

Module 3: An introduction to meta-analysis

Learning objectives


- Be able to explain the difference between meta-analyses and systematic reviews
- Understand that meta-analysis is a two stage-process involving (a) computing summary statistics for each trial, (b) averaging the summary statistics
- Understand why simply adding up data from individual studies is inappropriate
- Understand that a full analysis also involves proper consideration of the consistency of trial results
- Be able to name and explain the main characteristics of a forest plot

Systematic reviews and meta-analyses

In Module 1 we summarised the process of preparing a systematic review. Part of that process is to calculate the results of each study identified by the reviewer, and then to calculate an average of those results – *if appropriate* – in a meta-analysis. Systematic reviews do not *have* to have a meta-analysis – there are times when it is not appropriate or possible.

To represent this visually, the figure below shows that a meta-analysis may be part of a systematic review. A meta-analysis is also possible without doing a systematic review – you could just find a few studies and calculate a result, with no attempt to be systematic about how the particular studies were chosen.

A systematic review may have a statistical combination of studies (a meta-analysis) but it does not have to



One slight complication is that these two terms are often used interchangeably, particularly in North America. In this learning material, the term ‘systematic review’ will refer to the entire process of collecting, reviewing and presenting all available evidence, while the term ‘meta-analysis’ will refer to the statistical technique involved in extracting and combining data to produce a summary result.

What is a meta-analysis?

There are two stages in a meta-analysis: the results for each study are calculated, then a pooled average of those results is calculated

A meta-analysis is a two-stage process. The first stage is the extraction of data from each individual study and the calculation of a result for that study (the ‘point estimate’ or ‘summary statistic’), with an estimate of the chance variation we would expect with studies like that (the ‘confidence interval’).

The second stage involves deciding whether it is appropriate to calculate a pooled average result across studies and, if so, calculating and presenting such a result. Part of this process is to give greater weight to the results from studies which give us more information, because these are likely to be closer to the truth we are trying to estimate. We’ll come back to these topics in later modules.

The results of meta-analyses are often presented in a forest plot. Run through this slide show, which explains the parts of these plots.

See presentation at end of modules for an explanation of a forest plot

A meta-analysis does not just add up the numbers from the trials

There is more to meta-analysis than simply adding up the numbers from studies

One common criticism of meta-analysis is that it somehow simply adds together the results from quite different studies and calculates a summary statistic as if it is one big study. It would be wrong to do this, and this is not what a meta-analysis does. A meta-analysis looks at the results *within* each study, and then calculates a weighted average.

The reasons for this are explained in a later module. For now, it's enough to realise that if we just add up the numbers of people and events (such as deaths) from a number of trials, we effectively treat it as one big trial. In effect we will be comparing people in the treatment group of one trial with people in the control group of another trial. This comparison is not randomised, and it is likely that there will be some differences in the way the trials were carried out. This doesn't make sense when we have gone to a lot of trouble to find randomised comparisons, and it does not make logical sense to do this.

Is it sensible to calculate a pooled estimate?

With the right software, it's very easy to do a meta-analysis. In fact it's almost too easy. Before pressing the button to calculate a meta-analysis, it's important to ask whether it is sensible to do so.

There are two parts to making this decision – clinical and statistical. The first, clinical, part involves asking yourself whether the studies you have found really do all address the same question so that an average of their results would be sensible. There might be differences in the participants, interventions or outcomes that lead you to think that the treatment effect is really very different in the different studies. Here are some examples to think about:


- Would you expect pregnant women to respond differently from teenage boys to interventions designed to help them stop smoking?
- Would you expect the effect of acupuncture to vary depending on who did it?
- Would you expect the outcome of treatments for depression to be similar at 3 days and 3 months?


There is always an element of judgement in these decisions. But if you think there are good reasons why you would expect the effects of an intervention to differ substantially between studies, you should not pool the results.

The second, statistical, way of thinking about the consistency of the results in the studies included in a review is to look for big differences between the results of the trials. This will be covered later, in the module on Heterogeneity. For now, it's enough to know that you should look for studies where the results don't seem to fit, and then investigate possible reasons for this.

Comparison: Any steroid administered in any dose against no steroid
Outcome: Death at end of follow up period

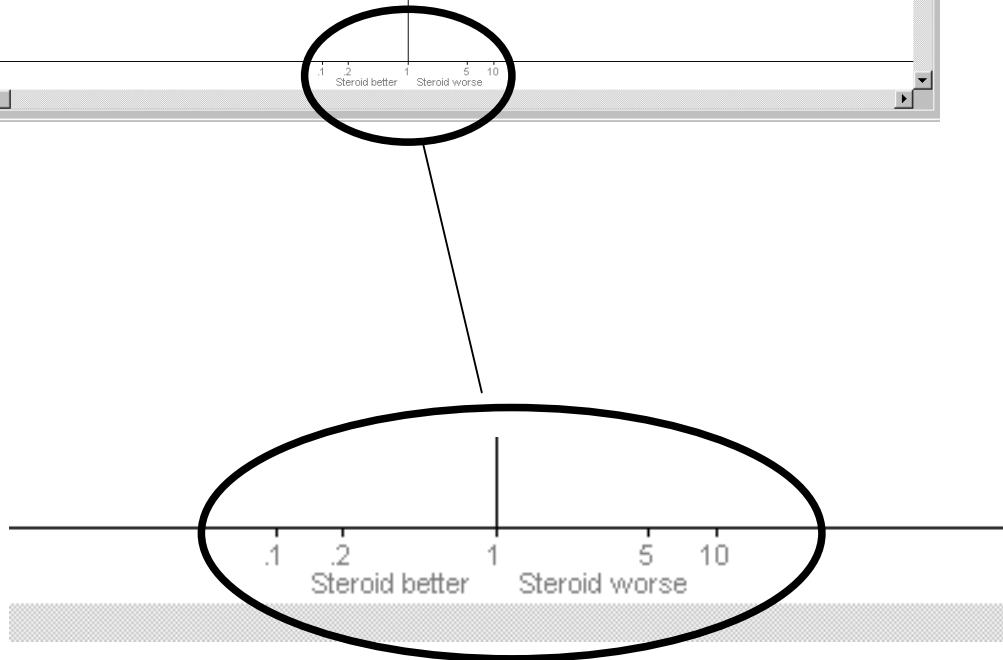
there's a label to tell you what the comparison is and what the outcome of interest is






At the bottom there's a horizontal line. This is the scale measuring the treatment effect. Here the outcome is death and towards the left the scale is less than one, meaning the treatment has made death less likely.

Take care to read what the labels say - things to the left do not always mean the treatment is better than the control.





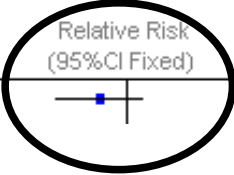
The vertical line in the middle is where the treatment and control have the same effect - there is no difference between the two

Study	Expt n/N	Ctrl n/N	Relative Risk (95%CI Fixed)	Weight %	RR (95%CI Fixed)
Alexander 1972	16 / 55	22 / 55		6.6	0.73 [0.43,1.23]


For each study there is an id

The data for each trial are here, divided into the experimental and control groups

This is the % weight given to this study in the pooled analysis

Study	Expt n/N	Ctrl n/N	Relative Risk (95%CI Fixed)	Weight %	RR (95%CI Fixed)
Alexander 1972	16 / 55	22 / 55		6.6	0.73 [0.43,1.23]

The data shown in the graph are also given numerically



The label above the graph tells you what statistic has been used

- Each study is given a blob, placed where the data measure the effect.
- The size of the blob is proportional to the % weight
- The horizontal line is called a confidence interval and is a measure of how we think the result of this study might vary with the play of chance.
- The wider the horizontal line is, the less confident we are of the observed effect.

Total (95%CI)

410 / 1194


301 / 925



100.0

Chi-square 18.11 (df=14) Z=0.68

The pooled analysis is given a diamond shape where the widest bit in the middle is located at the calculated best guess (point estimate), and the horizontal width is the confidence interval



Note on interpretation

If the confidence interval crosses the line of no effect, this is equivalent to saying that we have found no statistically significant difference in the effects of the two interventions