Research Synthesis Methods (2026), **0**: 1–14
doi:10.1017/rsm.2026.10127

Research Synthesis Methods

TUTORIAL

Meta-analysis with JASP, Part I: Classical approaches

František Bartoš[1], Eric-Jan Wagenmakers[1] and Wolfgang Viechtbauer[2]

[1]Department of Psychological Methods, University of Amsterdam, Netherlands
[2]Department of Psychiatry and Neuropsychology, Maastricht University, Netherlands

Corresponding author: František Bartoš; Email: f.bartos96@gmail.com

Received: 12 September 2025 UTC; **Revised:** 5 August 2026 UTC; **Accepted:** 13 September 2026 UTC

Keywords: graphical user interface; meta-regression; multilevel; multivariate; open-source software; effect size

Abstract

Meta-analyses play a crucial part in empirical science, enabling researchers to synthesize evidence across studies and draw more precise and generalizable conclusions. Despite their importance, access to advanced meta-analytic methodology is often limited to scientists and students with considerable expertise in computer programming. To lower the barrier for adoption, we have developed the Meta-Analysis module in JASP (https://jasp-stats.org/), a free and open-source software for statistical analyses. The module offers standard and advanced meta-analytic techniques through an easy-to-use graphical user interface, allowing researchers with diverse technical backgrounds to conduct state-of-the-art analyses. This manuscript presents an overview of the meta-analytic tools implemented in the module and showcases how JASP supports a meta-analytic practice that is rigorous, relevant, and reproducible. Tutorial videos accompany the examples presented in this manuscript.

Highlights

What is already known?

- Meta-analysis is central to empirical science.
- Advanced methods are often out of reach for researchers and students with limited programming skills.

What is new?

- We implemented a comprehensive set of standard and advanced meta-analytic methods through the Meta-Analysis module in JASP, an open-source statistical software program.

Potential impact for RSM readers

- The JASP graphical user interface enables users across technical backgrounds to run state-of-the-art meta-analyses with only a couple of mouse actions.
- This manuscript provides an overview of the meta-analysis tools available in the Meta-Analysis module.

1. Introduction

Meta-analyses are central to modern empirical science. In parallel, meta-analytic methodology has advanced rapidly, with many methods disseminated as packages in the R programming language[1] (see https://cran.r-project.org/view=MetaAnalysis). While this ecosystem is convenient for programming-savvy methodologists, many applied scientists (and their students) prefer not to analyze data using a programming language.

This article was awarded Open Data and Open Materials badges for transparent practices. See the Data availability statement for details.

© The Author(s), 2026. Published by Cambridge University Press on behalf of The Society for Research Synthesis Methodology. This is an Open Access article, distributed under the terms of the Creative Commons Attribution licence (https://creativecommons.org/licenses/by/4.0), which permits unrestricted re-use, distribution and reproduction, provided the original article is properly cited.

https://doi.org/10.1017/rsm.2026.10127 Published online by Cambridge University Press

As an alternative, applied scientists often turn to "point-and-click" statistical software with a graphical user interface (GUI). For meta-analysis, prominent options include both proprietary (e.g., CMA[2] and RevMan[3]) and open-source (e.g., MetaWin,[4] MetaEssentials,[5] and OpenMEE[6]) software. However, these packages rarely provide advanced capabilities, such as multilevel and multivariate models with complex dependency structures[7,8] or location-scale meta-regression.[9] Many of these packages also lack comprehensive diagnostics and integration with interpretation tools, such as estimated marginal means (EMMs), contrasts, and publication-ready visualizations.

In order to overcome these limitations, we have developed the JASP Meta-Analysis module; this module provides a GUI for the `metafor`[10] `R` package enhanced with additional post-processing, integrated with the free and open-source statistical program JASP.[11] As such, the module provides a user-friendly GUI to state-of-the-art meta-analytic tools. This manuscript focuses on the implemented statistical procedures rather than on detailed guidance about when to use them; for that, we refer readers to the cited literature and meta-analytic textbooks.[12–15]

In the following sections, we first briefly introduce JASP. Second, we demonstrate the meta-analytic methodology implemented in the Meta-Analysis module through three examples: Example 1 features the Bacillus Calmette–Guerin (BCG) vaccine and covers effect size computation, funnel plots, a random-effects meta-analysis, forest plots, and subgroup analysis; Example 2 features writing-to-learn interventions and illustrates meta-regression, bubble plots, estimated marginal means and contrasts, and location-scale models; finally, Example 3 features recidivism and mental health and demonstrates multilevel models, variance–covariance construction for multivariate models, and cluster-robust standard errors. Along the way, we highlight additional features and options available in the software.

The Meta-Analysis module also includes other functionality not treated here (e.g., meta-analyses of prediction model performance, publication bias adjustment, and meta-analytic SEM). A tutorial-style introduction of the Bayesian functionality is presented in Part II.[16] JASP is available at https://jasp-stats.org/download/; source code can be found at https://github.com/jasp-stats/. Annotated JASP files are available on the OSF project page at https://osf.io/mb56t/. Tutorial videos are available on the JASP YouTube channel at https://www.youtube.com/playlist?list=PLEIybnA0SEf0.

2. Introduction to JASP

JASP is a free and open-source statistical software package whose development started in 2013.[17] With over 100,000 downloads per month, JASP is used in the curricula of 389 universities from 77 different countries.[18] JASP is primarily developed at the University of Amsterdam, with growing contributions from other institutions across the world (https://jasp-stats.org/community-institutional-members/).

Key features of JASP include a user-friendly point-and-click interface, dynamically updated results, progressive disclosure, translations into many languages, Open Science Framework (OSF) integration, cross-platform support (Windows, macOS, and Linux), an annotated data library, fully integrated AI support, and APA-style output with figures and tables exportable to Word or LATE X. The analysis output can be annotated and exported as an .html or .pdf file. The analysis itself can be saved as a .JASP file, which allows users to share the analyses. The analysis files are reproducible in the same version of JASP (in case of analyses with sampling/bootstrap, a seed needs to be set), and when uploaded to OSF, the results can be directly previewed online. JASP runs locally, does not collect or send user information, and uses a sandboxed `R` environment, which guarantees that data always stay on one's computer (in contrast to online-hosted applications).

Many of the analysis methods in JASP are accompanied by annotated example datasets and extensive documentation, accessible via on-hover tooltips or the help files linked through the "(i)" button in each analysis. Feature requests, bug reports, and questions are handled on the active GitHub issue tracker (https://github.com/jasp-stats/jasp-issues).

To replicate the examples that follow, you may download the current version of JASP from https://jasp-stats.org/download/ and enable the Meta-Analysis module from the module panel using the

https://doi.org/10.1017/rsm.2026.10127 Published online by Cambridge University Press

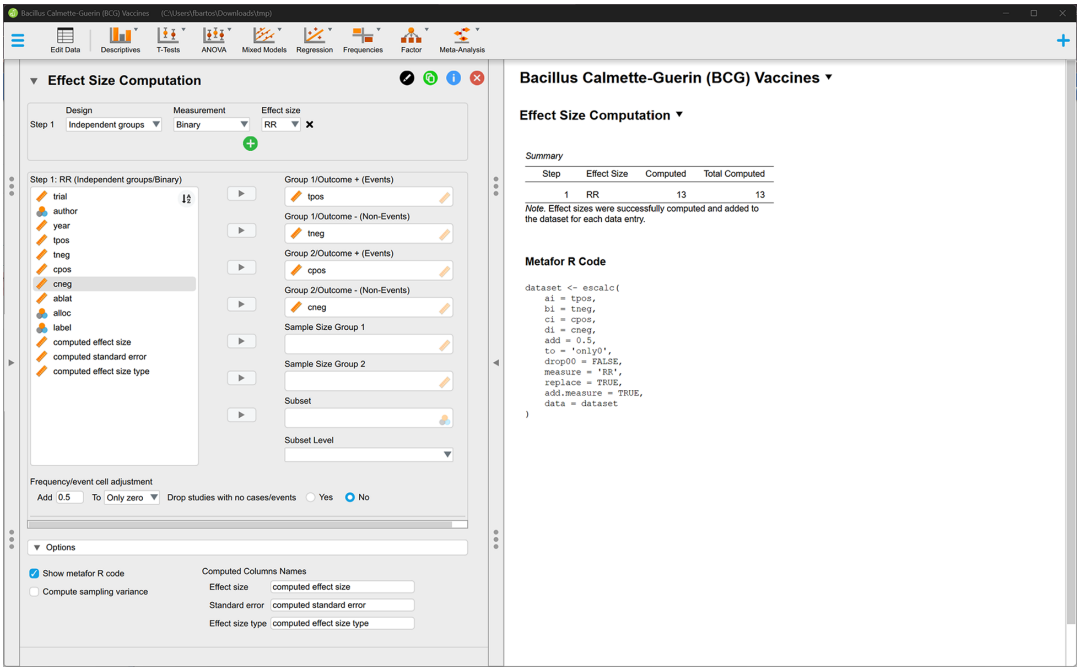


Figure 1. *Effect size computation menu in JASP.*

"+" button in the top right corner of the application. To use the most up-to-date version of the Meta-Analysis module, you might need to update the module via the Module Store that opens under the same menu. Raw datasets and annotated JASP files for the examples below are included in the JASP library and at https://osf.io/mb56t/. Datasets can be loaded into JASP via the main menu in the top left corner of the application ("≡" button) and the "Open" option. JASP can import many different data formats, ranging from common spreadsheet formats like .csv and .xlsx to SPSS, Stata, SAS, Minitab, and R-specific data formats. For a general overview of JASP, see Wagenmakers et al.[19] and Ly et al.[20]

3. Example 1: Bacillus Calmette–Guerin vaccine

The first example features the famous BCG vaccine dataset distributed with the `metadat` R package.[21] Based on the study of Colditz et al.,[22] the dataset features the counts of positive and negative tuberculosis cases for control (non-vaccinated) and experimental (vaccinated) groups across 13 studies, along with absolute latitude and the treatment allocation procedure.

3.1. Effect size computation

We begin by loading the csv dataset (available from https://osf.io/mb56t/) into JASP via the "≡" button in the top left corner of the application and the "Open" option. Once the dataset is loaded into JASP, we proceed by computing effect size estimates and their standard errors from the case counts. This is done in the "Effect Size Computation" menu underneath the "Meta-Analysis" module heading (Figure 1). In the left input panel, the "Design," "Measurement," and "Effect size" dropdowns determine which effect size is computed, and the variable inputs (e.g., "Group 1/Outcome +" and "Group 1/Outcome -") map dataset columns to the calculations. The right output panel then shows how many effect sizes were computed successfully. Notice that the output panel dynamically updates as the options are selected; "run" or "execute" commands are usually not needed in JASP.

https://doi.org/10.1017/rsm.2026.10127 Published online by Cambridge University Press

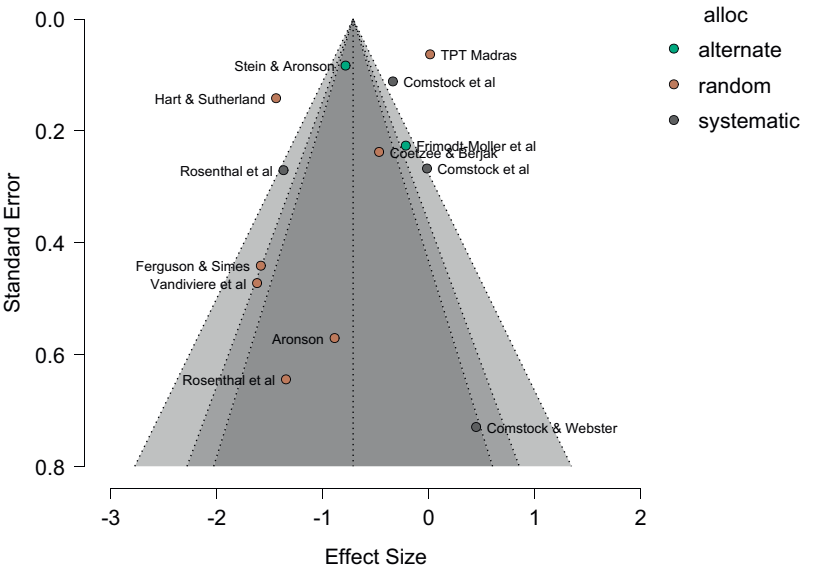


Figure 2. *Funnel plot created in JASP.*

Behind the scenes, "Effect Size Computation" calls the `metafor::escalc` function. As such, the analysis supports a wide range of effect sizes suitable for many different settings. The corresponding R code can be generated using the "Show metafor R code" option in the "Options" section. Furthermore, the analysis allows for the chaining of multiple calculations with the green "+" button and can restrict computations to subsets of rows via the corresponding "Subset" variable input. Details about the available effect sizes and the required variable input are available through the "(i)" help button in the top-right of the input section.

3.2. Funnel plot

Next, the data are visualized with a funnel plot. Funnel plots are a useful tool for visual inspection of the data before proceeding with the analyses; for instance, funnel plots can help identify odd patterns, incorrectly extracted information (e.g., standard errors mistaken for standard deviations), and are used to examine the data for "small-study effects."

The funnel plot is created using the "Funnel Plot" menu. The analysis requires specification of the computed "Effect Size" and "Effect Size Standard Error" variable inputs. By default, JASP shows a scatterplot of the effect sizes against their standard errors with a funnel that reflects the sampling distribution centered at the random-effects estimate (Figure 2; a second funnel under the null hypothesis can be created via the "$\mathcal{H}_0$" checkbox). In the "Estimates Mapping" subsection, the "Study Label" option is set to the "author" variable and the "Color" option is set to the "allocation" variable, which extends the funnel plot with the author labels and allocation color coding. The figure can be exported in various file formats via the small triangle just above the plot.

Additional options include incorporating heterogeneity into the sampling distribution, displaying power enhancement ("power bands"),[23] and creating separate funnels for subsets of data. Related statistical tests are in the "Publication Bias/Sensitivity Analyses" section, including funnel plot asymmetry tests,[24,25] trim-and-fill,[26] and fail-safe N variants.[27–29]

3.3. Meta-analysis

We proceed by estimating a random-effects model via the "Meta-Analysis" menu in the "Classical" part of the "Meta-Analysis." The user interface contains several sections separating it into easy-to-navigate

https://doi.org/10.1017/rsm.2026.10127 Published online by Cambridge University Press

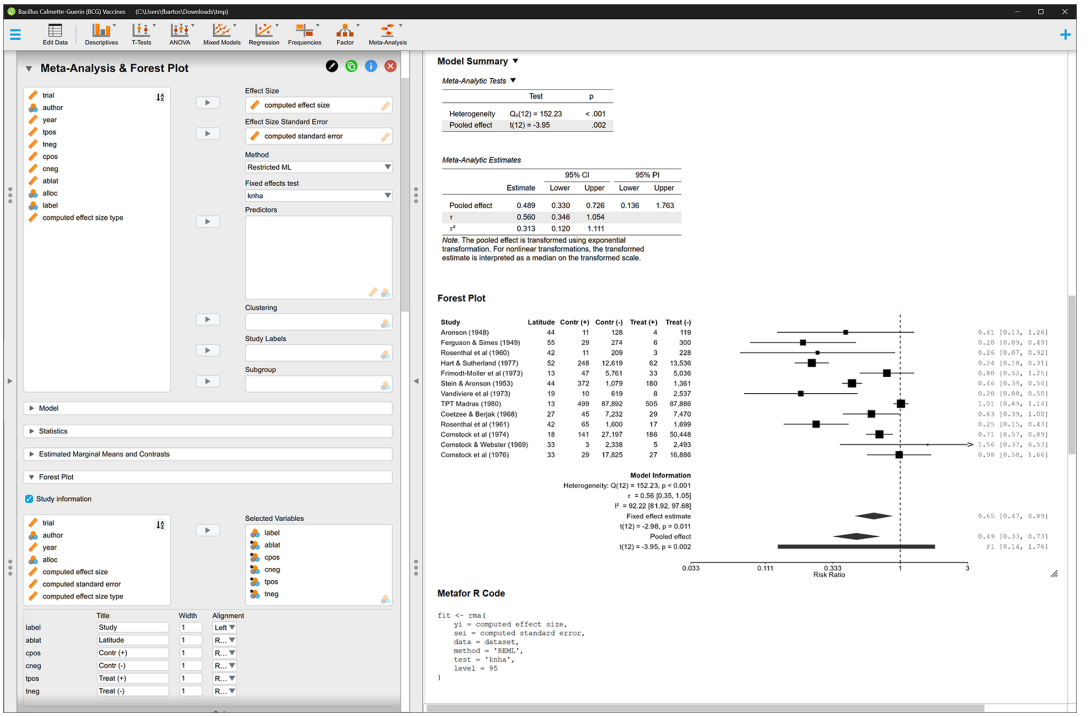


Figure 3. *Random-effects meta-analysis with forest plot.*

option groups. In the pre-opened data input section, we specify the required "Effect Size" and "Effect Size Standard Error" variable inputs. Default settings follow `metafor` recommendations, including the "REML" random-effects estimator and the "Knapp and Hartung" adjustment for the fixed-effects test.[30]

With the variable input set as in the left panel of Figure 3, the "Meta-Analytic Tests" and "Meta-Analytic Estimates" tables appear in the "Model Summary" output section (right output panel of Figure 3). In this example, the tests reject the null hypothesis that the pooled effect is zero and the null hypothesis of no between-study heterogeneity. Estimates are by default reported on the input scale (i.e., the log risk ratio). To present pooled estimates and confidence/prediction intervals on the risk ratio scale instead, we set the "Exponential" option in the "Transform effect size" dropdown located in the "Statistics" section. This transformation propagates to all effect size estimates in the forest plot, bubble plots, and EMMs.

Options within the "Statistics" section allow users to obtain additional information about the fitted model, such as relative heterogeneity measures I^2 and H^2, fit measures (e.g., log likelihood, AIC, and BIC), and confidence and prediction intervals of different levels. In addition, the "Diagnostics" section presents comprehensive model diagnostics, including casewise diagnostics,[31,32] a profile likelihood plot,[33] the Baujat plot,[34] and the residual funnel plot. Finally, the "Advanced" section contains additional settings that allow advanced users to fully control the `metafor::rma.uni` function by changing the optimization settings, fixing between-study heterogeneity, extending the `R` function call directly, and applying permutation tests.[35,36] Importantly, the "Advanced" section also allows the `R` code that corresponds to the fitted model to be generated via the "Generate metafor R code" option. The resulting output, shown at the bottom of Figure 3, can be especially helpful when teaching users of various levels or when attempting a graceful transition to `R`.

3.4. Forest plot

Forest plots are the most commonly used visualization in meta-analyses and can be created in the "Forest Plot" section. The forest plot can include three components that are toggled independently:

https://doi.org/10.1017/rsm.2026.10127 Published online by Cambridge University Press

"Study information," "Estimated Marginal Means," and "Model Information." To produce the forest plot shown in Figure 3, we select the "Study information" component and specify "Selected Variables" to define study information to be shown on the left side of the study information panel (including information about column titles, centering, etc.). We also select the "Model Information" panel and specify the desired information to be included in the panel (e.g., heterogeneity estimates, fixed-effect estimate, etc.). Final styling, such as modifying the relative width of the left, middle, and right sides of the figure (automatic width adjustment is not always perfect), setting the *x*-axis label, or the addition of vertical line(s), is controlled in the "Settings" subsection.

Other notable options in the "Study Information" panel include the ability to show predicted effects based on the estimated model, and the ability to aggregate effect size estimates according to a grouping variable, which is especially useful when individual studies contribute multiple effect sizes. The "Settings" subsection controls the details of the output, such as the number of decimals, whether hypothesis tests and/or estimates are reported on the right or left side of the figure, mapping variables to the effect size estimates' color and shape, and additional behavior when a subgroup analysis is specified.

3.5. Subgroup analysis

The first example is completed by performing a subgroup analysis based on the type of treatment allocation. The subgroup analysis is specified via the "Subgroup" variable input. Assigning the allocation variable into the variable input stratifies all previous outputs (tests, estimates, and forest plot) based on the treatment allocation. Furthermore, a new row called "Subgroup differences" is added to the "Meta-Analytic Tests" table with results of a null hypothesis of no differences in the pooled effect across the subgroups. Here, we cannot reject the null hypothesis of no difference between the subgroups ($p = 0.361$, not shown).

In the case of binary or person-time outcomes, the Mantel–Haenszel[37] and Peto[38] methods with equivalent input and output options are available under the "Mantel–Haenszel/Peto" menu of the Meta-Analysis module.

4. Example 2: Writing-to-learn interventions

The second example involves a dataset on the effectiveness of writing-to-learn interventions,[39] distributed with the `metadat` R package.[21] The dataset contains standardized mean differences and sampling variances[1] comparing academic achievement (e.g., final grades and exam scores) of students who received instructions with increased emphasis on writing tasks (i.e., the experimental group) versus students who received conventional instructions (i.e., the control group) from 48 studies, together with several covariates. Since the effect size estimates and their standard errors are readily available from the dataset, we can skip the effect size calculation step and directly proceed to the analysis. For illustration, we focus on two moderators: the continuous intervention length (in weeks) and a categorical indicator of whether participants received feedback.

4.1. Meta-regression

Meta-regression examines how study-level predictors (moderators, such as the length of the intervention) relate to the effect size while adjusting for other predictors (e.g., whether participants received feedback). In contrast to subgroup analysis, common between-study heterogeneity is estimated from the full dataset after covariate adjustment (unless, as showcased later, a location-scale model is applied).

After setting the required "Effect Size" and "Effect Size Standard Error" variables in the "Meta-Analysis" menu, the meta-regression specification is completed by assigning the intervention "length" and "feedback" variables to the "Predictors" variable input. Continuous and categorical moderators

[1] Users can transform the sampling variance into the standard error in the JASP data editor using the square-root transformation.

https://doi.org/10.1017/rsm.2026.10127 Published online by Cambridge University Press

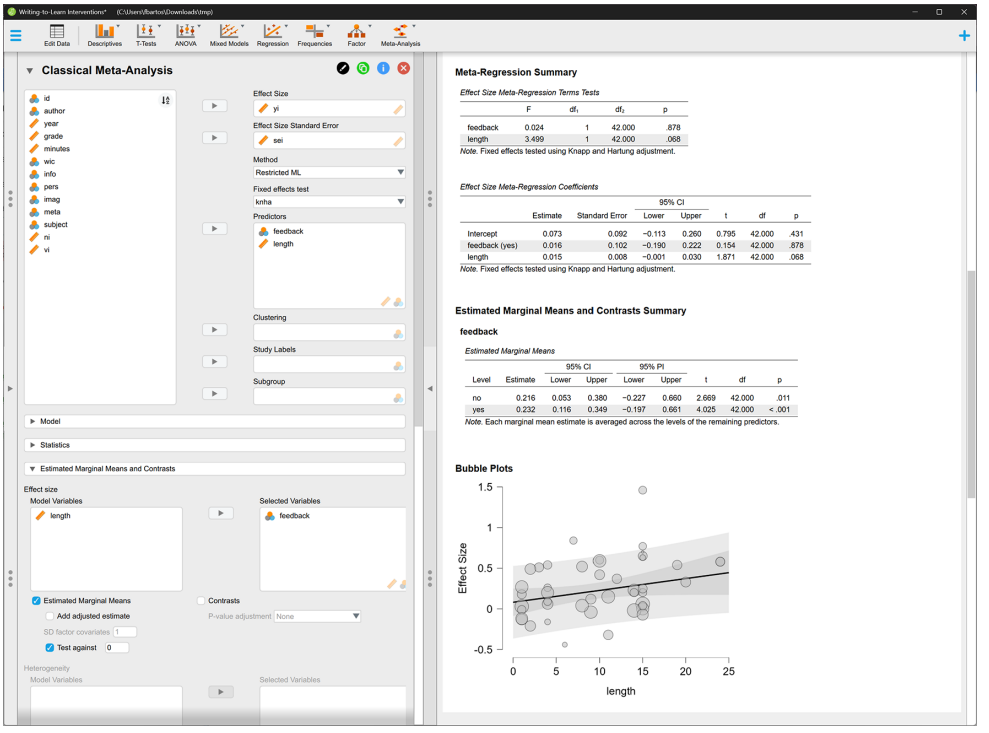


Figure 4. *Meta-regression with estimated marginal means and bubble plot.*

are handled differently by the analysis, and therefore it is important to verify the variable icons: three intersecting circles indicate a categorical variable, whereas the ruler icon indicates a continuous variable. The complete input for our example is shown on the left side of Figure 4.

The corresponding output shows a "Meta-Regression Summary" output section that contains the "Meta-Regression Terms Tests" table. For categorical predictors, the term-specific omnibus test evaluates whether all $p - 1$ coefficients equal zero (i.e., a null hypothesis of no differences among the moderator levels). For continuous predictors, the test corresponds to a test of the slope coefficient (i.e., a null hypothesis of zero slope). In our example, neither null hypothesis is rejected at the $\alpha = 0.05$ level. The "Meta-Regression Coefficients" table then lists the individual meta-regression coefficients.

The omnibus moderation test of all moderators simultaneously is added to the main "Model Summary" output section (not visible in Figure 4). Moderator interaction terms can be added via the "Model" section by selecting multiple "Available Components" while holding the shift key and assigning them to the "Model Terms" variable input.

Missing values in moderators, effect sizes, or standard errors result in list-wise deletion (complete-case analysis). The number of omitted observations (three in our example) is reported below the "Meta-Analytic Estimates" table in the "Model Summary" output section. The JASP Team is currently working on implementing multiple imputation via the `mice` R package[40] across all analyses in JASP.

4.2. Estimated marginal means and contrasts

The interpretation of meta-regression can be aided by the "Estimated Marginal Means and Contrasts" section. EMMs are the model-predicted effects at each level of the moderator of interest, with every other moderator held fixed at a representative value: a continuous moderator is set to its mean, and a categorical moderator is averaged over its levels (each level weighted equally). For example, the EMMs

https://doi.org/10.1017/rsm.2026.10127 Published online by Cambridge University Press

for the categorical moderator "feedback" are the predicted effects at its two levels ("no" and "yes"), each evaluated at the mean of the continuous moderator "length." If a second categorical moderator were present, for example, "subject" with levels "Math," "Social Studies," and "Biology," then the "feedback" EMMs would still be evaluated at the mean of "length," but would additionally be averaged across the three "subject" levels, so that they reflect no single subject in particular.

Returning to the current example with "feedback" and "length" moderators, Figure 4 shows the "Estimated Marginal Means" output table for the "feedback" moderator at levels "no" and "yes." Furthermore, the EMMs allow users to test whether the predicted effect for each level differs from zero (or any other value of interest). The level-specific estimates can be tested against a null hypothesis of no effect via the "Test against" option with the following value input set to "0" in the left input panel. We find that the estimated effect size at both levels of the moderator is significantly different from zero.

Pairwise differences between the EMMs can be obtained by selecting the "Contrasts" option. Compared to raw meta-regression coefficients, EMMs and their contrasts incorporate moderator interactions[2] and avoid re-leveling of categorical variables. EMMs can also be shown in the "Forest plot" via its corresponding subsection. When a continuous moderator is added to the "Selected Variables" input field, predictions are produced at ±1 standard deviation from the moderator's mean.

The "Estimated Marginal Means and Contrasts" section also provides an "Adjusted estimate" option which calculates the effect for the *unweighted* average across the moderator levels. This differs from the pooled effect in the "Model Summary" output, which is based on the *weighted* average across the moderator levels.

4.3. Bubble plot

The "Bubble Plot" section (not shown in Figure 4) allows users to visualize moderator values against observed effect size estimates as bubbles, with larger bubbles indicating higher precision, together with the estimated meta-regression trend. The bottom-right panel of Figure 4 shows a bubble plot of intervention length with a positive, although not statistically significant, increase in the effect size estimate in studies with longer interventions. The trend reflects predictions over the "Selected Variable" input, averaging across the remaining moderators as in EMMs. Specifying "Separate Lines" and "Separate Plots" variables allows users to visualize predicted trends at different levels of other moderators, which is particularly helpful in the presence of interaction effects. The darker band around the regression line depicts the confidence interval for the mean effect; the lighter band depicts the prediction interval for true study effects.

The JASP bubble plots also enable visualization of categorical moderators; instead of a regression line, vertical "box plots" with confidence and prediction intervals are displayed alongside the jittered effect size estimate bubbles.

4.4. Location-scale models

We conclude the second example by extending the analysis to a location-scale model.[9] Location-scale models relax the assumption of a common between-study heterogeneity by including a meta-regression model for the heterogeneity parameter. In other words, standard meta-analysis and meta-regression models assume that the between-study heterogeneity τ is constant. Location-scale models allow the between-study heterogeneity to vary as a function of the moderators. The scale part of the model (with the location part corresponding to the standard regression on the effect size) is specified in the "Heterogeneity Model (Scale)" subsection of the "Model" section.

The previously estimated "location" meta-regression can be extended to a location-scale meta-regression by assigning the study length moderator to the corresponding ("Heterogeneity Model

[2]Note that the `metafor` function `pairmat` does not compute EMMs, as it does not average across the levels of the other terms.

https://doi.org/10.1017/rsm.2026.10127 Published online by Cambridge University Press

(Scale)") "Model Terms" variable input. The previous output is extended with appropriate "Meta-regression Terms Tests" and "Meta-regression Coefficients Table" for the scale (Heterogeneity) term. Although the moderation effect on heterogeneity is not statistically significant at the $\alpha = 0.05$ level, the corresponding bubble plot would show slightly wider prediction intervals for longer interventions.

As with the pooled effect (computed at the weighted average of observed moderator values), the heterogeneity estimates from the location-scale model are also computed at the weighted average of observed moderator values. The adjusted heterogeneity at the unweighted average of moderator levels can be obtained with the "Adjusted estimate" option in the "Heterogeneity" subsection of the "Estimated Marginal Means and Contrasts" section. This section also provides estimated marginal means and contrasts for the heterogeneity regression itself, parallel to those for the effect size regression.

5. Example 3: Recidivism and mental health

The third example concerns studies on the association between recidivism and mental health.[41] This dataset again comes bundled with the `metadat` R package.[21] The dataset contains standardized mean differences and sampling variances comparing recidivism of delinquent juveniles with and without mental health disorders across 17 studies. Most studies contributes several effect size estimates for different types of delinquent behavior measured for the same participants, which induces within-study dependence. Accordingly, this dataset allows a demonstration of multilevel models, multivariate models, and cluster-robust standard errors, available from the "Meta-Analysis (Multilevel/Multivariate)" menu of the Meta-Analysis module. These three adjustments are all driven by the same study grouping, but they enter the analysis through three different inputs and address the dependence in three different ways: the multilevel model (the "Level 1" and "Level 2" inputs of the "Random Effects/Model Components" section) captures the similarity among the true effects estimated within a study, the multivariate model (the "Cluster" input of the "Effect Size Variance–Covariance Matrix" section) captures the correlation among the sampling errors of estimates that share participants, and cluster-robust standard errors (the "Clustering" input of the data input section) guard the standard errors against misspecification of that assumed dependence structure. Assigning the study variable through three separate inputs gives the same grouping a different role. Previously described functionality, such as meta-regression, estimated marginal means and contrasts, bubble plots, etc., can be applied to this example as well (apart from scale moderation).

5.1. Multilevel models

If individual effect size estimates are not independent, a simple random-effects model can inflate the precision of the effect size estimate, resulting in overly optimistic *p*-values and poor confidence interval coverage. One source of dependence is that estimates within the same study tend to be more similar than estimates across studies (e.g., due to shared populations, manipulations, or measures). For instance, in the current example, each study provides multiple estimates of recidivism; a study with samples from a population with an extensive criminal history will lead to systematically different recidivism estimates than a study with samples from a population with a short criminal history. Multilevel models can account for this dependence by introducing separate heterogeneity components for each "level" (i.e., study-level and estimate-level) to account for the estimates-within-studies nesting dependency.[8,42]

After assigning the usual "Effect Size" and "Effect Size Standard Error" variables, the multilevel model is specified by setting the "Nested (multilevel)" component in the "Random Effects/Model Components" section. In the "Specification" subsection depicted in the left panel of Figure 5, we define "Level 1" as the higher nesting level (study) and "Level 2" as the lower nesting level (estimates). After completing the model specification this way, the output is then extended with a "Random Effects/Model Structure Summary" section shown in the upper right output panel of Figure 5, which reports the study-level (τ_{study}) and estimate-level ($\tau_{\text{study/esid}}$) heterogeneity estimates; note that these estimates are

https://doi.org/10.1017/rsm.2026.10127 Published online by Cambridge University Press

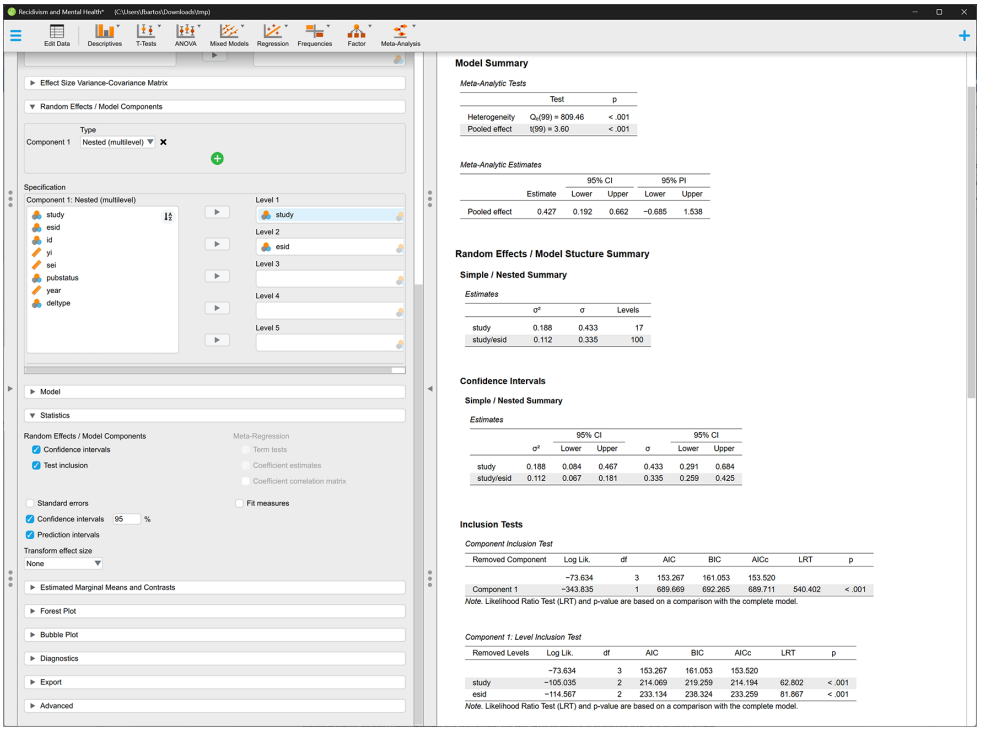


Figure 5. *Multilevel meta-analysis and heterogeneity components inclusion tests.*

no longer summarized in the "Model Summary" section. Confidence intervals for the heterogeneity components, often requiring additional extensive calculation, can be requested via the "Confidence intervals" option in the "Random Effects/Model Components" subsection of the "Statistics" section.

The "Test inclusion" option in the "Statistics" subsection assesses whether the specified random-effects structure improves model fit, and it does so at two levels of granularity. A model may contain one or more random-effects components (the terms listed in the "Random Effects/Model Components" section), and a nested component may in turn span several levels (here, "study" and "estimate"). At a coarse level of granularity, the "Component Inclusion Test" table compares the full model against a model with one entire component removed (or their combinations); when several components are specified, each is dropped in turn so that its contribution can be judged on its own, which is where this test is most informative. In the present example, only a single nested component is specified, so removing it leaves a fixed-effect model. The test rejects the null hypothesis of equivalent fit, suggesting that a random-effects structure is preferred over a model without one.

At a finer level of granularity, and only when a nested component spans multiple levels, the "Component 1: Level Inclusion Test" table compares the full model against models in which individual levels (or their combinations) are dropped. In our example, dropping either the study-level or the estimate-level significantly worsens the fit, so both nesting levels are warranted: the two-level structure is preferred over the simpler single-level alternatives that retain only estimate-level or only study-level random effects.

The "Random Effects/Model Structure Summary" section allows the specification of a wide range of random effects available through the `metafor::rma.mv` function, including random slopes, structured, autoregressive, and spatial random effects, as well as phylogenetic models, some of which are showcased in the JASP Data Library or under the "Specifying Random Effects" heading at https://wviechtb.github.io/metafor/reference/rma.mv.

https://doi.org/10.1017/rsm.2026.10127 Published online by Cambridge University Press

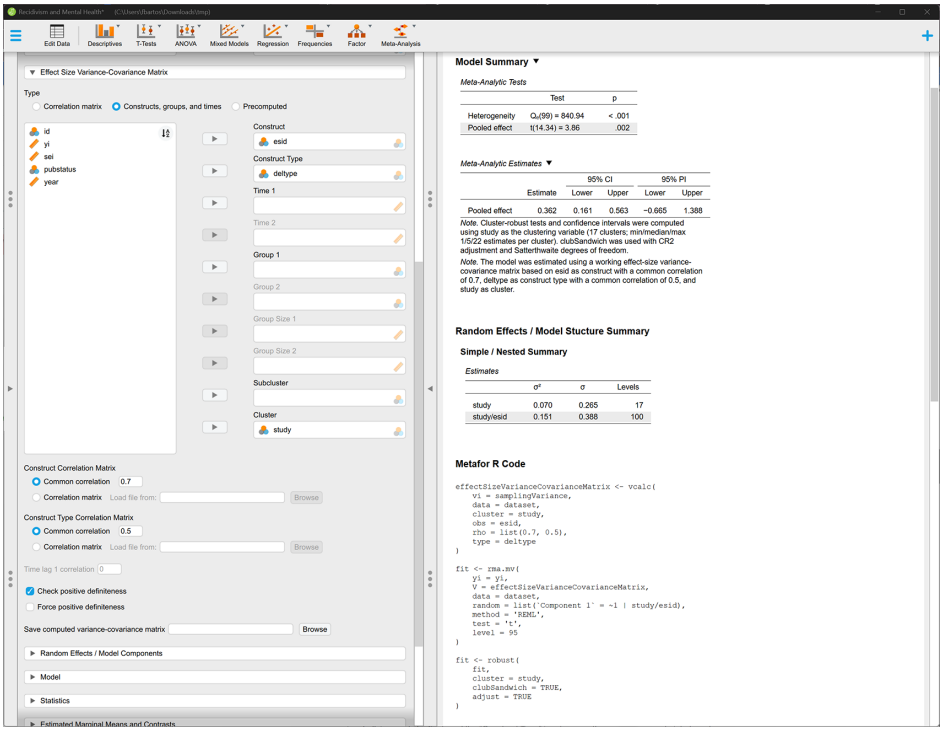


Figure 6. *Multilevel and multivariate meta-analysis with cluster-robust standard errors.*

5.2. Multivariate models

Another source of dependence is that repeated assessments on the same participants (across time points or measures) produce correlated estimates that do not contribute fully independent information. For instance, the current example features measures of recidivism in different types of delinquent behavior (e.g., aggression, bullying, and stealing)[3] from the same groups of participants. Ignoring this multivariate dependency can again result in overly optimistic *p*-values and poor confidence interval coverage.[43,44]

Multivariate models are generally used to account for this particular dependence. For their application, these models require the specification of the between-estimate dependencies via the effect-size variance–covariance matrix. This can be done in the "Effect Size Variance–Covariance Matrix" section (see the left input panel of Figure 6), which uses the `metafor::vcalc` function in the background.

In our example, the variance–covariance matrix is built from three variable inputs. The "Cluster" input identifies sets of estimates that may be correlated; estimates in different clusters are assumed independent, and here each study forms its own cluster because studies report on different participants. Within a cluster, the "Construct" input identifies the individual measurements (recidivism in different delinquent behaviors, such as aggression, bullying, and stealing), and the "Construct Type" input groups those measurements into broader categories (here, overt and covert behavior). Specifying both inputs is useful when measurements of the same type are expected to correlate more strongly than measurements of different types: aggression (overt) and bullying (overt) recidivism should correlate more strongly than aggression (overt) and stealing (covert) recidivism. The strength of these correlations is supplied by the user rather than estimated from the data through the "Construct Correlation Matrix" option

[3]The exact recidivism measures are not reported in the original article: we input generic delinquent behavior types for the clarity of the illustration.

https://doi.org/10.1017/rsm.2026.10127 Published online by Cambridge University Press

for measurements sharing a "Construct Type" ($r = 0.70$ in our example) and the "Construct Type Correlation Matrix" option for measurements of different "Construct Type" ($r = 0.50$).

Once these inputs are set, the output updates to reflect the multivariate adjustment (the right panel of Figure 6, which also includes the cluster-robust standard errors described next). In our example, the adjustment reduces the study-level heterogeneity estimate but increases the estimate-level heterogeneity estimate: part of the within-study dependence is now attributed to the assumed correlation between sampling errors, so less of it needs to be absorbed by the study-level random-effects variances.

The specification of the variance–covariance matrix can be challenging and error-prone, and therefore we recommend users consult the `metafor::vcalc` documentation at https://wviechtb.github.io/metafor/reference/vcalc.html and then use "Show Metafor R Code" in the "Advanced" section to verify the matrix specification. The interface also supports per-level correlation matrices for all "Construct" or "Construct Type" pairs (see the JASP Help File for formatting) and allows dependencies from repeated measures across time or from a shared comparison group via the corresponding variable inputs. Alternatively, correlations between two variables can be supplied directly via the "Correlation matrix" option, or the entire variance–covariance structure can be provided via the "Precomputed" option. Finally, a computed variance–covariance matrix can be saved using the "Save computed variance–covariance matrix" option, manually edited, and reloaded via the "Precomputed" option.

5.3. Cluster-robust standard errors

The example concludes with an adjustment that yields cluster-robust standard errors.[45] Cluster-robust standard errors are obtained by assigning the study indicator to the "Clustering" input in the data input section (the first, pre-opened section), a separate input from the "Cluster" variable used to build the variance–covariance matrix. After the "Clustering" variable is set, the output in the right panel of Figure 6 incorporates the adjustment. With the multilevel, multivariate, and cluster-robust adjustments all in place, a statistically significant difference in recidivism between juveniles with and without mental health disorders remains: the pooled effect size estimate is $d = 0.36$, 95% CI [0.16, 0.56], $t(14.34) = 3.86$, $p = 0.002$.

By default, cluster-robust standard errors use the small-sample adjustment from the `clubSandwich`[46] `R` package. The cluster-robust standard errors settings can be further modified in the "Advanced" section.

6. Concluding comments

This manuscript summarizes the functionality of the classical Meta-Analysis module in JASP. The software is free, open-source, and makes state-of-the-art meta-analytic methodology available within a couple of mouse clicks. This means that analysts, teachers, and students can focus their efforts on choosing the appropriate methodology and interpreting its results, rather than devoting considerable time to writing and debugging relatively complex analysis scripts. JASP also generates `R` code snippets for the underlying analytic functionality, which can facilitate teaching and learning. Finally, the reproducible and easy-to-annotate analysis files directly render on the OSF, further facilitating reproducibility and data-sharing practices.

This manuscript covered effect size computation, funnel plots, random-effects meta-analysis, forest plots, subgroup analysis, meta-regression, estimated marginal means and contrasts, bubble plots, location-scale models, multilevel models, multivariate models, and cluster-robust standard errors, built using the `metafor`[10] `R` package. The Meta-Analysis module also contains additional tools not discussed here, including Bayesian meta-analytic methodology mirroring the classical features (presented in Part II[16]), publication bias adjustments,[47] meta-analytic generalized mixed-effects models,[48] meta-analyses of prediction model performance,[49] meta-analytic SEM and SEM-based meta-analysis,[50] penalized meta-analysis,[51] effect size aggregation, and risk-of-bias visualizations.[52] A full list of the currently

https://doi.org/10.1017/rsm.2026.10127 Published online by Cambridge University Press

implemented features is available at https://github.com/jasp-stats/jaspMetaAnalysis/. Features such as network meta-analysis[53] are scheduled for later implementation.

By providing classical meta-analytic capabilities through a user-friendly interface, the JASP Meta-Analysis module makes meta-analytic methods available to a wider audience and, hopefully, facilitates the adoption of state-of-the-art methodology across a broad range of scientific disciplines.

Acknowledgements. F.B. is grateful to the JASP Team for reviewing the Meta-Analysis module pull requests and developing new features necessary to implement the module.

Author contributions. Conceptualization: F.B.; Data curation: F.B. and W.V.; Data visualization: F.B.; Methodology: F.B. and W.V.; Writing—original draft: F.B. All authors approved the final submitted draft.

Competing interest statement. F.B. and E.-J.W. declare their involvement in the open-source software package JASP (jasp-stats.org), a non-commercial, publicly funded effort to make Bayesian statistics accessible to a broader group of researchers and students. In addition, F.B. and E.-J.W. are involved in JASP Services BV, a company that supports organizations who wish to adopt JASP, with a focus on applications in quality control. The remaining author declares no competing interests exist.

Data availability statement. JASP can be downloaded from https://jasp-stats.org/download/ with the source code available at https://github.com/jasp-stats/. Annotated example analysis files are available at https://osf.io/mb56t/. Tutorial videos are available on the JASP YouTube channel at https://www.youtube.com/playlist?list=PLEIybnA0SEf0.

Funding statement. This work is supported by ERC grant 101201482, BAYESIAN BOOST, awarded to E.-J.W. Views and opinions expressed are those of the authors only. Open access funding provided by University of Amsterdam.

Generative AI. ChatGPT was used for sentence-level editing and proofreading.

Ethical standards. The research meets all ethical guidelines, including adherence to the legal requirements of the study country.

References

[1] R Core Team. *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing; 2021. https://www.R-project.org/.
[2] Borenstein M, Hedges L, Higgins J, Rothstein H. *Comprehensive Meta-Analysis Version 4*. Biostat; 2022.
[3] Review Manager (RewMan). Version 7.12.0. 2024. https://documentation.cochrane.org/revman-kb/cite-revman-in-reference-list-110242006.html.
[4] Rosenberg MS. MetaWin 3: open-source software for meta-analysis. *Front Bioinform*. 2024;4: 1305969.
[5] Suurmond R, van Rhee H, Hak T. Introduction, comparison and validation of meta-essentials: a free and simple tool for meta-analysis. *Res Synth Methods*. 2017;8: 537–553.
[6] Wallace BC, Lajeunesse MJ, Dietz G, et al. OpenMEE: intuitive, open-source software for meta analysis in ecology and evolutionary biology. *Methods Ecol Evol*. 2016;8(8): 941–947. https://doi.org/10.1111/2041-210X.12708.
[7] Olkin I, Gleser L. Stochastically dependent effect sizes. In: Cooper H, Hedges LV, Valentine JC, eds. *The Handbook of Research Synthesis and Meta-Analysis*. 2nd ed. Russell Sage Foundation; 2009: 357–376.
[8] Konstantopoulos S. Fixed effects and variance components estimation in three-level meta-analysis. *Res Synth Methods*. 2011;2: 61–76.
[9] Viechtbauer W, López-López JA. Location-scale models for meta-analysis. *Res Synth Methods*. 2022;13: 697–715.
[10] Viechtbauer W. Conducting meta-analyses in R with the metafor package. *J Stat Softw*. 2010;36: 1–48.
[11] JASP Team. JASP (version 0.95) [computer software]. 2025. https://jasp-stats.org/.
[12] Borenstein M, Hedges LV, Higgins JP, Rothstein HR. *Introduction to Meta-Analysis*. John Wiley & Sons; 2009.
[13] Schmid CH, Stijnen T, White I. *Handbook of Meta-Analysis*. CRC Press; 2020.
[14] Borenstein M, Egger M, Higgins J, Smith G. *Systematic Reviews in Health Research: Meta-Analysis in Context*. John Wiley & Sons; 2022.
[15] Harrer M, Cuijpers P, Furukawa TA, Ebert DD. Meta regression. In: *Doing Meta-Analysis with R: A Hands-on Guide*. 1st ed. Boca Raton, FL and London: Chapman & Hall/CRC Press; 2021. ISBN 978-0-367-61007-4.
[16] Bartoš F, Wagenmakers EJ. Meta-analysis with JASP, Part II: Bayesian approaches. Preprint, 2025. https://doi.org/10.48550/arXiv.2509.09850.
[17] Love J, Selker R, Marsman M, et al. JASP: graphical statistical software for common statistical designs. *J Stat Softw*. 2019;88: 1–17.
[18] Field AP, van Doorn J, Wagenmakers EJ. *Discovering Statistics Using JASP*. Sage; 2025.
[19] Wagenmakers EJ, Kucharský Š, van Doorn J, van den Bergh, D. Accessible and sustainable statistics with JASP. 2023. https://doi.org/10.31234/osf.io/ud2vj.
[20] Ly A, van den Bergh D, Bartoš F, Wagenmakers EJ. Bayesian inference with JASP. *ISBA Bull.* 2021;28: 7–15.

[21] Viechtbauer W, White T, Noble D, Senior A, Hamilton WK. metadat: meta-analysis datasets. R package version 1.4-0. 2025. https://doi.org/10.32614/CRAN.package.metadat.
[22] Colditz GA, Brewer TF, Berkey CS, et al. Efficacy of BCG vaccine in the prevention of tuberculosis: meta-analysis of the published literature. *JAMA J Am Med Assoc*. 1994;271: 698–702.
[23] Kossmeier M, Tran US, Voracek M. Power-enhanced funnel plots for meta-analysis. *Z Psychol.* 2020;228(1): 43–49.
[24] Egger M, Smith GD, Schneider M, Minder C. Bias in meta-analysis detected by a simple, graphical test. *BMJ*. 1997;315: 629–634.
[25] Begg CB, Mazumdar M. Operating characteristics of a rank correlation test for publication bias. *Biometrics*. 1994;50(4): 1088–1101.
[26] Duval S, Tweedie R. Trim and fill: a simple funnel-plot–based method of testing and adjusting for publication bias in meta-analysis. *Biometrics*. 2000;56: 455–463.
[27] Rosenthal R. The file drawer problem and tolerance for null results. *Psychol Bull*. 1979;86: 638–641.
[28] Orwin RG. A fail-safe N for effect size in meta-analysis. *J Educ Stat*. 1983;8: 157–159.
[29] Rosenberg MS. The file-drawer problem revisited: a general weighted method for calculating fail-safe numbers in meta-analysis. *Evolution*. 2005;59: 464–468.
[30] Knapp G, Hartung J. Improved tests for a random effects meta-regression with a single covariate. *Stat Med*. 2003;22: 2693–2710.
[31] Viechtbauer W, Cheung MWL. Outlier and influence diagnostics for meta-analysis. *Res Synth Methods*. 2010;1: 112–125.
[32] Viechtbauer W. Model checking in meta-analysis. In: Schmid CH, Stijnen T, White IR, eds. *Handbook of Meta-Analysis*. CRC Press; 2021: 219–254. https://doi.org/10.1201/9781315119403.
[33] Raue A, Kreutz C, Maiwald T, et al. Structural and practical identifiability analysis of partially observed dynamical models by exploiting the profile likelihood. *Bioinformatics*. 2009;25: 1923–1929.
[34] Baujat B, Mahe C, Pignon JP, Hill C. A graphical method for exploring heterogeneity in meta-analyses: application to a meta-analysis of 65 trials. *Stat Med*. 2002;21: 2641–2652.
[35] Follmann DA, Proschan MA. Valid inference in random effects meta-analysis. *Biometrics*. 1999;55: 732–737.
[36] Good PI. *Permutation, Parametric, and Bootstrap Tests of Hypotheses*. 3rd ed. Springer; 2009.
[37] Mantel N, Haenszel W. Statistical aspects of the analysis of data from retrospective studies of disease. *J Natl Cancer Inst*. 1959;22: 719–748.
[38] Yusuf S, Peto R, Lewis J, Collins R, Sleight P. Beta blockade during and after myocardial infarction: an overview of the randomized trials. *Progr Cardiovasc Dis*. 1985;27: 335–371.
[39] Bangert-Drowns RL, Hurley MM, Wilkinson B. The effects of school-based writing-to-learn interventions on academic achievement: a meta-analysis. *Rev Educ Res*. 2004;74: 29–58.
[40] van Buuren S, Groothuis-Oudshoorn K. mice: multivariate imputation by chained equations in R. *J Stat Softw*. 2011;45: 1–67.
[41] Assink M, van der Put CE, Hoeve M, de Vries SLA, Stams GJJM, Oort FJ. Risk factors for persistent delinquent behavior among juveniles: a meta-analytic review. *Clin Psychol Rev*. 2015;42: 47–61.
[42] Nakagawa S, Santos ESA. Methodological issues and advances in biological meta-analysis. *Evol Ecol*. 2012;26: 1253–1274.
[43] Kalaian HA, Raudenbush SW. A multivariate mixed linear model for meta-analysis. *Psychol Methods*. 1996;1: 227–235.
[44] van Houwelingen HC, Arends LR, Stijnen T. Advanced methods in meta-analysis: multivariate approach and meta-regression. *Stat Med*. 2002;21: 589–624.
[45] Tipton E, Pustejovsky JE. Small-sample adjustments for tests of moderators and model fit using robust variance estimation in meta-regression. *J Educ Behav Stat*. 2015;40: 604–634.
[46] Pustejovsky JE. clubSandwich: cluster-robust (sandwich) variance estimators with small-sample corrections. R package version 0.6.1. 2025. https://doi.org/10.32614/CRAN.package.clubSandwich.
[47] Bartoš F, Maier M, Quintana D, Wagenmakers EJ. Adjusting for publication bias in JASP and R: selection models, PET-PEESE, and robust Bayesian meta-analysis. *Adv Methods Pract Psychol Sci*. 2022;5: 1–19.
[48] Turner RM, Omar RZ, Yang M, Goldstein H, Thompson SG. A multilevel model framework for meta-analysis of clinical trials with binary outcomes. *Stat Med*. 2000;19: 3417–3432.
[49] Debray TP, Damen JA, Snell KIE, et al. A framework for meta-analysis of prediction model studies with binary and time-to-event outcomes. *Stat Methods Med Res*. 2019;28: 2768–2786.
[50] Cheung MWL. metaSEM: an R package for meta-analysis using structural equation modeling. *Front Psychol*. 2015;5:1521. https://doi.org/10.3389/fpsyg.2014.01521.
[51] Van Lissa CJ, van Erp S, Clapper EB. Selecting relevant moderators with Bayesian regularized meta-regression. *Res Synth Methods*. 2023;14: 301–322.
[52] McGuinness LA. robvis: An R package and web application for visualising risk-of-bias assessments. Version 0.3.1. 2019. https://CRAN.R-project.org/package=robvis.
[53] Dias S, Ades AE, Welton NJ, Jansen JP, Sutton AJ. *Network Meta-Analysis for Decision-Making*. John Wiley & Sons; 2018.

Cite this article: Bartoš F, Wagenmakers E.-J., Viechtbauer W. Meta-analysis with JASP, Part I: Classical approaches. *Research Synthesis Methods*. 2026;0: 1–14. https://doi.org/10.1017/rsm.2026.10127

https://doi.org/10.1017/rsm.2026.10127 Published online by Cambridge University Press