0.0001
Random effects model  -15.6 [-18.1; -13.2] -12.3 < 0.0001

Quantifying heterogeneity:
tau^2 = 2.4374; H = 1.05 [1; 1.35]; I^2 = 8.9% [0%; 45.3%]

Test of heterogeneity:
     Q d.f.  p-value
 17.57   16   0.3496

Details on meta-analytical method:
- Inverse variance method
- DerSimonian-Laird estimator for tau^2
```

Fig. 2.2 Output from meta-analysis of the bronchoconstriction meta-analysis [37]. The output starts with a table of the included studies. For each study, the mean difference (MD) with 95 % confidence interval is given, along with weights used for fixed effect and random effects model. There are 17 studies in the example. Next, the results of fixed effect and random effects model are presented with 95 % confidence intervals, z statistic and p-value. Heterogeneity is quantified by the estimated between-study variance τ^2, H and I^2, see Sects. 2.3 and 2.4, and tested using Cochran's Q statistic, see Eq. (2.12). There is not much heterogeneity present in this example. The output ends with details of the methods used, e.g. how τ^2 was estimated, see Sect. 2.3.1

This table is also shown in Fig. 1.4 on the right side of the forest plot. The column `MD` is the mean difference of the response (maximum change in FEV_1 as a percentage) between the Nedocromil sodium and placebo group. Next comes a 95 % confidence interval for this difference, calculated based on (2.3). The next two columns are the

weights given to the study under the fixed effect (`%W(fixed)`) and random effects model (`%W(random)`).

The weight of study 1 (`Boner 1988`) in the fixed effect meta-analysis is given by the inverse of the variance (2.2) which can be calculated as

$$1 \Big/ \left(\frac{13.85^2}{13} + \frac{21.46^2}{13} \right) = 1/50.18108 = 0.01992783.$$

The percentage weight of study 1 (`Boner 1988`) in the fixed effect meta-analysis reported in Figs. 1.4 and 2.2 is

$$100 \cdot \frac{w_1}{\sum_{i=1}^{17} w_i} = 100 \cdot \frac{0.01992783}{0.7079028} = 2.82\,\%.$$

We could also use R to calculate these values:

```
> mc1$w.fixed[1]
[1] 0.01992783
> sum(mc1$w.fixed)
[1] 0.7079028
> round(100*mc1$w.fixed[1] / sum(mc1$w.fixed), 2)
[1] 2.82
```

After reporting the number of studies combined in meta-analysis, fixed effect estimate $\hat{\theta}_F$, random effects estimate $\hat{\theta}_R$ (see Sect. 2.3) and their 95 % confidence intervals, z and p-values are given in Fig. 2.2. Next come the measures for heterogeneity and a test for heterogeneity (see Sect. 2.4). Finally a note indicates that the "Inverse variance method" has been used. This is in fact the only method for continuous data; but with binary data (see Chap. 3) we shall see there are other alternatives.

A forest plot is shown in Fig. 2.3 which has been produced by the R command

```
> forest(mc1, comb.random=FALSE, xlab=
+        "Difference in mean response (intervention - control)
+ units: maximum % fall in FEV1",
+        xlim=c(-50,10), xlab.pos=-20, smlab.pos=-20)
```

Note the use of the `xlab` option to label the x-axis, and in particular how a line break in the input text creates a line break in the axis label on the graph. The option `xlim=c(-50,10)` is used to specify that the limits of the x-axis are between -50 and 10. The options `xlab.pos` and `smlab.pos` specify the centre of the label on x-axis and the summary measure at the top of the figure; otherwise these texts would be centred around 0.

Note, the meta-analysis could have also been done using the `metagen` function which is the primary function in R package **meta** to conduct a meta-analysis based on the generic inverse variance method.

```
> # 1. Apply generic inverse variance method
> mc1.gen <- metagen(mc1$TE, mc1$seTE, sm="MD")
```

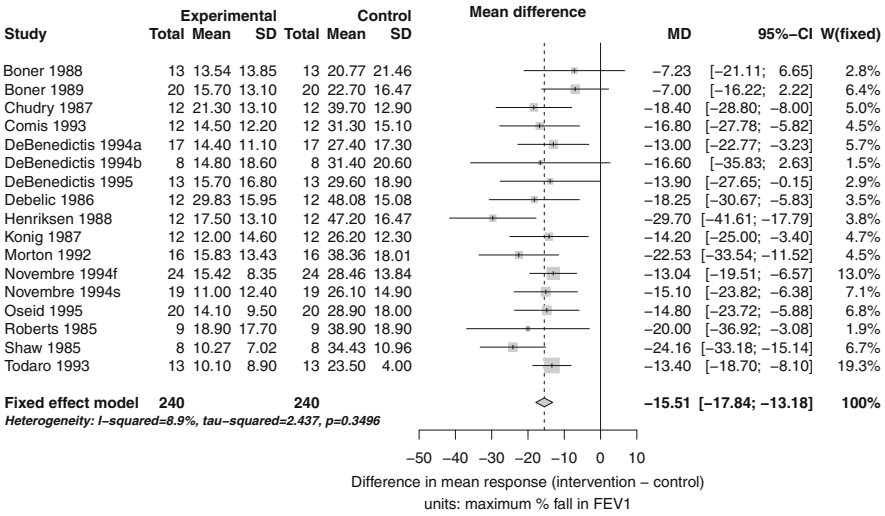

Fig. 2.3 Forest plot for the bronchoconstriction meta-analysis [37]. For details, see text

```
> # 2. Same result
> mc1.gen <- metagen(TE, seTE, data=mc1, sm="MD")
> # 3. Print results for fixed effect and random effects method
> c(mc1$TE.fixed, mc1$TE.random)
[1] -15.51403 -15.64357
> c(mc1.gen$TE.fixed, mc1.gen$TE.random)
[1] -15.51403 -15.64357
```

In steps 1 and 2, the generic inverse variance method is applied using the `metagen` function; we use the list elements `mc1$TE` (treatment effect) and `mc1$seTE` (standard error) as inputs to the `metagen` function. Output of resulting object `mc1.gen` is identical to results using the `metacont` function as exemplified in step 3 for the fixed effect and random effects estimates. Applying the `metagen` function in this way seems rather artificial, however, as we will see in Sect. 2.6 this function can be used to conduct a meta-analysis for other outcomes. □

Following RevMan 5, the following quantities are used to estimate the standardised mean difference in the fixed effect model:

$$\hat{\theta}_k = \hat{g}_k$$
$$\hat{\sigma}_k^2 = \widehat{\mathrm{Var}}\,(\hat{g}_k)$$

with $\hat{g}_k$ and $\widehat{\mathrm{Var}}\,(\hat{g}_k)$ defined in (2.4) and (2.5). These quantities are utilised in formulae (2.8)–(2.10) to calculate the fixed effect estimate of the standardised mean difference.