episummer
@columbia

Meta-Analysis – Concepts and Applications

Intensive course

COURSE DURATION

This is an on-line, asynchronous course and material will be available from June 1 – 30, 2024

INSTRUCTOR

Michael Borenstein, PhD

Director, Biostat, Inc.
Associate Professor, Albert Einstein College of Medicine (1990-2001)
Director of Biostatistics, Hillside Hospital, Long Island Jewish Medical Center (1982-2002)
Biostat100@GMail.com

Michael Borenstein is the co-author (with Larry Hedges, Julian Higgins, and Hannah Rothstein) of the text *Introduction to Meta Analysis* (Wiley, 2009, 2021). He is also the author of the text *Common Mistakes in Meta-Analysis and How to Avoid Them*. A PDF of this book may be downloaded here for free. Dr. Borenstein has published numerous papers on meta-analysis and contributed chapters to texts on meta-analysis and systematic reviews. He is the primary developer of the software Comprehensive Meta Analysis (CMA) which will be used in this course. This software was developed with funding from the National Institutes of Health in the United States, in collaboration with experts from the US, the UK, and Australia. It was initially published in the year 2000 and is now in its fourth major release. The software is used by tens of thousands of researchers throughout the world. Dr. Borenstein has been the primary investigator on numerous grants from the National Institutes of Health (NIH) to develop methods and software for conducting and teaching meta- analysis. He serves as a reviewer for meta-analyses submitted for publication. He is a founding member of the Society for Research Synthesis Methodology, and served as its president in 2017-2018. Dr. Borenstein has been teaching workshops on meta-analysis since 2005 at numerous venues including the National Institutes of Health, the Centers for Disease Control, the Federal Drug Administration, and the Karolinska Institute, as well as scores of workshops open to the public.

Testimonials are available here. Many of these cite the instructor's ability to take complex concepts and present them in a way that makes them seem intuitive.

COURSE DESCRIPTION

Meta-Analysis is the process of synthesizing data from multiple studies to yield a more complete understanding of the effect. It is the core of evidence-based medicine and epidemiology. The goal of this course is to teach students how to perform a meta-analysis, to interpret the results properly, and to present the results in a manner that is clear and intuitive. The techniques that we discuss in these workshops are relevant to any field where a meta-analysis might be employed. Participants typically

work in such fields as medicine, epidemiology, health sciences, biostatistics, nursing, veterinary medicine, psychology, education, criminal justice, ecology, marketing, among others.

Basic course (20 hours) June 1-15

In the basic course we explain how to conduct and report the results for a meta-analysis. In an analysis where the effect size is consistent across studies the goal will be to identify that common effect size and consider the clinical implications of the findings. In an analysis where the effect size varies across studies the goal will be to identify the mean effect size and also to quantify the variation in effects and consider the clinical implications of this variation: Is it the case that the treatment's effect varies from moderate to large; or from minor to moderate; or is the treatment helpful in some cases but harmful in others?

Advanced course (20 hours) June 16-30

In the advanced course we revisit the same datasets we had used in the basic course. Now, rather than simply quantify the heterogeneity in effects, we try to identify factors that may explain that heterogeneity. Specifically, we use subgroup analysis (analogous to ANOVA) to compare the effect size in different sets of studies. And we use meta-regression (analogous to multiple regression) to identify the unique impact of any covariates on the effect size. We also explore other issues in meta-analysis such as publication bias; limitations of the random-effects model; and options for working with a small number of studies.

Intensive course (40 hours) June 1-30

The intensive course is simply the basic course plus the advanced course. The basic course runs from June 1-15 and the advanced course runs from June 16-30.

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.

PREREQUISITES

Basic course

This workshop is intended for researchers, clinicians, educators, statisticians, graduate students, and anyone with an interest in understanding or performing meta-analyses. Our approach is primarily conceptual rather than mathematical, and participants with only one or two statistics classes in their background should be able to fully understand the materials.

Advanced course

The main prerequisite for the advanced course is the basic course. If you register for the intensive course, you will take the basic course first, followed immediately by the advanced course.

The first two modules in the advanced course deal with subgroup analysis and meta-regression. To fully understand these modules, you will need to be familiar with analysis of variance and multiple regression in primary studies.

Is this course too basic for me?

Probably not. While the course is geared to researchers, statisticians who attend our workshops invariably report that they found the course well worth their time. The overwhelming majority of meta-analyses contain some serious mistakes. You will learn what these mistakes are and how to avoid them in your own work. 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.

Is the course too advanced for me?

Probably not. The course assumes only basic knowledge of research techniques and no knowledge of meta-analysis. The course is intended for researchers rather than statisticians. I focus on a conceptual approach to each issue and use practical examples from published analyses. Statistical formulas are typically addressed separately and explained in a way that focuses on the logic, rather than the math.

I have additional questions

Please e-mail me at Biostat100@GMail.com

Can I see a sample video?

To get a sense of the course, click here to watch a sample video. This should give you a sense of how the course works.

Can I ask questions about my own data?

The zoom sessions are intended to clarify the issues covered in the videos and other general issues in meta-analysis. I would be happy to discuss your data as it relates to these issues, and to answer any questions of general interest based on your data. However, we cannot spend much time on issues that are unique to a specific dataset and not of general interest.

TECHNICAL REQUIREMENTS

Students will need access to a computer with internet access. To run the exercises students will need to download and install the software Comprehensive Meta-Analysis (CMA). The software runs on Windows. Students using a MAC will need to install a Windows emulator such as Parallels. All students will be given a free license to use the software for the duration of the course.

COURSE LEARNING OBJECTIVES

Basic course (20 hours June 1-15)

The objective of this course is to introduce students to the logic of a basic meta-analysis and teach them how to perform a meta-analysis, critique a meta-analysis, and avoid common mistakes in meta-analysis. By the end of the course students will understand

- The goals of a meta-analysis
- How to choose a statistical model
- How to choose an effect-size index
- How to enter data for a simple meta-analysis
- How to estimate the mean effect size
- How to quantify and understand heterogeneity in effects
- How to report the results of the analysis
- How to create forest plots
- How to create plots that show the distribution of true effects
- How to avoid common mistakes in all these areas

Advanced course (20 hours June 15-30)

The objective of this course is to teach students advanced issues in meta-analysis. By the end of the course the student will understand

- How to use subgroup analyses to compare the impact of a treatment in sets of studies that enrolled different populations or employed different variants of an intervention (analogous to ANOVA in a primary study)
- How to use meta-regression to assess the unique impact of continuous or categorical covariates on the effect size (analogous to multiple-regression in a primary study)
- How to assess the potential impact of publication bias on the analysis
- What to do when there are only a small number of studies in the analysis
- How to avoid common mistakes in all these areas

Intensive course (40 hours June 1-30)

The intensive course is simply the basic course plus the advanced course. The basic course runs from June 1-15 and the advanced course runs from June 16-30.

COURSE READINGS

There are no required readings, but the following are strongly recommended:

Introduction to Meta-Analysis (2nd Ed) Michael Borenstein, Larry V. Hedges, Julian PT Higgins, Hannah Rothstein. Wiley (2021). Available on Amazon

Common Mistakes in Meta-Analysis and How to Avoid Them. Michael Borenstein, Biostat (2019). Available on Amazon. To download a free PDF of the book click here`

COURSE STRUCTURE

Online videos

The course is centered about a series of video tutorials. There are roughly 3 hours of videos for each module. Watch these toward the beginning of each module's time period.

Group discussions

A key part of the class is the opportunity for group discussions with the instructor. These will be scheduled several times a week for the duration of the class. We will also record these and make them available to all students. For students in Europe or Australia I will schedule discussions as needed.

Exercises

For all modules there are PDFs that offer the opportunity to work through an analysis on your own, with step-by-step instructions. If you have any questions, we can discuss these during the zoom calls

Readings

For most modules there are optional readings that provide additional details on the topic.

Software

We will be using the software Comprehensive Meta-Analysis, or CMA. This software was developed with funding from the National Institutes of Health and is used by tens of thousands of researchers throughout the world. It incorporates many features that make ideal for teaching including the ability to present results in intuitive ways and to generate detailed text reports that explain the meaning of the statistics. Full disclosure: The instructor has a commercial interest in CMA.

All students will be given a free license for the duration of the course.

While we will be using this software for the examples, the course is primarily about meta-analysis rather than this specific software, and can be applied regardless of what software you employ in the future.

We will also use ChatGPT, which is free, for some exercises when discussing common mistakes in meta-analysis.

COURSE SCHEDULE

Module 1. Basic concepts + Meta-Analysis where effect size is consistent across studies June 1-4

Videos and zoom discussions

Canon Analysis

We start with a simple meta-analysis to compare the relative utility of two treatments for preventing cardiovascular events. I use this to outline the elements of a meta-analysis that we will be exploring in later modules. I also use this to show how we can perform a simple analysis from start to finish, including generating a high-resolution plot and writing a report.

Tamiflu Symptom Relief

This is an analysis of randomized controlled trials (RCTs) that compared the duration of flu symptoms in patients treated with Tamiflu to those treated with a placebo. This is the second example where the effect size is consistent across studies.

How a meta-analysis works.

I use a series of fictional studies to show what happens as we add studies to a meta-analysis. When the effect size is consistent across studies, we focus on the common effect size. When the effect size varies

across studies we estimate the mean effect size, but we also need to estimate the dispersion in effects and consider the implications of this dispersion for the utility of the intervention.

Fixed-effect vs. random-effects

Every meta-analysis must be based on a statistical model. The model tells us how the studies were sampled and how we can generalize from them to other studies or populations. We discuss how to select a model, and also how to avoid common mistakes related to this issue.

Learning Objectives

At the conclusion of this module students will be able to:

- Perform a meta-analysis where the effect-size is consistent across studies
- Understand how to choose a statistical model
- Report the results for a meta-analysis where the effect size is consistent
- Identify and avoid common mistakes related to choosing a statistical model

Module 2 Basic concepts + Meta-Analysis where effect size varies across studies June 5-8

Videos and zoom discussions

Effect size indices

The effect size is the unit of currency in a meta-analysis. We compute an effect size for each study. Then, we pool these values to estimate the common (or mean) effect size, and the dispersion in effects. We discuss how to choose an effect size index, and how to understand and explain the meaning of the effect size.

Effect size vs p-values

In primary studies that compare outcomes in two groups, there are two general approaches that researchers apply. One is to test the null hypothesis of no effect and report a p-value. The other is to estimate the effect size, and report that effect size along with a confidence interval. In this module I show why we almost always want to focus on the effect size approach. Then I extend this to meta-analysis, where it is imperative that we work with the effect size from each study rather than the p-value.

What studies to include

In any meta-analysis we can choose to work with a narrowly defined population and a specific variant of the intervention. In this case, we assume that the effect size will be reasonably consistent across studies, and our goal will be to estimate this effect size. Alternatively, we can choose to include an array of populations and/or variants of the intervention. In this case, our goal will be to assess the dispersion in effects and possibly to see what moderators are associated with this dispersion. We explain how to make these decisions and how to map them to the inclusion/exclusion criteria.

Heterogeneity

In most meta-analyses, the effect size varies from study to study. It's important to understand how much the effect size varies, and to consider the clinical or substantive implications of this variation. In

this module, I start by reviewing how we think about heterogeneity in a primary study. Then, I show that the same ideas apply in a meta-analysis.

The ADHD analysis

This is an analysis of seventeen randomized controlled trials (RCTs) that assessed the impact of methylphenidate on adults with ADHD (attention deficit hyperactivity disorder). This is an opportunity to develop a feel (in practice) for the concepts we earlier discussed in the abstract.

Learning Objectives

At the conclusion of this module students will be able to:

- Understand how to choose an effect-size index
- Understand how to determine what kinds of studies to include in an analysis
- Understand how to quantify heterogeneity in a meta-analysis
- Identify and avoid common mistakes

Module 3 How to perform and report a meta-analysis using continuous and dichotomous outcomes June 9-12

The I^2 statistic

The vast majority of meta-analyses use the I^2 statistic to quantify heterogeneity. Many use this statistic to quantify heterogeneity as being low, moderate or high. While this practice is ubiquitous it is nevertheless incorrect. I^2 is a proportion, not an absolute amount. It tells us what proportion of the variance is due to differences in the true effects rather than sampling error. It does not tell us how much variance there is. Classifications of heterogeneity based on I^2 are uninformative at best, and sometimes misleading. I explain this in detail, and then discuss the prediction interval, the statistic that provides the information that researchers believe (incorrectly) is being reported by I^2.

Since many students find this idea so hard to accept, I include a video clip from one of my in-person workshops where Julian Higgins (the co-creator of I^2) was kind enough to speak and confirm that he completely agrees with this point.

We will work through a series of analyses from start to finish. This will provide concrete examples of the topics introduced earlier and an opportunity for students to become comfortable performing these analyses on their own. The datasets include the following (among others)

The ADHD analysis

This is an analysis of studies where adults with ADHD (attention deficit hyperactivity disorder) were randomized to received either methylphenidate or a placebo. The studies looked at the impact of the treatment on the patients' cognitive functioning. We will use this analysis to discuss how to estimate the mean effect size and the dispersion in effects working with continuous outcomes.

The hypericum analysis

This is an analysis of studies where patients with severe depression were randomized randomized to received either hypericum or a placebo. The studies looked at the impact of the treatment on the

depression. We will use this analysis to discuss how to estimate the mean effect size and the dispersion in effects working with dichotomous outcomes.

Learning Objectives

In this module students will gain experience in all steps related to an analysis including

- How to enter data
- How to run the analysis
- How to understand and report the results
- How to properly quantify heterogeneity in effects
- How to create a high-resolution plot
- How to perform sensitivity analyses
- To avoid mistakes related to heterogeneity

Module 4 How to perform and report a meta-analysis using incidence, correlation, and risk difference June 13-15

Overview

We will work through a series of analyses from start to finish. This will provide concrete examples of the topics introduced earlier and an opportunity for students to become comfortable performing these analyses on their own. The datasets include the following (among others)

The Tocilizumab analysis

This is an analysis of observational studies of patients who had been hospitalized for Covid-19. In each study some patients were treated with tocilizumab and others were not. The analysis looked at the relationship between treatment and risk of death. The paper reported that patients treated with the drug were more likely to survive and that this relationship was statistically significant. However, we will reanalyze the data and show that the truth is more complicated. While the drug was associated with increased survival on average, it was actually associated with an increased risk of death in roughly 35% of studies. This analysis provides an example of why it is important to quantify heterogeneity properly and take account of the heterogeneity when assessing the potential utility of an intervention,

The mitral-valve analysis

While most meta-analyses work with studies that compared two groups, we can also use meta-analysis to work with studies that looked at the risk of an event in one group – such as incidence or prevalence. The mitral valve analysis is one example. This analysis looks at the risk of death following mitral-valve surgery in elderly patients.

Learning Objectives

In this module students will gain experience in all steps related to an analysis including

- How to enter data
- How to run the analysis
- How to understand and report the results

- How to properly quantify heterogeneity in effects
- How to create a high-resolution plot
- How to perform sensitivity analyses
- To avoid mistakes related to heterogeneity

Module 5 Subgroup Analyses – Comparing the effect size in different sets of studies June 16-19

Overview

In the basic course we introduced a number of meta-analyses where the effect size varied across studies, and we learned how to estimate the mean effect size and the heterogeneity in effects. In the advanced course we return to those data sets and learn how to explain some of that heterogeneity using subgroup analysis and meta-regression (analogous to ANOVA and multiple regression). We also learn how to assess the potential impact of publication bias in these analyses.

Comparing subgroups

In a primary study we may use a t-test or a one-way analysis of variance to compare the mean score in two or more groups of people. In a meta-analysis we may use a subgroup analysis to compare the mean effect size in two or more subgroups of studies.

- We will revisit the ADHD analysis to look at the relationship between dose and patient type
- We will revisit the Mitral-valve analysis to compare the risk for studies that employed different variants of the procedure
- We will revisit the hypericum analysis to compare the effect size for different types of patients

Learning Objectives

In this module students will learn

- How to run a subgroups analysis
- How to select a statistical model for a subgroups analysis
- How to estimate Tau-squared in the presence of subgroups
- That subgroup comparisons are observational
- How to report the results of a subgroup analysis
- How to avoid common mistakes when comparing subgroups

Module 6 Meta-regression – Assessing the relationship of covariates with effect size June 20-23

Meta-regression

In a primary study we may use multiple regression to assess the relationship between covariates and scores. In a meta-analysis we may use meta-regression to assess the relationship between covariates and effect sizes

- We will revisit the ADHD analysis to look at the relationship between dose and effect sizes
- We will revisit the Mitral-valve analysis to look for covariates related to the risk of death
- We will revisit the hypericum analysis to look for covariates related to the effect size.

In each case we will isolate the unique impact of specific factors while controlling for potential confounds. We will discuss the meaning of T2 and prediction intervals in the presence of covariates. We will discuss the meaning and limitations of the R2 analog. We will discuss how to plot the impact of continuous and categorical predictors, and interactions

Learning Objectives

In this module students will learn

- How to perform a meta-regression in CMA
- How to work with continuous covariates
- How to work with categorical covariates
- How to work with sets of covariates
- How to understand the I-sq analog
- How to understand the R-sq analog
- How to plot results of a meta-regression
- How to assess the unique impact of covariates
- How to use the Knapp-Hartung adjustment in meta-regression
- How to avoid common mistakes in meta-regression

Module 7 Working with complex data sets June 24-26

Complex data sets

A meta-analysis may include studies that report data for complex data sets. This may include data for two or more independent subgroups. It may include data for two or more outcomes or timepoints based on the same people. It may include studies that reported data for two or more treatments vs. a common control group.

Learning Objectives

In this module students will learn

- How to work with studies that present data for two or more independent subgroups
- How to work with studies that present data for two or more outcomes based on the same subjects
- How to work with studies that present data for two or more timepoints based on the same subjects
- How to work with studies that present data for studies that reported data for two or more treatments vs. a common control group.

Module 8 Publication bias and other issues June 27-30

Publication bias

Most meta-analyses are based on studies pulled from the literature. Since studies are more likely to be published if they report statistically significant results, there is a tendency for the literature to incorporate a biased subset of studies. It follows that a meta-analysis based on these studies will tend to overestimate the magnitude of the effect. We discuss how to address this problem.

Limitations of the random-effects model

While meta-analysis provides an exceptional tool, it is imperative that we understand the limitations of this tool. We will discuss the limitations of the random-effects model. In particular, we will discuss limitations that apply when we are working with a small number of studies, and how to deal with these cases.

Learning Objectives

In this module students will learn

- How to assess the potential impact of publication bias
- How to generalize from the studies in the analysis to a wider universe of studies
- How to work with analyses when there are only a small number of studies
- How to use the Knapp-Hartung adjustment
- When it makes sense to perform a meta-analysis
- How to avoid common mistakes associated with these issues