Northwestern Medicine®
Feinberg School of Medicine

A gentle introduction to meta-analysis

Applied Statistics Seminar

April 19, 2022

Jacob M. Schauer, Department of Preventive Medicine– Biostatistics, Northwestern University

What is this talk about?

Meta-analysis and statistics, right?

- What are effect sizes, and why should we care?
- Why would we combine effect sizes and how?
- What are some things to consider when combining effect sizes?
- Complexifying issues in analysis

- Some examples with the R package metafor.
  - Also weightr for one example
  - Feel free to code along while I talk!

What are the statistics of meta-analysis?

- We have already identified relevant studies and data within studies.
  - Dr. Muhammad's *Statistically Speaking* talk.

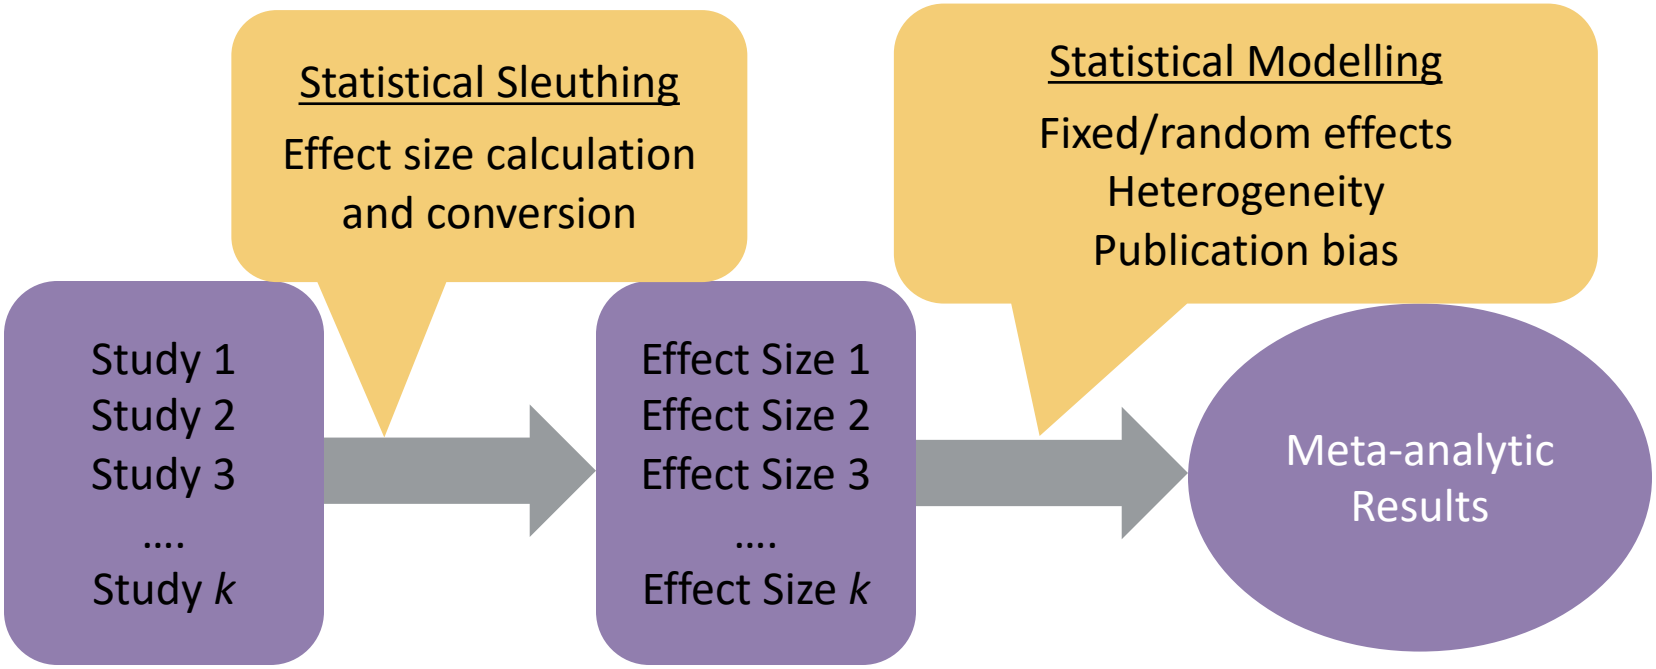

Example

Dagostino, 1998

9 studies examining the impact of antihistamines on runny nose severity for the common cold

- Outcome: Change in runny nose severity after 2 days
  - 4 different scales (0-3, 0-4, 0-8, 0-10)
- 2 different drugs: chlorpheniramine and doxylamine
- Some studies find statistically significant effects, some don't
- One study finds a negative effect

How do we make sense of all of these?

Northwestern Medicine®
Feinberg School of Medicine

Effect Sizes

What are they? Can we combine them?

What is an effect size?

Some statistical considerations for a single study

- **Estimand θ**
  - "True" effect
  - Parameter
- **Estimates *T***
  - Function of the data
- **Variance σ^2**
  - Standard error: σ
  - Sampling or estimation error variance that decreases with sample size
- Confidence/credible intervals

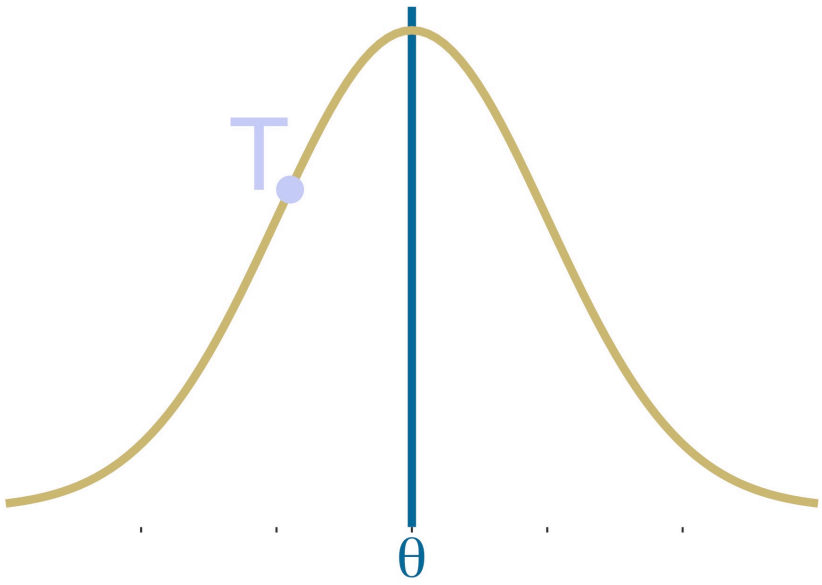

What is an effect size?

Some statistical considerations for multiple studies

For $i = 1, \ldots, k$

- **Estimands θ_i**
  - "True" effects
  - Parameters
- **Estimates T_i**
  - Functions of the data
- **Variances σ_i^2**
  - Standard errors: σ_i
  - Sampling or estimation error variance that decrease with sample size
- Confidence/credible intervals

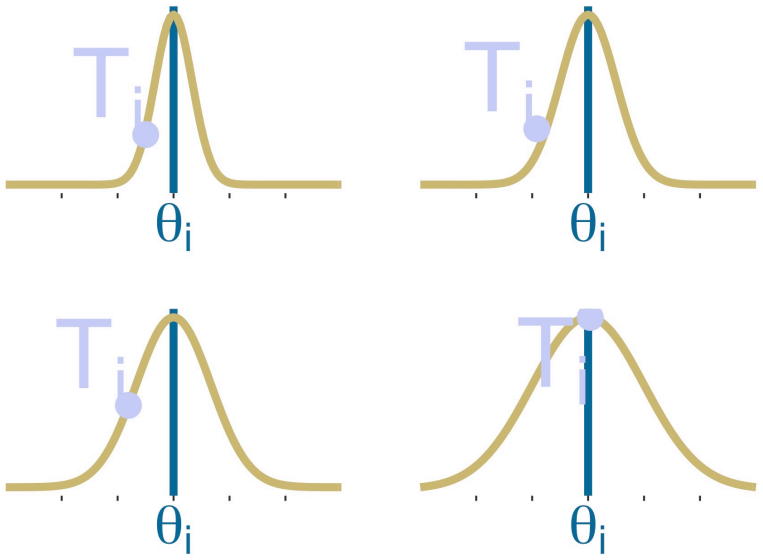

What we talk about when we talk about effect sizes

Some statistical considerations for multiple studies

- Estimands θ_i (*effect size parameter*)
- Estimates T_i (*effect size estimate*)
- The scale of estimands and estimates (*effect size index*)
  - Consider a two-armed study (**T**reatment vs. **C**ontrol)
    - T-C mean difference: $\mu_T - \mu_C$
    - T-C standardized mean difference (Cohen's *d* is "scale-free"…kind of): $(\mu_T - \mu_C)/\varsigma$
    - T-C odds ratio (log transform), risk difference, …
  - Correlation coefficient (arctan transform)

Where do effect sizes come from?

- To run a meta-analysis we need both the effect estimate T_i and variance σ_i^2.
- To compute an effect size estimates and variances, we need data:
  - Raw data (unlikely for every or even most studies)
  - **Summary statistics**
    - Often reported in primary research

Effect size calculation is not always trivial

Example: Cohen's *d* in a 2-armed RCT

- $Y_{iT} \sim N(\mu_T, \varsigma^2)$ and $Y_{iC} \sim N(\mu_C, \varsigma^2)$
  - $i = 1, \ldots n_T$ and $i = 1, \ldots, n_C$
  - $n = n_T + n_C$
- Cohen's $d = (\mu_T - \mu_C)/\varsigma$
- Estimate (Glass, 1976)
- *Bias correction (Hedges' g):* $g = \frac{\Gamma(n-2/2)}{\sqrt{\frac{n-2}{2}}\Gamma(n-3/2)} \frac{\bar{Y}_T - \bar{Y}_C}{\sqrt{\frac{(n_T-1)s_T^2+(n_C-1)s_C^2}{n-2}}}$
- *Approximate bias correction:* $g \approx \frac{4n-12}{4n-9} \frac{\bar{Y}_T - \bar{Y}_C}{\sqrt{\frac{(n_T-1)s_T^2+(n_C-1)s_C^2}{n-2}}}$
- Required to compute:
  - Treatment and control means
  - Treatment and control sample sizes
  - Treatment and control standard deviations (or some pooled SD)

Effect size calculation is not always trivial

Example: Cohen's *d* in a 2-armed RCT

- *Bias correction (Hedges' g):* $g = \frac{\Gamma(n-2/2)}{\sqrt{\frac{n-2}{2}}\Gamma(n-3/2)} \frac{\bar{Y}_T - \bar{Y}_C}{\sqrt{\frac{(n_T-1)s_T^2+(n_C-1)s_C^2}{n-2}}}$
- *Approximate bias correction:* $g \approx \frac{4n-12}{4n-9} \frac{\bar{Y}_T - \bar{Y}_C}{\sqrt{\frac{(n_T-1)s_T^2+(n_C-1)s_C^2}{n-2}}}$
- It is possible to obtain the pooled standard deviation via test statistics:
  - $t = \frac{\bar{Y}_T - \bar{Y}_C}{\sqrt{s^2\left(\frac{1}{n_T}+\frac{1}{n_C}\right)}} \cong \frac{\bar{Y}_T - \bar{Y}_C}{\sqrt{\frac{(n_T-1)s_T^2+(n_C-1)s_C^2}{n-2}}} \frac{1}{\sqrt{\left(\frac{1}{n_T}+\frac{1}{n_C}\right)}}$
- It is possible to get (approximate) pooled SD from the SE:
  - $SE \cong \sqrt{s^2\left(\frac{1}{n_T}+\frac{1}{n_C}\right)}$

Effect size calculation is not always trivial

Example: (log) odds ratio in a 2-armed RCT

- $Y_T \sim B(n_T, \pi_T)$ and $Y_C \sim B(n_C, \pi_C)$
  - $n_T + n_C = n$
- Odds ratio $\lambda = \frac{\pi_T/(1-\pi_T)}{\pi_C/(1-\pi_C)}$
- Estimate $\frac{Y_T/(n_T-Y_T)}{Y_C/(n_C-Y_C)}$
  - $z = log\left(\frac{Y_T/(n_T-Y_T)}{Y_C/(n_C-Y_C)}\right)$
- Asymptotically, $\log(z) \sim N(\log(\lambda), \frac{1}{Y_T} + \frac{1}{n_T-Y_T} + \frac{1}{Y_C} + \frac{1}{n_C-Y_C})$

A quick look at effect size calculations in R

Introduction to metafor

- Dagostino (1998) impact of antihistamines on runny nose severity for the common cold.

```
library(metafor)
dag_es <- escalc(
    measure = "SMD", # "OR", "RR", etc.
    m1i = mt, # treatment means
    m2i = mc, # control means
    sd1i = sdt, # treatment SDs
    sd2i = sdc, # control SDs
    n1i = nt, # treatment group sample size
    n2i = nc # control group sample size
    data = metafor::dat.dagostino1998 %>%
filter(outcome == "rnic2")
)
```

What do we need to know about effect sizes?

- Effect parameters should be conceptually similar enough to consider jointly.
- Effect size indices need to be the same across studies.
  - What effect size index makes sense?
    - In our example, outcomes pertain to the same construct (runny nose), but are on different scales (e.g., 0-3, 0-4, 0-8, or 0-10).
    - We can put them on similar scales via Cohen's *d*.
  - It is often possible (with some normality assumptions) to convert from one scale to another:
    - *d* <–> log(odds ratio) <–> correlation
    - It may not be conceptually appropriate to change scales even if it is technically feasible.

How do we visualize data in a meta-analysis?

Forest plot

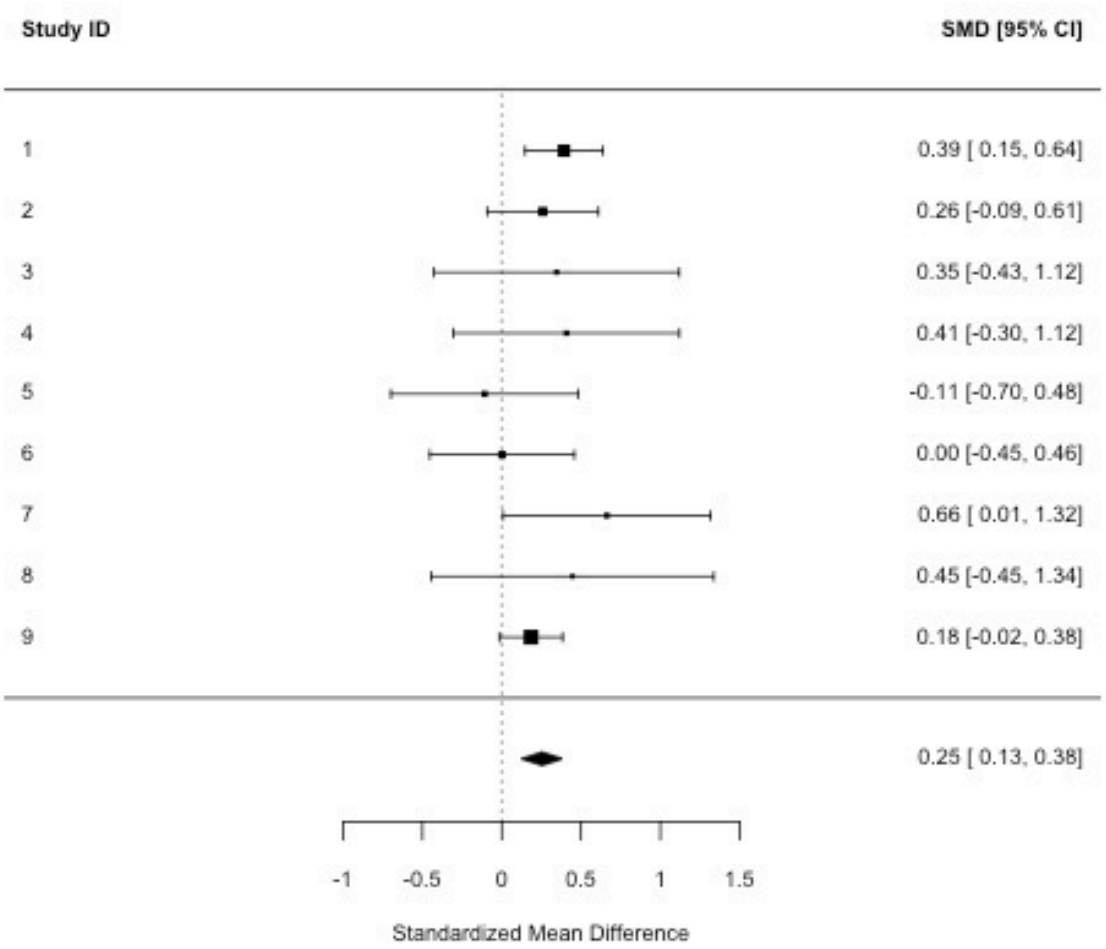

How do we visualize data in a meta-analysis?

Forest plot

```
# quick meta-analysis fit, we'll come back to this later
remod <- rma(yi = yi, vi = vi, data = dag_es, method = "PM", knha = TRUE)
# make a forest plot
forest(remod, cex=.75, header="Study ID",
       mlab="", slab = dag_es$study)
```

Effect sizes

Summary

- We need effect estimates and variances (or SEs)
- They need to be on the same scale
  - We can often convert between effect size scales
- We should start by visualizing with a forest plot

Northwestern Medicine®
Feinberg School of Medicine

Combining Effect Sizes

Goals of analyses, and basics for estimation

What is an effect size?

Some statistical considerations for multiple studies

- **Estimands θ_i**
  - "True" effects
  - Parameters
- **Estimates T_i**
  - Functions of the data
- **Variances σ_i^2**
  - Standard errors: σ_i
  - Sampling or estimation error variance that decrease with sample size
- Confidence/credible intervals

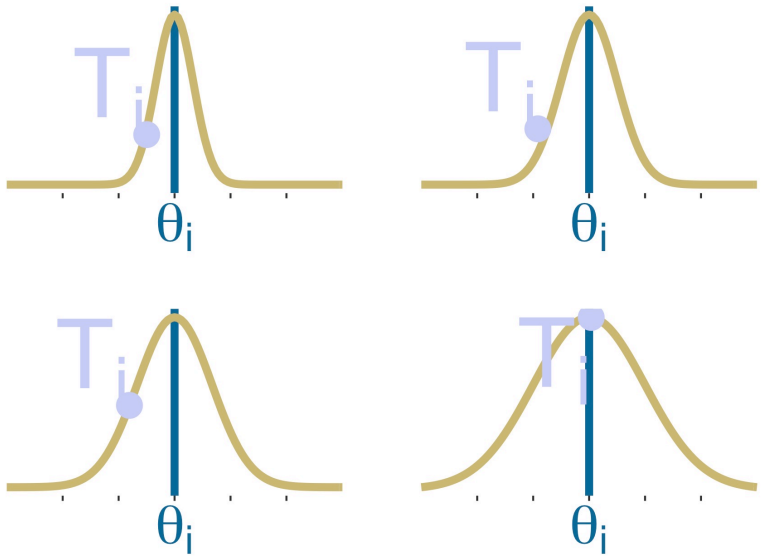

Assumptions about the studies/effects will govern if we do a ***fixed-*** or ***random-effects*** meta-analysis

Fixed-effects meta-analysis

Strong assumptions ahead!

- Early statistical theory in the 1980s focused on fixed-effects models:
  - $\theta_1 = \theta_2 = \ldots = \theta_k = \theta$
- Inferential goal: Estimate θ and report SE/CI, etc.

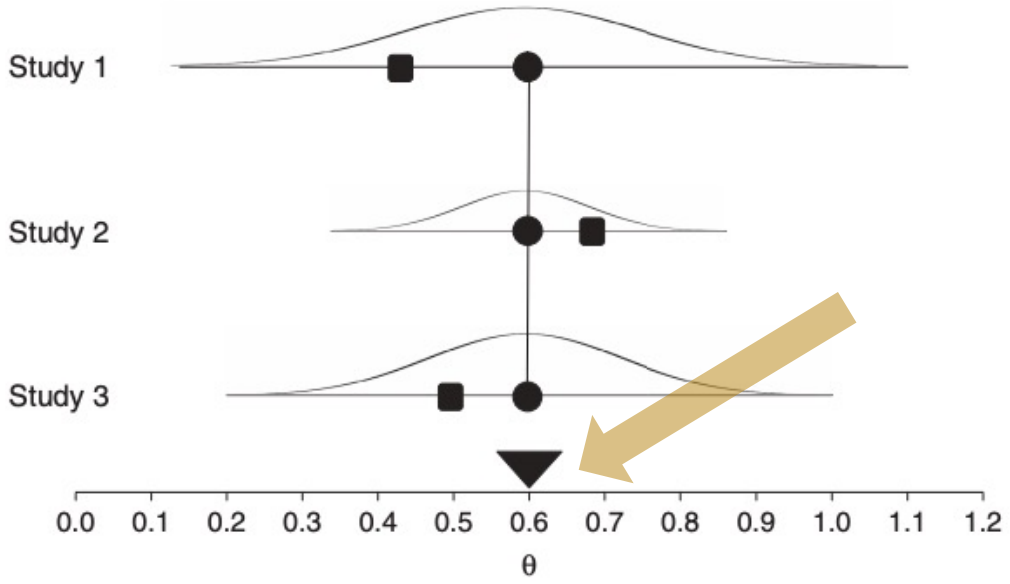

Fixed-effects meta-analysis

Strong assumptions ahead!

- Early statistical theory in the 1980s focused on fixed-effects models:
  - $\theta_1 = \theta_2 = \ldots = \theta_k = \theta$
- Assumes that studies are identical enough to produce identical effects.
  - Evidence from direct replications suggests we can't always do this even if we're explicitly trying to do so.

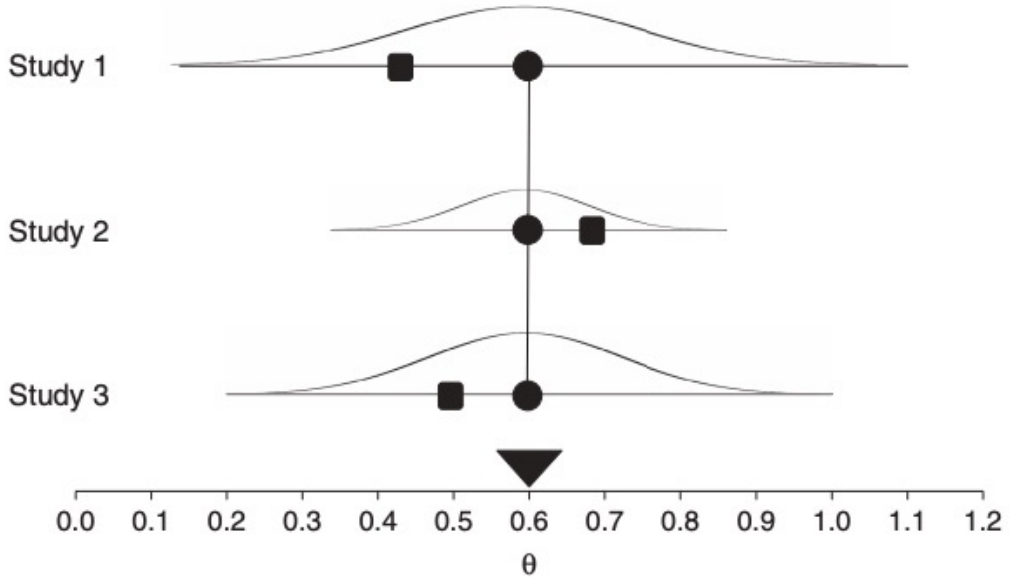

Random-effects meta-analysis

Weaker assumptions

- Assumes $\theta_i \neq \theta_j$, instead the θ_i vary randomly:
  - $\theta_i \sim N(\mu, \tau^2)$
    - Need not be normal, but it's a common assumption.
- Assumes that studies are a random sample from some population.

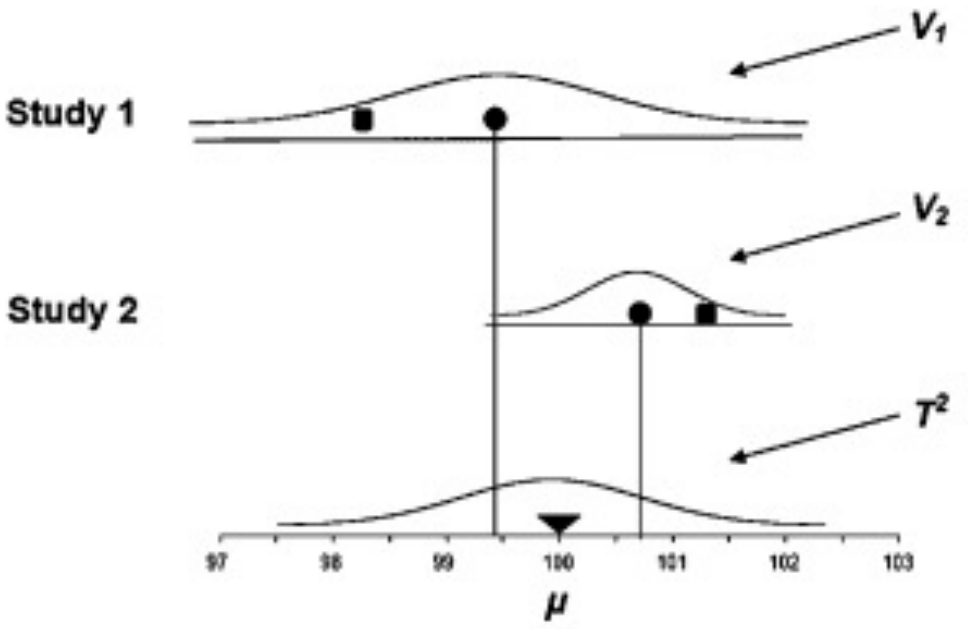

Random-effects meta-analysis

Weaker assumptions

- Assumes $\theta_i \neq \theta_j$, instead the θ_i vary randomly:
  - $\theta_i \sim N(\mu, \tau^2)$
    - Need not be normal, but it's a common assumption.
- Inferential goal: Estimate μ, τ^2 and report SE/CI, etc.
  - Report intervals likely to contain future values of θ_i (prediction interval)

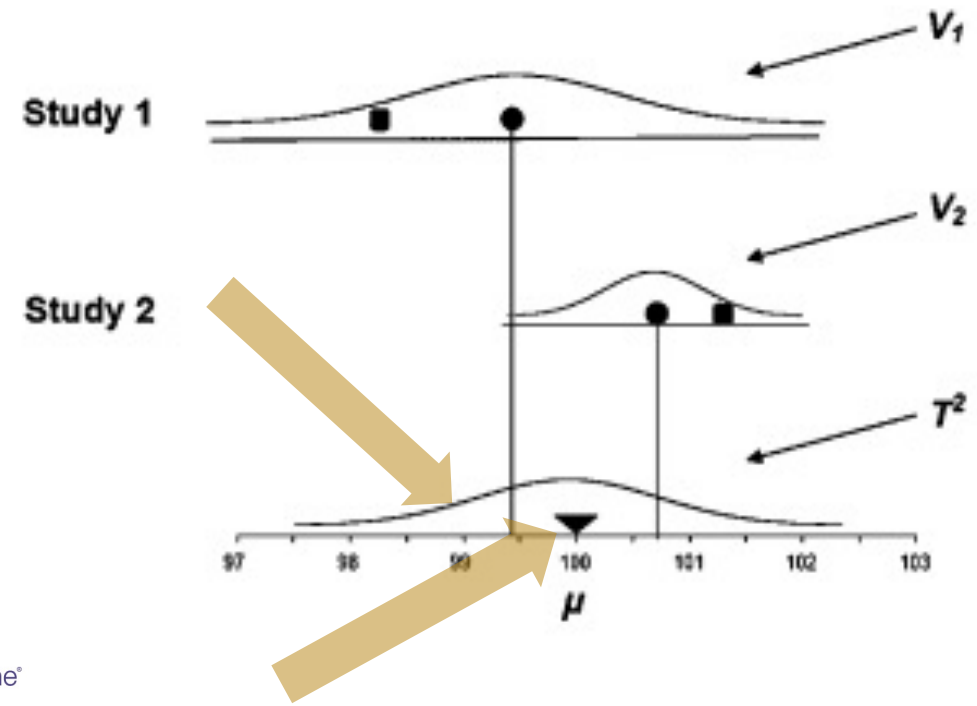

Comparing fixed- and random-effects analyses

Fixed-effects analysis

- Estimate mean effect
  - Assumes common underlying effect across all studies
- One source of variation
  - Within-study (sampling) variation

Random-effects analysis

- Estimate mean effect
  - Mean of a distribution of effect parameters
  - Prediction interval for future effects
- Two sources of variation
  - Within-study (sampling) variation
  - Between-study variation
- Estimate between-study variation

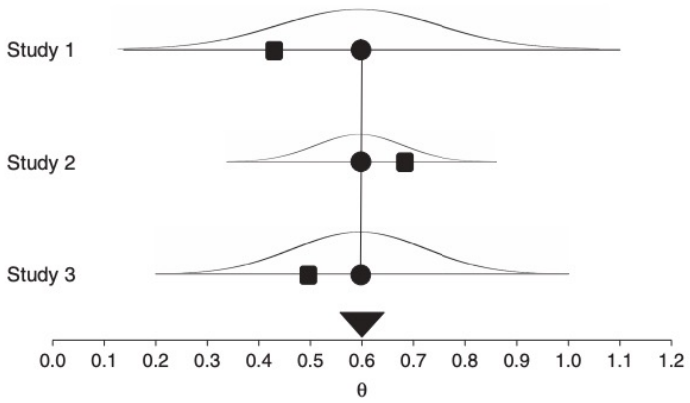

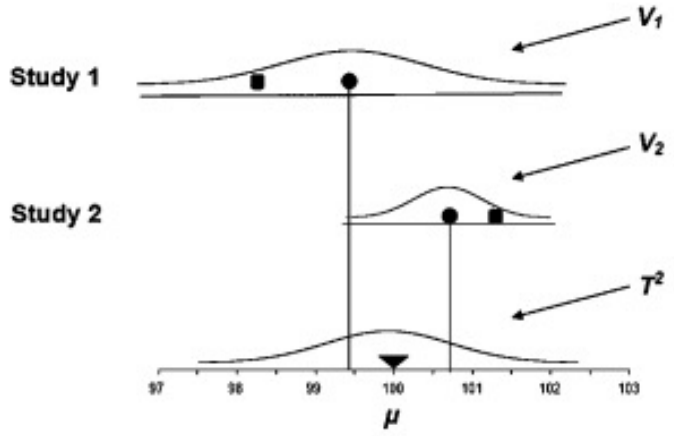

Should I use fixed- or random-effects models?

- The *Q*-test for heterogeneity tests H_0: $\theta_1 = \theta_2 = \ldots = \theta_k$
  - $Q = \sum_{i=1}^{k} \frac{\left(T_i - \sum_{i=1}^{k} T_i / \sigma_i^2\right)^2}{\sigma_i^2} \sim \chi^2_{k-1}$
- However, the *Q*-test has low power unless there are a large number of effects ($k > 50$-80).
- Unless there is a large # of effects, *Q* is not advised for discerning between model specification.
- Instead, choice should be consistent with beliefs about the studies
  - My default: random-effects
  - Caveat: there needs to be enough studies to estimate the between-study variation ($k > 5$-10)

Meta-analysis model

$$T_i \sim N(\mu, \tau^2 + \sigma_i^2)$$

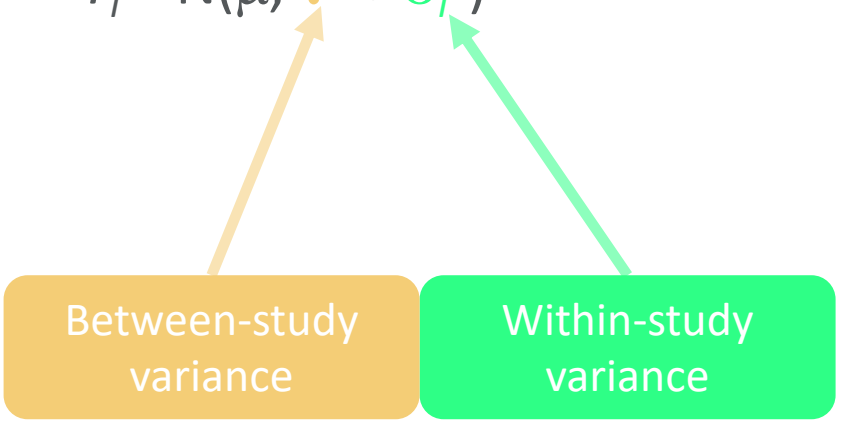

- Target of inference is the distribution of the effect parameters characterized by μ and τ^2

Meta-analysis model

$$T_i = \mu + r_i + e_i \text{ where } r_i \sim N(0, \tau^2) \text{ and } e_i \sim N(0\ \sigma_i^2)$$

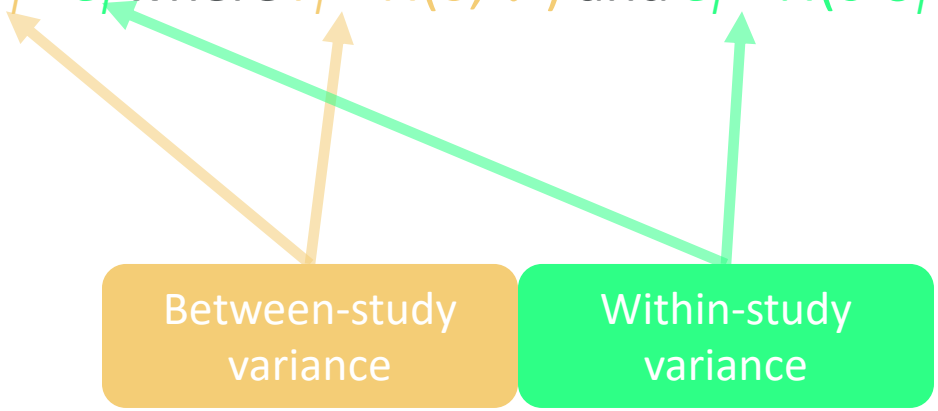

- Target of inference is the distribution of the effect parameters characterized by μ and τ^2

Meta-analysis model

$T_i \sim N(\mu, \tau^2 + \sigma_i^2)$ $T_i = \mu + r_i + e_i$ where $r_i \sim N(0, \tau^2)$ and $e_i \sim N(0\ \sigma_i^2)$

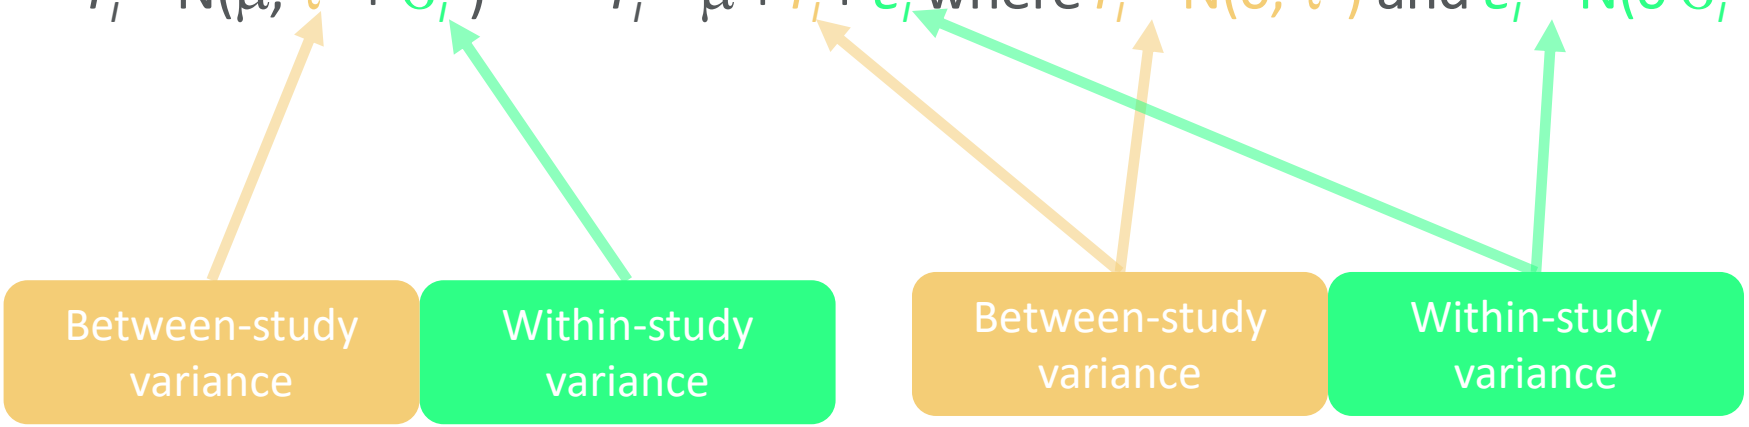

- Target of inference is the distribution of the effect parameters characterized by μ and τ^2

Meta-analysis model

- $T_i \sim N(\mu, \tau^2 + \sigma_i^2)$ $\quad\quad$ $T_i = \mu + r_i + e_i$ where $r_i \sim N(0, \tau^2)$ and $e_i \sim N(0\ \sigma_i^2)$
- UMVUE (and MLE) of μ
  - $\bar{T}_\cdot = \frac{\sum_{i=1}^{k} w_i T_i}{\sum_{i=1}^{k} w_i}$ where $w_i = \frac{1}{\tau^2 + \sigma_i^2}$
  - $\bar{T}_\cdot$ is asymptotically normal with variance $V[\bar{T}_\cdot] = \frac{1}{\sum_{i=1}^{k} w_i}$
    - NHST H_0: $\mu = 0$
    - 95% CI for μ
- No UMVUE for τ^2
  - REML
  - Moment estimators: DerSimonian-Laird, **Paule-Mandel**, etc.

Meta-analysis model

Estimation

- $T_i \sim N(\mu, \tau^2 + \sigma_i^2)$
- $\bar{T}_. = \frac{\sum_{i=1}^{k} w_i^* T_i}{\sum_{i=1}^{k} w_i^*}$ where $w_i^* = \frac{1}{\hat{\tau}^2 + \sigma_i^2}$
  - $\bar{T}_.$ is asymptotically normal with variance $\hat{V}[\bar{T}_.] = \frac{1}{\sum_{i=1}^{k} w_i^*}$
    - NHST H_0: $\mu = 0$
      - Use a **Knapp-Hartung correction** (like a t-test)
    - 95% CI for μ
- Use $\bar{T}_.$ and $\hat{\tau}^2$ to make inferences about the *distribution* of future θ_i
  - 95% prediction interval

Weighting

Mean effects are estimated using

$$w_i^* = \frac{1}{\hat{\tau}^2 + \sigma_i^2}$$

More weight goes to T_i with smaller σ_i^2

In a FE model, $\tau^2 = 0$, so $w_i^* = \frac{1}{\sigma_i^2}$

- Larger variation in weights
- Mean pulled harder toward some T_i

In a RE model, $w_i^* = \frac{1}{\hat{\tau}^2 + \sigma_i^2}$

- Less variation in weights than FE
- Mean pulled less strongly

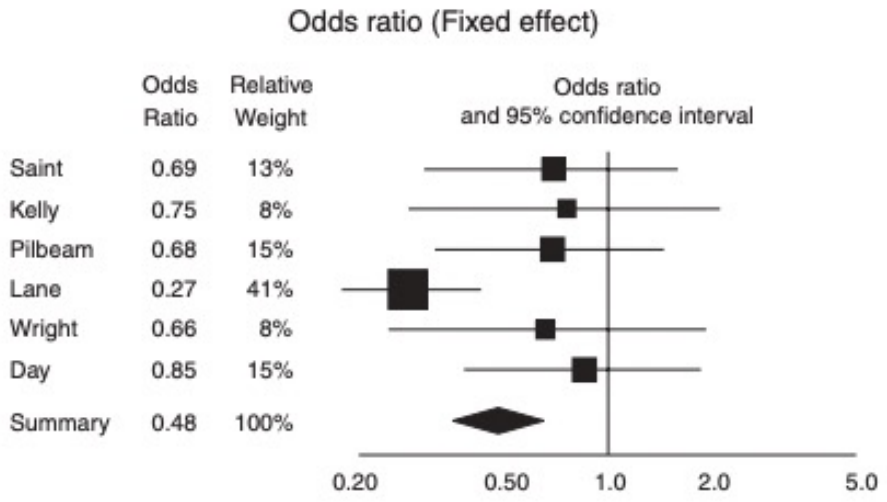

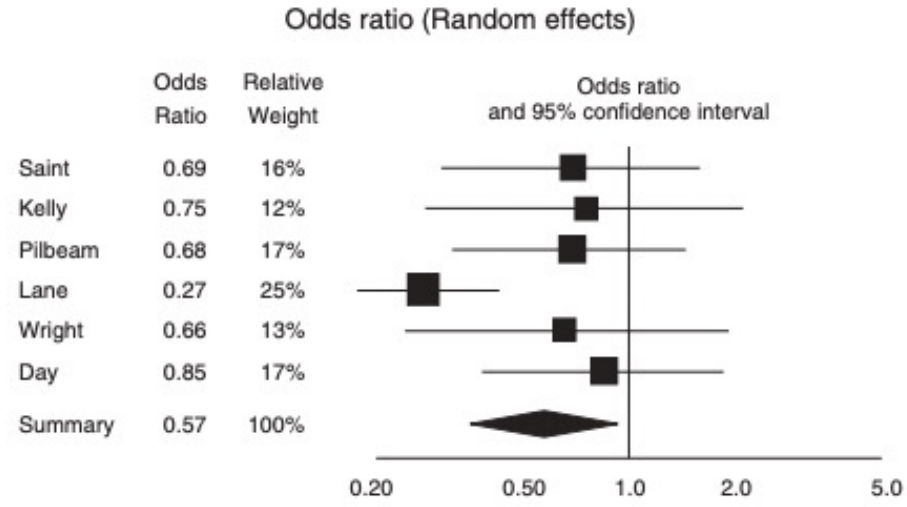

A note on inference for between-study variance

- The scale of τ^2 depends on the scale of the θ_i
  - Alternatively, we can quantify τ^2 in a manner that is "scale-free" relative to the "typical" within-study variance σ^2
    - H^2 estimates $1 + \frac{\tau^2}{\sigma^2} = \frac{\tau^2+\sigma^2}{\sigma^2} = \frac{total\ variation}{within-study\ varation}$
    - I^2 estimates $\frac{\tau^2}{\tau^2+\sigma^2} = \frac{between-study\ variation}{total\ varation}$
  - I^2 values > 30-40% are often considered "meaningful"
  - H^2 values > 1.33-1.75 are considered "large" or "meaningful"

A typical meta-anlysis

- Estimate of the mean effect μ
  - SE, CI
  - NHST that $\mu = 0$
    - Use KNHA adjustment!
- Estimate of the variance
  - I^2
  - H^2
- Prediction Interval

```
remod <- rma(
    yi = yi, # effect estimates
    vi = vi, # variances of effect estimates
    data = dag_es,  # dataset
    method = "PM",  # use the Paule-Mandel RE model
    knha = TRUE # small-sample adjustment for tests
)
summary(remod) # view results
predict(remod) # get prediction interval
```

A typical meta-anlysis

Results

- Estimate of the mean effect μ
  - SE, CI
  - NHST that $\mu = 0$
    - Use KNHA adjustment!
- Estimate of the variance
  - I^2
  - H^2
- Prediction Interval

Random-Effects Model (k = 9; tau^2 estimator: PM)

logLik	deviance	AIC	BIC	AICc
1.1942	6.2132	1.6117	2.0061	3.6117

tau^2 (estimated amount of total heterogeneity): 0 (SE = 0.0181)

tau (square root of estimated tau^2 value): 0

I^2 (total heterogeneity / total variability): 0.00%

H^2 (total variability / sampling variability): 1.00

Test for Heterogeneity:

Q(df = 8) = 6.2132, p-val = 0.6234

A typical meta-anlysis

Results

- Estimate of the mean effect μ
  - SE, CI
  - NHST that $\mu = 0$
    - Use KNHA adjustment!
- Estimate of the variance
  - I^2
  - H^2
- Prediction Interval

Model Results:

estimate	se	tval	df	pval	ci.lb	ci.ub	
0.2539	0.0558	4.5480	8	0.0019	0.1252	0.3827	**

Signif. codes: 0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1

A typical meta-anlysis

Results

- Estimate of the mean effect μ
  - SE, CI
  - NHST that $\mu = 0$
    - Use KNHA adjustment!
- Estimate of the variance
  - I^2
  - H^2
- Prediction Interval

pred	se	ci.lb	ci.ub	pi.lb	pi.ub
0.2539	0.0558	0.1252	0.3827	0.1252	0.3827

Summarize everything with a forest plot

```
# make a forest plot
forest(remod, cex=.75, header="Study ID",
      mlab="", slab = dag_es$study)
# add text with Q-value, dfs, p-value, and I^2 statistic
text(-16, -1, pos=4, cex=0.75,
      bquote(paste("RE Model (Q = ",
                        .(formatC(remod$QE, digits=2, format="f")),
                  ", df = ", .(remod$k - remod$p),
                  ", p = ", .(formatC(remod$QEp, digits=2, format="f")),
                  "; ", I^2, " = ",
                              .(formatC(remod$I2, digits=1, format="f")), "%)")))
```

A typical meta-analysis

Results

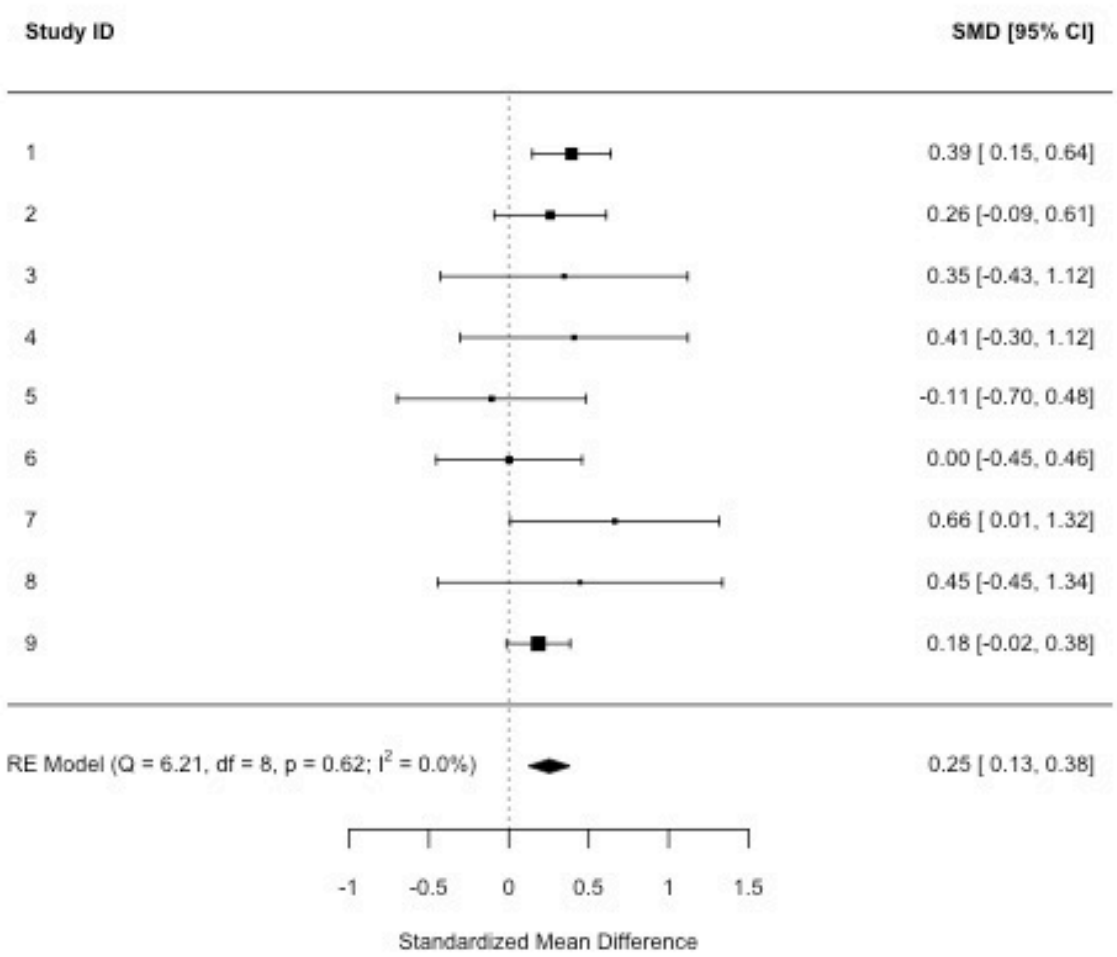

Northwestern Medicine®
Feinberg School of Medicine

Selection and publication bias

Beware the published record

Are our effect sizes "representative"?

Rosenthal's File Drawer problem

- Systematic reviews are often dominated by published research.

- Are we only seeing a subset of relevant effect sizes?
  - Selective reporting within studies
    - "We reported the contrasts for which we found significant results"
  - Selective reporting of entire studies
    - "We didn't get a significant result so we didn't feel the need to publish"

Can we tell if we're missing "null" results?

Funnel Plots

- Studies with high statistical power are unlikely to have null results (assuming effects are nonzero).
- Studies with low statistical power are more likely to have null results.
- Studies with low statistical power tend to have higher within-study variation.

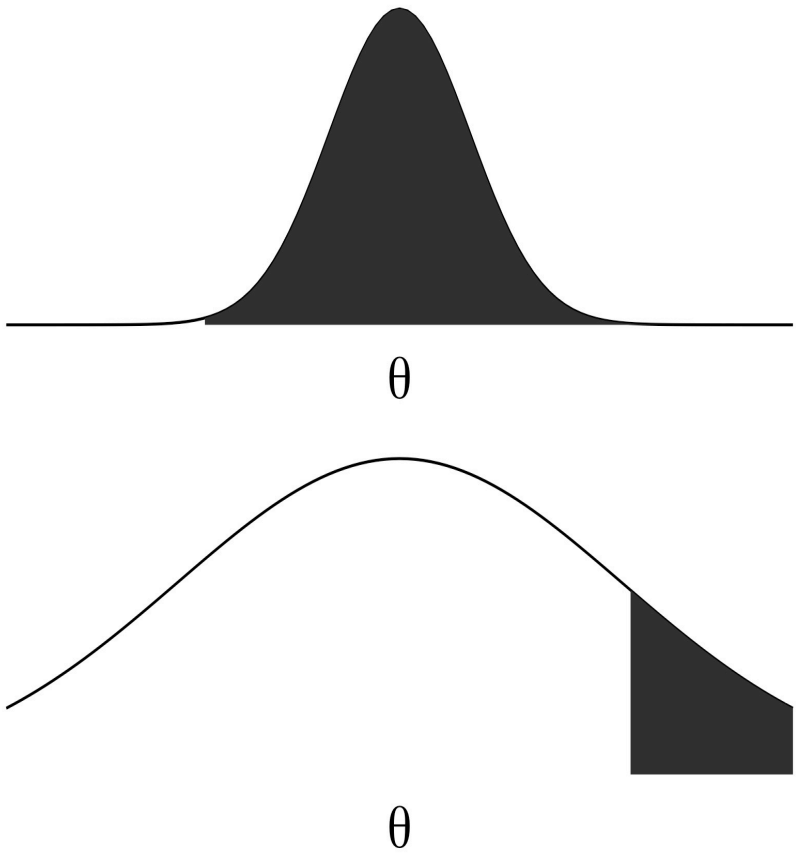

Can we tell if we're missing "null" results?

Funnel Plots

- Studies with high statistical power are unlikely to have null results (assuming effects are nonzero).
- Studies with low statistical power are more likely to have null results.
- Studies with low statistical power tend to have higher within-study variation.

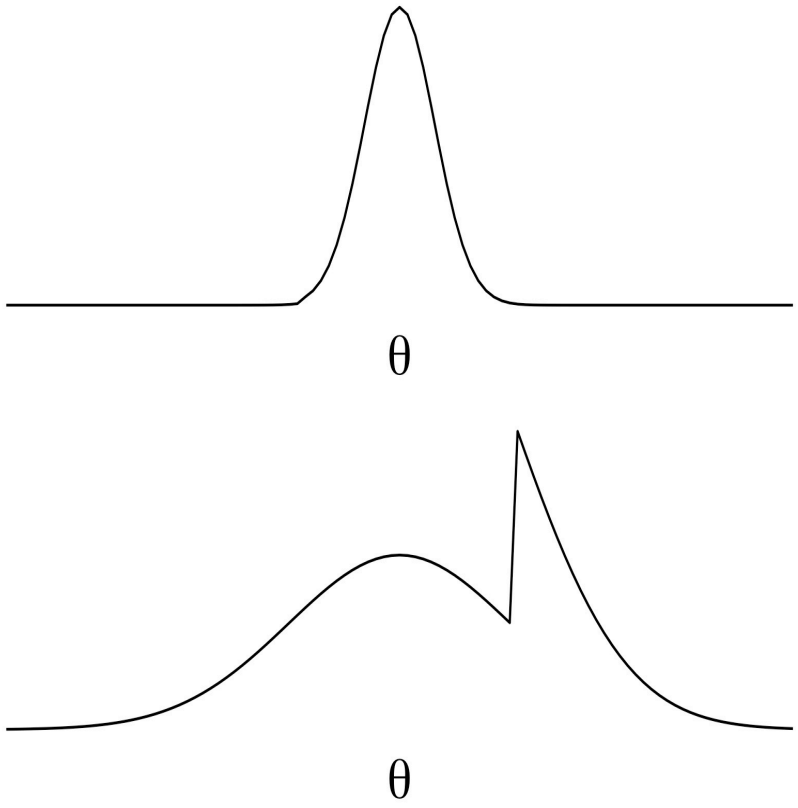

Funnel plots

funnel(remod)

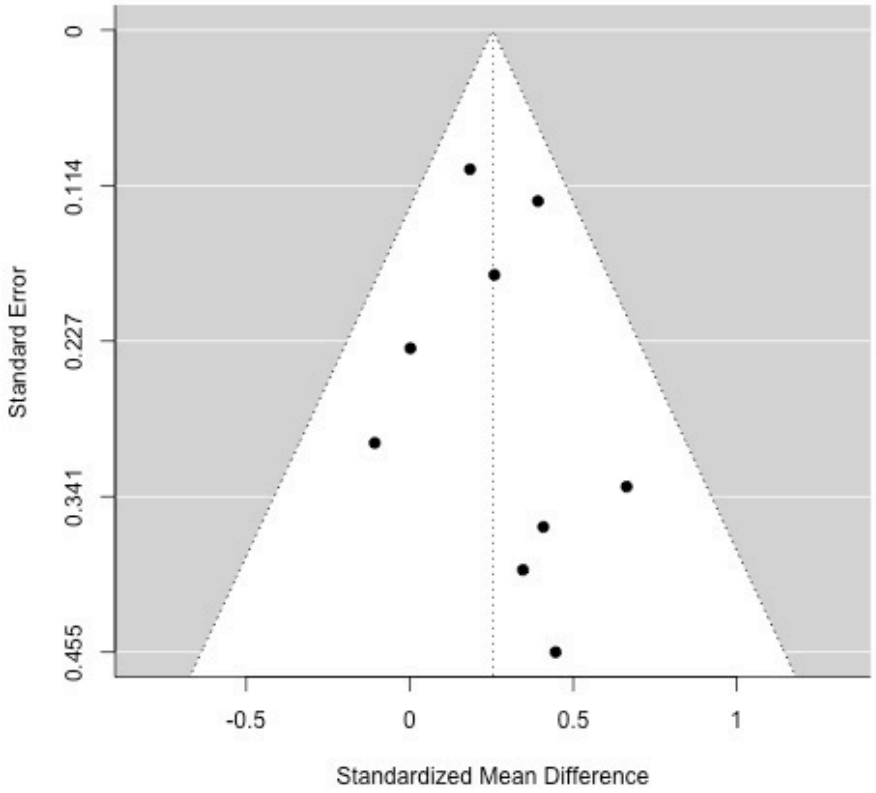

Tests for funnel plot asymmetry

- Egger's test
  1. Fit the model $T_i = \beta_0 + \beta_1 \sigma_i$
  2. Test H_0: $\beta_1 = 0$

Tests can also regress T_i on σ_i^2 or look at the rank correlation between T_i & σ_i

```
regtest(remod)
Regression Test for Funnel Plot Asymmetry

Model:     mixed-effects meta-regression model
Predictor: standard error

Test for Funnel Plot Asymmetry: t = 0.3521, df = 7, p = 0.7351
Limit Estimate (as sei -> 0):   b = 0.2160 (CI: -0.0750, 0.5069)
```

Adjustments

- Trim and fill
- Selection models (likelihood based approach)
- Can be seen as outright corrections to biased parameter estimates, or as sensitivity analyses.

- Avoid p-curve, PEESE, and PET-PEESE

Trim and fill

- Trim: remove some of the oversampled significant results
- Fill: impute “missing” nonsignficant results

Iterate

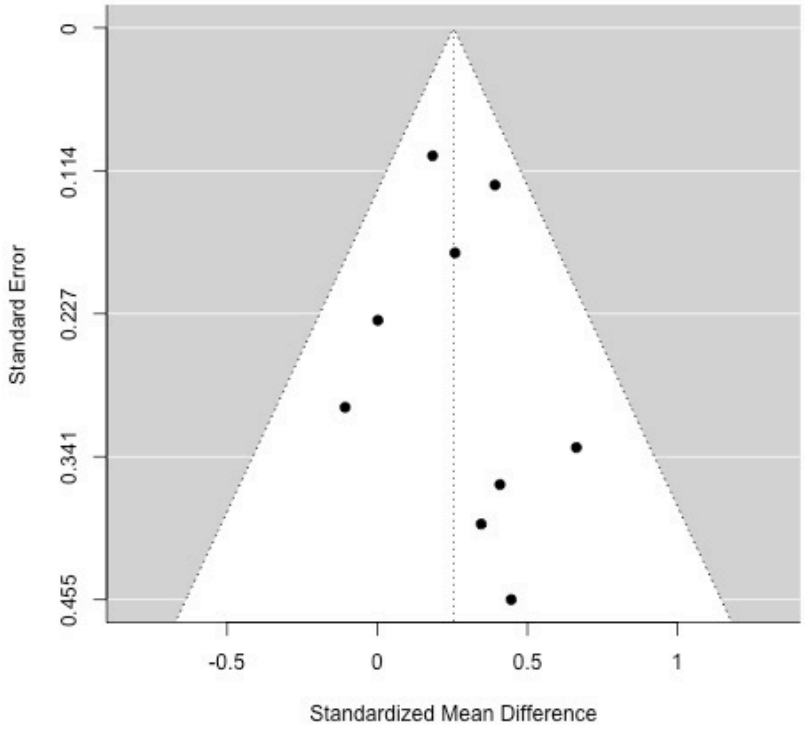

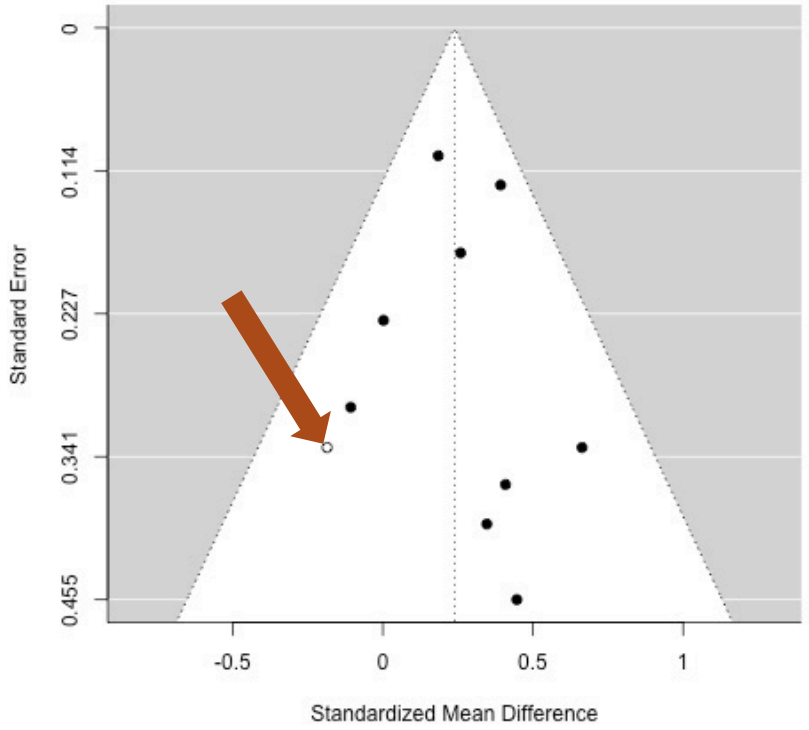

Selection models

- Likelihood based approach
  - Assume T_i are unconditionally normal
  - T_i is observed ($R_i = 1$) with some probability π given its *p*-value is <0.05 (3P model)
  - $p(T_i|R_i = 1) \propto$

$$\phi\left(T_i;\ \mu, \tau^2 + \sigma_i^2\right)\mathbf{1}\left\{\frac{|T_i|}{\sigma_i} \geq 1.96\right\} + \pi\phi\left(T_i;\ \mu, \tau^2 + \sigma_i^2\right)\mathbf{1}\left\{\frac{|T_i|}{\sigma_i} < 1.96\right\}$$

  - Likelihood-based estimates for μ, τ^2 and π

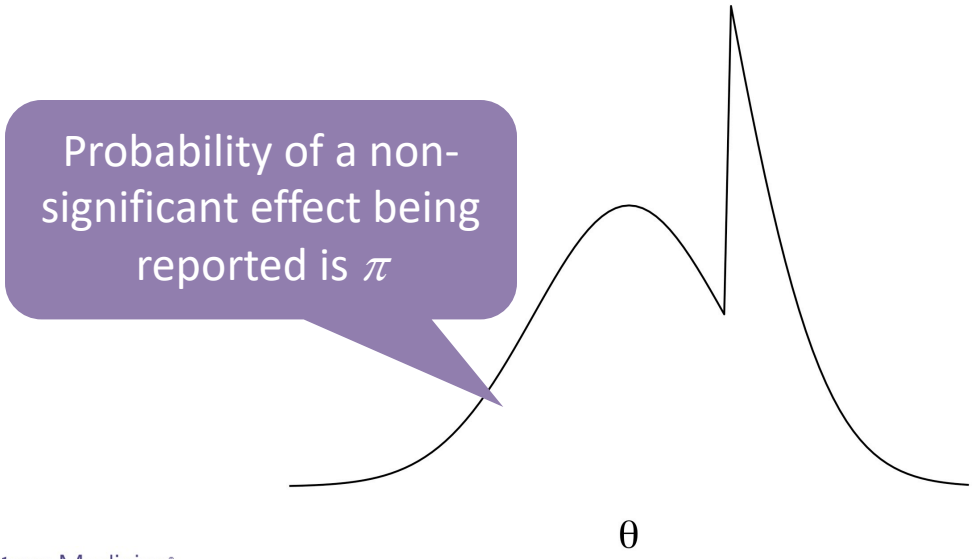

Example

Trim and fill

```
tf <- trimfill(remod)
summary(tf)
predict(tf)
funnel(tf)
```

```
Estimated number of missing studies on the left side: 1 (SE = 2.1192)

tau^2 (estimated amount of total heterogeneity): 0 (SE = 0.0183)
tau (square root of estimated tau^2 value):      0
I^2 (total heterogeneity / total variability):   0.00%
H^2 (total variability / sampling variability):  1.00

Test for Heterogeneity:
Q(df = 9) = 7.8900, p-val = 0.5453

Model Results:
estimate   se        zval      pval      ci.lb      ci.ub
0.2386     0.0622    3.8338    0.0001    0.1166     0.3606  ***
```

Example

Selection model

```
weightr::weightfunct(
    estimate = dag_es$yi,
    vi =  dag_es$vi,
    steps = c(0.05, 1)
)
```

Adjusted Model (k = 9):

tau^2 (estimated amount of total heterogeneity): 0.0000
(SE = 0.0195)
tau (square root of estimated tau^2 value): 0.0000

Test for Heterogeneity:
Q(df = 8) = 6.2132, p-val = 0.7184092

Model Results:

	estimate	std.error	z-stat	p-val	ci.lb	ci.ub
Intercept	0.2553	0.08518	2.9969	0.00273	0.08832	0.4222
0.05 < p < 1	1.0259	1.09665	0.9355	0.3495	-1.12346	3.1753

Likelihood Ratio Test:
X^2(df = 1) = 0.0005748692, p-val = 0.98087

Summary of publication bias

- If it makes sense, conduct and report assessments of publication bias (funnel plots, Egger's or Begg's test)
- If there appears to be some publication bias, conduct and report adjustments
  - Trim-and-fill
  - Selection models (for larger *k*)
  - Avoid p-curve, PEESE, and PET-PEESE

Northwestern Medicine®
Feinberg School of Medicine

Meta-regression

It's just like regular regression…sort of

Summarizing conditional distributions

Are effects parameters related to observed covariates?

- Meta-regression concerns the relationship between effect sizes and observed covariates
  - How was a treatment implemented?
  - Where did the study take place?
  - On whom?
- Used to answer important questions:
  - What is the treatment effect in populations >65 years-old?
  - Does dosage matter for treatment effects?
  - Is the correlation stronger in some countries, but not others?
- Referred to sometimes as “subgroup analysis” or “meta-regression”

Meta-regression model

$\theta_i = \mathbf{X}\beta + r_i$ $\qquad$ $T_i = \mathbf{X}\beta + r_i + e_i$

- where $r_i \sim N(0, \tau^2)$ and $e_i \sim N(0, \sigma_i^2)$

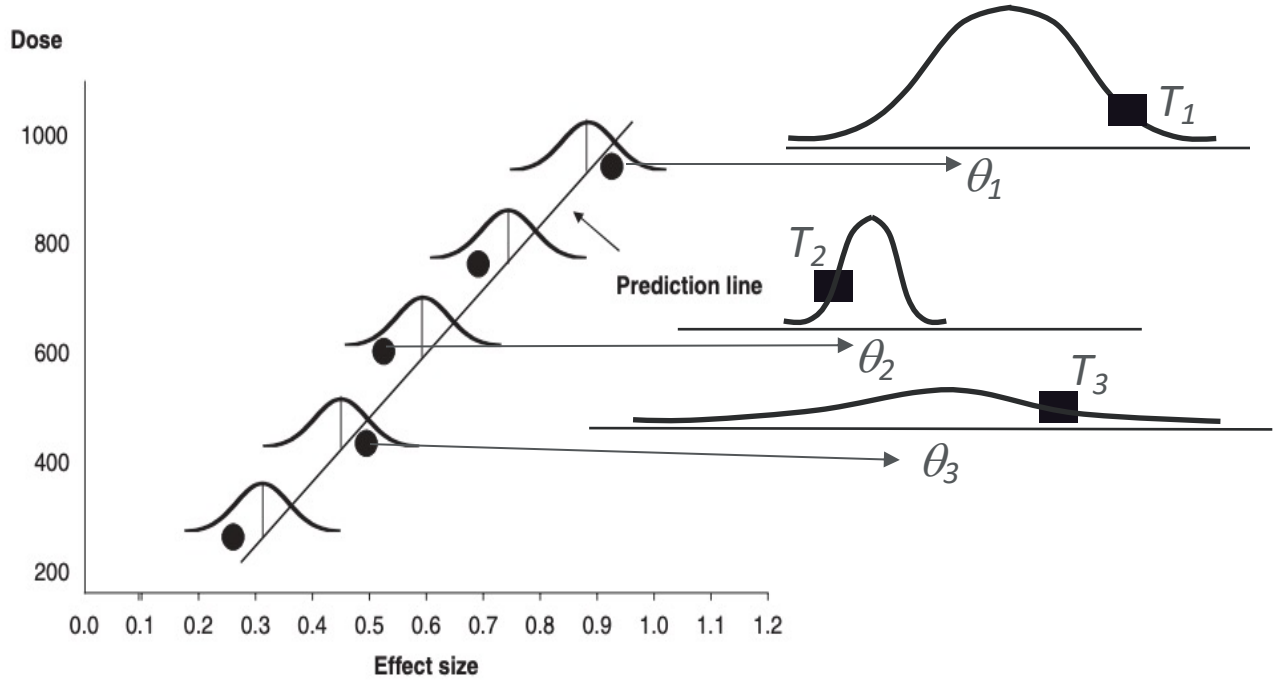

Estimation of meta-regression

- $\boldsymbol{T} \sim N(\boldsymbol{X}\beta, \boldsymbol{W}^{-1})$ where $\boldsymbol{W} = \text{diag}(\tau^2 + \sigma_i^2)$
- Basic WLS estimate: $\hat{\beta} = (\boldsymbol{X}'\boldsymbol{W}\boldsymbol{X})^{-1}\boldsymbol{X}'\boldsymbol{W}\boldsymbol{T}$
- In practice, use a moment-based estimator for τ^2 and plug-in to estimate of β
  - REML, Paule-Mandel, DerSimonian-Laird

Weighted least squares for meta-regression

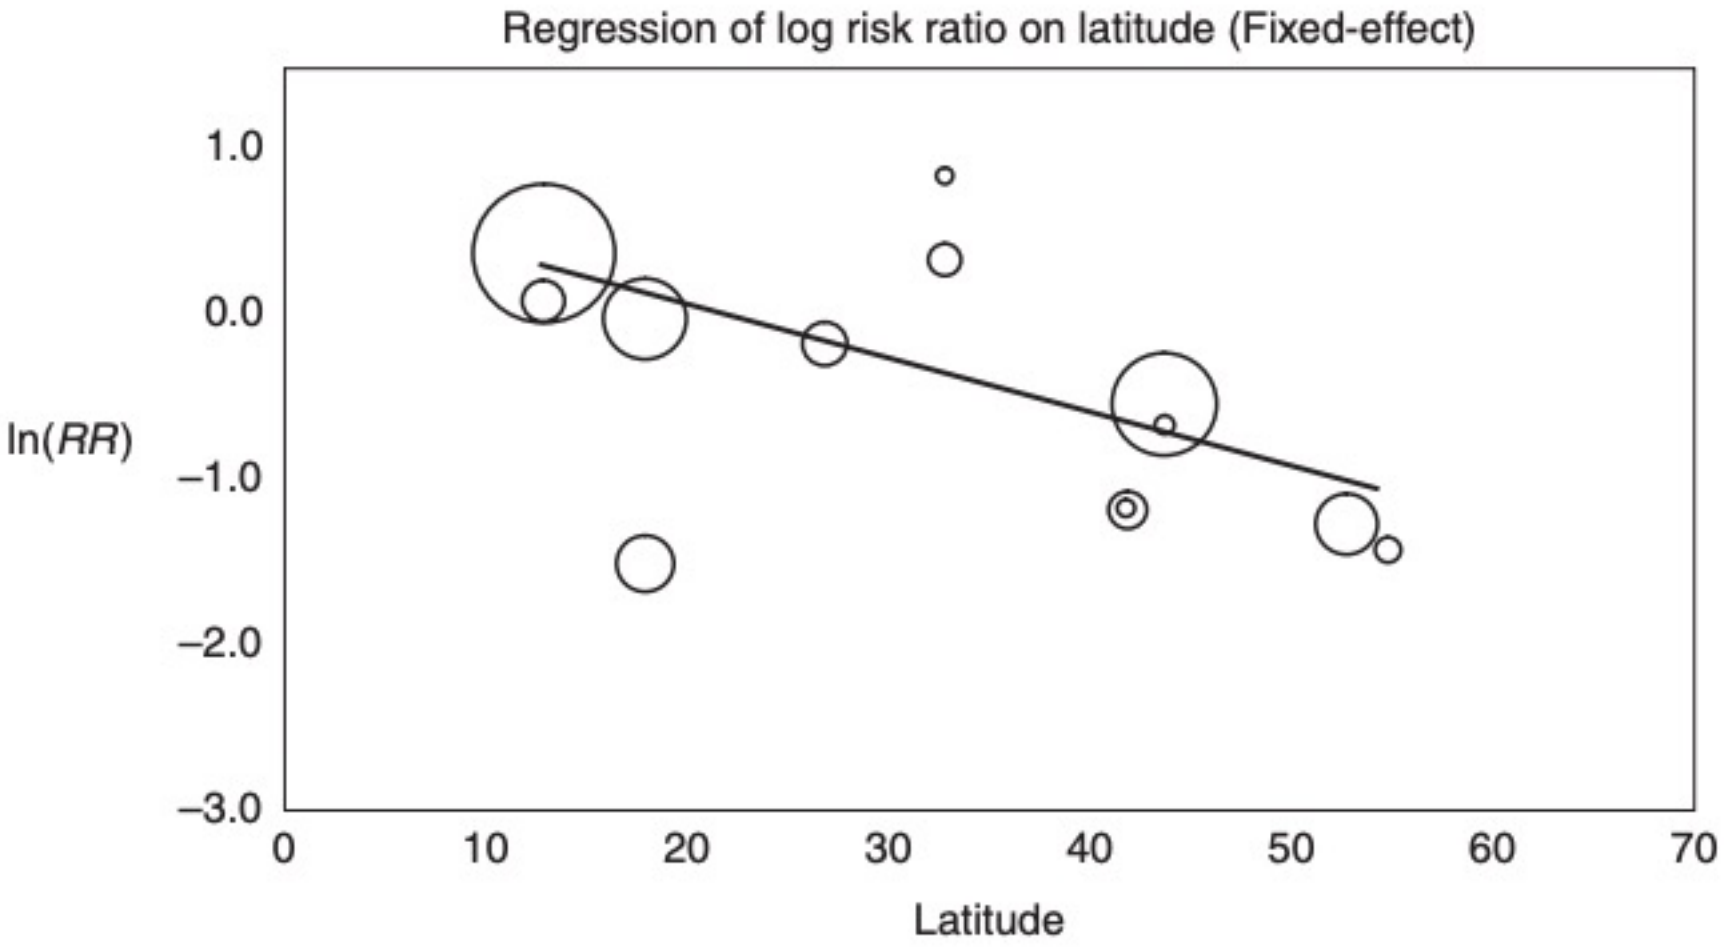

Inference for meta-regression

- Point estimates and SEs: $\hat{\beta}$ is consistent with variance $(\boldsymbol{X}'\boldsymbol{W}\boldsymbol{X})^{-1}$
- Omnibus test that all coefficients are 0
- Tests for individual coefficients
  - Knapp-Hartung corrections!
- Heterogeneity estimates (including I^2, H^2)
- Prediction intervals (given $\boldsymbol{X}$)

Example

Curtis 1998: Plant group and time of exposure

```
remod_mr <- rma(
    yi = yi,
    vi = vi,
    mods = ~ drug,
    data = dag_es,
    method = "PM",
    knha = TRUE
)
summary(remod_mr)
regtest(remod_mr)
```

```
weightr::weightfunc(
    estimate = dag_es$yi,
    vi =  dag_es$vi,
    steps = c(0.05, 1)
)
 # Cannot do trim-and-fill for
meta-regression
```

Example

Curtis 1998: Plant group and time of exposure

Mixed-Effects Model (k = 9; tau^2 estimator: PM)

tau^2 (estimated amount of residual heterogeneity): 0 (SE = 0.0193)
tau (square root of estimated tau^2 value): 0
I^2 (residual heterogeneity / unaccounted variability): 0.00%
H^2 (unaccounted variability / sampling variability): 1.00
R^2 (amount of heterogeneity accounted for): 0.00%

Test for Residual Heterogeneity:
QE(df = 7) = 4.6875, p-val = 0.6980

Example

Curtis 1998: Plant group and time of exposure

Test of Moderators (coefficient 2):

F(df1 = 1, df2 = 7) = 2.2783, p-val = 0.1749

Model Results:

	estimate	se	tval	df	pval	ci.lb	ci.ub	
intrcpt	0.3482	0.0812	4.2900	7	0.0036	0.1563	0.5401	**
drugdoxylamine	-0.1592	0.1055	-1.5094	7	0.1749	-0.4087	0.0902	

Northwestern Medicine®
Feinberg School of Medicine

Summary

Some points to consider

- Most meta-analyses involve
  - Mean effect estimate + inference
    - Weighted averages
    - Knapp-Hartung corrections
  - Heterogeneity estimate + inference
    - τ^2 , H^2, I^2,
  - Checks of funnel plots/publication bias
    - Egger's test
  - Publication bias corrections
    - Trim-and-fill
    - Selection weighting
  - Meta-regression models
- Dependent effect sizes
  - Model within- and between-study correlations
  - Robust variance estimation
- Missing data
  - FIML
  - Imputation

Northwestern Medicine®
Feinberg School of Medicine

Thank you!

Resources

- Metafor
- Introduction to Meta-Analysis
- Handbook of Research Synthesis