# THREE-DAY META-ANALYSIS COURSE

## WHAT WE COVER

This three-day workshop follows the same approach we employed in the best-selling text, *Introduction to Meta-Analysis* (Wiley, 2009).

The first day we deal primarily with conceptual issues in meta-analysis. We start by explaining what a meta-analysis is and how it should be performed. We introduce the two common statistical models (fixed-effect and random-effects), explain how to choose between them, and how the choice of a model affects the results. We introduce participants to the computer program Comprehensive Meta-Analysis (CMA), which automates all steps in a meta-analysis, from computing effect sizes to creating a publication-quality forest plot.

The second day we focus on more advanced issues, addressing these conceptually and also from a practical "How to" perspective. These issues include using subgroups-analysis to compare the effect size in different subgroups of studies, and using meta-regression to assess the relationship between covariates and effect size.

The third day we discuss working with complex data sets. These include studies that report effects for more than one subgroup of people, studies that report results for more than one outcome, and studies that report results at more than one time-point. It also includes studies that compare two or more treatment groups to a common control group. We also discuss publication bias, the question of when it makes sense to perform a meta-analysis, and how to use a cumulative meta-analysis.

At the conclusion of each topic we discuss common mistakes related to that topic, and how to avoid them. Throughout the course we focus on concepts rather than formulas. Researchers and clinicians with relatively little statistical background should be able to follow all materials without any problem. At the same time, statisticians are likely to have a series of "Aha" moments and walk away with a thorough grasp of the statistics underlying the concepts.

All participants will understand how to read and critique a meta-analysis, to perform a meta-analysis using the CMA software, and to report meta-analysis.

## WHAT WE DO NOT COVER

A systematic review is a lengthy process that includes formulating a research problem, searching the literature, deciding which studies to include in the synthesis, and then performing the statistical analysis, called the meta-analysis. This workshop focuses only on this last step, the meta-analysis.

# AGENDA

## DAY-1

### What is a meta-analysis?

Meta-analysis refers to the statistical procedures for synthesizing data from a set of studies. We start with the summary effect (such as the mean difference, risk ratio, or correlation) from each study, and the meta-analysis enables us to view these effects in context. If the effect size is consistent, we can report that the treatment effect is robust, and can also estimate the effect size with better precision than we could with a single study. If the effect size varies, we can describe the variation and may also be able to explain it.

### The shift from narrative reviews to systematic reviews

Prior to the 1990s, the primary mechanism for synthesizing data from multiple studies was the narrative review, where an expert would study the literature, critique and summarize the results, and come to a conclusion. Beginning in the early 1990s meta-analysis began to replace the narrative review, and it now dominates the field of research synthesis in such fields in medicine, education, ecology, psychology, business, criminal justice, and others. We will discuss the key differences between meta-analysis and narrative reviews, how researchers' thinking about these differences has evolved, and why meta-analysis is now generally embraced as the gold standard.

### Fixed-effect and random-effects models

In a meta-analysis we want to assign a weight to each study, with more precise studies getting more weight and less precise studies getting less weight. It turns out that the definition of "precise" depends on our understanding of the sampling frame. The fixed-effect model is appropriate in some cases, while the random-effects model is appropriate in others. We explain the conceptual and practical differences between the models. We also discuss common (and serious) mistakes in selecting a model, and how to avoid these mistakes.

## Effect sizes for means, binary, and correlational data

The effect size (or treatment effect) is the unit of currency in a meta-analysis. This could be a mean difference, a risk ratio, a correlation, or another index, but the basic idea is the always the same. We compute an effect size for each study, and then use the meta-analysis to synthesize these values. We also need to compute the variance of each effect size, since this affects the weight that is allocated to each study in the synthesis. We will discuss several effect sizes, and show how to compute these using Excel™ and CMA™.

## Performing a meta-analysis using CMA

We will show how to use CMA to perform a simple meta-analysis. Steps include entering the data, running the analysis, studying the results, creating a forest plot, and exporting the plot to PowerPoint© and to Word©. We will also use the educational features of CMA to illustrate conceptual issues in meta-analysis. For example, CMA is able to display the study weights under the fixed-effect and random-effects models, and how these weights affect the results.

## Avoiding common mistakes, Part 1

Over a period of years we have reviewed hundreds of meta-analyses and have compiled a list of mistakes that people make on a fairly regular basis. At the conclusion of every session we'll discuss common mistakes related to that session. This will help you to avoid these mistakes, and also to defend your analysis if the person reviewing your paper should raise questions about the procedures. Some of the mistakes we'll discuss on Day-1 include mistakes in the selection of a statistical model, and mistakes in interpreting the mean effect and its confidence interval.

# DAY-2

## Heterogeneity in treatment effects

We will explain how to quantify and to understand heterogeneity in effect sizes. We explain this conceptually, and then show how the concept leads to a series of distinct indices, each with a different meaning and purpose.

## Comparing the effect size in different groups of studies

Suppose we are working with 20 studies that tested the impact of a drug. Ten of the studies employed a standard dose of the drug while the others employed a high dose. Was the impact greater in one group then other? Or, suppose we are working with 20 studies that tested the impact of an intervention for students. Some of the studies ran the intervention for eight weeks while the others ran it for sixteen weeks. Was the longer intervention more effective? In a primary study we would address these kinds of questions by using a t-test or analysis of variance, and we can apply similar procedures in meta-analysis. We'll show how to perform these kinds of analyses using CMA.

## Using regression to assess the impact of continuous moderators on effect size

In a primary study we use multiple-regression to assess the relationship between covariates and outcome. This technique is simple enough to be applied with a single predictor, but also allows us to work with sets of predictors, and to assess the impact of one set with another set held constant. We will show how the same techniques can be employed with meta-analysis, and how to use CMA for this purpose.

## Avoiding common mistakes, Part 2

Some of the mistakes we'll discuss on Day-1 include mistakes in reporting heterogeneity. For example, the "common" wisdom is that the I-squared index reflects the amount of dispersion in treatment effects, with certain values (typically 25%, 50%, 80%) taken to mean small, medium, and large

dispersion. In fact, I-squared is NOT a measure of heterogeneity. We will discuss what I-squared does measure. We will also discuss indices that do reflect the dispersion in treatment effects, and explain how to use these.

## DAY-3

### Working with studies that report effects for two or more independent subgroups

Sometimes studies report the impact of a treatment separately for two (or more) independent samples – for example, the impact of an intervention for males and also for females. We explain the options for working with this kind of data, and how to implement these options in CMA.

### Working with studies that report effects for two or more outcomes

Suppose studies report the impact of an intervention on both math and reading scores. In the analysis we may want to assess the impact for each outcome separately, or we may want to assess the impact on a composite score called "Academic achievement". We explain the options for working with this kind of data, and how to implement these options in CMA.

### Working with studies that compare several treatment groups to a common control group

Suppose studies report the impact of an intervention on both math and reading scores. In the analysis we may want to assess the impact for each outcome separately, or we may want to assess the impact on a composite score called "Academic achievement". We explain the options for working with this kind of data, and how to implement these options in CMA.

## Important conceptual issues in meta-analysis

When does it make sense to perform a meta-analysis? How many studies do we need? How similar do the studies need to be? What is the role of a cumulative meta-analysis? What if the meta-analysis appears to conflict with a large-scale trial? We discuss these and other issues that arise frequently and are not well-understood.

## Criticisms of meta-analysis

In 1990 Charles Mann published a paper about the future of meta-analysis. While many researchers were convinced that meta-analysis would eventually serve as the basis for evidence-based practice, others were less enthusiastic. One person compared meta-analysis to alchemy and another compared it to terrorism. In 1993, an editorial in the New England Journal of Medicine suggested that meta-analyses are so likely to be flawed that it would be preferable to stick with narrative reviews. It's important to understand the criticisms outlined in these papers (and others), both to rebut the objections that are not valid and also to learn from the criticisms that are valid.

## Resources for meta-analysis

We provide a synopsis of various texts. Some of these show how to work with meta-analysis in a specific substantive area such as medicine, social science, or ecology. Others are intended as general texts or as handbooks. We also discuss a number of web sites and professional groups that serve as resources for systematic reviews and meta-analysis.

## Avoiding common mistakes, Part 3

The objective of this session is similar to the corresponding sessions on Days 1 and 2. We will discuss mistakes to avoid when working with complex data sets. We will discuss some of the key mistakes people make when discussing publication bias. We will discuss the argument that a meta-analysis should not be performed when the effect size varies across studies.

# LEARNING OUTCOMES

## AFTER DAY-1, YOU SHOULD UNDERSTAND:

- What a meta-analysis is, and how to perform one
- The key differences between meta-analysis and narrative reviews
- That the goal of a meta-analysis is to synthesize, and not simply combine, effect sizes
- The difference between fixed-effect and random-effects models
- How to use a forest-plot to understand and report a meta-analysis
- How to use CMA to compute effect sizes, perform a simple analysis, and create forest plots
- Common mistakes in meta-analysis, and how to avoid them
- Mistakes in choosing between fixed-effect and random-effects models
- Mistakes in understanding why a meta-analysis appears to conflict with a clinical trial
- Mistakes in using Vote-counting
- Mistakes in the goals of meta-analysis

## AFTER DAY-2, YOU SHOULD UNDERSTAND:

- How to quantify and interpret heterogeneity
- How to compare the effect size in subgroups of studies
- How to use regression to assess the relationship between covariates and effect size
- Common mistakes in meta-analysis, and how to avoid them
- Mistakes in interpreting indices of heterogeneity
- Mistakes in choosing between fixed-effect and random-effects models for subgroups-analysis and meta-regression

## AFTER DAY-3, YOU SHOULD UNDERSTAND:

- How to work with studies that report effects for two (or more) independent subgroups
- How to work with studies that report effects for two (or more) outcomes or time-points
- How to work with studies that compare two (or more) treatments to a common control group
- How to decide whether or not it makes sense to perform a meta-analysis
- How to assess the potential impact of publication bias
- How to perform a meta-analysis using studies that employed different designs (matched groups vs. independent groups), formats (some reported means, others reported t-tests) or outcomes (some worked with means, others with risks).
- How to use CMA to perform all of these analyses
- Common mistakes in meta-analysis, and how to avoid them
- Mistakes in working with multiple outcomes from the same sample
- Mistakes in interpreting publication bias