

Description

- The `meta` command performs meta-analysis. In a nutshell, you can do the following:
1. Compute or specify effect sizes; see `[META] meta esize` and `[META] meta set`.
  2. Summarize meta-analysis data; see `[META] meta summarize` and `[META] meta forestplot`.
  3. Examine heterogeneity and perform meta-regression; see `[META] meta galbraithplot`, `[META] meta labbeplot`, and `[META] meta regress`.
  4. Explore small-study effects and publication bias; see `[META] meta funnelplot`, `[META] meta bias`, and `[META] meta trimfill`.
  5. Perform multivariate meta-regression; see `[META] meta mvregress`.
  6. Perform multilevel meta-regression; see `[META] meta meregress` and `[META] meta multilevel`.

For software-free introduction to meta-analysis, see `[META] Intro`.

Declare, update, and describe `meta data`

<code>meta data</code>	Declare meta-analysis data
<code>meta esize</code>	Compute effect sizes and declare meta data
<code>meta set</code>	Declare meta data using precalculated effect sizes
<code>meta update</code>	Update current settings of meta data
<code>meta query</code>	Describe current settings of meta data
<code>meta clear</code>	Clear current settings of meta data

Summarize meta data by using a table

<code>meta summarize</code>	Summarize meta-analysis data
<code>meta summarize, subgroup()</code>	Perform subgroup meta-analysis
<code>meta summarize, cumulative()</code>	Perform cumulative meta-analysis
<code>meta summarize, leaveoneout</code>	Perform leave-one-out meta-analysis

Summarize meta data by using a forest plot

<code>meta forestplot</code>	Produce meta-analysis forest plots
<code>meta forestplot, subgroup()</code>	Produce subgroup meta-analysis forest plots
<code>meta forestplot, cumulative()</code>	Produce cumulative meta-analysis forest plots
<code>meta forestplot, leaveoneout</code>	Produce leave-one-out meta-analysis forest plots

Explore heterogeneity and perform meta-regression

<code>meta galbraithplot</code>	Produce Galbraith plots
<code>meta labbeplot</code>	Produce L'Abbé plots for two-group comparison of binary outcomes
<code>meta regress</code>	Fit meta-regression
<code>estat bubbleplot</code>	Produce bubble plots after meta-regression

Explore and address small-study effects (funnel-plot asymmetry, publication bias)

<code>meta funnelplot</code>	Produce funnel plots
<code>meta funnelplot, contours()</code>	Produce contour-enhanced funnel plots
<code>meta bias</code>	Test for small-study effects or funnel-plot asymmetry
<code>meta trimfill</code>	Perform trim-and-fill analysis of publication bias

Perform multivariate meta-regression

<code>meta mvregress</code>	Fit multivariate meta-regression
<code>estat heterogeneity (mv)</code>	Assess heterogeneity in multivariate meta-regression

`meta mvregress` does not require your dataset to be `meta set`.

Perform multilevel meta-regression

<code>meta meregress</code>	Fit multilevel meta-regression
<code>meta multilevel</code>	Fit random-intercepts multilevel meta-regression
<code>estat heterogeneity (me)</code>	Assess heterogeneity in multilevel meta-regression

`meta meregress` and `meta multilevel` do not require your dataset to be `meta set`.

## Remarks and examples

This entry describes Stata's suite of commands, `meta`, for performing meta-analysis. For a software-free introduction to meta-analysis, see [\[META\] Intro](#).

Remarks are presented under the following headings:

*Introduction to meta-analysis using Stata*

*Example datasets*

*Effects of teacher expectancy on pupil IQ (pupiliq.dta)*

*Effect of streptokinase after a myocardial infarction (strepto.dta)*

*Efficacy of BCG vaccine against tuberculosis (bcg.dta)*

*Effectiveness of nonsteroidal anti-inflammatory drugs (nsaids.dta)*

*Treatment of moderate periodontal disease (periodontal.dta)*

*Tour of meta-analysis commands*

*Prepare your data for meta-analysis in Stata*

*Basic meta-analysis summary*

*Subgroup meta-analysis*

*Cumulative meta-analysis*

*Heterogeneity: Galbraith plot, meta-regression, and bubble plot*

*Funnel plots for exploring small-study effects*

*Testing for small-study effects*

*Trim-and-fill analysis for addressing publication bias*

*Multivariate meta-regression*

*Multilevel meta-regression*

## Introduction to meta-analysis using Stata

Stata's `meta` command offers full support for meta-analysis from computing various effect sizes and producing basic meta-analytic summary and forest plots to accounting for between-study heterogeneity and potential publication bias. Random-effects, common-effect, and fixed-effects meta-analyses are supported.

Standard effect sizes, such as the log odds-ratio for a two-group comparison of binary outcomes, Hedges's  $g$  for a two-group comparison of continuous outcomes, the Fisher's  $z$ -transformed correlation for correlation data, or the Freeman–Tukey-transformed proportion for estimating a single proportion (prevalence), may be computed using the `meta esize` command; see [\[META\] meta esize](#). Generic (pre-calculated) effect sizes may be specified by using the `meta set` command; see [\[META\] meta set](#).

`meta esize` and `meta set` are part of the meta-analysis declaration step, which is the first step of meta-analysis in Stata. During this step, you specify the main information about your meta-analysis such as the study-specific effect sizes and their corresponding standard errors and the meta-analysis model and method. This information is then automatically used by all subsequent `meta` commands for the duration of your meta-analysis session. You can use `meta update` to easily update some of the specified information during the session; see [\[META\] meta update](#). And you can use `meta query` to remind yourself about the current [meta settings](#) at any point of your meta-analysis; see [\[META\] meta update](#). For more information about the declaration step, see [\[META\] meta data](#). Also see [Prepare your data for meta-analysis in Stata](#).

Random-effects, common-effect, and fixed-effects meta-analysis models are supported. You can specify them during the declaration step and use the same model throughout your meta-analysis or you can specify a different model temporarily with any of the `meta` commands. You can also switch to a different model for the rest of your meta-analysis by using `meta update`. See [Declaring a meta-analysis model](#) in [\[META\] meta data](#) for details.

Traditionally, meta-analysis literature and software used the term “fixed-effect model” (notice singular effect) to refer to the model that assumes a common effect for all studies. To avoid potential confusion with the term “fixed-effects model” (notice plural effects), which is commonly used in various disciplines to refer to the model whose effects vary from one group to another, we adopted the terminology from [Rice, Higgins, and Lumley \(2018\)](#) of the “common-effect model”. This terminology is also reflected in the option names for specifying the corresponding models with meta commands: `common` specifies a common-effect model and `fixed` specifies a fixed-effects model. (Similarly, `random` specifies a random-effects model.) Although the overall effect-size estimates from the common-effect and fixed-effects models are computationally identical, their interpretation is different. We provide the two options to emphasize this difference and to encourage proper interpretation of the final results given the specified model. See [common-effect versus fixed-effects models](#) in [\[META\] meta data](#) and [Meta-analysis models](#) in [\[META\] Intro](#) for more information.

Depending on the chosen meta-analysis model, various estimation methods are available: inverse-variance and Mantel–Haenszel for the common-effect and fixed-effects models and seven different estimators for the between-study variance parameter for the random-effects model. See [Declaring a meta-analysis estimation method](#) in [\[META\] meta data](#).

Also see [Default meta-analysis model and method](#) in [\[META\] meta data](#) for the default model and method used by the meta commands.

Results of a basic meta-analysis can be summarized numerically in a table by using `meta summarize` (see [\[META\] meta summarize](#)) or graphically by using forest plots; see [\[META\] meta forestplot](#). See [Basic meta-analysis summary](#).

To evaluate the trends in the estimates of the overall effect sizes, you can use the `cumulative()` option with `meta summarize` or `meta forestplot` to perform cumulative meta-analysis. See [Cumulative meta-analysis](#).

In the presence of [subgroup heterogeneity](#), you can use the `subgroup()` option with `meta summarize` or `meta forestplot` to perform single or multiple subgroup analyses. See [Subgroup meta-analysis](#).

Heterogeneity can also be explored by fitting meta-regression using the `meta regress` command; see [\[META\] meta regress](#). After meta-regression, you can produce bubble plots (see [\[META\] estat bubbleplot](#)) and perform other postestimation analysis (see [\[META\] meta regress postestimation](#)). Also see [Heterogeneity: Galbraith plot, meta-regression, and bubble plot](#).

In addition to forest plots, you can also visually explore heterogeneity using `meta galbraithplot`, which works with any type of data (see [\[META\] meta galbraithplot](#)), and `meta labbeplot`, which works with a two-group comparison of binary outcomes (see [\[META\] meta labbeplot](#)).

Publication bias, or more accurately, [small-study effects](#) or funnel-plot asymmetry, may be explored graphically via standard or contour-enhanced funnel plots (see [\[META\] meta funnelplot](#)). Regression-based and other tests for detecting small-study effects are available with the `meta bias` command; see [\[META\] meta bias](#). The trim-and-fill method for assessing the potential impact of publication bias on the meta-analysis results is implemented in the `meta trimfill` command; see [\[META\] meta trimfill](#). See [Funnel plots for exploring small-study effects](#), [Testing for small-study effects](#), and [Trim-and-fill analysis for addressing publication bias](#).

Multivariate meta-regression can be fit via `meta mvregress` (see [\[META\] meta mvregress](#)). After multivariate meta-regression, you can explore heterogeneity using `estat heterogeneity` (see [\[META\] estat heterogeneity \(mv\)](#)) and conduct other postestimation analysis (see [\[META\] meta mvregress postestimation](#)).

Multilevel meta-regression can be fit via `meta meregress` (see [META] [meta meregress](#)) or `meta multilevel` ([META] [meta multilevel](#)). After multilevel meta-regression, you can explore multilevel heterogeneity using `estat heterogeneity` (see [META] [estat heterogeneity \(me\)](#)) and conduct other postestimation analysis (see [META] [meta me postestimation](#)).

## Example datasets

We present several datasets that we will use throughout the documentation to demonstrate the meta suite. Feel free to skip over this section to [Tour of meta-analysis commands](#) and come back to it later for specific examples.

Example datasets are presented under the following headings:

*Effects of teacher expectancy on pupil IQ (pupiliq.dta)*  
*Effect of streptokinase after a myocardial infarction (strepto.dta)*  
*Efficacy of BCG vaccine against tuberculosis (bcg.dta)*  
*Effectiveness of nonsteroidal anti-inflammatory drugs (nsaids.dta)*  
*Treatment of moderate periodontal disease (periodontal.dta)*

### Effects of teacher expectancy on pupil IQ (pupiliq.dta)

This example describes a well-known study of [Rosenthal and Jacobson \(1968\)](#) that found the so-called Pygmalion effect, in which expectations of teachers affected outcomes of their students. A group of students was tested and then divided randomly into experimentals and controls. The division may have been random, but the teachers were told that the students identified as experimentals were likely to show dramatic intellectual growth. A few months later, a test was administered again to the entire group of students. The experimentals outperformed the controls.

Subsequent researchers attempted to replicate the results, but many did not find the hypothesized effect.

[Raudenbush \(1984\)](#) did a meta-analysis of 19 studies and hypothesized that the Pygmalion effect might be mitigated by how long the teachers had worked with the students before being told about the non-existent higher expectations for the randomly selected subsample of students. We explore this hypothesis in [Subgroup meta-analysis](#).

The data are saved in `pupiliq.dta`. Below, we describe some of the variables that will be used in later analyses.

```
. use https://www.stata-press.com/data/r19/pupiliq
(Effects of teacher expectancy on pupil IQ)

. describe

Contains data from https://www.stata-press.com/data/r19/pupiliq.dta
Observations:      19              Effects of teacher expectancy
                                on pupil IQ
                                24 Apr 2024 08:28
Variables:         14              (_dta has notes)
```

Variable name	Storage type	Display format	Value label	Variable label
study	byte	%9.0g		Study number
author	str20	%20s		Author
year	int	%9.0g		Publication year
nexper	int	%9.0g		Sample size in experimental group
ncontrol	int	%9.0g		Sample size in control group
stdmdiff	double	%9.0g		Standardized difference in means
weeks	byte	%9.0g		Weeks of prior teacher-student contact
catweek	byte	%9.0g	catwk	Weeks of prior contact (categorical)
week1	byte	%9.0g	catweek1	Prior teacher-student contact > 1 week
se	double	%10.0g		Standard error of stdmdiff
se_c	float	%9.0g		se from Pubbias book, p.322
setting	byte	%8.0g	testtype	Test setting
tester	byte	%8.0g	tester	Tester (blind or aware)
studylbl	str26	%26s		Study label

Sorted by:

Variables `stdmdiff` and `se` contain the effect sizes (standardized mean differences between the experimental and control groups) and their standard errors, respectively. Variable `weeks` records the number of weeks of prior contact between the teacher and the students. Its dichotomized version, `week1`, records whether the teachers spent more than one week with the students (high-contact group, `week1` = 1) or one week and less (low-contact group, `week1` = 0) prior to the experiment.

We perform basic meta-analysis summary of this dataset in [Basic meta-analysis summary](#) and explore the between-study heterogeneity of the results with respect to the amount of the teacher–student contact in [Subgroup meta-analysis](#).

This dataset is also used in [Examples of using meta summarize](#) of [META] [meta summarize](#), example 5 of [META] [meta forestplot](#), example 8 of [META] [meta funnelplot](#), and [Examples of using meta bias](#) of [META] [meta bias](#).

See [example 1](#) for the declaration of the `pupiliq.dta`. You can also use its predeclared version, `pupiliqset.dta`.

Effect of streptokinase after a myocardial infarction (`strepto.dta`)

Streptokinase is a medication used to break down clots. In the case of myocardial infarction (heart attack), breaking down clots reduces damage to the heart muscle.

Lau et al. (1992) conducted a meta-analysis of 33 studies performed between 1959 and 1988. These studies were of heart attack patients who were randomly treated with streptokinase or a placebo. Lau et al. (1992) introduced cumulative meta-analysis to investigate the time when the effect of streptokinase became statistically significant. Studies were ordered by time, and as each was added to the analysis, standard meta-analysis was performed. See *Cumulative meta-analysis* for details.

The data are saved in `strepto.dta`.

```
. use https://www.stata-press.com/data/r19/strepto
(Effect of streptokinase after a myocardial infarction)

. describe

Contains data from https://www.stata-press.com/data/r19/strepto.dta
Observations:      33      Effect of streptokinase after a
                        myocardial infarction
Variables:         7      14 May 2024 18:24
                        (_dta has notes)
```

Variable name	Storage type	Display format	Value label	Variable label
study	str12	%12s		Study name
year	int	%10.0g		Publication year
ndeadt	int	%10.0g		Number of deaths in treatment group
nsurvt	int	%9.0g		Number of survivors in treatment group
ndeadc	int	%10.0g		Number of deaths in control group
nsurvc	int	%9.0g		Number of survivors in control group
studyplus	str13	%13s		Study label for cumulative MA

Sorted by:

The outcome of interest was death from myocardial infarction. Variables `ndeadt` and `nsurvt` contain the numbers of deaths and survivals, respectively, in the treatment group and `ndeadc` and `nsurvc` contain those in the control (placebo) group.

See [example 5](#) for the declaration of the `strepto.dta`. You can also use its predeclared version, `streptoset.dta`.

### Efficacy of BCG vaccine against tuberculosis (`bcg.dta`)

BCG vaccine is a vaccine used to prevent tuberculosis (TB). The vaccine is used worldwide. Efficacy has been reported to vary. Colditz et al. (1994) performed meta-analysis on the efficacy using 13 studies—all randomized trials—published between 1948 and 1980. The dataset, shown below, has been studied by, among others, Berkey et al. (1995), who hypothesized that the latitude of the study location might explain the variations in efficacy. We explore this via meta-regression in *Heterogeneity: Galbraith plot, meta-regression, and bubble plot*.

The data are saved in `bcg.dta`. Below, we describe some of the variables we will use in future analyses.

```
. use https://www.stata-press.com/data/r19/bcg
(Efficacy of BCG vaccine against tuberculosis)

. describe

Contains data from https://www.stata-press.com/data/r19/bcg.dta
Observations:      13      Efficacy of BCG vaccine against
                        tuberculosis
Variables:         11      1 May 2024 14:40
                        (_dta has notes)
```

Variable name	Storage type	Display format	Value label	Variable label
trial	byte	%9.0g		Trial number
trialloc	str14	%14s		Trial location
author	str21	%21s		Author
year	int	%9.0g		Publication year
npost	int	%9.0g		Number of TB positive cases in treated group
nnegt	long	%9.0g		Number of TB negative cases in treated group
nposc	int	%9.0g		Number of TB positive cases in control group
nnegc	long	%9.0g		Number of TB negative cases in control group
latitude	byte	%9.0g		Absolute latitude of the study location (in degrees)
alloc	byte	%10.0g	alloc	Method of treatment allocation
study1bl	str27	%27s		Study label

Sorted by: trial

Variables `npost` and `nnegt` contain the numbers of positive and negative TB cases, respectively, in the treatment group (vaccinated group) and `nposc` and `nnegc` contain those in the control group. Variable `latitude` records the latitude of the study location, which is a potential moderator for the vaccine efficacy. Studies are identified by `study1bl`, which records the names of the authors and the year of the publication for each study.

This dataset is also used in [example 3](#) of [\[META\] meta data](#), [Examples of using meta forestplot](#) of [\[META\] meta forestplot](#), [example 1](#) of [\[META\] meta galbraithplot](#), [example 1](#) of [\[META\] meta labbeplot](#), [Examples of using meta regress](#) of [\[META\] meta regress](#), [Remarks and examples](#) of [\[META\] meta regress postestimation](#), and [Examples of using estat bubbleplot](#) of [\[META\] estat bubbleplot](#).

See [example 7](#) for the declaration of the `bcg.dta`. You can also use its predeclared version, `bcgset.dta`.

Effectiveness of nonsteroidal anti-inflammatory drugs (`nsaids.dta`)

Strains and sprains cause pain, and nonsteroidal anti-inflammatory drugs (NSAIDS) are used to treat it. How well do they work? People who study such things define success as a 50-plus percent reduction in pain. [Moore et al. \(1998\)](#) performed meta-analysis of 37 randomized trials that looked into successful pain reduction via NSAIDS. Following their lead, we will explore publication bias or, more generally, small-study effects in these data. See [Funnel plots for exploring small-study effects](#), [Testing for small-study effects](#), and [Trim-and-fill analysis for addressing publication bias](#).



The data are saved in `nsaids.dta`.

```
. use https://www.stata-press.com/data/r19/nsaids
(Effectiveness of nonsteroidal anti-inflammatory drugs)
. describe
Contains data from https://www.stata-press.com/data/r19/nsaids.dta
Observations:      37      Effectiveness of nonsteroidal
                        anti-inflammatory drugs
Variables:         5      24 Apr 2024 17:09
                        (_dta has notes)
```

Variable name	Storage type	Display format	Value label	Variable label
study	byte	%8.0g		Study ID
nstreat	byte	%8.0g		Number of successes in the treatment arm
nftreat	byte	%9.0g		Number of failures in the treatment arm
nscontrol	byte	%8.0g		Number of successes in the control arm
nfcontrol	byte	%9.0g		Number of failures in the control arm

Sorted by:

Variables `nstreat` and `nftreat` contain the numbers of successes and failures, respectively, in the experimental group and `nscontrol` and `nfcontrol` contain those in the control group.

This dataset is also used in *Examples of using meta funnelplot* of [META] [meta funnelplot](#) and *example 3* of [META] [meta bias](#).

See [example 11](#) for the declaration of the `nsaids.dta`. You can also use its predeclared version, `nsaidsset.dta`.

**Treatment of moderate periodontal disease (`periodontal.dta`)**

Periodontal disease is the inflammation of the gum that may destroy the bone supporting the teeth. [Antczak-Bouckoms et al. \(1993\)](#) investigated five randomized controlled trials that explored the impact of two procedures, surgical and nonsurgical, on treating periodontal disease. This dataset was also analyzed by [Berkey et al. \(1998\)](#). Subjects’ mouths were split into sections. These sections were randomly allocated to the two treatment procedures. For each patient, at least one section was treated surgically and at least one other section was treated nonsurgically. Two outcomes (effect sizes) of interest were mean improvements from baseline (pretreatment) in probing depth (`y1`) and attachment level (`y2`) around the teeth.

The data are saved in `periodontal.dta`.

```
. use https://www.stata-press.com/data/r19/periodontal
(Treatment of moderate periodontal disease)

. describe
Contains data from https://www.stata-press.com/data/r19/periodontal.dta
Observations:      5      Treatment of moderate
                        periodontal disease
Variables:         9      13 Jan 2025 18:11
                        (_dta has notes)
```

Variable name	Storage type	Display format	Value label	Variable label
trial	str23	%23s		Trial label
pubyear	byte	%9.0g		Publication year centered at 1983
y1	float	%6.2f		Mean improvement in probing depth (mm)
y2	float	%6.2f		Mean improvement in attachment level (mm)
v11	float	%6.4f		Variance of y1
v12	float	%6.4f		Covariance of y1 and y2
v22	float	%6.4f		Variance of y2
s1	double	%10.0g		Standard error of y1
s2	double	%10.0g		Standard error of y2

Sorted by:

Other variables of interest that will be used in [example 15](#) are the year of publication (`pubyear`) and three variables defining the within-study covariance matrix for each study: `v11`, `v12`, and `v22`.

This dataset is also used in [Examples of using meta mvregress](#) of [\[META\] meta mvregress](#).

Tour of meta-analysis commands

In this section, we provide a tour of Stata’s meta-analysis (`meta`) commands with applications to several real-world datasets. We demonstrate the basic meta-analysis summary and a forest plot and explore heterogeneity via subgroup analysis using the [pupil IQ dataset](#). We then demonstrate cumulative meta-analysis using the [streptokinase dataset](#). We continue with more heterogeneity analyses of the [BCG dataset](#). Finally, we explore and address publication bias for the [NSAIDS dataset](#).

Examples are presented under the following headings:

- [Prepare your data for meta-analysis in Stata](#)
- [Basic meta-analysis summary](#)
- [Subgroup meta-analysis](#)
- [Cumulative meta-analysis](#)
- [Heterogeneity: Galbraith plot, meta-regression, and bubble plot](#)
- [Funnel plots for exploring small-study effects](#)
- [Testing for small-study effects](#)
- [Trim-and-fill analysis for addressing publication bias](#)
- [Multivariate meta-regression](#)
- [Multilevel meta-regression](#)

Prepare your data for meta-analysis in Stata

The first step of meta-analysis in Stata is to declare your data as `meta data`. During this step, we specify the main information needed for meta-analysis such as effect sizes and their standard errors. We declare this information once by using either `meta set` or `meta esize`, and it is then used by all `meta`

commands. If needed, we can update our initial settings throughout the meta-analysis session by using `meta update`. The declaration step helps minimize potential mistakes and typing; see [META] `meta data` for details.

► Example 1: Set up your data for meta-analysis in Stata

Consider the pupil IQ dataset described in *Effects of teacher expectancy on pupil IQ (pupiliq.dta)*.

```
. use https://www.stata-press.com/data/r19/pupiliq
(Effects of teacher expectancy on pupil IQ)
. describe studylbl stdmdiff se week1
```

Variable name	Storage type	Display format	Value label	Variable label
studylbl	str26	%26s		Study label
stdmdiff	double	%9.0g		Standardized difference in means
se	double	%10.0g		Standard error of stdmdiff
week1	byte	%9.0g	catweek1	Prior teacher-student contact > 1 week

First, we prepare our data for use with meta commands. The dataset contains precomputed effect sizes, standardized mean differences stored in variable `stdmdiff`, and their standard errors stored in variable `se`. We will use `meta set` to declare these data. (If we needed to compute the individual effect sizes and their standard errors from the available `summary data`, we would have used [META] `meta esize`.)

We specify the effect sizes `stdmdiff` and their standard errors `se` with `meta set`. We also specify the variable that contains the study labels in the `studylabel()` option and the effect-size label in the `eslabel()` option. These are optional but useful for displaying the study and effect-size labels instead of generic study numbers and the generic label `Effect size`.

```
. meta set stdmdiff se, studylabel(studylbl) eslabel(Std. mean diff.)
Meta-analysis setting information

Study information
  No. of studies: 19
  Study label: studylbl
  Study size: N/A
Effect size
  Type: <generic>
  Label: Std. mean diff.
  Variable: stdmdiff
Precision
  Std. err.: se
  CI: [_meta_cil, _meta_ciu]
  CI level: 95%
Model and method
  Model: Random effects
  Method: REML
```

The header reports that there are  $K = 19$  studies in the meta-analysis and which variables contain the study labels, the effect sizes, and the standard errors. The output also shows that we will be using the random-effects model with the REML estimation method for our meta-analysis. This can be changed by specifying options with either `meta set` or the `meta` command of interest; see *Declaring a meta-analysis model* in [META] `meta data`.

meta.set creates some system variables beginning with `_meta_` and stores some data characteristics. For example, the system variables `_meta_cil` and `_meta_ciu` store the lower and upper limits of the CIs for the effect sizes. See [System variables](#) in [\[META\] meta data](#) for details.

See [\[META\] meta set](#) for more information about the command. o



Basic meta-analysis summary

In this section, we focus on basic meta-analysis summary by using [\[META\] meta summarize](#) and [\[META\] meta forestplot](#). See [Introduction of \[META\] meta summarize](#) and [Overview of \[META\] meta forestplot](#) for an overview of the meta-analysis summary and forest plots.

Example 2: Meta-analysis summary

Continuing with [example 1](#), we use [meta summarize](#) to combine the studies and estimate the overall effect size.

```
. meta summarize
Effect-size label: Std. mean diff.
Effect size: stdmdiff
Std. err.: se
Study label: studylbl

Meta-analysis summary
Random-effects model
Method: REML

Number of studies = 19
Heterogeneity:
    tau2 = 0.0188
    I2 (%) = 41.84
    H2 = 1.72
```

Effect size: Std. mean diff.					
	Study	Effect size	[95% conf. interval]		% weight
	Rosenthal et al., 1974	0.030	-0.215	0.275	7.74
	Conn et al., 1968	0.120	-0.168	0.408	6.60
	Jose & Cody, 1971	-0.140	-0.467	0.187	5.71
	Pellegrini & Hicks, 1972	1.180	0.449	1.911	1.69
	Pellegrini & Hicks, 1972	0.260	-0.463	0.983	1.72
	Evans & Rosenthal, 1969	-0.060	-0.262	0.142	9.06
	Fielder et al., 1971	-0.020	-0.222	0.182	9.06
	Claiborn, 1969	-0.320	-0.751	0.111	3.97
	Kester, 1969	0.270	-0.051	0.591	5.84
	Maxwell, 1970	0.800	0.308	1.292	3.26
	Carter, 1970	0.540	-0.052	1.132	2.42
	Flowers, 1966	0.180	-0.257	0.617	3.89
	Keshock, 1970	-0.020	-0.586	0.546	2.61
	Henrikson, 1970	0.230	-0.338	0.798	2.59
	Fine, 1972	-0.180	-0.492	0.132	6.05
	Grieger, 1970	-0.060	-0.387	0.267	5.71
	Rosenthal & Jacobson, 1968	0.300	0.028	0.572	6.99
	Fleming & Anttonen, 1971	0.070	-0.114	0.254	9.64
	Ginsburg, 1970	-0.070	-0.411	0.271	5.43
	theta	0.084	-0.018	0.185	

```
Test of theta = 0: z = 1.62
Test of homogeneity: Q = chi2(18) = 35.83

Prob > |z| = 0.1052
Prob > Q = 0.0074
```

The output from the standard meta-analysis summary includes heterogeneity statistics, the individual and overall effect sizes, and other information. The estimate of the overall effect size  $\theta$  is reported at the bottom of the table and labeled as `theta`. It is computed as the weighted average of study-specific effect sizes (standardized mean differences in our example). For these data, the overall estimate is 0.084 with a 95% CI of  $[-0.018, 0.185]$ . The significance test of  $H_0: \theta = 0$  is reported below the table and has a  $p$ -value of 0.1052, which suggests that the overall effect size is not statistically significantly different from zero.

We should be careful with our inferential conclusions about  $\theta$  because of the presence of between-study heterogeneity, as indicated, for instance, by the homogeneity test of  $H_0: \theta_1 = \theta_2 = \dots = \theta_{19} = \theta$  reported following the significance test. Its  $Q$  test statistic is 35.83 with a  $p$ -value of 0.0074, from which we can infer that there is significant heterogeneity between the individual studies.

The presence of heterogeneity among studies can be inferred also from the heterogeneity statistics reported in the header. For instance,  $I^2 = 41.84$  indicates that about 42% of the variability in the effect-size estimates is due to the differences between studies. The between-study heterogeneity must be addressed before final meta-analytic conclusions; see [Subgroup meta-analysis](#).

The table also reports the study-specific effect-sizes and their corresponding 95% CIs, but this information can be suppressed, if desired, by specifying the `nostudies` option.


See [\[META\] meta summarize](#) for details.



### ► Example 3: Forest plot

The results of meta-analysis are commonly displayed graphically using a forest plot. Continuing with [example 2](#), we can use `meta forestplot` to produce a meta-analysis forest plot for the pupil IQ data.

```
. meta forestplot
  Effect-size label: Std. mean diff.
    Effect size: stdmdiff
      Std. err.: se
    Study label: studylbl
```



Random-effects REML model

We obtain the same meta-analysis summary as with `meta summarize` in [example 2](#), but it is now displayed on a graph. In addition to the estimated values, the effect sizes are displayed graphically as blue squares centered at their estimates with areas proportional to the study weights and with horizontal lines or whiskers that represent the length of the corresponding CIs. The overall effect size is displayed as a green diamond with its width corresponding to the respective CI. (Notice that only the width and not the height of the diamond is relevant for the overall effect size.)

A forest plot provides an easy way to visually explore the agreement between the study-specific effect sizes and how close they are to the overall effect size. We can also spot the studies with large weights more easily by simply looking at the studies with large squares. In our example, the presence of between-study heterogeneity is evident—there are several studies whose effect-size estimates are very different from the overall estimate, and there are studies whose CIs do not even overlap.

See [\[META\] meta forestplot](#) for details.

## Subgroup meta-analysis

In [example 2](#) and [example 3](#), we established the presence of between-study heterogeneity in the pupil IQ dataset. Sometimes, the differences between studies may be explained by study-level covariates available in the data. When these covariates are categorical, we can perform meta-analysis separately for each category, which is known as subgroup meta-analysis; see [Subgroup meta-analysis](#) of [\[META\] Intro](#).

### ➤ Example 4: Subgroup meta-analysis

[Raudenbush \(1984\)](#) suspected that the amount of time the teachers spent with students before the experiment could impact their susceptibility to researchers' test results about children's intellectual abilities. If so, we would expect the effect sizes to be negatively associated with the amount of contact.

Continuing with [example 2](#), we see that the dataset contains a binary variable `week1` that records whether the teachers spend more than one week with children (high-contact group) or one week and less (low-contact group). Let's perform meta-analysis separately for each group. Under Raudenbush's hypothesis, we should expect to see larger effect sizes in the low-contact group and smaller effect sizes in the high-contact group.

We use the `subgroup()` option with `meta summarize` to perform a separate analysis for each group of `week1`.

```
. meta summarize, subgroup(week1)
Effect-size label: Std. mean diff.
Effect size: stdmdiff
Std. err.: se
Study label: studylbl
```

```
Subgroup meta-analysis summary
Random-effects model
Method: REML
Group: week1
```

Number of studies = 19

Effect size: Std. mean diff.

Study	Effect size	[95% conf. interval]	% weight	
Group: <= 1 week				
Pellegrini & Hicks, 1972	1.180	0.449	1.911	1.69
Pellegrini & Hicks, 1972	0.260	-0.463	0.983	1.72
Kester, 1969	0.270	-0.051	0.591	5.84
Maxwell, 1970	0.800	0.308	1.292	3.26
Carter, 1970	0.540	-0.052	1.132	2.42
Flowers, 1966	0.180	-0.257	0.617	3.89
Keshock, 1970	-0.020	-0.586	0.546	2.61
Rosenthal & Jacobson, 1968	0.300	0.028	0.572	6.99
theta	0.373	0.189	0.557	
Group: > 1 week				
Rosenthal et al., 1974	0.030	-0.215	0.275	7.74
Conn et al., 1968	0.120	-0.168	0.408	6.60
Jose & Cody, 1971	-0.140	-0.467	0.187	5.71
Evans & Rosenthal, 1969	-0.060	-0.262	0.142	9.06
Fielder et al., 1971	-0.020	-0.222	0.182	9.06
Claiborn, 1969	-0.320	-0.751	0.111	3.97
Henrikson, 1970	0.230	-0.338	0.798	2.59
Fine, 1972	-0.180	-0.492	0.132	6.05
Grieger, 1970	-0.060	-0.387	0.267	5.71
Fleming & Anttonen, 1971	0.070	-0.114	0.254	9.64
Ginsburg, 1970	-0.070	-0.411	0.271	5.43
theta	-0.021	-0.102	0.059	
Overall				
theta	0.084	-0.018	0.185	

Heterogeneity summary

Group	df	Q	P > Q	tau2	% I2	H2
<= 1 week	7	11.20	0.130	0.015	22.40	1.29
> 1 week	10	6.40	0.780	0.000	0.00	1.00
Overall	18	35.83	0.007	0.019	41.84	1.72

Test of group differences:  $Q_b = \chi^2(1) = 14.77$

Prob >  $Q_b = 0.000$




Indeed, if we look at the overall effect-size estimates for each group, the low-contact group has a larger estimate of 0.373 with a 95% CI of [0.189, 0.557], which suggests a statistically significant effect in this group, whereas the high-contact group has a smaller estimate of  $-0.021$  with a 95% CI of  $[-0.102, 0.059]$ , which suggests that the effect in this group is not different from 0 at a 5% significance level. Clearly, the amount of teacher contact with students has an impact on the meta-analysis results.

If we look at the heterogeneity summary reported following the main table, we will see that heterogeneity is reduced within each group. It is essentially nonexistent in the high-contact group and is much smaller (for instance,  $I^2 = 22\%$  versus the earlier  $I^2 = 42\%$ ) in the low-contact group.

The test of group differences (with  $Q_b = 14.77$  and the corresponding  $p$ -value of 0.000) reported at the bottom of the output also indicates that the group-specific overall effect sizes are statistically different.

We can also present the results of our subgroup analysis graphically by using the `subgroup()` option with `meta forest`:

```
. meta forestplot, subgroup(week1)
Effect-size label: Std. mean diff.
Effect size: stdmdiff
Std. err.: se
Study label: studylbl
```



It appears that stratifying our meta-analysis on the amount of prior contact between students and teachers explained most of the variability in the magnitudes of the effect sizes, at least in the high-contact group. ◀

When interpreting results from subgroup analysis, we should be mindful that the results are based on fewer studies and thus may not be as precise, in general.

See [META] [meta summarize](#) and [META] [meta forestplot](#).

## Cumulative meta-analysis

Cumulative meta-analysis performs multiple meta-analyses by accumulating studies one at a time after ordering them with respect to a variable of interest. This analysis is useful to monitor the trend in the estimates of the overall effect sizes with respect to some factor. For instance, it may be used to detect the time when the effect size of interest became significant.

### ➤ Example 5: Computing log odds-ratios using meta esize

Consider the streptokinase dataset described in [Effect of streptokinase after a myocardial infarction \(strepto.dta\)](#).

```
. use https://www.stata-press.com/data/r19/strepto, clear
(Effect of streptokinase after a myocardial infarction)

. describe
Contains data from https://www.stata-press.com/data/r19/strepto.dta
Observations:      33      Effect of streptokinase after a
                        myocardial infarction
Variables:          7      14 May 2024 18:24
                        (_dta has notes)
```

Variable name	Storage type	Display format	Value label	Variable label
study	str12	%12s		Study name
year	int	%10.0g		Publication year
ndeadt	int	%10.0g		Number of deaths in treatment group
nsurvt	int	%9.0g		Number of survivors in treatment group
ndeadc	int	%10.0g		Number of deaths in control group
nsurvc	int	%9.0g		Number of survivors in control group
studyplus	str13	%13s		Study label for cumulative MA

Sorted by:

As in [example 1](#), first we prepare our data for use with meta commands. Our dataset contains the summary data that represent the study-specific  $2 \times 2$  tables. The variables `ndeadt`, `nsurvt`, `ndeadc`, and `nsurvc` record the numbers of deaths and survivors in the treatment and control groups.

[Lau et al. \(1992\)](#) considered an odds ratio as the effect size of interest for these data. For odds ratios, meta-analysis is performed in the log metric. We can use `meta esize` to compute study-specific log odds-ratios and their corresponding standard errors and declare them for the subsequent meta-analysis. To compute log odds-ratios, we specify the four variables containing table cell counts with `meta esize`. As with `meta set` in [example 1](#), we specify the study labels in the `studylabel()` option with `meta esize`.

```
. meta esize ndeadt nsurvt ndeadc nsurvc, studylabel(studyplus) common
Meta-analysis setting information

Study information
  No. of studies: 33
  Study label: studyplus
  Study size: _meta_studysize
  Summary data: ndeadt nsurvt ndeadc nsurvc
  Effect size
    Type: lnoratio
    Label: Log odds-ratio
    Variable: _meta_es
  Zero-cells adj.: 0.5, only0
  Precision
    Std. err.: _meta_se
    CI: [_meta_cil, _meta_ciu]
    CI level: 95%
  Model and method
    Model: Common effect
    Method: Mantel-Haenszel
```

`meta esize` reports that there are 33 trials and that the computed effect size is log odds-ratio. This is the default effect size with a two-group comparison of binary outcomes. You can specify other effect sizes in the `esize()` option, which include a log risk-ratio, risk difference, and log Peto's odds-ratio. (After the declaration, you can use `meta update` to change the effect size more easily without having to respecify your summary data variables; see [\[META\] meta update](#).)

[Lau et al. \(1992\)](#) used a common-effect model with the Mantel–Haenszel method to perform their cumulative meta-analysis. We will follow their approach. Thus, we also specified the `common` option with `meta esize`. The command reported that the assumed meta-analysis model is a common-effect model. The Mantel–Haenszel estimation method is the default method for log odds-ratios under a common-effect model.


◀

## ► Example 6: Cumulative meta-analysis

After the data declaration in [example 5](#), we are ready to perform the cumulative meta-analysis. [Lau et al. \(1992\)](#) used cumulative meta-analysis to investigate the trends in the effect of the streptokinase drug used to prevent death after a myocardial infarction. We replicate their analysis below by producing a cumulative meta-analysis plot over the years for these data. Also see [Borenstein, Hedges, Higgins, and Rothstein \(2009\)](#) for the analysis of these data.

We use the `meta forestplot` command with the `cumulative()` option. We use the `or` option to display odds ratios instead of the default log odds-ratios. To match figure 1 in [Lau et al. \(1992\)](#) more closely, we also specify the `crop(0.5 .)` option to crop the lower CI limits and log odds-ratios estimates that are smaller than 0.5.

```
. meta forestplot, cumulative(year) or crop(0.5 .)
Effect-size label: Log odds-ratio
Effect size: _meta_es
Std. err.: _meta_se
Study label: studyplus
```



The cumulative meta-analysis forest plot displays the overall effect-size estimates and the corresponding CIs computed for the first study, for the first two studies, for the first three studies, and so on. The point estimates are represented by green circles, and the CIs are represented by the CI lines. The change in style and color of the plotted markers emphasizes that the (cumulative) overall effect sizes and not the study-specific effect sizes are being plotted.

The “+” sign in front of the study label we used for this analysis (variable `studyplus`) indicates that each subsequent study is being added to the previous ones for each analysis. In addition to the ordered values of the specified variable of interest (year in our example), the plot also displays the  $p$ -values corresponding to the tests of significance of the computed overall effect sizes.

For example, the cumulative odds ratio in the fourth row marked as +European 2 is 0.70 with a 95% CI of [0.52, 0.95] and a  $p$ -value of 0.023. So, based on the first four trials, the overall odds of death is roughly 30% less in the treatment group (treated with streptokinase) compared with the placebo group.

Notice that the first two odds-ratio estimates (and their lower CI limits) are smaller than 0.5. Because we used the `crop(0.5 .)` option, their values are not displayed on the graph. Instead, the arrowheads are displayed at the lower ends of the CI lines to indicate that the lower limits and the effect-size estimates are smaller than 0.5.

Borenstein, Hedges, Higgins, and Rothstein (2009) states that with the inclusion of additional trials in the cumulative meta-analysis, the overall effect sizes become more uniform because the chance of any new trial reporting a drastically different overall effect size is low. Also, the CIs become more narrow because the precision increases as more data become available.

If we look back at the plot, we will notice that starting from 1977, the overall effect size becomes (and stays) highly significant over the next decade of additional trials. Lau et al. (1992) and Borenstein et al. (2009, chap. 42) noted that if cumulative meta-analysis was used at that time to monitor the accumulated evidence from the trials, perhaps, the benefits from streptokinase could have been adopted in practice as early as 1977.

We can also obtain the same results as above but in a table by using `meta summarize`.

```
. meta summarize, cumulative(year) or
Effect-size label: Log odds-ratio
Effect size: _meta_es
Std. err.: _meta_se
Study label: studyplus
Cumulative meta-analysis summary          Number of studies =    33
Common-effect model
Method: Mantel-Haenszel
Order variable: year
```

Study	Odds ratio	[95% conf. interval]	p-value	year	
Fletcher	0.159	0.015	1.732	0.131	1959
+Dewar	0.345	0.104	1.141	0.081	1963
+European 1	0.952	0.514	1.760	0.874	1969
+European 2	0.702	0.517	0.951	0.023	1971
+Heikinheimo	0.776	0.589	1.023	0.072	1971
+Italian	0.806	0.624	1.040	0.097	1971
+Australian 1	0.796	0.632	1.004	0.054	1973
+Frankfurt 2	0.740	0.594	0.921	0.007	1973
+NHLBI SMIT	0.765	0.616	0.950	0.015	1974
+Frank	0.770	0.623	0.953	0.016	1975
+Valere	0.781	0.635	0.962	0.020	1975
+Klein	0.792	0.644	0.974	0.027	1976
+UK-Collab	0.809	0.670	0.979	0.029	1976
+Austrian	0.762	0.641	0.906	0.002	1977
+Australian 2	0.751	0.636	0.887	0.001	1977
+Lasierra	0.746	0.632	0.881	0.001	1977
+N Ger Collab	0.797	0.683	0.930	0.004	1977
+Witchitz	0.797	0.683	0.928	0.004	1977
+European 3	0.781	0.673	0.906	0.001	1979
+ISAM	0.793	0.690	0.910	0.001	1986
+GISSI-1	0.801	0.734	0.874	0.000	1986
+Olson	0.800	0.733	0.873	0.000	1986
+Baroffio	0.796	0.730	0.869	0.000	1986
+Schreiber	0.795	0.729	0.867	0.000	1986
+Cribier	0.795	0.729	0.868	0.000	1986
+Sainsous	0.794	0.728	0.866	0.000	1986
+Durand	0.793	0.727	0.865	0.000	1987
+White	0.787	0.721	0.858	0.000	1987
+Bassand	0.785	0.721	0.856	0.000	1987
+Vlay	0.785	0.720	0.856	0.000	1988
+Kennedy	0.783	0.718	0.853	0.000	1988
+ISIS-2	0.766	0.718	0.817	0.000	1988
+Wisenberg	0.765	0.717	0.816	0.000	1988

See [META] [meta summarize](#) and [META] [meta forestplot](#).



**Heterogeneity: Galbraith plot, meta-regression, and bubble plot**

The Galbraith plot ([Galbraith 1988](#)) is mainly used to assess heterogeneity of the studies and detect potential outliers. It may also be an alternative to forest plots for summarizing meta-analysis results, especially when there are many studies. See [META] [meta galbraithplot](#).

Meta-regression performs a weighted linear regression of effect sizes on moderators; see [META] [meta regress](#). With one moderator, the relationship between the effect sizes and the moderator may be further explored via a bubble plot after meta-regression; see [META] [estat bubbleplot](#).

In this section, we will demonstrate how to use Galbraith plots, meta-regression, and bubble plots to assess heterogeneity and examine relationships between effects sizes and moderators.

### ► Example 7: Computing log risk-ratios using meta esize

Consider the BCG dataset described in *Efficacy of BCG vaccine against tuberculosis (bcg.dta)*.

```
. use https://www.stata-press.com/data/r19/bcg, clear
(Efficacy of BCG vaccine against tuberculosis)
```

```
. describe studylbl npost nnegt nposc nnegc latitude
```

Variable name	Storage type	Display format	Value label	Variable label
studylbl	str27	%27s		Study label
npost	int	%9.0g		Number of TB positive cases in treated group
nnegt	long	%9.0g		Number of TB negative cases in treated group
nposc	int	%9.0g		Number of TB positive cases in control group
nnegc	long	%9.0g		Number of TB negative cases in control group
latitude	byte	%9.0g		Absolute latitude of the study location (in degrees)

As in [example 5](#), this dataset also records summary data for a two-group comparison of binary outcomes, so we will again use `meta esize` to compute our effect sizes.

In this example, our effect size of interest is a risk ratio. Just like with odds ratios, the meta-analysis of risk ratios is performed in the log metric, so we will be computing log risk-ratios.

```
. meta esize npost nnegt nposc nnegc, esize(lnrratio) studylabel(studylbl)
```

```
Meta-analysis setting information
```

```
Study information
```

```
No. of studies: 13
```

```
Study label: studylbl
```

```
Study size: _meta_studysize
```

```
Summary data: npost nnegt nposc nnegc
```

```
Effect size
```

```
Type: lnrratio
```

```
Label: Log risk-ratio
```

```
Variable: _meta_es
```

```
Zero-cells adj.: None; no zero cells
```

```
Precision
```

```
Std. err.: _meta_se
```

```
CI: [_meta_cil, _meta_ciu]
```

```
CI level: 95%
```

```
Model and method
```

```
Model: Random effects
```

```
Method: REML
```

Our specification of `meta esize` is similar to that from [example 5](#), except here we specify the `esize(lnrratio)` option to compute log risk-ratios instead of the default log odds-ratios.



The output indicates that there are  $K = 13$  studies in the meta-analysis and the default random-effects meta-analysis model (with the REML estimation method) will be used.


◀

Let's investigate the presence of heterogeneity in these data. For the purpose of illustration, we will do this using a Galbraith plot; see [META] [meta galbraithplot](#).

### ► Example 8: Galbraith plot

We use `meta galbraithplot` to produce a Galbraith plot for the BCG data.

```
. meta galbraithplot
Effect-size label: Log risk-ratio
Effect size: _meta_es
Std. err.: _meta_se
Model: Common effect
Method: Inverse-variance
```



The blue circles form a scatterplot of the study-specific standardized log risk-ratios against study precisions. Studies that are close to the  $y$  axis have low precision. Precision of studies increases as you move toward the right on the  $x$  axis.

The reference black line ( $y = 0$ ) represents the “no-effect” line. If a circle is above the reference line, the risk in the treatment group is higher than the risk in the control group for that study. Conversely, if a circle is below the line, the risk in the treatment group is lower than the risk in the control group.

The red line is the regression line through the origin. The slope of this line equals the estimate of the overall effect size. In the absence of substantial heterogeneity, we expect around 95% of the studies to lie within the 95% CI region (shaded area). In our example, there are 6 (out of 13) trials that are outside the CI region. We should suspect the presence of heterogeneity in these data, and we will investigate the reasons behind it in [example 9](#). For more interpretation of the above Galbraith plot, see [META] [meta galbraithplot](#).

◀

We have established that there is heterogeneity among the studies. Let's explore this further using meta-regression.

## ► Example 9: Meta-regression

As we discussed in *Subgroup meta-analysis*, when effect sizes vary greatly between different subgroups, one can perform separate meta-analysis on each subgroup to account for the between-study heterogeneity. But what if there is an association between the effect sizes and other study-level covariates or moderators that may be continuous? Meta-regression addresses this problem. Its goal is to investigate whether the differences between the effect sizes can be explained by one or more moderators. See *Introduction* of [META] *meta regress*.

The efficacy of the BCG vaccine against TB may depend on many factors such as the presence of environmental mycobacteria that provides some immunity to TB. [Berkey et al. \(1995\)](#) suggested that the distance of a study from the equator (the absolute latitude) may be used as a proxy for the presence of environmental mycobacteria and perhaps explain the lower efficacy of the BCG vaccine against TB in some studies. [Borenstein et al. \(2009\)](#) also commented that, in hotter climates, the vaccine may lose potency and certain bacteria necessary for the vaccine to work well are less likely to survive with more exposure to sunlight.

Following [Berkey et al. \(1995\)](#), we will explore these observations by using *meta regress* with the centered latitude as the moderator.

First, we generate a new variable, `latitude_c`, that is the mean-centered version of `latitude`. The mean value of `latitude`, 33.46, can be thought of as the latitude of the city of Atlanta in the United States or the city of Beirut in Lebanon.

```
. summarize latitude, meanonly
. generate double latitude_c = latitude - r(mean)
. label variable latitude_c "Mean-centered latitude"
```

We then fit meta-regression with `latitude_c` as the moderator.

```
. meta regress latitude_c
    Effect-size label: Log risk-ratio
      Effect size: _meta_es
      Std. err.: _meta_se

Random-effects meta-regression               Number of obs =          13
Method: REML                               Residual heterogeneity:
                                          tau2 = .07635
                                          I2 (%) = 68.39
                                          H2 = 3.16
                                          R-squared (%) = 75.63
                                          Wald chi2(1) = 16.36
                                          Prob > chi2 = 0.0001
```

<code>_meta_es</code>	Coefficient	Std. err.	z	P> z	[95% conf. interval]	
<code>latitude_c</code>	-.0291017	.0071953	-4.04	0.000	-.0432043	-.0149991
<code>_cons</code>	-.7223204	.1076535	-6.71	0.000	-.9333174	-.5113234

```
Test of residual homogeneity: Q_res = chi2(11) = 30.73    Prob > Q_res = 0.0012
```

The regression coefficient for `latitude_c` is  $-0.0291$ , which means that every one degree of latitude corresponds to a decrease of 0.0291 units in the log risk-ratio. In other words, the vaccine appears to work better in colder climates.

The proportion of between-study variance explained by the covariates can be assessed via the  $R^2$  statistic. Here roughly 76% of the between-study variance is explained by the covariate `latitude_c`. From the value of  $I^2$  in the output, roughly 68% of the residual variation is due to heterogeneity, which may potentially be explained by other covariates, with the other 32% due to the within-study sampling variability.

The test statistic for residual homogeneity,  $Q_{\text{res}}$ , is 30.73 with a  $p$ -value of 0.0012, so the null hypothesis of no residual heterogeneity is rejected, which is consistent with the reported residual heterogeneity summaries.


See [META] [meta regress](#) for more examples.



## ► Example 10: Bubble plot

Whenever there is one continuous covariate in the meta-regression, we may explore the relationship between the effect sizes and that covariate via a bubble plot using the `estat bubbleplot` command. Continuing with [example 9](#), we explore the relationship between the log risk-ratios and `latitude_c`.

```
. estat bubbleplot
```



The bubble plot is a scatterplot of effect sizes and covariate values. Each study is represented by a circle with the size of the circle proportional to the effect-size precision,  $1/\hat{\sigma}_j^2$ . The fitted line (predicted log risk-ratios) is also plotted on the graph.

The log risk-ratio for the BCG vaccine decreases as the distance from the equator increases. The plot also reveals a few outlying studies that require more thorough investigation. We continue exploring this model in [META] [meta regress postestimation](#).

See [META] [estat bubbleplot](#).



## Funnel plots for exploring small-study effects

A funnel plot (Light and Pillemer 1984) plots study-specific effect sizes against measures of study precision such as standard errors. This plot is commonly used to explore publication bias or, more precisely, small-study effects. Small-study effects (Sterne, Gavaghan, and Egger 2000) arise when smaller studies tend to report different results such as larger effect-size estimates than larger studies. In the absence of small-study effects, the shape of the plot should resemble a symmetric inverted funnel.

Publication bias arises when smaller studies with nonsignificant findings are being suppressed from publication. It is one of the more common reasons for the presence of small-study effects, which leads to the asymmetry of the funnel plot. Another common reason for the asymmetry in the funnel plot is the presence of between-study heterogeneity.

See [Introduction in \[META\] meta funnelplot](#) for details.

### ► Example 11: Funnel plot

Let’s explore the funnel-plot asymmetry for the NSAIDS dataset described in [Effectiveness of nonsteroidal anti-inflammatory drugs \(nsaids.dta\)](#).

```
. use https://www.stata-press.com/data/r19/nsaids, clear
(Effectiveness of nonsteroidal anti-inflammatory drugs)
. describe
Contains data from https://www.stata-press.com/data/r19/nsaids.dta
Observations:      37      Effectiveness of nonsteroidal
                        anti-inflammatory drugs
Variables:         5      24 Apr 2024 17:09
                        (_dta has notes)
```

Variable name	Storage type	Display format	Value label	Variable label
study	byte	%8.0g		Study ID
nstreat	byte	%8.0g		Number of successes in the treatment arm
nftreat	byte	%9.0g		Number of failures in the treatment arm
nscontrol	byte	%8.0g		Number of successes in the control arm
nfcontrol	byte	%9.0g		Number of failures in the control arm

Sorted by:

As before, our first step is to declare our data. `nsaids.dta` records summary data for a two-group comparison of binary outcomes, so we will again use `meta esize` to compute our effect sizes as in [example 5](#) and [example 7](#).

Our effect size of interest is an odds ratio, so we can use the default specification of `meta esize`.


```
. meta esize nstreat-nfcontrol
Meta-analysis setting information
Study information
  No. of studies: 37
  Study label: Generic
  Study size: _meta_studysize
  Summary data: nstreat nftreat nscontrol nfcontrol
  Effect size
    Type: lnoratio
    Label: Log odds-ratio
    Variable: _meta_es
  Zero-cells adj.: 0.5, only0
  Precision
    Std. err.: _meta_se
    CI: [_meta_cil, _meta_ciu]
    CI level: 95%
Model and method
  Model: Random effects
  Method: REML
```

In the above, instead of listing all four variables with `meta esize` as we did in previous examples, we use one of the varlist shortcuts (see [\[U\] 11.4 varname and varlists](#)) to include all variables between `nstreat` and `nfcontrol`. We could do this because our variables appear in the dataset in the same order they need to be listed with `meta esize`: numbers of successes and failures in the treatment group followed by those in the control group.

There are  $K = 37$  trials in this dataset. We will continue using the default random-effects meta-analysis model with the REML estimation method.

We use `meta funnelplot` to produce a funnel plot for the NSAIDS data.

```
. meta funnelplot
Effect-size label: Log odds-ratio
Effect size: _meta_es
Std. err.: _meta_se
Model: Common effect
Method: Inverse-variance
```



On a funnel plot, the more precise trials (with smaller standard errors) are displayed at the top of the funnel, and the less precise ones (with larger standard errors) are displayed at the bottom. The red reference line is plotted at the estimate of the overall effect size, the overall log odds-ratio in our example. In the absence of small-study effects, we would expect the points to be scattered around the reference line with the effect sizes from smaller studies varying more around the line than those from larger studies, forming the shape of an inverted funnel.

In our plot, there is an empty space in the bottom left corner. This suggests that the smaller trials with log odds-ratio estimates close to zero may be missing from the meta-analysis.

See [\[META\] meta funnelplot](#) for more examples.




## ► Example 12: Contour-enhanced funnel plot

The asymmetry is evident in the funnel plot from [example 11](#), but we do not know the cause for this asymmetry. The asymmetry can be the result of publication bias or may be because of other reasons. The so-called contour-enhanced funnel plots can help determine whether the asymmetry of the funnel plot is because of publication bias. The contour lines that correspond to certain levels of statistical significance (1%, 5%, and 10%) of tests of individual effects are overlaid on the funnel plot. Generally, publication bias is suspect when smaller studies are missing in the nonsignificant regions.

Let's add the 1%, 5%, and 10% significance contours to our funnel plot by specifying them in the `contours()` option.

```
. meta funnelplot, contours(1 5 10)
Effect-size label: Log odds-ratio
Effect size: _meta_es
Std. err.: _meta_se
Model: Common effect
Method: Inverse-variance
```



From this plot, we can see that the reported effects of almost all smaller trials (those at the bottom of the funnel) are statistically significant at a 5% level and less. On the other hand, a fair number of the larger trials (at the top of the funnel) reported nonsignificant results. For the funnel plot to look symmetric with respect to the reference line, we should have observed some trials in the middle and the bottom of the darkest region (with  $p$ -values larger than 10%). This suggests that we are missing some of the smaller trials with nonsignificant results, which would be consistent with the presence of publication bias.

There is also a chance that the funnel-plot asymmetry is induced by the between-study heterogeneity. Using a random-effects model and investigating the study-level covariates that may account for the heterogeneity should also be considered when exploring the funnel-plot asymmetry.

Also see [example 5](#) of [\[META\] meta funnelplot](#) for more details about this example.



## Testing for small-study effects

We can test for the presence of small-study effects or, technically, the asymmetry in the funnel plot more formally by using, for example, one of the regression-based tests. The main idea behind these tests is to determine whether there is a statistically significant association between the effect sizes and their measures of precision such as effect-size standard errors.

See [Introduction](#) in [\[META\] meta bias](#) for details.

### ► Example 13: Harbord's regression-based test

In [example 11](#), we investigated the funnel-plot asymmetry visually. Let's check for it more formally by using the [meta bias](#) command. We will use the Harbord regression-based test ([Harbord, Egger, and Sterne 2006](#)), which is often recommended when the effect size of interest is an odds ratio (or log odds-ratio).

To perform this test, we specify the `harbord` option with `meta bias`.

```
. meta bias, harbord
      Effect-size label: Log odds-ratio
      Effect size: _meta_es
      Std. err.: _meta_se

Regression-based Harbord test for small-study effects
Random-effects model
Method: REML

H0: beta1 = 0; no small-study effects
      beta1 =      3.03
      SE of beta1 =    0.741
           z =      4.09
      Prob > |z| =    0.0000
```

The test uses a type of weighted regression that explores the relationship between the effect sizes and their precision. The slope in that regression, labeled as `beta1` in the output, describes the asymmetry of the funnel plot and represents the magnitude of the small-study effects. The further it is from zero, the more asymmetry is present in the funnel plot.

`meta bias` reports the  $z$ -test statistic of 4.09 with a  $p$ -value less than 0.0000 for the test of  $H_0: \beta_1=0$  assuming a random-effects model with the REML estimation method. We have statistically significant evidence to reject the null hypothesis of the funnel-plot symmetry.

See [\[META\] meta bias](#).



### Trim-and-fill analysis for addressing publication bias

When the presence of publication bias is suspected, it is important to explore its impact on the final meta-analysis results. The trim-and-fill method of [Duval and Tweedie \(2000a, 2000b\)](#) provides a way to evaluate the impact of publication bias on the results. The idea of the method is to estimate the number of studies potentially missing because of publication bias, impute these studies, and use the observed and imputed studies to obtain the overall estimate of the effect size. This estimate can then be compared with the estimate obtained using only the observed studies. For details, see [Introduction](#) in [\[META\] meta trimfill](#).

#### ► Example 14: Trim-and-fill analysis

From [example 12](#) and [example 13](#), we suspect the presence of publication bias in the meta-analysis of the NSAIDS data. Let’s use the trim-and-fill method to investigate the impact of potentially missing studies on the estimate of the overall log odds-ratio.

We use the `meta trimfill` command. We specify the `eform` option (synonym for `or` when the computed effect sizes are log odds-ratios) to report the results as odds ratios instead of the default log odds-ratios. We also draw a contour-enhanced funnel plot that contains both the observed and imputed studies.

```
. meta trimfill, eform funnel(contours(1 5 10))
    Effect-size label: Log odds-ratio
      Effect size: _meta_es
        Std. err.: _meta_se


Nonparametric trim-and-fill analysis of publication bias
Linear estimator, imputing on the left

Iteration                                Number of studies =    47
  Model: Random-effects                   observed =    37
  Method: REML                           imputed =    10

Pooling
  Model: Random-effects
  Method: REML
```

Studies	Odds ratio	[95% conf. interval]	
Observed	3.752	2.805	5.018
Observed + Imputed	2.815	2.067	3.832





`meta trimfill` reports that 10 hypothetical studies are estimated to be missing. When 10 studies are imputed and added to the meta-analysis, the overall odds ratio reduces from 3.752 (based on 37 observed studies) to 2.815 (based on 47 observed and imputed studies). This suggests that the treatment benefit as reported in the literature may be larger than it would be in the absence of publication bias.

From the funnel plot, almost all the imputed studies fall in the darkest-gray region corresponding to a  $p$ -value of more than 10%. This further supports the conclusion that the small-study effect is most likely because of publication bias.

See [META] [meta trimfill](#).



## Multivariate meta-regression

Multivariate meta-regression is a multivariate statistical technique used to investigate reasons behind between-study heterogeneity of multiple dependent effect sizes. The technique explores whether there are associations between the effect sizes and other study-level covariates or moderators. You can think of multivariate meta-regression as an extension of meta-regression in univariate meta-analysis to the multivariate setting.

### ► Example 15: Multivariate meta-regression

In this example, we will use the periodontal disease dataset described in [Treatment of moderate periodontal disease \(periodontal.dta\)](#) to explore whether the moderator `pubyear` can explain some of the between-study heterogeneity of the two dependent effect-size variables `y1` and `y2`. We will perform a random-effects multivariate meta-regression using the `meta mvregress` command. Unlike other `meta` commands that are designed for standard meta-analysis, the `meta mvregress` command does not require your dataset to be declared as meta data.

```
. meta mvregress y1 y2 = pubyear, wcovvariables(v11 v12 v22)
```

Performing EM optimization ...

Performing gradient-based optimization:

Iteration 0: Log restricted-likelihood = -3.5544446  
 Iteration 1: Log restricted-likelihood = -3.5402086  
 Iteration 2: Log restricted-likelihood = -3.5399568  
 Iteration 3: Log restricted-likelihood = -3.5399567

Multivariate random-effects meta-regression	Number of obs	=	10
Method: REML	Number of studies	=	5
	Obs per study:		
	min	=	2
	avg	=	2.0
	max	=	2
	Wald chi2(2)	=	0.40
	Prob > chi2	=	0.8197

Log restricted-likelihood = -3.5399567

	Coefficient	Std. err.	z	P> z	[95% conf. interval]
y1					
pubyear	.0048615	.0218511	0.22	0.824	-.0379658 .0476888
_cons	.3587569	.07345	4.88	0.000	.2147975 .5027163
y2					
pubyear	-.0115367	.0299635	-0.39	0.700	-.070264 .0471907
_cons	-.3357368	.0979979	-3.43	0.001	-.5278091 -.1436645

Test of homogeneity: Q\_M = chi2(6) = 125.76                      Prob > Q\_M = 0.0000

Random-effects parameters	Estimate
Unstructured:	
sd(y1)	.1429917
sd(y2)	.2021314
corr(y1,y2)	.561385

The output shows information about the optimization algorithm, the iteration log, and the model (random-effects) and method (REML) used for estimation. It also displays the number of studies,  $K = 5$ , and the total number of observations on the outcomes,  $N = 10$ , which is equal to  $Kd$  because no observations are missing. The minimum, maximum, and average numbers of observations per study are also reported. Because there were no missing observations, all of these numbers are identical and are equal to 2. The Wald statistic,  $\chi^2 = 0.4$ , tests the joint hypothesis that the coefficients of pubyear for outcomes y1 and y2 are equal to 0.

The first table displays the fixed-effects coefficients for each dependent (outcome) variable. The coefficients of pubyear for outcomes y1 and y2 are not significant ( $p = 0.824$  and  $p = 0.7$ , respectively), so it does not appear that pubyear explains much of the between-study heterogeneity of effect sizes y1 and y2. In fact, the multivariate Cochran’s homogeneity test strongly suggests the presence of a between-study heterogeneity even after accounting for pubyear:  $Q_M = 125.76$  with a  $p < 0.0001$ .

The second table displays the random-effects parameters, which are used to compute an estimate of the between-study covariance matrix  $\Sigma$ . For details, see [META] meta mvregress.

After you fit your model, you can use estat heterogeneity to assess the residual heterogeneity in your model. To conduct other postestimation analysis, see [META] meta mvregress postestimation.



### Multilevel meta-regression

Multilevel meta-regression is a statistical technique used to study the relationship between potentially dependent effect sizes and covariates. The dependence among the effect sizes stems from a hierarchical or multilevel structure that is assumed present in the data. The standard random-effects meta-analysis can be viewed as a two-level meta-analysis model with studies as level-2 groups and subjects within studies as level-1 observations. When the term “multilevel meta-analysis” is used in the literature, it typically refers to models that incorporate more than two levels of hierarchy.

#### ► Example 16: Multilevel meta-regression

Møller and Mousseau (2015) conducted a meta-analysis to investigate the effect of radiation from Chernobyl on mutation rates across different taxonomic groups (`taxon`) and species (`species`). The relation between radiation and mutation rates was quantified by Pearson’s product-moment correlation coefficient (`correlation`). Study labels are stored in variable `studylbl`. A key feature of this dataset is that most studies contributed more than one observed effect size. Therefore, the effect sizes, identified by variable `id` (level 2), can be seen as nested within `studylbl` (level 3). The original dataset had 45 studies reporting 172 effect sizes corresponding to 8 different taxonomic groups. Here we focus only on the radiation effect on mutation rates for the largest two taxonomic groups in the dataset: mammals and plants. This leaves us with 42 studies reporting 158 effect sizes. We first describe the variables that will be used in our model:

```
. use https://www.stata-press.com/data/r19/chernobyl
(Effect of radiation from Chernobyl on mutation rates)
. describe studylbl - taxon
```

Variable name	Storage type	Display format	Value label	Variable label
studylbl	str29	%29s		Study label
id	int	%9.0g		Effect-size ID
z	double	%10.0g		Fisher’s z-transformed correlations
var	float	%9.0g		Variance of Fisher’s z-transformed correlations
taxon	byte	%9.0g	taxon1	Taxonomic group

Variables `z` and `var` store Fisher’s  $z$ -transformed correlation values and their variances. This transformed metric is typically used for estimation when pooling correlations; see [example 10](#) of [\[META\] meta summarize](#) for details about Fisher’s  $z$ -transformed correlations and their asymptotic standard-errors computation.

Because multiple effect sizes are nested within each study, we fit the three-level random-intercepts model

$$z_{jk} = \beta_1 I(\text{taxon}_j = \text{mammals}) + \beta_2 I(\text{taxon}_j = \text{plants}) + u_j^{(3)} + u_{jk}^{(2)} + \epsilon_{jk}, \quad j = 1, \dots, 42$$

where  $u_j^{(3)} \sim N(0, \tau_3^2)$ ,  $u_{jk}^{(2)} \sim N(0, \tau_2^2)$ , and  $\epsilon_{jk} \sim N(0, \text{var}_{jk})$ .  $I(\text{taxon}_j = \text{mammals})$  and  $I(\text{taxon}_j = \text{plants})$  are indicator variables for the mammals and plants taxonomic groups, respectively. You can think of the above model as a form of multilevel subgroup analysis.

We will perform a multilevel meta-regression using the `meta mereregress` command. Unlike other `meta` commands that are designed for standard meta-analysis, the `meta mereregress` command does not require your dataset to be declared as `meta data`.

```
. meta meregress z ibn.taxon, noconstant || studylbl: || id:, esvarvariable(var)
Performing EM optimization ...
Performing gradient-based optimization:
Iteration 0: Log restricted-likelihood = -321.43393 (not concave)
Iteration 1: Log restricted-likelihood = -187.61078
Iteration 2: Log restricted-likelihood = -183.46436
Iteration 3: Log restricted-likelihood = -181.8944
Iteration 4: Log restricted-likelihood = -181.83596
Iteration 5: Log restricted-likelihood = -181.83585
Iteration 6: Log restricted-likelihood = -181.83585
Computing standard errors ...
Multilevel REML meta-regression                                Number of obs =    158
    Grouping information
```

Group variable	No. of groups	Observations per group		
		Minimum	Average	Maximum
studylbl	42	1	3.8	22
id	158	1	1.0	1

```
Log restricted-likelihood = -181.83585                                Wald chi2(2) = 130.11
                                                                    Prob > chi2   = 0.0000
```

z	Coefficient	Std. err.	z	P> z	[95% conf. interval]	
taxon						
Mammals	.6622741	.1066936	6.21	0.000	.4531586	.8713897
Plants	1.031014	.1077358	9.57	0.000	.8198556	1.242172

```
Test of homogeneity: Q_M = chi2(156) = 1.0e+05                                Prob > Q_M = 0.0000
```

Random-effects parameters	Estimate
studylbl: Identity	
sd(_cons)	.2427429
id: Identity	
sd(_cons)	.7406531

In the syntax, we wrote `z ibn.taxon, noconstant` to specify the response (`z`) and the fixed-effects part of the model. The `esvarvariable(var)` option specifies the variable (`var` in our case) that stores the effect-size variances (sampling variances). The `|| studylbl: || id:` portion of the syntax adds to the model the random intercepts (the  $u_j^{(3)}$ 's and  $u_{jk}^{(2)}$ 's) at the respective `studylbl` and `id` levels. The order in which the levels are specified (from left to right) is important—`meta meregress` assumes that `id` is nested within `studylbl`. Because the above model is a random-intercepts three-level meta-regression (that is, a model without random slopes), it could have also been fit by using the `meta multilevel` command (`[META] meta multilevel`), which provides a simpler syntax for models with only random intercepts:

```
. meta multilevel z ibn.taxon, noconstant relevels(studylbl id) esvarvariable(var)
```

The output shows information about the optimization algorithm, the iteration log, and the estimation method (REML). It also displays the total number of effect sizes,  $n = 158$ . The minimum, maximum, and average numbers of observations per group at each hierarchical level are also reported. The Wald statistic,  $\chi^2 = 130.11$ , tests the joint hypothesis that Fisher's  $z$ -values for mammals and plants are equal to 0.

The second table displays the fixed-effects coefficients. Both overall effect sizes for mammals and plants are different from 0. The interpretation of the results, however, is easier in the natural correlation-coefficient metric, which we can compute using the inverse transformation:

$$\text{rho} = \frac{\exp(2z) - 1}{\exp(2z) + 1} = \text{tanh}(z)$$

For example, you may obtain the value of the correlation coefficient corresponding to mammals and its confidence interval as follows:

```
. display tanh(e(b) [1,1])
.57987485

. display "[" tanh(r(table) ["l1",1]) " , " tanh(r(table) ["u1",1]) "]"
[.42449189 , .70207952]
```

The multilevel Cochran's homogeneity test strongly suggests the presence of heterogeneity among the effect sizes even after partitioning the data by taxonomic groups ( $p < 0.0001$ ).

The third table displays the random-effects parameters, which are estimates of the level-3 and level-2 random-effects standard deviations,  $\tau_3$  and  $\tau_2$ , respectively. For details, see [META] [meta meregress](#).

After you fit your model, you can use [estat heterogeneity](#) to assess the multilevel heterogeneity in your model. To conduct other postestimation analysis, see [META] [meta me postestimation](#).

◀

## Acknowledgments

Previous and still ongoing work on meta-analysis in Stata influenced the design of meta. We gratefully acknowledge the contributions of the Stata developers who wrote the community-contributed commands. We thank Jonathan Sterne of the University of Bristol, Roger Harbord of the Met Office, Ross Harris of Public Health England, Thomas Steichen (retired) of RJRT, Mike Bradburn of the University of Sheffield, Jon Deeks of the University of Birmingham, and Doug Altman (1948–2018) of the University of Oxford, for `metan`; Evangelos Kontopantelis and David Reeves, both of the University of Manchester, for `metaan` and `ipdforest`; Roger Harbord, Julian Higgins of the University of Bristol, and Stephen Sharp of the MRC Epidemiology Unit, University of Cambridge, for `metareg`; Jonathan Sterne and Ross Harris for `metacum`; Jonathan Sterne and Roger Harbord for `metafunnel`; Tom Palmer of the University of Bristol, Alex Sutton of the University of Leicester, Santiago Moreno of HEOR Global, and Jaime Peters of the University of Exeter for `confunnel`; Roger Harbord, Ross Harris, and Jonathan Sterne for `metabias`; Thomas Steichen for `metatrim`; Roger Harbord and Penny Whiting of the University of Bristol for `metandi`; Ian White of the MRC Clinical Trials Unit at UCL for `mvmeta` and `network`; and David Fisher of the MRC Clinical Trials Unit at UCL for `ipdmetan` and `admetan`, and many more.

We also thank the editors, Tom Palmer and Jonathan Sterne, of the Stata Press book *Meta-Analysis in Stata: An Updated Collection from the Stata Journal* and the authors of the articles therein for providing valuable information about meta-analysis, in addition to developing the meta-analysis software.

## References

- Antczak-Bouckoms, A., K. Joshipura, E. Burdick, and J. F. Camilla Tulloch. 1993. Meta-analysis of surgical versus non-surgical methods of treatment for periodontal disease. *Journal of Clinical Periodontology* 20: 259–268. <https://doi.org/10.1111/j.1600-051X.1993.tb00355.x>.
- Berkey, C. S., D. C. Hoaglin, F. Mosteller, and G. A. Colditz. 1995. A random-effects regression model for meta-analysis. *Statistics in Medicine* 14: 395–411. <https://doi.org/10.1002/sim.4780140406>.
- . 1998. Meta-analysis of multiple outcomes by regression with random effects. *Statistics in Medicine* 17: 2537–2550. [https://doi.org/10.1002/\(SICI\)1097-0258\(19981130\)17:22%3C2537::AID-SIM953%3E3.0.CO;2-C](https://doi.org/10.1002/(SICI)1097-0258(19981130)17:22%3C2537::AID-SIM953%3E3.0.CO;2-C).
- Borenstein, M., L. V. Hedges, J. P. T. Higgins, and H. R. Rothstein. 2009. *Introduction to Meta-Analysis*. Chichester, UK: Wiley.
- Colditz, G. A., T. F. Brewer, C. S. Berkey, M. E. Wilson, E. Burdick, H. V. Fineberg, and F. Mosteller. 1994. Efficacy of BCG vaccine in the prevention of tuberculosis: Meta-analysis of the published literature. *Journal of the American Medical Association* 271: 698–702. <https://doi.org/10.1001/jama.1994.03510330076038>.
- Duval, S., and R. L. Tweedie. 2000a. Trim and fill: A simple funnel-plot-based method of testing and adjusting for publication bias in meta-analysis. *Biometrics* 56: 455–463. <https://doi.org/10.1111/j.0006-341X.2000.00455.x>.
- . 2000b. A nonparametric “trim and fill” method of accounting for publication bias in meta-analysis. *Journal of the American Statistical Association* 95: 89–98. <https://doi.org/10.1080/01621459.2000.10473905>.
- Galbraith, R. F. 1988. A note on graphical representation of estimated odds ratios from several clinical trials. *Statistics in Medicine* 7: 889–894. <https://doi.org/10.1002/sim.4780070807>.
- Harbord, R. M., M. Egger, and J. A. C. Sterne. 2006. A modified test for small-study effects in meta-analyses of controlled trials with binary endpoints. *Statistics in Medicine* 25: 3443–3457. <https://doi.org/10.1002/sim.2380>.
- Lau, J., E. M. Antman, J. Jimenez-Silva, B. Kupelnick, F. Mosteller, and T. C. Chalmers. 1992. Cumulative meta-analysis of therapeutic trials for myocardial infarction. *New England Journal of Medicine* 327: 248–254. <https://doi.org/10.1056/NEJM199207233270406>.
- Light, R. J., and D. B. Pillemer. 1984. *Summing Up: The Science of Reviewing Research*. Cambridge, MA: Harvard University Press. <https://doi.org/10.2307/j.ctvk12px9>.
- Møller, A. P., and T. A. Mousseau. 2015. Strong effects of ionizing radiation from Chernobyl on mutation rates. *Scientific Reports* 5: 8363. <https://doi.org/10.1038/srep08363>.
- Moore, R. A., M. R. Tramèr, D. Carroll, P. J. Wiffen, and H. J. McQuay. 1998. Quantitative systematic review of topically applied non-steroidal anti-inflammatory drugs. *BMJ* 316: 333–338. <https://doi.org/10.1136/bmj.316.7128.333>.
- Raudenbush, S. W. 1984. Magnitude of teacher expectancy effects on pupil IQ as a function of the credibility of expectancy induction: A synthesis of findings from 18 experiments. *Journal of Educational Psychology* 76: 85–97. <https://doi.org/10.1037/0022-0663.76.1.85>.
- Rice, K., J. P. T. Higgins, and T. S. Lumley. 2018. A re-evaluation of fixed effect(s) meta-analysis. *Journal of the Royal Statistical Society, A ser.*, 181: 205–227. <https://doi.org/10.1111/rssa.12275>.
- Rosenthal, R., and L. Jacobson. 1968. Pygmalion in the classroom. *Urban Review* 3: 16–20. <https://doi.org/10.1007/BF02322211>.
- Sterne, J. A. C., D. Gavaghan, and M. Egger. 2000. Publication and related bias in meta-analysis: Power of statistical tests and prevalence in the literature. *Journal of Clinical Epidemiology* 53: 1119–1129. [https://doi.org/10.1016/S0895-4356\(00\)00242-0](https://doi.org/10.1016/S0895-4356(00)00242-0).

## Also see

[META] **Intro** — Introduction to meta-analysis

[META] **Glossary**

Stata, Stata Press, and Mata are registered trademarks of StataCorp LLC. Stata and Stata Press are registered trademarks with the World Intellectual Property Organization of the United Nations. StataNow and NetCourseNow are trademarks of StataCorp LLC. Other brand and product names are registered trademarks or trademarks of their respective companies. Copyright © 1985–2025 StataCorp LLC, College Station, TX, USA. All rights reserved.



For suggested citations, see the FAQ on [citing Stata documentation](#).