

meta funnelplot — Funnel plots

Description	Quick start	Menu	Syntax
Options	Remarks and examples	Stored results	Methods and formulas
References	Also see		

Description

`meta funnelplot` produces funnel plots, which are used to explore the presence of [small-study effects](#) often associated with publication bias. A funnel plot is a scatterplot of study-specific effect sizes on the x axis against the measures of study precision such as standard errors and inverse standard errors on the y axis. In the absence of small-study effects, the plot should look symmetrical. `meta funnelplot` can also draw contour-enhanced funnel plots, which are useful for investigating whether the plot asymmetry can be attributed to publication bias.

Quick start

Construct a funnel plot for `meta data`, which was declared by either `meta set` or `meta esize`

```
meta funnelplot
```

Specify 1%, 5%, and 10% significance contours to produce a contour-enhanced funnel plot

```
meta funnelplot, contours(1 5 10)
```

Same as above, but base the significance contours on a one-sided lower-tailed z test, and request separate plots for each group of variable `groupvar`

```
meta funnelplot, contours(1 5 10, lower) by(groupvar)
```

Specify the inverse standard error as the precision metric on the y axis

```
meta funnelplot, metric(invse)
```

Menu

Statistics > Meta-analysis

Syntax

Construct a funnel plot

```
meta funnelplot [if] [in] [, level(#) options]
```

Construct a contour-enhanced funnel plot

```
meta funnelplot [if] [in], contours(contourspec) [options]
```

options

Description

<i>options</i>	Description
Model	
random [<i>remethod</i>]	random-effects meta-analysis
common [<i>cefemethod</i>]	common-effect meta-analysis
fixed [<i>cefemethod</i>]	fixed-effects meta-analysis
Options	
by (<i>varlist</i> , ...)	construct a separate plot for each group formed by <i>varlist</i>
metric (<i>metric</i>)	specify <i>y</i> -axis metric; default is metric (<i>se</i>)
n (#)	evaluate CI lines or significance contours at # points; default is n (300)
[no] metashow	display or suppress meta settings in the output
graph_options	affect rendition of overall funnel plot

[collect](#) is allowed; see [U] [11.1.10 Prefix commands](#).

remethod

Description

reml	restricted maximum likelihood; the default
mle	maximum likelihood
ebayes	empirical Bayes
dlaird	DerSimonian–Laird
sjonkman	Sidik–Jonkman
hedges	Hedges
hschmidt	Hunter–Schmidt

cefemethod

Description

mhaenszel	Mantel–Haenszel
invvariance	inverse variance
ivariance	synonym for invvariance

<i>graph_options</i>	Description
ES line <code>esopts</code> (<i>line_options</i>)	affect rendition of estimated effect-size line
CI plot * <code>ciopts</code> (<i>ciopts</i>)	affect rendition of the confidence intervals
Add plots <code>addplot</code> (<i>plot</i>)	add other plots to the funnel plot
Y axis, X axis, Titles, Legend, Overall <code>twoway_options</code>	any options documented in [G-3] <i>twoway_options</i>

* `ciopts`(*ciopts*) is not available for a contour-enhanced funnel plot.

Options

Main

`contours`(*contourspec*) specifies that a contour-enhanced funnel plot be plotted instead of the default standard funnel plot; see *Contour-enhanced funnel plots*. This option may not be combined with options `ciopts`() and `level`().

contourspec is `numlist`[, lower upper lines *graph_options*]. *numlist* specifies the levels of significance (as percentages) and may contain no more than 8 integer values between 1 and 50.

`lower` and `upper` specify that the significance contours be based on one-sided lower- or upper-tailed z tests of individual effect sizes. In other words, the studies in the shaded area of a specific contour c are considered not statistically significant based on one-sided lower- or upper-tailed z tests with $\alpha = c/100$. By default, the contours correspond to the two-sided z tests.

`lines` specifies that only the contours lines be plotted. That is, no shaded regions will be displayed.

graph_options are any of the options documented in [G-3] *area_options* except `recast`() or, if option `lines` is specified, any of the options documented in [G-3] *line_options*.

Model

Options `random`(), `common`(), and `fixed`() specify a meta-analysis model to use when estimating the overall effect size. For historical reasons, the default is `common`(*invariance*), regardless of the global model declared by `meta set` or `meta esize`. Specify one of these options with `meta funnelplot` to override this default. Options `random`(), `common`(), and `fixed`() may not be combined. Also see *Meta-analysis models* in [META] **Intro**.

`random` and `random`(*remethod*) specify that a random-effects model be assumed for meta-analysis; see *Random-effects model* in [META] **Intro**.

remethod specifies the type of estimator for the between-study variance τ^2 . *remethod* is one of `reml`, `mle`, `ebayes`, `dlaird`, `sjonkman`, `hedges`, or `hschmidt`. `random` is a synonym for `random`(`reml`). See *Options* in [META] **meta esize** for more information.

`common` and `common`(*cefemethod*) specify that a common-effect model be assumed for meta-analysis; see *Common-effect (“fixed-effect”) model* in [META] **Intro**. Also see the *discussion* in [META] **meta data** about common-effect versus fixed-effects models.

common implies common(mhaenszel) for effect sizes lnratio, lnrratio, and rdiff and common(invvariance) for all other effect sizes. common(mhaenszel) is supported only with effect sizes lnratio, lnrratio, and rdiff.

cefemethod is one of mhaenszel or invvariance (synonym ivariance). See [Options](#) in [META] [meta esize](#) for more information.

fixed and fixed(cefemethod) specify that a fixed-effects model be assumed for meta-analysis; see [Fixed-effects model](#) in [META] [Intro](#). Also see the [discussion](#) in [META] [meta data](#) about fixed-effects versus common-effect models.

fixed implies fixed(mhaenszel) for effect sizes lnratio, lnrratio, and rdiff and fixed(invvariance) for all other effect sizes. fixed(mhaenszel) is supported only with effect sizes lnratio, lnrratio, and rdiff.

cefemethod is one of mhaenszel or invvariance (synonym ivariance); see [Options](#) in [META] [meta esize](#) for more information.

Options

by(*varlist* [, *byopts*]) specifies that a separate plot for each group defined by *varlist* be produced. *byopts* are any of the options documented in [G-3] [by_option](#). by() is useful to explore publication bias in the presence of between-study heterogeneity induced by a set of categorical variables. These variables must then be specified in the by() option.

metric(*metric*) specifies the precision metric on the *y* axis. *metric* is one of se, invse, var, invvar, n, or invn. When *metric* is one of n or invn, no CIs or significance contours are plotted. The default is metric(se).

se specifies that the standard error, $\hat{\sigma}_j$, be used as the precision metric.

invse specifies that the inverse of the standard error, $1/\hat{\sigma}_j$, be used as the precision metric.

var specifies that the variance, $\hat{\sigma}_j^2$, be used as the precision metric.

invvar specifies that the inverse of the variance, $1/\hat{\sigma}_j^2$, be used as the precision metric.

n specifies that the sample size, n_j , be used as the precision metric.

invn specifies that the inverse of the sample size, $1/n_j$, be used as the precision metric.

level(#) specifies the confidence level, as a percentage, for confidence intervals. The default is as declared for the meta-analysis session; see [Declaring a confidence level for meta-analysis](#) in [META] [meta data](#). Also see option level() in [META] [meta set](#). This option may not be combined with option contours().

n(#) specifies the number of points at which to evaluate the CIs or, if option contours() is specified, significance contours. The default is n(300).

metashow and nometashow display or suppress the meta setting information. By default, this information is displayed at the top of the output. You can also specify nometashow with [meta update](#) to suppress the meta setting output for the entire meta-analysis session.

ES line

esopts(*line_options*) affects the rendition of the line that plots the estimated overall effect size; see [G-3] [line_options](#).

CI plot

`ciopts(ciopts)` affects the rendition of the (pseudo) CI lines in a funnel plot. *ciopts* are any of the options documented in [G-3] [line_options](#) and option `recast(rarea)` as described in [G-3] [advanced_options](#). This option may not be combined with option `contours()`.

Add plots

`addplot(plot)` allows adding more `graph twoway` plots to the graph; see [G-3] [addplot_option](#).

Y axis, X axis, Titles, Legend, Overall

twoway_options are any of the options documented in [G-3] [twoway_options](#). These include options for titling the graph (see [G-3] [title_options](#)) and for saving the graph to disk (see [G-3] [saving_option](#)).

Remarks and examples

stata.com

Remarks are presented under the following headings:

- [Introduction](#)
- [Funnel plots](#)
- [Contour-enhanced funnel plots](#)
- [Using meta funnelplot](#)
- [Examples of using meta funnelplot](#)

Introduction

A funnel plot is used to visually explore “small-study effects”. The term small-study effects (Sterne, Gavaghan, and Egger 2000) is used in meta-analysis to describe the cases when the results of smaller studies differ systematically from the results of larger studies. For instance, smaller studies often report larger effect sizes than the larger studies. One of the reasons for the presence of small-study effects is publication bias, also referred to as reporting bias.

For more formal testing of small-study effects, see [META] [meta bias](#). To assess the impact of publication bias on the results, see [META] [meta trimfill](#).

Also see [Publication bias](#) of [META] [Intro](#) for information about publication bias.

Funnel plots

The funnel plot (Light and Pillemer 1984) is a scatterplot of the study-specific effect sizes against measures of study precision. This plot is commonly used to explore publication bias. In the absence of publication bias, the shape of the scatterplot should resemble a symmetric (inverted) funnel.

In a funnel plot, the effect sizes, $\hat{\theta}_j$'s, from individual studies are plotted on the x axis, and measures of study precision such as standard errors, $\hat{\sigma}_j$'s, or sample sizes, n_j 's, are plotted on the y axis (Sterne and Harbord 2016). The line corresponding to the estimated overall effect size is also plotted on a funnel plot. In addition to standard errors and sample sizes, other choices for metrics on the y axis include inverse standard errors, $1/\hat{\sigma}_j$'s; variances, $\hat{\sigma}_j^2$'s; inverse variances, $1/\hat{\sigma}_j^2$'s; and inverse sample sizes, $1/n_j$'s. Sterne and Egger (2001) studied various metrics and found that the standard error metric performed well in many cases.

In the absence of publication bias (and between-study heterogeneity), the studies should be distributed symmetrically about the overall effect size because the sampling error is random. Also, the effect-size estimates from the smaller studies will be more variable than those from the larger studies. Thus, the scatter will be wider at the base of the plot creating, in the absence of bias, a symmetrical funnel shape or, more precisely, a symmetrical inverted funnel shape. When the statistically nonsignificant results of smaller studies are not published (and thus not included in the meta-analysis), an asymmetrical shape of the funnel plot may be observed. In this case, the estimate of the overall effect size will overestimate the true effect size. See [Sterne, Becker, and Egger \(2005\)](#) for details. Also see [Examples of using meta funnelplot](#) for examples of funnel plots.

[Sutton \(2009\)](#) states that when the y -axis metric is one of standard error, variance, or their inverses, a $(1 - \alpha) \times 100\%$ CI can be formed around the overall estimate. This CI can provide more formal interpretation of the plot. But the author suggests that caution be used when interpreting these CIs because they are formed around the estimate of the overall effect size that may be affected by publication bias. This is one of the reasons why the funnel-plot CIs are often referred to as pseudo CIs.

In general, there may be many reasons for an asymmetric funnel plot such as the choice of the plotted effect size ([Sterne et al. 2011](#)), the presence of a moderator correlated with the study effect and study size ([Peters et al. 2008](#)), or simply chance. One of the more common reasons, however, is the presence of substantial between-study heterogeneity ([Sterne, Gavaghan, and Egger 2000](#)).

The between-study heterogeneity, if present, must be addressed before the exploration of publication bias. For instance, if there are study-level covariates that explain the differences between the studies, their influence can distort a funnel plot if they are not accounted for in the main meta-analysis ([Sutton 2009](#)). Suppose that during our subgroup meta-analysis (see option `subgroup()` in `[META] meta summarize`), we identified a categorical variable that explains most of the heterogeneity between the studies. The exploration of the publication bias should then be performed separately for each group. That is, a separate funnel plot should be constructed for each subgroup. In the case of a continuous variable, some authors suggest constructing a funnel plot based on the residuals, $\hat{\theta}_j - \mathbf{x}_j \hat{\beta}$, on the x axis against their standard errors on the y axis, where the residuals are obtained from a meta-regression that uses this continuous variable as the moderator; see `[META] meta regress` and `[META] meta regress postestimation`.

Contour-enhanced funnel plots

[Peters et al. \(2008\)](#) (also see [Palmer et al. \[2016\]](#)) suggest that contour lines of statistical significance (or significance contours) be added to the funnel plot. These “contour-enhanced” funnel plots are useful for determining whether the funnel-plot asymmetry is potentially caused by publication bias or is perhaps due to other reasons. The contour lines that correspond to certain significance levels ($\alpha = 0.01, 0.05, 0.1$, etc.) of tests of zero effect sizes are overlaid on the funnel plot. Publication bias is suspect if there are studies, especially smaller studies, that are missing in the nonsignificant regions. Otherwise, other reasons may explain the presence of the funnel-plot asymmetry.

Using meta funnelplot

`meta funnelplot` produces funnel plots. By default, it plots effect sizes against the standard errors, but you can specify a different metric in the `metric()` option. For historical reasons, the default meta-analysis model is the common-effect model with the inverse-variance estimation method. You can specify one of `random()`, `common()`, or `fixed()` to use a different model or method.

By default, the CIs are displayed, which correspond to the confidence level as declared by `meta set` or `meta esize`. You can specify a different level in the `level()` option. You can also specify the `ciopts()` option to modify the default look of the CI lines.

Instead of the CIs, you can request a contour-enhanced funnel plot by specifying the desired levels of significance (as a percentage) in the `contours()` option. The default significance contours are based on two-sided significance tests of individual effect sizes. You can use `lower` or `upper` within `contours()` to specify that the significance contours be based on the corresponding one-sided tests. You can also specify `lines` within `contours()` to recast the contours to be displayed as lines instead of shaded area plots.

You can use the `by(varlist)` option to produce separate funnel plots for each group defined by `varlist`. This option is useful after a subgroup analysis (see option `subgroup()` in [META] `meta summarize`). If a subgroup analysis identified a categorical variable that explains some of the between-study variability, that variable must be specified in the `by()` option when using `meta funnelplot` to explore publication bias.

You can also change the default look of the effect-size line by specifying the `esopts()` option.

In the next section, we describe some of the uses of `meta funnelplot`.

Examples of using meta funnelplot

Recall the NSAIDS data of 37 trials on the effectiveness and safety of topical nonsteroidal anti-inflammatory drugs for acute pain described in *Effectiveness of nonsteroidal anti-inflammatory drugs (nsaids.dta)* of [META] `meta`. In this section, we use its declared version and focus on the demonstration of various options of `meta funnelplot` and explanation of its output.


```
. use https://www.stata-press.com/data/r18/nsaidsset
(Effectiveness of nonsteroidal anti-inflammatory drugs; set with -meta esize-)
. meta query, short
-> meta esize nstreat nftreat nscontrol nfcontrol
    Effect-size label: Log odds-ratio
    Effect-size type: lnoratio
    Effect size: _meta_es
    Std. err.: _meta_se
    Model: Random effects
    Method: REML
```

`meta query, short` reminds us about the main settings of the declaration step. Our data were declared by using `meta esize` with variables `nstreat`, `nftreat`, `nscontrol`, and `nfcontrol` representing the summary data from 2×2 tables, which record the numbers of successes and failures in the treatment and control arms. The computed effect sizes are log odds-ratios; their values and standard errors are stored in the respective `system variables` `_meta_es` and `_meta_se`. The declared meta-analysis model is the default random-effects model with the REML estimation method.

► Example 1: Funnel plot

To produce a funnel plot for the declared NSAIDS data, we simply type

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



The scatterplot of log odds-ratios against their standard errors is produced. The estimated effect-size line and the corresponding pseudo 95% CIs are also plotted. The funnel plot is clearly asymmetric with smaller, less precise studies—studies with larger standard errors—reporting larger effect sizes than the more precise studies. This may suggest the presence of publication bias. The plotted pseudo CI lines are not genuine CI limits, but they provide some insight into the spread of the observed effect sizes about the estimate of the overall effect size. In the absence of publication bias and heterogeneity, we would expect the majority of studies to be randomly scattered within the CI region resembling an inverted funnel shape.


Notice that although the declared meta-analysis model was the random-effects model with the REML estimation method, the default model used by `meta funnelplot` was the common-effect model with the inverse-variance method, as is indicated in the brief output of meta settings reported by the command. This is the model traditionally used with funnel plots in the literature. The reported model and method are used to compute the estimate of the overall effect size, the overall log odds-ratio in our example, which is depicted by the reference (red) effect-size line.

► Example 2: Random-effects REML model

In [example 1](#), we pointed out that, for historical reasons, `meta funnelplot` uses the common-effect model with the inverse-variance method by default instead of the declared ones.

If desired, we can obtain the results assuming the declared random-effects model with the REML estimation method by specifying the `random` option.

```
. meta funnelplot, random
Effect-size label: Log odds-ratio
Effect size: _meta_es
Std. err.: _meta_se
Model: Random effects
Method: REML
```



From the output of meta settings, a random-effects model with the REML method is now used to estimate the overall log odds-ratio. Our conclusion about the existence of potential publication bias remains the same.

For brevity, let's suppress the meta setting information from the output of `meta funnelplot` for the rest of the analysis. We can do this by specifying the `nometashow` option with `meta update` (see [\[META\] meta update](#)).


```
. quietly meta update, nometashow
```

▷ Example 3: Inverse-standard-error metric

Continuing with [example 1](#), we can use a different precision metric on the y axis. Using different metrics may reveal different characteristics of the observed studies. For example, for the standard error metric, more weight is given to the smaller studies, which are more susceptible to publication bias. For the inverse-standard-error metric, more weight is given to the larger studies.

Below, we specify the inverse-standard-error metric by using the `metric(invse)` option.

```
. meta funnelplot, metric(invse)
```



Because our y -axis metric is now $1/\hat{\sigma}_j$, the shape of the plotted CI lines is a hyperbola. The interpretation, however, is similar. We still want to see the majority of studies be concentrated within the regions defined by this hyperbola.


In this metric, the focus is on larger studies with the smaller studies compressed at the bottom. We can see that the asymmetry does not appear to be present for larger studies with, say, $1/\hat{\sigma}_j > 2$. But the asymmetry is present for the smaller studies.

◀

▷ Example 4: The ylabel() option

Continuing with [example 3](#), we can improve the look of the plot by restricting the y -axis values to the observed range. We can do this by specifying the `ylabel(1 2 3)` option:

```
. meta funnelplot, metric(invse) ylabel(1 2 3)
```




◀

▷ Example 5: Contour-enhanced funnel plot

In [example 12](#) of [\[META\] meta](#), we considered contour-enhanced funnel plots for the NSAIDS data to help us identify whether the asymmetry observed in [example 1](#) is because of publication bias or perhaps some other reasons. Let's obtain the 1%, 5%, and 10% contour-enhanced funnel plots by specifying the `contours(1 5 10)` option:

```
. meta funnelplot, contours(1 5 10)
```



The plotted contour regions define the regions of statistical significance (or nonsignificance) of the individual effect sizes θ_j . That is, if you consider the significance test of $H_0: \theta_j = 0$ for a study j , the contour regions represent the critical regions of such tests for all studies at the specified significance level such as 1%, 5%, and 10% levels in our example. Thus, if a study falls outside a

certain region, we have statistical evidence to reject the null hypothesis of no effect at the significance level corresponding to that region.

In our example, for studies in the white region, the null hypothesis of no effect can be rejected at the 1% significance level. That is, the significance tests for these studies would have p -values less than 0.01 or 1%. For studies in the light-gray region, the p -values would be between 1% and 5%. For studies in the darker-gray region, the p -values would be between 5% and 10%. And for studies in the darkest-gray region, the p -values would be larger than 10%.

The plot clearly shows that almost all smaller studies report a statistically significant result, favoring the treatment, either at the 1% or 5% level. On the other hand, some of the larger (more precise) studies (in the darkest-gray region) report nonsignificant results. The hypothetical missing studies—the studies that would make the scatterplot look symmetric with respect to the solid red vertical line—appear to fall in the darkest-gray region corresponding to a p -value of more than 10%. Because we are “missing” small studies in a region of statistical nonsignificance, this suggests that the observed asymmetry in the funnel plot is likely because of publication bias.

Also see [example 14](#) of [\[META\] meta](#).



► Example 6: The `addplot()` option

Continuing with [example 5](#), let's use the `addplot()` option to add the pseudo 95% CIs around the effect-size line to our contour-enhanced plot. These CIs are computed as follows: $\hat{\theta}_{IV} \pm 1.96 \times y$, where y represents a standard error of $\hat{\theta}_{IV}$, varying within the range of observed standard errors. Here 1.96 corresponds to the $z_{0.975}$ critical value.

We first obtain the estimated value of the overall effect size, which is available in the `r(theta)` scalar from the previous run of `meta funnelplot`. We store it in the `theta` scalar.

```
. display r(theta)
1.0604829
. scalar theta = r(theta)
```

The estimate of the overall effect size may also be obtained from `meta summarize` (see [\[META\] meta summarize](#)):

```
. meta summarize, common(invvariance) nostudies
Meta-analysis summary                               Number of studies =    37
Common-effect model
Method: Inverse-variance
      theta: Overall Log odds-ratio
```

	Estimate	Std. err.	z	P> z	[95% conf. interval]
theta	1.060483	.0758709	13.98	0.000	.9117788 1.209187

```
. display r(theta)
1.0604829
```


The CI lines can be constructed by using `twoway`'s function command (see [G-2] [graph twoway function](#)):

```
. twoway function theta-1.96*x, horizontal range(0 1.6) ||
>      function theta+1.96*x, horizontal range(0 1.6)
```

In the above, `x` plays the role of `y` in the earlier expression for CIs. We used the `horizontal` option to interchange the roles of `y` and `x` in the function because $\hat{\theta}_j$ appears on the `x` axis and standard errors on the `y` axis in the funnel plot. We also specified the `range()` option so that the range of the plotted function matches the observed range for the standard errors.

We use the above specification in the `addplot()` option with `meta funnelplot`. Because `addplot()` implies a twoway plot, we can omit `twoway` within `addplot()`. We also specify several other options to improve the look of the graph, which we describe later.

```
. local opts horizontal range(0 1.6) lpattern(dash) lcolor("red")
> legend(order(1 2 3 4 5 6) label(6 "95% pseudo CI"))
. meta funnel, contours(1 5 10)
> addplot(function theta-1.96*x, 'opts' || function theta+1.96*x, 'opts')
```



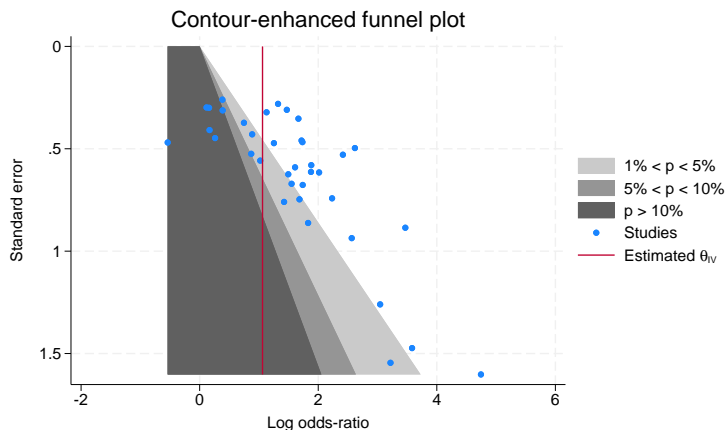
We changed the color and pattern of the CI lines by using options `lpattern()` and `lcolor()`. We used `legend()`'s suboption `order()` to display only the first six keys for the legend to avoid duplicate keys for the two CI plots. We also provided a more descriptive label for the CI legend key. Also, to be more precise, we could have replaced 1.96 in the above with `invnormal(.975)`, which computes the corresponding quantile of the standard normal distribution.

◀

► Example 7: Upper one-sided significance contours

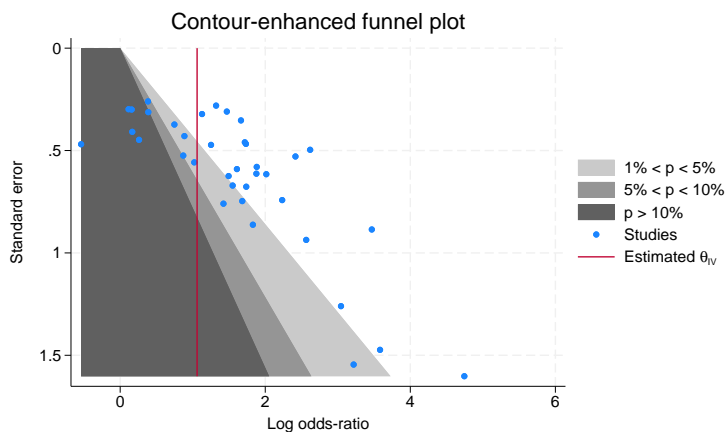
Continuing with [example 5](#), let's produce a contour-enhanced funnel plot based on one-sided significance tests instead of the default two-sided tests. We can specify `lower` or `upper` within `contours()` to produce a funnel plot with contours based on lower or upper one-sided tests. Below, we specify the upper suboption:

```
. meta funnelplot, contours(1 5 10, upper)
```




The interpretation of a one-sided contour-enhanced funnel plot is similar to that of a two-sided one. Studies that fall in the upper-tailed region (the white region to the right) are statistically significant at the 1% level based on a one-sided upper-tailed test. The white space on the left is uninformative, and we can suppress it by disallowing the x axis to extend to -2 . This may be done by specifying `xlabel(0(2)6)` (see [G-3] [axis_label_options](#)).

```
. meta funnelplot, contours(1 5 10, upper) xlabel(0(2)6)
```



If we want to suppress all the extra white space around the plot, we can specify `plotregion(margin(zero))` (see [G-3] *region_options*).

```
. meta funnelplot, contours(1 5 10, upper) xlabel(0(2)6)
> plotregion(margin(zero))
```



◀


► Example 8: Group-specific funnel plots: option `by()`

In [example 9](#) of [META] [meta summarize](#), we performed subgroup analysis for the pupil IQ data to account for the between-study heterogeneity induced by the binary covariate `week1`, which records whether teachers had prior contact with students for more than 1 week or for 1 week or less. See [Effects of teacher expectancy on pupil IQ \(pupliq.dta\)](#) of [META] [meta](#) for details about the data.

Let's check for publication bias in these data, accounting for the between-study heterogeneity. We can do this by looking at a funnel plot separately for each category of `week1`. We will use `meta funnel`'s option `by()` to produce such plots.

We use a previously declared version of the dataset.

```
. use https://www.stata-press.com/data/r18/pupiliqset, clear
(Effects of teacher expectancy on pupil IQ; set with -meta set-)
. meta funnelplot, by(week1)
Effect-size label: Std. mean diff.
Effect size: stdmdiff
Std. err.: se
Model: Common effect
Method: Inverse-variance
```



The above graph shows funnel plots of the subgroups for prior contact of one week or less and more than one week, respectively. These funnels are centered on different effect-size values, but there is little evidence of asymmetry in either plot. We should be careful with our interpretation, however, because we have only a few studies in each plot.

◀

Stored results

`meta funnelplot` stores the following in `r()`:

Scalars

<code>r(theta)</code>	estimated overall effect size
<code>r(xmin)</code>	minimum abscissa of scatter points
<code>r(xmax)</code>	maximum abscissa of scatter points
<code>r(ymin)</code>	minimum ordinate of scatter points
<code>r(ymax)</code>	maximum ordinate of scatter points

Macros

<code>r(model)</code>	meta-analysis model
<code>r(method)</code>	meta-analysis estimation method
<code>r(metric)</code>	metric for the y axis
<code>r(contours)</code>	significance levels of contours

Methods and formulas

`meta funnel` produces a scatterplot of individual effect sizes, $\hat{\theta}_j$'s, on the x axis and measures of study precision on the y axis. The supported measures of precision are the standard errors, $\hat{\sigma}_j$'s (default); inverse standard errors, $1/\hat{\sigma}_j$'s; variances, $\hat{\sigma}_j^2$'s; inverse variances, $1/\hat{\sigma}_j^2$'s; sample sizes, n_j 's; and inverse sample sizes, $1/n_j$'s.

The effect-size reference (solid red) line is plotted at the estimate of the overall effect size. By default, the overall effect size is estimated assuming a common-effect model with the inverse-variance method, but this can be changed by specifying one of `random()`, `common()`, or `fixed()`.

The pseudo $(1 - \alpha) \times 100\%$ CI lines, plotted by default, are constructed as $\hat{\theta} \pm z_{1-\alpha/2} f(y)$, where $\hat{\theta}$ is the estimate of the overall effect size, $z_{1-\alpha/2}$ is the $(1 - \alpha/2)$ th quantile of the standard normal distribution, and $f(y)$ plays the role of the varying standard error of $\hat{\theta}$ as a function of the y -axis values and depends on the chosen metric. When standard errors, $\hat{\sigma}_j$'s, are plotted on the y axis,

$$f(y) = y$$

and the CI curves form straight lines. When variances, $\hat{\sigma}_j^2$'s, are plotted on the y axis,

$$f(y) = \sqrt{y}$$

and the CI curves form a parabola. When inverse standard deviations, $1/\hat{\sigma}_j$'s, or inverse variances, $1/\hat{\sigma}_j^2$, are plotted on the y axis

$$f(y) = \frac{1}{y} \quad \text{and} \quad f(y) = \frac{1}{\sqrt{y}}$$

and the CI curves form a hyperbola.

When the `contours()` option is specified, the contour region corresponding to a significance level α (specified as a percentage in `contours()`) for a two-sided test is defined as the set of points (x, y) ,

$$\left\{ (x, y) : \left| \frac{x}{f(y)} \right| \leq z_{1-\alpha/2} \right\}$$

where $f(y)$ depends on the chosen metric and is defined as before. For upper one-sided tests, the contour region is defined as

$$\left\{ (x, y) : \frac{x}{f(y)} \geq z_{1-\alpha} \right\}$$

and for lower one-sided tests, it is defined as

$$\left\{ (x, y) : \frac{x}{f(y)} \leq z_{\alpha} \right\}$$

The `n(#)` option specifies how many evaluation points are used to construct the CI lines or, when the `contours()` option is specified, the significance contours. By default, 300 points are used for y for CI lines and for each of x and y for contours.

References

- Light, R. J., and D. B. Pillemer. 1984. *Summing Up: The Science of Reviewing Research*. Cambridge, MA: Harvard University Press.
- Palmer, T. M., J. L. Peters, A. J. Sutton, and S. G. Moreno. 2016. Contour-enhanced funnel plots for meta-analysis. In *Meta-Analysis in Stata: An Updated Collection from the Stata Journal*, ed. T. M. Palmer and J. A. C. Sterne, 2nd ed., 139–152. College Station, TX: Stata Press.
- Peters, J. L., A. J. Sutton, D. R. Jones, K. R. Abrams, and L. Rushton. 2008. Contour-enhanced meta-analysis funnel plots help distinguish publication bias from other causes of asymmetry. *Journal of Clinical Epidemiology* 61: 991–996. <https://doi.org/10.1016/j.jclinepi.2007.11.010>.
- Sterne, J. A. C., B. J. Becker, and M. Egger. 2005. The funnel plot. In *Publication Bias in Meta-Analysis: Prevention, Assessment and Adjustments*, ed. H. R. Rothstein, A. J. Sutton, and M. Borenstein, 75–98. Chichester, UK: Wiley.
- Sterne, J. A. C., and M. Egger. 2001. Funnel plots for detecting bias in meta-analysis: Guidelines on choice of axis. *Journal of Clinical Epidemiology* 54: 1046–1055. [https://doi.org/10.1016/s0895-4356\(01\)00377-8](https://doi.org/10.1016/s0895-4356(01)00377-8).
- 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).
- Sterne, J. A. C., and R. M. Harbord. 2016. Funnel plots in meta-analysis. In *Meta-Analysis in Stata: An Updated Collection from the Stata Journal*, ed. T. M. Palmer and J. A. C. Sterne, 2nd ed., 124–138. College Station, TX: Stata Press.
- Sterne, J. A. C., A. J. Sutton, J. P. A. Ioannidis, N. Terrin, D. R. Jones, J. Lau, J. R. Carpenter, G. Rücker, R. M. Harbord, C. H. Schmid, J. Tetzlaff, J. J. Deeks, J. L. Peters, P. Macaskill, G. Schwarzer, S. Duval, D. G. Altman, D. Moher, and J. P. T. Higgins. 2011. Recommendations for examining and interpreting funnel plot asymmetry in meta-analyses of randomised controlled trials. *BMJ* 343: d4002. <https://doi.org/10.1136/bmj.d4002>.
- Sutton, A. J. 2009. Publication bias. In *The Handbook of Research Synthesis and Meta-Analysis*, ed. H. Cooper, L. V. Hedges, and J. C. Valentine, 2nd ed., 435–452. New York: Russell Sage Foundation.

Also see

- [META] **meta bias** — Tests for small-study effects in meta-analysis
- [META] **meta data** — Declare meta-analysis data
- [META] **meta trimfill** — Nonparametric trim-and-fill analysis of publication bias
- [META] **meta** — Introduction to meta
- [META] **Glossary**
- [META] **Intro** — Introduction to meta-analysis

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. Other brand and product names are registered trademarks or trademarks of their respective companies. Copyright © 1985–2023 StataCorp LLC, College Station, TX, USA. All rights reserved.

