

Chapter 455

Meta-Analysis of Means

Introduction

This module performs a meta-analysis on a set of two-group, continuous-scale, studies. These studies have a treatment group and a control group. Each study's result may be summarized by the sample size, mean, and standard deviation of each of the two groups. The program provides a complete set of numeric reports and plots to allow the investigation and presentation of the studies. The plots include the *forest plot*, *radial plot*, and *L'Abbe plot*. Both fixed-, and random-, effects models are available for analysis.

Meta-Analysis refers to methods for the systematic review of a set of individual studies with the aim to combine their results. Meta-analysis has become popular for a number of reasons:

1. The adoption of evidence-based medicine which requires that all reliable information is considered.
2. The desire to avoid narrative reviews which are often misleading.
3. The desire to interpret the large number of studies that may have been conducted about a specific treatment.
4. The desire to increase the statistical power of the results by combining many small-size studies.

The goals of meta-analysis may be summarized as follows. A meta-analysis seeks to systematically review all pertinent evidence, provide quantitative summaries, integrate results across studies, and provide an overall interpretation of these studies.

We have found many books and articles on meta-analysis. In this chapter, we briefly summarize the information in Sutton et al (2000) and Thompson (1998). Refer to those sources for more details about how to conduct a meta-analysis.

Treatment Effects

Suppose you have obtained the results for k studies, labeled $i = 1, \dots, k$. Each study consists of a treatment group (T) and a control group (C). The results of each study are summarized by six statistics:

- n_{Ti} the number of subjects in the treatment group.
- n_{Ci} the number of subjects in the control group.
- \bar{x}_{Ti} the sample mean of the treatment group which estimates the treatment mean μ_{Ti} .
- \bar{x}_{Ci} the sample mean of the control group which estimates the control mean μ_{Ci} .
- s_{Ti} the sample standard deviation of the treatment group.
- s_{Ci} the sample standard deviation of the control group.

Meta-Analysis of Means

Mean Difference

The scales (e.g. blood pressure, pulse rate, volume, etc.) of all studies must be the same. If the logarithm has been taken in one study, it must be taken in all studies. You cannot combined studies with different scales using this program!

The measure of treatment effect for study i is

$$\theta_i = \mu_{Ti} - \mu_{Ci}$$

which is estimated by

$$\hat{\theta}_i = \bar{x}_{Ti} - \bar{x}_{Ci}$$

The standard deviation of the difference is

$$V(\hat{\theta}_i) = \sigma_i^2 \left(\frac{1}{n_{Ti}} + \frac{1}{n_{Ci}} \right)$$

The value of σ^2 is estimated by the pooled sample standard deviation given by

$$s_i^2 = \frac{(n_{Ti} - 1)s_{Ti}^2 + (n_{Ci} - 1)s_{Ci}^2}{n_{Ti} + n_{Ci}}$$

so that

$$\hat{V}(\hat{\theta}_i) = s_i^2 \left(\frac{1}{n_{Ti}} + \frac{1}{n_{Ci}} \right)$$

Defining the Study Parameters

Confidence intervals based on the t distribution may be defined for θ_i in the usual manner.

$$\hat{\theta}_i \pm t_{n_{Ti}+n_{Ci}-2, 1-\alpha/2} \sqrt{\hat{V}(\hat{\theta}_i)}$$

It will be useful in the sequel to make the following definition of the weights.

$$v_i = \hat{V}(\hat{\theta}_i)$$

$$w_i = 1/v_i$$

Hypothesis Tests

Several hypothesis tests have be developed to test the various hypotheses that may be of interest. These will be defined next.

Overall Null Hypothesis

Two statistical tests have been devised to test the overall null hypothesis that all treatment effects are zero. The null hypothesis is written

$$H_0 : \theta_i = 0 \quad i = 1, \dots, k$$

Meta-Analysis of Means

Nondirectional Test

The nondirectional alternative hypothesis that at least one $\theta_i \neq 0$ may be tested by comparing the quantity

$$X_{ND} = \sum_{i=1}^k w_i \hat{\theta}_i^2$$

with a χ_k^2 distribution.

Directional Test

A test of the more interesting directional alternative hypothesis that $\theta_i = \theta \neq 0$ for all i may be tested by comparing the quantity

$$X_D = \frac{\left(\sum_{i=1}^k w_i \hat{\theta}_i \right)^2}{\sum_{i=1}^k w_i}$$

with a χ_1^2 distribution. Note that this tests the hypothesis that all effects are equal to the same nonzero quantity.

Effect-Equality (Heterogeneity) Test

When the overall null hypothesis is rejected, the next step is to test whether all effects are equal, that is, whether the effects are homogeneous. Specifically, the hypothesis is

$$H_0 : \theta_i = \theta \quad i = 1, \dots, k$$

versus the alternative that at least one effect is different, that is, that the effects are heterogeneous. This may also be interpreted as a test of the study-by-treatment interaction.

This hypothesis is tested using Cochran's Q test which is given by

$$Q = \sum_{i=1}^k w_i (\hat{\theta}_i - \hat{\theta})^2$$

where

$$\hat{\theta} = \frac{\sum_{i=1}^k w_i \hat{\theta}_i}{\sum_{i=1}^k w_i}$$

The test is conducted by comparing Q to a χ_{k-1}^2 distribution.

Fixed versus Random Effects Combined Confidence Interval

If the effects are assumed to be equal (homogeneous), either through testing or from other considerations, a *fixed effects model* may be used to construct a combined confidence interval. However, if the effects are heterogeneous, a *random effects model* should be used to construct the combined confidence interval.

Fixed Effects Model

The fixed effects model assumes homogeneity of study results. That is, it assumes that $\theta_i = \theta$ for all i . This assumption may not be realistic when combining studies with different patient pools, protocols, follow-up strategies, doses, durations, etc.

If the fixed effects model is adopted, the *inverse variance-weighted* method as described by Sutton (2000) page 58 is used to calculate the confidence interval for θ . The formulas used are

$$\hat{\theta} \pm z_{1-\alpha/2} \sqrt{\hat{V}(\hat{\theta})}$$

where $z_{1-\alpha/2}$ is the appropriate percentage point from the standardized normal distribution and

$$\hat{\theta} = \frac{\sum_{i=1}^k w_i \hat{\theta}_i}{\sum_{i=1}^k w_i}$$

$$\hat{V}(\hat{\theta}) = \frac{1}{\sum_{i=1}^k w_i}$$

Random Effects Model

The random effects model assumes that the individual θ_i come from a random distribution with fixed mean $\bar{\theta}$ and variance σ^2 . Sutton (2000) page 74 presents the formulas necessary to conduct a random effects analysis using the *weighted* method. The formulas used are

$$\hat{\bar{\theta}} \pm z_{1-\alpha/2} \sqrt{\hat{V}(\hat{\bar{\theta}})}$$

where $z_{1-\alpha/2}$ is the appropriate percentage point from the standardized normal distribution and

$$\hat{\bar{\theta}} = \frac{\sum_{i=1}^k \bar{w}_i \hat{\theta}_i}{\sum_{i=1}^k \bar{w}_i}$$

$$\hat{V}(\hat{\bar{\theta}}) = \frac{1}{\sum_{i=1}^k \bar{w}_i}$$

$$\bar{w}_i = \frac{1}{\frac{1}{w_i} + \hat{\tau}^2}$$

$$\hat{\tau}^2 = \begin{cases} \frac{Q - k + 1}{U} & \text{if } Q > k - 1 \\ 0 & \text{otherwise} \end{cases}$$

Meta-Analysis of Means

$$Q = \sum_{i=1}^k w_i (\hat{\theta}_i - \hat{\theta})^2$$

$$U = (k-1) \left(\bar{w} - \frac{s_w^2}{k\bar{w}} \right)$$

$$s_w^2 = \frac{1}{k-1} \left(\sum_{i=1}^k w_i^2 - k\bar{w}^2 \right)$$

$$\bar{w} = \frac{1}{k} \left(\sum_{i=1}^k w_i \right)$$

Graphical Displays

A number of plots have been devised to display the information in a meta-analysis. These include the forest plot, the radial plot, and the L'Abbe plot. More will be said about each of these plots in the Output section.

Data Structure

The data are entered into a dataset using one row per study. Six variables are required to hold the sample size, mean, and standard deviation of each study. In addition to these, an additional variable is usually used to hold a short (3 or 4 character) label. Another variable may be used to hold a grouping variable.

As an example, we will use data referred to in Sutton (2000) page 30 as the dental dataset. This dataset reviews nine randomized clinical trials that were conducted to study the effects of sodium fluoride (NaF) with sodium monofluorophosphate (SMFP). These nine studies were all on the same continuous scale, so their results could be analyzed using the meta-analysis techniques presented in this chapter. These data are contained in the SUTTON30 database. You should load this database to see how the data are arranged.

Procedure Options

This section describes the options available in this procedure.

Variables Tab

The options on this screen control the variables that are used in the analysis.

Variables

Treatment N Variable

Specify the variable containing the sample size for the treatment group. Each row of data represents a separate study.

Treatment Mean Variable

Specify the variable containing the mean of the treatment group. Each row of data represents a separate study.

Meta-Analysis of Means

Treatment S.D. Variable

Specify the variable containing the standard deviation (not the standard error) of the treatment group. Each row of data represents a separate study.

Control N Variable

Specify the variable containing the sample size for the control group. Each row of data represents a separate study.

Control Mean Variable

Specify the variable containing the mean of the control group. Each row of data represents a separate study.

Control S.D. Variable

Specify the variable containing the standard deviation (not the standard error) of the control group. Each row of data represents a separate study.

Variables – Optional Variables

Label Variable

Specify an optional variable containing a label for each study (row) in the database. This label should be short (< 8 letters) so that it can fit on the plots.

Group Variable

Specify an optional variable containing a group identification value. Each unique value of this variable will receive its own plotting symbol on the forest plots. Some reports are sorted by these group values.

Combine Studies Method

Combine Studies Using

Specify the method used to combine treatment effects.

Use the **Fixed Effects** method when you do not want to account for the variation between studies.

Use the **Random Effects** method when you want to account for the variation between studies as well as the variation within the studies.

Reports Tab

The options on this screen control the appearance of the reports.

Select Reports

Summary Report - Outcome Detail Reports

Indicate whether to display the corresponding report.

Alpha Level

This setting controls the confidence coefficient used in the confidence limits. Note that $100 \times (1 - \alpha)\%$ confidence limits will be calculated. This must be a value between 0 and 0.5. The most common choice is 0.05, which results in 95% confidence intervals.

Meta-Analysis of Means

Report Options

Show Notes

Indicate whether to show the notes at the end of reports. Although these notes are helpful at first, they may tend to clutter the output. This option lets you omit them.

Precision

Specify the precision of numbers in the report. A single-precision number will show seven-place accuracy, while a double-precision number will show thirteen-place accuracy. Note that the reports are formatted for single precision. If you select double precision, some numbers may run into others. Also note that all calculations are performed in double precision regardless of which option you select here. Single precision is for reporting purposes only.

Variable Names

This option lets you select whether to display only variable names, variable labels, or both.

Report Options – Decimal Places

Probability Values - T Values

This setting controls the number of digits to the right of the decimal place that are displayed when showing this item.

Plots Tab

The options on this panel control the inclusion and the appearance of the forest plot, radial plot and L'Abbe plot.

Select Plots

Forest Plot – L'Abbe Plot

Indicate whether to display the corresponding plot. Click the plot format button to change the plot settings.

Storage Tab

These options let you specify if, and where on the database, various statistics are stored.

Warning: Any data already in these columns are replaced by the new data. Be careful not to specify columns that contain important data.

Data Storage Options

Storage Option

This option controls whether the values indicated below are stored on the dataset when the procedure is run.

- **Do not store data**
No data are stored even if they are checked.
- **Store in empty columns only**
The values are stored in empty columns only. Columns containing data are not used for data storage, so no data can be lost.

Meta-Analysis of Means

- **Store in designated columns**

Beginning at the *First Storage Variable*, the values are stored in this column and those to the right. If a column contains data, the data are replaced by the storage values. Care must be used with this option because it cannot be undone.

Store First Item In

The first item is stored in this column. Each additional item that is checked is stored in the columns immediately to the right of this variable.

Leave this value blank if you want the data storage to begin in the first blank column on the right-hand side of the data.

Warning: any existing data in these columns is automatically replaced, so be careful.

Data Storage Options – Select Items to Store with the Dataset

Mean Difference - Weights

Indicate whether to store these row-by-row values, beginning at the column indicated by the *Store First Item In* option.

Example 1 – Meta-Analysis of Means

This section presents an example of how to analyze the data contained in the Sutton30 dataset. This dataset contains data for nine randomized clinical trials that were conducted to study the effect of fluoridation. The NaF data represent the control group and the SMFP represent the treated group.

You may follow along here by making the appropriate entries or load the completed template **Example 1** by clicking on Open Example Template from the File menu of the Meta-Analysis of Means window.

1 Open the Sutton30 dataset.

- From the File menu of the NCSS Data window, select **Open Example Data**.
- Click on the file **Sutton30.NCSS**.
- Click **Open**.

2 Open the Meta-Analysis of Means window.

- Using the Analysis menu or the Procedure Navigator, find and select the **Meta-Analysis of Means** procedure.
- On the menus, select **File**, then **New Template**. This will fill the procedure with the default template.

3 Select the variables.

- Select the **Variables tab**.
- Set the **Treatment N Variables** to **SMFPN**.
- Set the **Treatment Mean Variables** to **SMFPMean**.
- Set the **Treatment S.D. Variables** to **SMFPSD**.
- Set the **Control N Variables** to **NaFN**.
- Set the **Control Mean Variables** to **NaFMean**.
- Set the **Control S.D. Variables** to **NaFSD**.
- Set the **Label Variable** to **Study**.

Meta-Analysis of Means

4 Specify the reports.

- Select the **Reports tab**.
- Check the **Summary Report** option box.
- Check the **Heterogeneity Tests** option box.
- Check the **Outcome Detail Reports** option box.
- On the **Plots** tab, check the **Forest Plot** option box.
- Check **Radial Plot** option box.
- Check the **L'Abbe Plot** option box.

5 Run the procedure.

- From the Run menu, select **Run Procedure**. Alternatively, just click the green Run button.

Run Summary Section

| Parameter | Value | Parameter | Value |
|-------------------------|----------|-----------------------|---------|
| N Treatment Variable | SMFPN | N Control Variable | NaFN |
| Mean Treatment Variable | SMFPMean | Mean Control Variable | NaFMean |
| SD Treatment Variable | SMFPSD | SD Control Variable | NaFSD |
| Group Variable | None | Number Groups | 1 |
| Row Label Variable | Study | Rows Processed | 9 |

This report records the variables that were used and the number of rows that were processed.

Numeric Summary Section

| Study | NT/NC | Mean T | Mean C | Difference | SD T | SD C |
|-------------------|-----------|---------|---------|------------|--------|--------|
| S1 | 113/134 | 6.8200 | 5.9600 | 0.8600 | 4.7200 | 4.2400 |
| S2 | 151/175 | 5.0700 | 4.7400 | 0.3300 | 5.3800 | 4.6400 |
| S3 | 140/137 | 2.5100 | 2.0400 | 0.4700 | 3.2200 | 2.5900 |
| S4 | 179/184 | 3.2000 | 2.7000 | 0.5000 | 2.4600 | 2.3200 |
| S5 | 169/174 | 5.8100 | 6.0900 | -0.2800 | 5.1400 | 4.8600 |
| S6 | 736/754 | 4.7600 | 4.7200 | 0.0400 | 5.2900 | 5.3300 |
| S7 | 209/209 | 10.9000 | 10.1000 | 0.8000 | 7.9000 | 8.1000 |
| S8 | 1122/1151 | 3.0100 | 2.8200 | 0.1900 | 3.3200 | 3.0500 |
| S9 | 673/679 | 4.3700 | 3.8800 | 0.4900 | 5.3700 | 4.8500 |
| [Combined] | | | | | | |
| Average | | | | 0.2835 | | |

Note: This report shows the input data for each study in the analysis. The 'Average' values are actually weighted averages with weights based on the effects model that was selected.

This report summarizes the input data. You should scan it for any mistakes. Note that the 'Average' line provides the estimated group average.

NT/NC

These are the count values that were read from the database.

Mean T

These are the input treatment means.

Mean C

These are the input control means.

Difference

These are the computed values of Mean T minus Mean C. These difference values are the effects of interest in the analysis.

Meta-Analysis of Means

SD T

These are the input treatment standard deviations.

SD C

These are the input control standard deviations.

Nondirectional Zero-Effect Test

| Rows | Outcome Measure | Chi-Square | DF | Prob Level |
|----------|-----------------|------------|----|------------|
| Combined | Mean Difference | 14.8417 | 9 | 0.0954 |

This reports the results of the nondirectional zero-effect chi-square test designed to test the null hypothesis that all treatment effects are zero. The null hypothesis is written

$$H_0: \theta_i = 0 \quad i = 1, \dots, k$$

The alternative hypothesis is that at least one $\theta_i \neq 0$, that is, at least one study had a statistically significant result.

Chi-Square

This is the computed chi-square value for this test. The formula was presented earlier.

DF

This is the degrees of freedom. For this test, the degrees of freedom is equal to the number of studies.

Prob Level

This is the significance level of the test. If this value is less than the nominal value of alpha (usually 0.05), the test is statistically significant and the alternative is concluded. If the value is larger than the specified value of alpha, no conclusion can be drawn other than that you do not have enough evidence to reject the null hypothesis.

Directional Zero-Effect Test

| Rows | Outcome Measure | Chi-Square | DF | Prob Level |
|----------|-----------------|------------|----|------------|
| Combined | Mean Difference | 9.4395 | 1 | 0.0021 |

This reports the results of the directional zero-effect chi-square test designed to test the overall null hypothesis that all treatment effects are zero. The null hypothesis is written

$$H_0: \theta_i = 0 \quad i = 1, \dots, k$$

The alternative hypothesis is that $\theta_i = \theta \neq 0$ for all i , that is, that all effects are equal to the same, non-zero value.

Chi-Square

This is the computed chi-square value for this test. The formula was presented earlier.

DF

This is the degrees of freedom. For this test, the degrees of freedom is equal one.

Prob Level

This is the significance level of the test. If this value is less than the specified value of alpha (usually 0.05), the test is statistically significant and the alternative is concluded. If the value is larger than the specified value of alpha, no conclusion can be drawn other than that you do not have enough evidence to reject the null hypothesis.

Meta-Analysis of Means

Effect-Equality (Heterogeneity) Test

| Treatment | Outcome Measure | Cochran's Q | DF | Prob Level |
|-----------|-----------------|-------------|----|------------|
| Combined | Mean Difference | 5.4022 | 8 | 0.7139 |

This reports the results of the effect-equality (homogeneity) test. This chi-square test was designed to test the null hypothesis that all treatment effects are equal. The null hypothesis is written

$$H_0: \theta_i = \theta \quad i = 1, \dots, k$$

The alternative is that at least one effect is different, that is, that the effects are heterogeneous. This may also be interpreted as a test of the study-by-treatment interaction. This test may help you determine whether to use a Fixed Effects model (used for homogeneous effects) or a Random Effects model (heterogeneous effects).

Cochran's Q

This is the computed chi-square value for Cochran's Q statistic. The formula was presented earlier.

DF

This is the degrees of freedom. For this test, the degrees of freedom is equal to the number of studies minus one..

Prob Level

This is the significance level of the test. If this value is less than the specified value of alpha (usually 0.05), the test is statistically significant and the alternative is concluded. If the value is larger than the specified value of alpha, no conclusion can be drawn other than that you do not have enough evidence to reject the null hypothesis.

Mean Difference Detail Section

| Study | Mean Difference | Standard Error | 95.0% Lower Confidence Limit | 95.0% Upper Confidence Limit | Percent Random Effects Weight |
|-------------------|-----------------|----------------|------------------------------|------------------------------|-------------------------------|
| S1 | 0.8600 | 0.5704 | -0.2579 | 1.9779 | 2.6172 |
| S2 | 0.3300 | 0.5549 | -0.7577 | 1.4177 | 2.7648 |
| S3 | 0.4700 | 0.3516 | -0.2191 | 1.1591 | 6.8887 |
| S4 | 0.5000 | 0.2509 | 0.0082 | 0.9918 | 13.5238 |
| S5 | -0.2800 | 0.5400 | -1.3384 | 0.7784 | 2.9199 |
| S6 | 0.0400 | 0.2752 | -0.4993 | 0.5793 | 11.2455 |
| S7 | 0.8000 | 0.7826 | -0.7340 | 2.3340 | 1.3900 |
| S8 | 0.1900 | 0.1337 | -0.0720 | 0.4520 | 47.6529 |
| S9 | 0.4900 | 0.2782 | -0.0554 | 1.0354 | 10.9974 |
| [Combined] | | | | | |
| Average | 0.2835 | 0.0923 | 0.1026 | 0.4643 | |

This report displays results for the mean difference outcome measure.

Confidence Limits


These are the lower and upper confidence limits (the formulas were given earlier in this chapter).

Weights

The last column gives the relative (percent) weight used in creating the weighted average. Using these values, you can decide how much influence each study has on the weighted average.

Meta-Analysis of Means


Forest Plot



This plot presents the results for each study on one plot. The size of the plot symbol is proportional to the sample size of the study. The points on the plot are sorted by the mean difference. The lines represent the confidence intervals about the mean differences. Note that the narrower the confidence limits, the better.

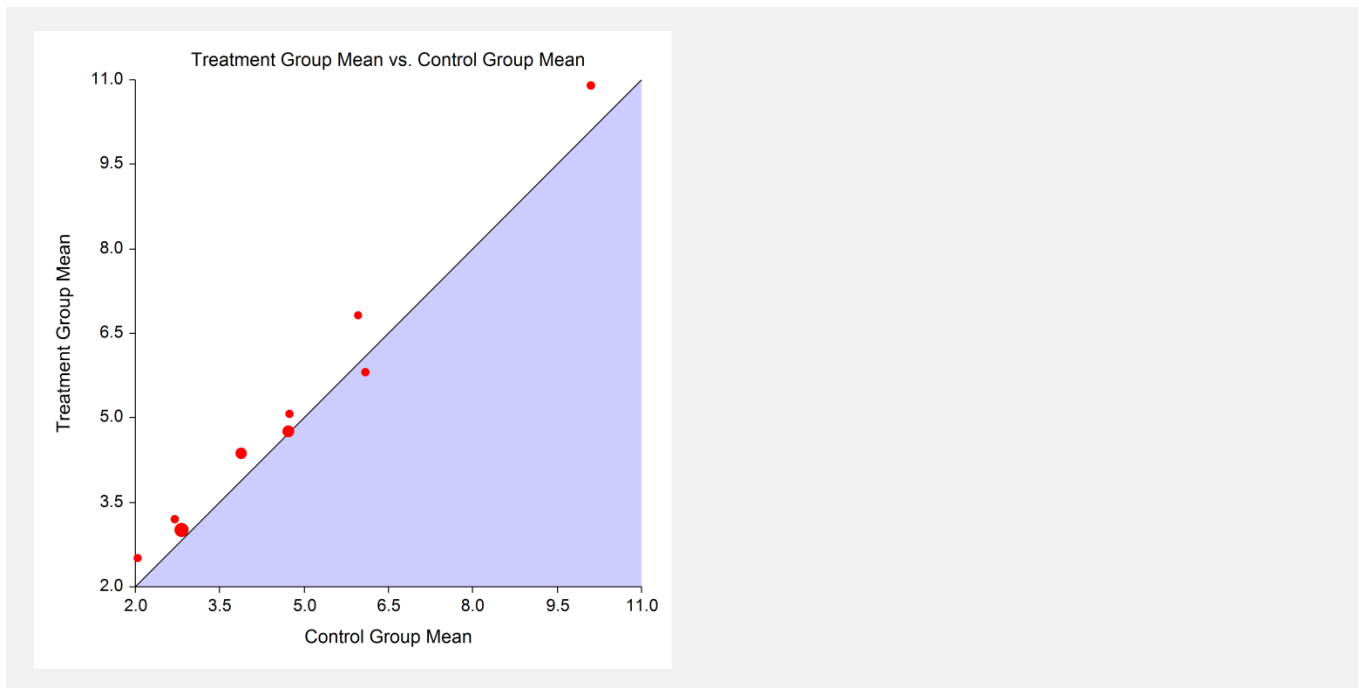
By studying this plot, you can determine the main conclusions that can be drawn from the set of studies. For example, you can determine how many studies were significant (the confidence limits do not intersect the vertical line at 0.0).

Radial Plot



The radial (or Galbraith) plot shows the z-statistic (outcome divided by standard error) on the vertical axis and a measure of weight on the horizontal axis. Studies that have the largest weight are closest to the Y axis. Studies within the limits are interpreted as homogeneous. Studies outside the limits may be outliers.

L'Abbe Plot



The L'Abbe plot displays the treatment mean on vertical axis versus the control mean on the horizontal axis. Homogenous studies will be arranged along the diagonal line. This plot is especially useful in determining if the relationship between the treatment group and the control group is the same for all values of the control group risk.