

Funnel Plots and Publication Bias in Meta-Analyses

I. The Basics of Funnel Plots and Publication Bias

Publication bias refers to the selective publication of statistically significant or positive study findings. A **funnel plot** is a traditional tool used to visually assess evidence of publication bias (or more aptly, small study effects, as we will discuss below).

Funnel plots typically plot the estimated effect size from each study (x-axis) against the estimated standard error for that study (y-axis). Each point on the plot represents a single study.

For example, let's implement a meta-analysis using the same data we used in the handout for Lab 11. For illustration, we will implement a fixed-effects meta-analysis.

```
library(metafor)
d = dat.bangertdrowns2004[ !is.na( dat.bangertdrowns2004$yi ), ]


m = rma(yi = d$yi,
        vi = d$vi,
        ni = d$ni,
        knha = TRUE,
        method = "FE")

summary(m)
```

```
##
## Fixed-Effects Model (k = 48)
##
##   logLik  deviance      AIC      BIC     AICc
## -26.2784  107.1061   54.5568   56.4280   54.6438
##
## I2 (total heterogeneity / total variability):  56.12%
## H2 (total variability / sampling variability):  2.28
##
## Test for Heterogeneity:
## Q(df = 47) = 107.1061, p-val < .0001
##
## Model Results:
##
## estimate      se      tval  df      pval   ci.lb   ci.ub
##   0.1656   0.0407   4.0739  47   0.0002   0.0838   0.2474   ***
##
## ---
## Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
```

The estimated mean effect obtained via meta-analysis is 0.1656. Now let's look at the funnel plot.

```
# funnel plot
funnel(m)
```



The Estimated Mean Obtained Via Meta-Analysis

The vertical dotted line which runs through the center of the triangular funnel represents the **estimated mean effect** (i.e., the pooled meta-analytic estimate of the effect across all studies, which we will denote $\hat{\mu}$). That is, in the funnel plot shown above, this line corresponds to the estimated meta-analytic mean of 0.1656 we obtained previously.

Note that, when constructing the funnel plot, this mean estimate is usually obtained using a **fixed-effects** (rather than a random-effects) meta-analysis, regardless of which method was used in the primary analysis. Why? A meta-analytic estimate of the mean which is based on a random-effects meta-analysis gives a **larger relative weight to smaller studies** than a fixed-effect meta-analysis would. As a result, compared to estimates obtained from a fixed-effects meta-analysis, meta-analytic estimates obtained from a random-effects meta-analysis may be more detrimentally affected by publication biases which affect small studies more than large studies. For the purposes of visually assessing evidence of publication bias in the funnel plot (e.g., an asymmetrical scattering of points around the vertical line representing the mean estimate), a fixed-effects meta-analytic mean is therefore usually preferred.

To understand why a random-effects meta-analysis give a larger relative weight to small studies than a fixed-effects meta-analysis, let's review the inverse variance weights that determine how much each study's estimate contributes to the meta-analytic mean estimate.

For a fixed-effects meta-analysis, each study i is weighted by the inverse of its estimated study-specific variance, $\hat{\sigma}_i^2$:

$$w_i = \frac{1}{\hat{\sigma}_i^2}$$

In a random-effects meta-analysis, each study i is weighted by the inverse of the *sum* of the estimated variance of the population effects across studies, $\hat{\tau}^2$, plus the estimated study-specific variance, $\hat{\sigma}_i^2$:

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

The **relative weight** that the i^{th} study receives (in either a fixed or random-effects meta-analysis) is defined as:

$$\text{Relative weight for study } i = \frac{w_i}{\sum_{i=1}^k w_i}$$

Where k is the number of studies included in the meta-analysis.

We said that a random-effects meta-analysis will give a larger relative weight to a small study than a fixed-effects meta-analysis would. To see why this is the case, let's imagine we conduct a meta-analysis with just $k = 3$ studies. The first study is a 'small' study which has an estimated variance of $\hat{\sigma}_1^2 = 0.83$. The second and third studies are 'large' studies which have estimated variances of $\hat{\sigma}_2^2 = 0.19$ and $\hat{\sigma}_3^2 = 0.22$, respectively. The weight for each of these three studies in a fixed-effects meta-analysis would then be:

$$\begin{aligned} \text{Weight for study 1} = w_1 &= \frac{1}{\hat{\sigma}_1^2} = \frac{1}{0.83} = 1.2048 \\ \text{Weight for study 2} = w_2 &= \frac{1}{\hat{\sigma}_2^2} = \frac{1}{0.19} = 5.2632 \\ \text{Weight for study 3} = w_3 &= \frac{1}{\hat{\sigma}_3^2} = \frac{1}{0.22} = 4.5454 \end{aligned}$$

We can now calculate the relative weight of the first study (i.e., the 'small' study) as follows:

$$\text{Relative weight for study 1} = \frac{w_1}{\sum_{i=1}^k w_i} = \frac{1.2048}{1.2048 + 5.2632 + 4.5454} = 0.1094$$

Now let's see how these weights would change if we performed a random-effects meta-analysis where $\hat{\tau}^2 = 0.21$. Note that, for these results to hold, $\hat{\tau}^2$ could be any number greater than 0 (i.e., as long as there is heterogeneity, these results hold).

The weight for each of the three studies in this random-effects meta-analysis would be:

$$\begin{aligned} \text{Weight for study 1} = w_1 &= \frac{1}{\hat{\tau}^2 + \hat{\sigma}_1^2} = \frac{1}{0.21 + 0.83} = 0.9615 \\ \text{Weight for study 2} = w_2 &= \frac{1}{\hat{\tau}^2 + \hat{\sigma}_2^2} = \frac{1}{0.21 + 0.19} = 2.500 \\ \text{Weight for study 3} = w_3 &= \frac{1}{\hat{\tau}^2 + \hat{\sigma}_3^2} = \frac{1}{0.21 + 0.22} = 2.3256 \end{aligned}$$

The relative weight of the first study is:

$$\text{Relative weight for study 1} = \frac{w_1}{\sum_{i=1}^k w_i} = \frac{0.9615}{0.9615 + 2.500 + 2.3256} = 0.1661$$

We have just shown that the relative weight of the 'small' study is approximately 10.9% when using a fixed-effects approach but is 16.6% when using a random-effects approach! As we noted above, **a summary estimate obtained from a random-effects meta-analysis provides greater relative weight to smaller studies.**

The Pseudo 95% Confidence Region

Now let's think a bit about the two diagonal dotted lines which form opposite sides of the funnel. These two lines form the bounds of a **pseudo 95% confidence region** around the mean effect estimate. Specifically, these lines represent the lower and upper bounds of the expected confidence interval calculated for each standard error value represented on the y-axis. The expected confidence interval bounds are calculated using the following formula:

$$\hat{\mu} \pm 1.96 \times SE$$

Where $\hat{\mu}$ is the estimated mean effect and SE is the standard error corresponding to a given value of the y-axis. To illustrate how these calculations are performed, let's investigate two simple examples from the meta-analysis we performed above. First, let's calculate the values that these lines take at the very top of the plot, where $SE = 0$. When $SE = 0$, the bounds of the pseudo 95% confidence region are:

$$0.1656 \pm 1.96 \times 0 = (0.1656, 0.1656)$$

That is, as we would expect based on the funnel plot, the two dotted diagonal lines are exactly equal to the estimated mean effect (the vertical dotted line) when $SE = 0$.


The standard errors increase as we move down the y-axis, and so does the width of the corresponding pseudo 95% confidence region. Near the bottom of the previous funnel plot, $SE = 0.515$ and the corresponding bounds of the pseudo 95% confidence interval at this SE are then:

$$0.1656 \pm 1.96 \times 0.515 = (-0.844, 1.175)$$

Alternative Funnel Plot Presentations

While this is the typical format for a funnel plot, other formats are possible. For example, we could plot the estimated effect size from each study on the x-axis against the **variance** of those estimates on the y-axis:

```
# funnel plot
funnel(m, yaxis="vi")
```



More details on the `funnel()` function and its arguments can be found at:

- https://www.metafor-project.org/doku.php/plots:funnel_plot_variations

For additional information on “standard” funnel plot presentation and why it became “standard,” see:

- <https://pubmed-ncbi-nlm-nih-gov.ezp-prod1.hul.harvard.edu/11576817/>

II. Limitations of Funnel Plot and Alternative Approaches

How might we use funnel plots to assess evidence of publication bias? If there is no publication bias, then at any horizontal cross-section of this plot, we might expect the point estimates from the different studies to be randomly and symmetrically scattered around the estimated meta-analytic mean, creating a symmetrical funnel shape.

However, funnel plots also have several important limitations. Funnel plots are helpful in detecting **publication bias that affects small studies more than large studies**. For this reason, funnel plots would be more aptly described as a visual method to assess evidence of small study effects. Why might a funnel plot fail to illustrate publication bias that affects all studies regardless of size? Well, if publication bias affected all studies regardless of size, then the entire funnel and the corresponding mean effect estimate might be shifted towards larger effect sizes. In this case, there may be no evidence of asymmetry in the distribution of the points around the mean estimated effect because the entire distribution has been shifted.

Funnel plots also **assume no heterogeneity** in the true underlying population effects. When there is in fact heterogeneity, this may lead to asymmetry in the funnel plot that is not due to publication bias (e.g., if the true treatment effect is larger in smaller studies, then the effect estimates from small studies may seem inflated or asymmetrical around the estimated mean). In addition, especially when the number of included studies is relatively small, funnel plots may illustrate asymmetry simply due to **random, stochastic variations**.

Finally, funnel plots may fail to detect publication bias that operates solely based on statistical significance testing.

To overcome some of the shortcomings of funnel plots, we might consider alternative approaches. For example, publication bias can be assessed using selection models. These models assume that publication bias results in the preferential publication of statistically significant, positive study results regardless of study size. For more on these approaches, see:

- McShanre, B.B., Bockenholt, U., & Hansen, K.T. Adjusting for publication bias in meta-analysis: An evaluation of selection methods and some cautionary notes. *Perspectives on Psychological Science*, 11(5), 730-749.
- Vevea, J. L., & Hedges, L. V. (1995). A general linear model for estimating effect size in the presence of publication bias. *Psychometrika*, 60(3), 419-435.

Sensitivity analyses may also be used to characterize how severe publication bias would need to be in order to explain away the effect (think an E-value for meta-analysis!). These approaches may also be used to provide a meta-analytic effect estimate that is corrected for ‘worst-case scenario’ publication bias. For more details, see:

- Mathur, M. B., & VanderWeele, T. J. (2020b). Sensitivity analysis for publication bias in meta-analyses. *Journal of the Royal Statistical Society: Series C (Applied Statistics)*, 69(5), 1091-1119.