Effect sizes reported in highly-cited emotion research compared to larger studies and meta-analyses addressing the same questions

Running head: Overestimation of effects in highly-cited emotion research

Ioana A. Cristea^{1,2}, Raluca Georgescu³, John P. A. Ioannidis^{2,4}

¹ Department of Brain and Behavioural Sciences, University of Pavia, Italy

² Meta-Research Innovation Center at Stanford (METRICS), Departments of Medicine, of Epidemiology and Population Health, of Biomedical Data Science, and of Statistics, Stanford University

³ Department of Clinical Psychology and Psychotherapy, Babes-Bolyai University, Cluj-Napoca, Romania

⁴ Meta-Research Innovation Center Berlin (METRIC-B), QUEST Center for Transforming Biomedical Research, Berlin Institute of Health, Charité – Universitätsmedizin Berlin

Word-count (not including abstract, tables, references and acknowledgements): 6283

Correspondence concerning this article should be addressed to Ioana A. Cristea, Department of Brain and Behavioral Sciences, University of Pavia, Piazza Botta 11, 27100 Pavia, Italy. E-mail: ioana.cristea@unipv.it

Abstract

We assessed whether the most highly-cited studies in emotion research reported larger effect sizes compared to meta-analyses and the largest studies on the same question. We screened all reports with at least 1000 citations and identified matching meta-analyses for 40 highly-cited observational and 25 highly-cited experimental studies. Observational studies had effects greater on average by 1.42-fold (95% CI 1.09 to 1.87) versus meta-analyses and 1.99-fold (95% CI 1.33 to 2.99) versus largest studies on the same questions. Experimental studies had fold-increases of 1.29 (95% CI 1.01 to 1.63) versus meta-analyses and 2.02 (95% CI 1.60 to 2.57) versus largest studies. There was substantial between-topic heterogeneity, more prominently for observational studies. Highly-cited studies uncommonly (12/65 topics, 18%) had the largest weight in meta-analyses. They were frequently (31/65 topics, 48%) the earliest published on the topic. Highly-cited studies may offer, on average, exaggerated estimates of effects in both observational and experimental designs.

Keywords: emotion; highly-cited; meta-analysis; meta-research; citation bias

Introduction

Highly-cited (HC) studies are often considered to be the most valued and influential scholarship, leading to an expectation that they should hopefully report the most accurate findings. However, meta-epidemiological investigations in some scientific fields have found that highly-cited studies may report overestimated effects relative to larger or better designed studies (Ioannidis, 2005; Tajika, Ogawa, Takeshima, Hayasaka, & Furukawa, 2015) or to meta-analyses on the same topic (Ioannidis & Panagiotou, 2011). Additionally, influential studies often produce substantially larger or contradictory effects relative to subsequent pre-registered replication attempts (Camerer et al., 2018; Klein et al., 2018; Open Science Collaboration, 2015; Wagenmakers et al., 2016).

Multiples sources of bias may contribute to effect size inflation (Fanelli, Costas, & Ioannidis, 2017; Ioannidis, Patsopoulos, & Rothstein, 2008). A major concern is that when research findings are incentivized to pass a prescribed threshold of statistical significance in order to be published ('publication bias') and research designs have suboptimal statistical power, published effect sizes (ES) are inflated on average (Bakker, van Dijk, & Wicherts, 2012; Button et al., 2013; Gelman, 2017; Ioannidis, 2008). Additionally, flexibility in analytical choices (Simmons, Nelson, & Simonsohn, 2011) can lead to a large 'vibration of effects' (i.e., the range of possible effects obtained for different analysis specifications estimating the same association) which, when combined with selective reporting, can lead to an upward bias for published effects (Patel, Burford, & Ioannidis, 2015; Steegen, Tuerlinckx, Gelman, & Vanpaemel, 2016). Finally, influential stakeholders within the scientific ecosystem, such as funders and journals, exert a preference for aesthetically appealing ('positive', 'clean', or 'novel') results (Nosek, Spies, & Motyl, 2012), which could lead to preferential citation ('citation bias') of studies that report

larger effects, compared to those reporting smaller or null effects (Cristea & Naudet, 2018; Gøtzsche, 1987; Greenberg, 2009).

Though systematic investigations of effect size inflation in highly-cited papers in the social and behavioural sciences are lacking, indirect evidence from replication studies suggests that effects reported by highly-cited studies may be exaggerated. For example, three large-scale studies have found that effects reported in multi-laboratory pre-registered replication attempts are on average 49-66% smaller than corresponding effects reported in previously published research (Camerer et al., 2016; Camerer et al., 2018; Open Science Collaboration, 2015). Many, but not all, of the included original studies were highly-cited, and they were all published in high-profile journals. Other multi-laboratory replication efforts specifically targeting influential psychology studies often report smaller (and sometimes null) effects relative to original studies (Klein et al., 2014; Klein et al., 2018; Wagenmakers et al., 2016). These replication efforts did not systematically target specifically the most highly-cited papers – even though some of the assessed work was highly-cited. Moreover, they have also focused predominantly on randomized experiments. However, there are many other studies that attract a lot of attention and citations, including diverse observational associations, biomarkers or predictive markers- and more. It would be important to assess whether highly-cited papers covering such a broad spectrum of designs have inflated effect sizes, and if so, how big the inflation is, compared with other studies on the same questions that do not get so many citations.

Hence, the goal of the present study was to investigate whether effect sizes reported in highly-cited emotion research are greater relative to larger studies and meta-analyses addressing the same questions. We focused on emotion research because it is a major topic domain in psychology with a breadth of content and research designs, covering both highly exploratory

basic research, as well as applied research with clinical implications. Our goal was to gauge the extent to which effects differed between highly-cited studies and summary effects from meta-analyses, as well the larger studies on the same topic. We also wanted to map the timing of publication of highly-cited studies, largest studies, and other studies on the same topic.

Methods:

We adopted the approach of previous similar investigations in clinical research (Ioannidis, 2005; Ioannidis & Panagiotou, 2011) and psychiatry (Tajika et al., 2015). The protocol of this study was registered prospectively on the Open Science Framework (<https://osf.io/wrnq9>). Changes to the protocol are detailed in the Supplement.

Identification and selection of target highly cited papers

The database Scopus was searched through October 8th 2019, using key-words generically related to “emotion”, “mood”, “anxiety” or “depression”, present in the title, abstract or keywords.

Eligible records reported on primary data that could be used for generating effect sizes in human participants, and that mentioned findings related to emotions in the abstract (even if these were peripheral to the goals of the study), had an experimental (i.e., randomized) or observational design, and had been cited at least 1000 times in Scopus as of the date of the search. Papers in which the abstract made no mention of emotion or focused exclusively on biomedical, molecular or other aspects not related to emotional disorders or conditions were excluded. However, papers that upon abstract inspection mentioned emotion, even if in a peripheral role (e.g., as one of many secondary outcomes, a component in a model) were included.

We also excluded: (1) meta-analyses and other papers using secondary data; (2) observational studies focused on prevalence; (3) studies describing the development or

subsequent validation of scales; (4) estimations of disease burden, such as the Global Burden of Disease (GBD).

One researcher (IAC) screened all records with at least 1000 citations by title and abstract and selected those mentioning emotion describing observational (including pre-post designs and non-randomized studies of various associations) and experimental (including all studies where participants were randomized to an intervention, or to different modalities of an independent variable) designs.

Identification and selection of meta-analyses

For each eligible observational or experimental highly cited record, we searched for the most recent meta-analysis including effect size data from any finding in the paper, provided it was related with emotion. In cases where the HC paper mentioned emotion but in a peripheral role, an eligible meta-analysis had to report ES related to the emotion finding and not to the paper's other findings.

Meta-analyses for each target paper were identified by downloading the most recent 2000 records citing the target paper as a searchable .csv file and using the “find” command in a document processor with the key word “meta-analy*” in the title, author- or index keywords. Citing records were screened starting with the most recent ones and moving downwards on the list. Whenever a potentially eligible meta-analysis was identified, the full-text was retrieved and manually searched to identify whether (1) the HC study was included and (2) an ES of interest from the HC study was reported. If these criteria were not satisfied, we moved down the list of citing papers chronologically until identifying another eligible meta-analysis. Meta-analyses that substituted the HC study with a larger study that encompassed it or with another publication on

the same sample were eligible. In these cases, we planned to recalculate the ES from the original report, if possible.

For eligible records that described more than one meta-analysis (i.e., reported more than one forest plots) including the target HC study, we chose the one with the highest number of studies, or, if there were ties in this regard, the one that appeared first in the text, provided it reported a finding related to emotion.

One researcher (IAC) searched citing records and identified meta-analyses.

Data extraction

For each matching meta-analysis, we coded information about publication year, meta-analysis model used (fixed or random), total number of included ES in the selected forest plot, ES measure (e.g., mean difference, standardized mean difference, correlation, odds ratio, risk ratio, hazard ratio), earliest study (by publication year) in the forest plot, ES and 95% CI for the HC and, respectively, largest study, summary ES and 95% CIs in the meta-analysis. When the meta-analysis reported different models of estimating ES, we preferred random effects. The largest study was defined as the study with the lowest standard error in the matching meta-analysis. To select the largest study, we relied on the following succession of information, if reported: (1) weights in the forest plots, followed by (2) standard errors/variance associated to individual ESs, followed by (3) recalculation of the 95% CI width for those individual ESs where the CI appeared visually smaller in the forest plot and, finally (4) study sample size. If more studies with the same weight or standard error were included, sample size was used to break the tie.

If more studies, including HC study, were the earliest in the forest plot (published in the same year), the HC was considered as the earliest.

When the forest plot only included graphic information, we attempted to contact the authors or use tools like WebPlotDigitizer (<https://automeris.io/WebPlotDigitizer/>) to reconstruct the data from the plots.

Outcomes

All outcomes were assessed separately for observational and experimental designs.

The primary outcome was the degree of agreement between the ES of the HC study and that of the summary ES of the matched meta-analysis; and between the ES of the HC study and the ES in the largest study in the matching meta-analysis. To this purpose, we calculated the ratios of odds ratios (RORs), as detailed in the data analysis section.

This outcome is reported both nominally, as the proportion (in percentages) of topics where the 95% CI of ROR includes 1, as well as statistically, as the meta-analytical aggregate across topics, separately for experimental and observational studies.

Secondary outcomes were the proportion (in percentages) of HC studies with ESs that differed by 2- ($\text{ROR} \geq 2$ or ≤ 0.5) or 4-fold ($\text{ROR} \geq 4$ or ≤ 0.25) from the ES in the matching meta-analysis and, respective, largest study in the meta-analysis.

Data analysis

Analyses were performed in Microsoft Excel for Mac version 16.43 and STATA/SE 16.1 for Mac (programs *admetan* and *metaeff*). Scatterplots were constructed in R (RStudio version 1.2.5033) using the lessR package (Gerbing, 2020).

For each identified meta-analysis, we extracted the ES and standard error (SE) or 95% CI reported for the HC study, the summary ES and the ES of the largest study in the meta-analysis. The preferred meta-analytic estimate was the odds ratio (OR). ESs were extracted as reported in the meta-analysis, without retrieving the primary studies. When meta-analyses reported estimates other than the OR, we employed standard procedures for converting estimates into ORs. Standardized mean differences (SMDs), including Hedges' g , were transformed to naturally logarithmic ORs (lnORs) using the Chinn transformation (Chinn, 2000). Correlation coefficients were first converted into SMDs (Polanin & Snilstveit, 2016) and then into ORs. For risk ratios (RRs) that cannot be converted into ORs without estimates of baseline risk, often not reported, we first checked whether or study level event data (e.g., 2X2 table) was reported. If yes, we extracted them and re-ran the meta-analysis with ESs expressed as ORs, using the authors' specified meta-analytic model. If neither baseline risk, nor event data were provided, remaining RRs were considered as ORs in the main analyses and excluded in sensitivity analyses. Similarly, hazard ratios (HRs) were assimilated to ORs. For continuous outcomes expressed as mean differences and standard errors or confidence intervals, we also re-ran the meta-analysis to produce standardized mean difference (SMDs).

For meta-analysis reporting data by subgroups, we took the pooled (i.e., across subgroups) estimate if available and the estimate in the largest subgroup including the HC study if the pooled estimate across all subgroups was not reported. For forest plots that included separate ESs from the HC or largest study (e.g., different subgroups or outcomes), we first pooled these distinct estimates under a fixed effects model and used that estimate for further analyses, while the summary estimate remained the one reported in the meta-analysis.

To assess the magnitude of the differences for each pair (HC versus summary estimate; HC versus largest study), we computed the ratio of ORs (ROR) using the Altman-Bland approach (Altman & Bland, 2003). Briefly, RORs were obtained by dividing the OR of the HC paper to the (a) summary ES of the meta-analysis, (b) the ES in the largest study.

In order to ensure coherence across studies, we used coining of effects. For experiments, coining was performed so that $ROR > 1$ implied that the intervention or experimental manipulation had more favorable results than control. For observational studies, exposures were coined so as to represent values over 1 for the HC study, so that a $ROR > 1$ meant that the ES in the HC study is larger than the one in the meta-analysis or larger study. For each comparison HC versus meta-analysis and respectively HC versus largest study, we noted whether RORs were statistically significant (i.e., 95% CI did not include 1) and whether estimates from the HC study differed by 2-fold or more ($ROR \geq 2$ or ≤ 0.5) or 4-fold more ($ROR \geq 4$ or ≤ 0.25) or more.

We also conducted meta-analyses of RORs separately for experimental and observational designs. We used a random-effects model with the Paule and Mandel estimator (Paule & Mandel, 1989), recommended for dichotomous outcomes in the presence of high heterogeneity (Veroniki et al., 2016). For comparisons with the largest study, cases where the largest study coincided with the HC study were excluded. Heterogeneity was assessed with the between-topic variance τ^2 , I^2 and its 95% confidence interval (CI) estimated using the Q-profile method (Viechtbauer, 2007). Since many clinical psychologists are accustomed to SMDs rather than ORs metrics for expressing effects, we also transformed summary RORs from the main analysis in differences of standardized mean differences (dSMDs), by applying the conversion formula described by Chinn (2000) using the natural logarithm ($\ln$) of the ROR. For the ROR of the HC study versus the summary estimated of the meta-analysis (MA), we have:

$$\ln(RORHCMA) = \ln\left(\frac{ORHC}{ORMA}\right) = \ln(ORHC) - \ln(ORMA) = 1.81 \times SMDHC -$$

$1.81 \times SMDMA = 1.81 dSMD$. Therefore the $dSMD = \ln(RORHCMA) \div 1.81$.

The standard errors (SE) can be computed by the same formula, $SE(dSMD) =$

$SE(\ln RORHCMA) \div 1.81$. We also reported estimates as dSMDs for the topics where the SMD was the effect measure used in the selected meta-analysis.

Sensitivity analyses were performed by repeating the main analyses (1) excluding studies for which HRs and RRs were considered as ORs; (2) limited to the topics where the HC was the earliest published on the topic; (3) limited to the topics where the largest study was been published later than the HC study; (4) restricted to HC studies where the abstract mentioned the outcome and exposure/intervention extracted from the matching meta-analysis. The analysis was added *post-hoc* to verify whether the finding selected for evaluation was considered central in the HC study. Finally, as we observed extremely large heterogeneity for observational designs, we added a series of exploratory sensitivity analyses for this cohort: (5) Limited to topics where both the exposure and outcome are clinical manifestations (e.g., depression, anxiety, insomnia), demographic variables (e.g., gender) or major life events (e.g., adverse events, childhood abuse); (6) Limited to topics where either exposure or outcome are non-clinical or surrogate measurements (e.g., genes, neuroimaging, cognitive tasks); (7) Limited to topics where the matching meta-analysis was lower variance or respectively higher variance compared to the median of the entire sample. For this analysis we calculated the median standard error of the log odds ratio for the entire cohort of meta-analyses and dichotomized the sample based on the median into topics where the standard error of the log odds ratio in the matching meta-analysis was below and respectively above the median.

Results

Selection of target highly cited papers and matching meta-analyses

The search produced 1 686 834 records, out of which 1183 had at least 1000 citations. From these, 187 studies were selected, 114 observational and 73 experimental (Figure 1 PRISMA flow diagram). Twenty-seven studies (14%) had over 2000 citations and, as per protocol and owing to Scopus limitations on download of citing records, we only screened through the first 2000 most recent citations until identifying a matching meta-analysis. This procedure failed to identify a matching meta-analysis for 19 of these 27.

We contacted authors of 3 meta-analyses where forest plots did not contain ES data or were incomplete (i.e., presented only a subgroup) and retrieved data for 2 meta-analysis. The remaining one (McKinnon, Yucel, Nazarov, & MacQueen, 2009) was excluded.

Therefore, we identified matching meta-analyses with study-level effects size data for 41/114 (36%) for observational studies (37 unique meta-analyses, 4 of them contained more than one HC study) and 25/73 (34%) for experimental studies (22 unique meta-analysis, 3 of them contained two different HC studies).

Characteristics of the sample (Table S1, Table S2)

The 41 observational studies were published between and 1972 and 2013, and citation counts ranged from 1001 to 5497 (median 1357, IQR 1087 to 1769). The 25 experimental studies spanned 1989 to 2006, with citation counts ranging from 1126 to 2374 (median 1426, IQR 1290 to 1723). Matching meta-analyses were published between 1998 and 2019 for observational studies and between 2014 and 2019 for experimental studies. 28/37 meta-analyses for

observational studies (76%) and 21/22 (95%) for experimental studies were published in the last 5 years (from 2015 onwards).

For observational studies, 17 meta-analyses used odds ratio (OR), 14 standardized mean differences (SMD) (5 Hedges' g), 7 risk ratio (RR), 1 hazard ratio (HR), 1 mean difference (MD) and 1 standardized beta coefficients from linear regression. In this last case, we did not have enough information to convert or recalculate the regression coefficient, and the meta-analysis (Martinez-Calderon, Flores-Cortes, Morales-Asencio, & Luque-Suarez, 2019) was excluded, leaving a total of 40 meta-analyses for quantitative synthesis.

There were a few special cases. In the case of one individual patient data meta-analysis (Culverhouse et al., 2018), the HC study was eligible for inclusion, but did not provide primary data. As per protocol, we used the estimates from the meta-analysis for the summary and the largest study estimates, and recalculated the effect for the HC study from the primary report, using coining so it would represent the same contrast as the meta-analysis. For another HC study (Regier et al., 1990), the meta-analysis (Lai, Cleary, Sitharthan, & Hunt, 2015) included data from a pooled analysis (Swendsen et al., 1998) combining the cohort reported in the HC study with other cohorts. The HC study did not report the sufficient data for ES calculation, but the meta-analysis included separate estimates for the HC study cohort (Epidemiologic Catchment Area), which we extracted. Another meta-analysis (Reising, Ttofi, Farrington, & Piquero, 2019) included a larger study that contained the HC study (Odgers et al., 2008). The HC study did not report sufficient data for ES calculation. We substituted its estimate with the one from the overlapping larger study (Odgers et al., 2008) reported in the meta-analysis. Since the original HC study only reported on males, and the overlapping study included separate estimates for

males and females, we only used the former. We conducted additional sensitivity analyses excluding all these special cases.

For experimental studies, 13 meta-analyses reported SMDs (6 Hedges' g), 4 ORs, 4 RR, 3 Pearson correlation coefficient (r), 1 HR. For one HC study describing the Enhancing Recovery in Coronary Heart Disease Patients (ENRICHD) randomized trial (Berkman et al., 2003), the corresponding meta-analysis (Richards et al., 2017) combined data from all trial publications. As per protocol, we used summary and largest study estimates from the meta-analysis, and recalculated the HC effects using the primary report.

Primary outcomes and meta-analysis

Observational studies (Table 1, Table 3, Figure 2)

Effect estimates were recalculated ($n=2$) or converted ($n=15$) for 16 meta-analyses. For 6 meta-analyses, RRs and HRs were assimilated to ORs in computing ratios of odds ratio (RORs). Effect estimates were coined for 8 meta-analyses. For 27/40 (67.5%) HC studies, we rated the abstract as describing the finding extracted from the matching meta-analysis. 25/40 (62.5%) HC studies were the earliest or conducted within 3 years of the earliest study in the meta-analysis.

In 27/40 HC studies (67.5%), estimates were nominally larger (i.e., $ROR > 1$) than the summary effect in the corresponding meta-analysis (Figure 2). In 12/40 HC studies (30%) effects were statistically significantly different from the summary estimate, with $ROR > 1$ in 10 cases. The difference was 2-fold or more for 15 (37.5%) pairs and 4-fold or more for 6 (15%). The summary ROR (Figure S1) across all topics was 1.42, 95% CI 1.09 to 1.87 with extremely high heterogeneity ($\tau^2 = 0.55$, $I^2 = 98\%$, 95% CI 95% to 99%). This was equivalent to a difference in SMDs (dSMDs) of 0.19, 95% CI 0.04 to 0.34 (Table 3). RORs were somewhat larger for topics

where the HC was the earliest study ($n=21$), 1.77 , 95% CI 1.07 to 2.94 , $\tau^2=1.07$, as well as those where the HC study abstract mentioned the exposure and outcome used in the meta-analysis ($n=27$), ROR of 1.60 , 95% CI 1.08 to 2.36 , $\tau^2=0.77$, but heterogeneity remained very large. For topics where the matched meta-analysis reported effects as SMDs ($n=15$), estimates were higher ($dSMD=0.45$, 95% CI 0.04 to 0.86), with extremely high heterogeneity, $\tau^2=1.66$.

Heterogeneity was significantly reduced in exploratory sensitivity analyses of topics where both exposure and outcome were clinical manifestations, demographic variables or major life events ($n=25$), ROR= 1.16 (95% CI 0.94 to 1.42), $\tau^2=0.17$. Heterogeneity was also contained in analyses ($n=20$) circumscribed to the meta-analyses with lower variance (i.e., under the median variance of the entire cohort of meta-analyses), ROR= 1.18 (95% CI 0.92 to 1.52), $\tau^2=0.26$.

Five HC studies (12.5%) were also the largest, leaving 35 for further ROR analyses. HC estimates were greater than those in the largest study in the meta-analysis in 29/35 (83%) cases (Figure 2). RORs comparing estimates were statistically significantly different from 1 in 17 cases (49%), from which $ROR>1$ in 13 cases. RORs were 2-fold or more for 17/35 cases (49%) and 4-fold or more in 10/35 (29%). The summary ROR (Figure S2) was 1.99 , 95% CI 1.33 to 2.99 , with extremely high heterogeneity ($\tau^2=1.17$, $I^2=98\%$, 95% CI 97% to 99%). This corresponded to a $dSMD$ of 0.38 , 95% CI 0.15 to 0.61 (Table 3). The summary ROR was similar in sensitivity analyses restricted to topics where the HC predated the larger study ($n=33$), ROR= 1.99 , 95% CI 1.29 to 3.07 , with similarly high heterogeneity $\tau^2=1.25$. In exploratory sensitivity analyses (Table 1) on topics where both exposure and outcome were clinical, summary RORs were reduced, and heterogeneity was more contained, ROR= 1.49 , 95% CI 0.97 to 2.3 , $\tau^2=0.81$.

Heterogeneity was substantially reduced in analyses limited to meta-analyses with lower variance, ROR= 1.71, 95% CI 1.22 to 2.41, $\tau^2= 0.40$.

Experimental studies (Table 2, Table 3, Figure 3)

Effect estimates were recalculated (n=4) or converted (n=16) for 20 meta-analyses, while for one meta-analysis the HR was considered equivalent to the OR. Effect estimates were coined for 8 meta-analyses. 15/25 (60%) HC studies were the earliest or conducted within 3 years of the earliest study in the meta-analysis. The abstract of 21/25 (84%) HC studies described the intervention and outcome used in the meta-analysis.

For 17/25 (68%) HC studies, estimates were nominally larger than summary estimates of matching meta-analyses (Figure 3). The ROR of the HC study versus the summary estimate was statistically significantly different from 1 in 6/25 cases (24%), 3 of which had ROR>1. The estimates from the HC study differed by 2-fold or more (i.e., $ROR \geq 2$ or ≤ 0.5) in 5 cases (20%), and by 4-fold or more in 1 (4%). The summary ROR (Figure S3) was 1.25, 95% CI 0.97 to 1.61, with substantial heterogeneity ($\tau^2= 0.25$, $I^2= 73\%$, 95% CI 53% to 87%). The ROR corresponded to a dSMD of 0.14, 95% CI 0.007 to 0.27). For topics where the matched meta-analyses reported effects as SMDs, summary estimates were higher, dSMD= 0.24, 95% CI 0.05 to 0.44. Sensitivity analyses limited to topics where the HC was the earliest study (n=10) resulted in a similar summary ROR of 1.33, 95% CI 1.08 to 1.64) with no between-topic heterogeneity ($\tau^2= 0$). Analyses of topics where the HC study abstracted mentioned the intervention and outcome extracted from the matching meta-analysis (n= 21) led to a similar summary ROR of 1.30, 95% CI 0.97 to 1.73, $\tau^2= 0.33$.

Seven HC studies (28%) were also the largest in the matching meta-analysis, leaving 18 studies for ROR analyses. Estimates from the HC study were nominally higher than those of the largest study in for all 18 studies (Figure 3), and for 6/18 (33%), RORs were statistically significantly different from 1. RORs were 2-fold or more in 7/18 (39%) and 4-fold or more in 1/18 (6%). The summary ROR (Figure S4) of the HC study versus the largest study was 1.81, 95% CI 1.39 to 2.36, with moderate heterogeneity ($\tau^2 = 0.09$, $I^2 = 29\%$, 95% CI 0% to 68%). One HC study post-dated the largest study. Analyses restricted to the cases where the HC study (n=17) predated the larger one yielded a larger summary ROR of 1.85, 95% CI 1.39 to 2.46, $\tau^2 = 0.10$.

Discussion

Reports that collect extreme numbers of citations can be very influential in shaping the scientific literature and often also inform crucial decisions about which research to conduct, publish or finance. Hence the validity of their claims is paramount. While there is no perfect means to evaluate validity, placing the results of HC studies against those of meta-analyses and of the largest studies on the same topic can offer valuable comparative insights. In a large field-wide survey of emotion research, we showed that HC studies report more prominent effects compared to meta-analyses and larger studies on the same topic. For observational designs, HC studies produced on average effects about 1.4-fold higher than those from meta-analyses, and almost 2-fold higher compared to the largest studies. For experimental designs, the average difference was around 1.3-fold for summary estimates and 2-fold in comparisons to the largest study. Translated in differences of standardized mean differences, an estimate more habitually used by clinical psychologists, HC observational studies produced effects higher by 0.19 compared to summary meta-analytic estimates and by 0.38 compared to the largest study. The differences were similar for experimental designs (0.14 compared to the summary estimate and 0.39 compared to the largest study).

These average differences need to be seen with great caution, because there was extremely prominent heterogeneity across topics. Heterogeneity was extremely high for observational designs, and more moderate for experimental studies. Heterogeneity was considerably reduced in exploratory sensitivity analyses limited to topics where the exposure and outcome were both clinical, demographic or major life events variables, as well as when considering the meta-analyses with the lower variance (i.e., below the median variance of the entire cohort of matching meta-analyses). In these analyses, differences between estimates were

reduced to around 1.2-fold compared to meta-analyses and 1.5-1.7 compared to larger studies and, in most cases, non-significant (the confidence interval around RORs included 1, albeit narrowly). Therefore, while on average HC studies may be expected to report larger effects, it is not possible to predict in advance for which topics this will be most pronounced and for which topics HC may not have larger effects at all. It is impossible to “correct” the effect estimates of a HC study by using some standard inflation factor.

We also examined the timing of publication of the HC reports as compared to the other studies and to the largest studies on the same topic. HC studies are sometimes the first ones on the topic and thus they would be the earliest published among the studies included in a meta-analysis. This pattern occurred in almost half (31/65, 48%) of the topics that we examined. However, in approximately 40% of cases, HC studies were published later or even substantially later (i.e., more than 3 years after). Their high citation profile may reflect early publication (“being the first”), some citation bias favoring extreme results, or a combination thereof. Relatedly, the HC study predated the largest study in about two-thirds of the pairs for observational designs and in all but one for experimental ones. Sensitivity analyses limited to topics where the HC studies were the first ones mirrored the main analyses. If anything, the summary ROR estimates became slightly larger when only these topics were considered, a pattern compatible with some influence of “being first”. However, the available data are too limited to exclude that this observation may reflect chance.

Finally, our approach of selecting a recent meta-analysis that used emotion-related estimates from the HC studies could have failed to capture the main outcomes of these studies that led to their high citation impact. To account for this possibility, we added a post-hoc sensitivity analysis restricted to instances where the selected meta-analytic comparison included

the outcome and exposure/intervention also mentioned in the abstract of the HC study. For most observational (65%) and experimental (85%) HC studies, this was indeed the case and this sensitivity analysis resulted in very similar results to the main analysis. Of course, the approach cannot fully guarantee we examined the principal finding of the HC study, and it is often impossible to single out only one particular finding from a complex study. However, given that abstracts describe what are considered by the authors to be the most noteworthy results, this approach could represent a useful proxy to identifying the principal findings. We did not assess the quality of the HC studies, as this would have posed significant challenges given the diversity of topics, designs, and scientific standards at the time of publication. Though study size is not a surrogate for quality, larger studies are more precise in estimating effects. In general, it was uncommon for HC studies to be also the largest ones.

Some of the HC studies had extremely large effects that also differed tremendously from the respective meta-analyses and largest studies. In the most conspicuous case (RORs of 287 and 297 respectively), Hariri et al. (2002) examined neuroimaging differences in amygdala activation in carriers of the short (1 or 2 copies) versus those of the long 5-HTTLPR allele. The paper collected 1769 citations and its standardized effects size in the matching meta-analysis (Munafò, Brown, & Hariri, 2008) was an incredible standardized mean difference of 3.74 (95% CI 2.51 to 4.97). In contrast, the summary effect in the meta-analysis was considerably smaller (SMD= 0.62, 95% CI 0.42 to 0.82), similarly to the largest study (Hariri et al., 2005) (SMD= 0.6, 95% CI 0.14 to 1.06). The true effect may actually be entirely null. The reason is that this HC study, as well as the other studies in the meta-analysis depend on a candidate gene approach, a design that has since been shown as notoriously unreliable (Ioannidis, Tarone, & McLaughlin, 2011), even more so in neuropsychiatric genetics (Duncan, Ostacher, & Ballon, 2019). Moreover,

neuroimaging studies are a classic example of a literature replete with small, underpowered studies with high analytical flexibility and often spurious results (Botvinik-Nezer et al., 2020; David et al., 2013; Szucs & Ioannidis, 2020). The proposed association with amygdala activation (Hariri et al., 2005; Hariri et al., 2002) would suggest a role of this genetic polymorphism in depression. However, very large, rigorous genome-wide association study have found absolutely no effect for this polymorphism (Border et al., 2019).

RORs were also very large (22 and 29) for a study (Klin, Jones, Schultz, Volkmar, & Cohen, 2002) examining differences in visual fixation patterns between autistic males and healthy controls while viewing social situations (1150 citations). In this case, the effects between the index study (SMD= -1.47, 95% CI -2.27 to -0.66) versus the corresponding meta-analysis (SMD=0.24, 95% CI 0.1 to 0.39) and largest study in it (SMD= 0.39, 95% CI 0.19 to 0.58) differed not just by magnitude, but also by direction. There were no such large outliers in the analyses on experimental studies. Overall, for observational designs, topics where the original meta-analyses used the SMD as the metric of choice, tended to have greater differences in effect size between the highly cited study and the respective meta-analysis or largest study. Several of the large outliers identified in neuroimaging or genetics belong to this category.

In selecting meta-analyses that included the index study, we focused on the most recent one that contained effect size data. Around 80% of the identified meta-analyses for observational studies and all but one for experimental studies were published in the last 5 years. The recency of selected meta-analyses makes it more likely that they included a larger number of publications. Also, the quality of reporting and analysis might also have improved with time (Page et al., 2016; Wen et al., 2008). However, we should caution that the “true” effects for the topic examined are unknown, and effect may genuinely differ across studies on the same topic due to

genuine differences rather than bias. Moreover, meta-analyses and even single large studies may also be biased. Random effects models for obtaining summary results are appropriate in situations where there is substantial heterogeneity, as is often the case in emotion research, but random effects estimates are also susceptible to biases such as small-study effects that might underlie publication bias (Sterne et al., 2011). On average, meta-analyses may be more biased than the largest studies. This would be entirely consistent with our observation that HC results seemed to be less inflated when the comparison was made against the summary effect of a meta-analysis than when it was made against the largest study.

Kvarven, Strømland, and Johannesson (2020) have employed a somewhat similar methodological approach to compare results from registered replications with meta-analyses testing the same hypotheses. The starting point were multi-laboratory registered replication studies in psychology, for which matching meta-analyses on the same hypothesis, as identified by the study authors, were searched. The authors retrieved meta-analyses with effect size data for 15/62 replication studies selected and used a Z-test to compare replication effects with summary meta-analysis estimates, either by a random effects model or with bias adjustment. Results indicated an almost 3-fold increase in summary meta-analysis estimates compared to replication studies, even when using methods to adjust for publication bias, suggesting that better designed studies where publication bias is avoided (as in the case of pre-registered replications) may provide the most accurate effect estimates. If one were to extrapolate from their findings to ours, it is possible that HC studies provide highly inflated results, more inflated than what a comparison against meta-analyses would suggest. Even the comparison against the largest available study may not fully capture the inflation of results, because these largest studies that we

used were not pre-registered. Therefore, they could also suffer from some selective reporting of analyses.

Of direct relevance to the present work, Kvarven et al. (2020) also compared estimates from the original studies, defined as the study that was the object of the replication project, with those in the selected meta-analyses, reporting a non-statistical mean difference of 0.10 for 14 pairs of original studies and meta-analyses. However, the pairs of replication studies and meta-analyses included mostly meta-analyses of small studies and such meta-analyses may also be unreliable and biased. Similarly, we showed that for meta-analyses with reduced variance, and hence lower uncertainty around the summary effects, differences between estimates from HC studies and summary ones were reduced and no longer significant (RORs close to 1). Conversely, for meta-analyses with higher variance and highly uncertain estimates, differences with HC studies were augmented. In an analysis of 200 meta-analyses published in *Psychological Bulletin*, the flagship journal for psychological science, Stanley, Carter, and Doucouliagos (2018) found that only a tiny proportion (<1%) of experimental studies are adequately powered, as compared to about a third of observational studies. Meta-analyses comprised entirely of underpowered studies may not be a good gold standard.

Our findings need to be qualified by important limitations. We were able to identify a matching meta-analysis containing effect size data for only a third of our sample of target papers. We considered meta-analyses as eligible if they included any emotion related finding from the target paper so as to avoid ranking findings in the original paper in terms of importance. Previous research has dealt with this problem by choosing a finding for which effect size data is reported in the abstract (Ioannidis & Panagiotou, 2011). However, we were concerned that most of the papers in social and behavioral sciences might simply present findings narratively, with absent or

incomplete data, especially in abstracts. Moreover, there is evidence that abstracts are frequently inconsistent with full reports (Li et al., 2017). Nonetheless, ancillary analyses restricted to findings that were mentioned in the abstract supported our main findings. Also, in the interest of consistency, when a matching meta-analysis included multiple forest plots, we chose the largest one, though it might not have used the most important finding from the highly-cited study. We were able to screen a maximum of 2000 citations for each target paper, due to the limitations of exporting data from Scopus. As we were mostly interested in research on human participants, more general terms like “fear” or “stress” were not used, as they would have rendered the search overtly non-specific. Finally, we cannot exclude the possibility that in some cases where ES was larger in the HC study, the HC study may have been more “correct” than the respective meta-analysis and the largest study on the topic. For instance, the HC study might had some particularly high-quality features and protection from bias that other studies did not, and biases might have eroded an otherwise genuinely large effect in the other studies. However, this does not seem to be the case in other fields where HC studies versus other evidence have been assessed.

Investigations of HC papers in the social and behavioural sciences have been limited and mostly restricted to surveying content and design (Price, Floyd, Fagan, & Smithson, 2011) or the availability and sharing of the data underlying their findings in HC papers (Hardwicke & Ioannidis, 2018). We add to this meta-research literature by demonstrating a pervasive systematic citation bias towards exaggerated effects across empirical studies in emotion research.

Acknowledgments.

METRICS is supported by a grant from the Laura and John Arnold Foundation. The work of JPAI is supported by an unrestricted gift from Sue and Bob O'Donnell. The funders had no role in study design, data collection and analysis, decision to publish, or preparation of the manuscript. The authors thank Tom E. Hardwicke for contributions to revising the study protocol.

Competing interests: None

Author contributions:

IAC and JPAI were responsible for study concept and design. IAC, RG, and JPAI were responsible for acquisition, analysis and interpretation of data. IAC was responsible for statistical analysis. JPAI was responsible for study supervision. IAC was responsible for drafting of the manuscript. RG and JPAI were responsible for critical revision of the manuscript for important intellectual content. All authors reviewed the final version of the manuscript.

Data and code availability statement:

The datasets generated and analyzed during the study, and the STATA and R scripts used for the data analysis and visualization are publicly available on the Open Science Framework:

DOI 10.17605/OSF.IO/4Y8FK

References:

- Altman, D. G., & Bland, J. M. (2003). Interaction revisited: the difference between two estimates. *BMJ*, 326(7382), 219. doi:10.1136/bmj.326.7382.219
- Bakker, M., van Dijk, A., & Wicherts, J. M. (2012). The Rules of the Game Called Psychological Science. *Perspect Psychol Sci*, 7(6), 543-554. doi:10.1177/1745691612459060
- Berkman, L. F., Blumenthal, J., Burg, M., Carney, R. M., Catellier, D., Cowan, M. J., . . . Schneiderman, N. (2003). Effects of treating depression and low perceived social support on clinical events after myocardial infarction: the Enhancing Recovery in Coronary Heart Disease Patients (ENRICHD) Randomized Trial. *Jama*, 289(23), 3106-3116. doi:10.1001/jama.289.23.3106
- Border, R., Johnson, E. C., Evans, L. M., Smolen, A., Berley, N., Sullivan, P. F., & Keller, M. C. (2019). No Support for Historical Candidate Gene or Candidate Gene-by-Interaction Hypotheses for Major Depression Across Multiple Large Samples. *Am J Psychiatry*, 176(5), 376-387. doi:10.1176/appi.ajp.2018.18070881
- Botvinik-Nezer, R., Holzmeister, F., Camerer, C. F., Dreber, A., Huber, J., Johannesson, M., . . . Schonberg, T. (2020). Variability in the analysis of a single neuroimaging dataset by many teams. *Nature*, 582(7810), 84-88. doi:10.1038/s41586-020-2314-9
- Button, K. S., Ioannidis, J. P. A., Mokrysz, C., Nosek, B. A., Flint, J., Robinson, E. S. J., & Munafò, M. R. (2013). Power failure: why small sample size undermines the reliability of neuroscience. *Nature Reviews Neuroscience*, 14(5), 365-376. doi:10.1038/nrn3475
- Camerer, C. F., Dreber, A., Forsell, E., Ho, T.-H., Huber, J., Johannesson, M., . . . Wu, H. (2016). Evaluating replicability of laboratory experiments in economics. *Science*, 351(6280), 1433-1436. doi:10.1126/science.aaf0918
- Camerer, C. F., Dreber, A., Holzmeister, F., Ho, T.-H., Huber, J., Johannesson, M., . . . Wu, H. (2018). Evaluating the replicability of social science experiments in Nature and Science between 2010 and 2015. *Nature Human Behaviour*, 2(9), 637-644. doi:10.1038/s41562-018-0399-z
- Chinn, S. (2000). A simple method for converting an odds ratio to effect size for use in meta-analysis. *Statistics in Medicine*, 19(22), 3127-3131. doi:10.1002/1097-0258(20001130)19:22<3127::Aid-sim784>3.0.Co;2-m
- Cristea, I. A., & Naudet, F. (2018). Defending psychiatry or defending the trivial effects of therapeutic interventions? A citation content analysis of an influential paper. *Epidemiol Psychiatr Sci*, 27(3), 230-239. doi:10.1017/s2045796017000750
- Culverhouse, R. C., Saccone, N. L., Horton, A. C., Ma, Y., Anstey, K. J., Banaschewski, T., . . . Bierut, L. J. (2018). Collaborative meta-analysis finds no evidence of a strong interaction between stress and 5-HTTLPR genotype contributing to the development of depression. *Mol Psychiatry*, 23(1), 133-142. doi:10.1038/mp.2017.44
- David, S. P., Ware, J. J., Chu, I. M., Loftus, P. D., Fusar-Poli, P., Radua, J., . . . Ioannidis, J. P. A. (2013). Potential Reporting Bias in fMRI Studies of the Brain. *PLoS One*, 8(7), e70104. doi:10.1371/journal.pone.0070104
- Duncan, L. E., Ostacher, M., & Ballon, J. (2019). How genome-wide association studies (GWAS) made traditional candidate gene studies obsolete. *Neuropsychopharmacology*, 44(9), 1518-1523. doi:10.1038/s41386-019-0389-5

- Fanelli, D., Costas, R., & Ioannidis, J. P. A. (2017). Meta-assessment of bias in science. *Proc Natl Acad Sci U S A*, 114(14), 3714-3719. doi:10.1073/pnas.1618569114
- Gelman, A. (2017). The Failure of Null Hypothesis Significance Testing When Studying Incremental Changes, and What to Do About It. *Personality and Social Psychology Bulletin*, 0(0), 0146167217729162. doi:10.1177/0146167217729162
- Gerbing, D. (2020). lessR: Less Code, More Results. R package (Version 3.9.8). Retrieved from <https://cran.r-project.org/package=lessR>
- Gøtzsche, P. C. (1987). Reference bias in reports of drug trials. *Br Med J (Clin Res Ed)*, 295(6599), 654-656.
- Greenberg, S. A. (2009). How citation distortions create unfounded authority: analysis of a citation network. *BMJ*, 339, b2680. doi:10.1136/bmj.b2680
- Hardwicke, T. E., & Ioannidis, J. P. A. (2018). Populating the Data Ark: An attempt to retrieve, preserve, and liberate data from the most highly-cited psychology and psychiatry articles. *PLoS One*, 13(8), e0201856. doi:10.1371/journal.pone.0201856
- Hariri, A. R., Drabant, E. M., Munoz, K. E., Kolachana, B. S., Mattay, V. S., Egan, M. F., & Weinberger, D. R. (2005). A susceptibility gene for affective disorders and the response of the human amygdala. *Arch Gen Psychiatry*, 62(2), 146-152. doi:10.1001/archpsyc.62.2.146
- Hariri, A. R., Mattay, V. S., Tessitore, A., Kolachana, B., Fera, F., Goldman, D., . . . Weinberger, D. R. (2002). Serotonin transporter genetic variation and the response of the human amygdala. *Science*, 297(5580), 400-403. doi:10.1126/science.1071829
- Ioannidis, J. P. A. (2005). Contradicted and initially stronger effects in highly cited clinical research. *Jama*, 294(2), 218-228. doi:10.1001/jama.294.2.218
- Ioannidis, J. P. A. (2008). Why most discovered true associations are inflated. *Epidemiology*, 19(5), 640-648. doi:10.1097/EDE.0b013e31818131e7
- Ioannidis, J. P. A., & Panagiotou, O. A. (2011). Comparison of effect sizes associated with biomarkers reported in highly cited individual articles and in subsequent meta-analyses. *Jama*, 305(21), 2200-2210. doi:10.1001/jama.2011.713
- Ioannidis, J. P. A., Patsopoulos, N. A., & Rothstein, H. R. (2008). Reasons or excuses for avoiding meta-analysis in forest plots. *BMJ*, 336(7658), 1413-1415. doi:10.1136/bmj.a117
- Ioannidis, J. P. A., Tarone, R., & McLaughlin, J. K. (2011). The false-positive to false-negative ratio in epidemiologic studies. *Epidemiology*, 22(4), 450-456. doi:10.1097/EDE.0b013e31821b506e
- Klein, R. A., Ratliff, K. A., Vianello, M., Adams Jr, R. B., Bahník, Š., Bernstein, M. J., . . . Nosek, B. A. (2014). Investigating variation in replicability: A “many labs” replication project. *Social Psychology*, 45(3), 142-152. doi:10.1027/1864-9335/a000178
- Klein, R. A., Vianello, M., Hasselman, F., Adams, B. G., Adams, R. B., Alper, S., . . . Nosek, B. A. (2018). Many Labs 2: Investigating Variation in Replicability Across Samples and Settings. *Advances in Methods and Practices in Psychological Science*, 1(4), 443-490. doi:10.1177/2515245918810225
- Klin, A., Jones, W., Schultz, R., Volkmar, F., & Cohen, D. (2002). Visual Fixation Patterns During Viewing of Naturalistic Social Situations as Predictors of Social Competence in Individuals With Autism. *Arch Gen Psychiatry*, 59(9), 809-816. doi:10.1001/archpsyc.59.9.809

- Kvarven, A., Strømmland, E., & Johannesson, M. (2020). Comparing meta-analyses and preregistered multiple-laboratory replication projects. *Nature Human Behaviour*, 4(4), 423-434. doi:10.1038/s41562-019-0787-z
- Lai, H. M., Cleary, M., Sitharthan, T., & Hunt, G. E. (2015). Prevalence of comorbid substance use, anxiety and mood disorders in epidemiological surveys, 1990-2014: A systematic review and meta-analysis. *Drug Alcohol Depend*, 154, 1-13. doi:10.1016/j.drugalcdep.2015.05.031
- Li, G., Abbade, L. P. F., Nwosu, I., Jin, Y., Leenus, A., Maaz, M., . . . Thabane, L. (2017). A scoping review of comparisons between abstracts and full reports in primary biomedical research. *BMC Med Res Methodol*, 17(1), 181. doi:10.1186/s12874-017-0459-5
- Martinez-Calderon, J., Flores-Cortes, M., Morales-Asencio, J. M., & Luque-Suarez, A. (2019). Pain-Related Fear, Pain Intensity and Function in Individuals With Chronic Musculoskeletal Pain: A Systematic Review and Meta-Analysis. *J Pain*, 20(12), 1394-1415. doi:10.1016/j.jpain.2019.04.009
- McKinnon, M. C., Yucel, K., Nazarov, A., & MacQueen, G. M. (2009). A meta-analysis examining clinical predictors of hippocampal volume in patients with major depressive disorder. *Journal of psychiatry & neuroscience : JPN*, 34(1), 41-54.
- Munafò, M. R., Brown, S. M., & Hariri, A. R. (2008). Serotonin transporter (5-HTTLPR) genotype and amygdala activation: a meta-analysis. *Biol Psychiatry*, 63(9), 852-857. doi:10.1016/j.biopsych.2007.08.016
- Nosek, B. A., Spies, J. R., & Motyl, M. (2012). Scientific Utopia: II. Restructuring Incentives and Practices to Promote Truth Over Publishability. *Perspect Psychol Sci*, 7(6), 615-631. doi:10.1177/1745691612459058
- Odgers, C. L., Moffitt, T. E., Broadbent, J. M., Dickson, N., Hancox, R. J., Harrington, H., . . . Caspi, A. (2008). Female and male antisocial trajectories: from childhood origins to adult outcomes. *Dev Psychopathol*, 20(2), 673-716. doi:10.1017/s0954579408000333
- Open Science Collaboration. (2015). Estimating the reproducibility of psychological science. *Science*, 349(6251), aac4716. doi:10.1126/science.aac4716
- Page, M. J., Shamseer, L., Altman, D. G., Tetzlaff, J., Sampson, M., Tricco, A. C., . . . Moher, D. (2016). Epidemiology and Reporting Characteristics of Systematic Reviews of Biomedical Research: A Cross-Sectional Study. *PLOS Medicine*, 13(5), e1002028. doi:10.1371/journal.pmed.1002028
- Patel, C. J., Burford, B., & Ioannidis, J. P. A. (2015). Assessment of vibration of effects due to model specification can demonstrate the instability of observational associations. *J Clin Epidemiol*, 68(9), 1046-1058. doi:10.1016/j.jclinepi.2015.05.029
- Paule, R. C., & Mandel, J. (1989). Consensus Values, Regressions, and Weighting Factors. *J Res Natl Inst Stand Technol*, 94(3), 197-203. doi:10.6028/jres.094.020
- Polanin, J. R., & Snilstveit, B. (2016). Converting between effect sizes. *Campbell Systematic Reviews*, 12(1), 1-13. doi:10.4073/cmpn.2016.3
- Price, K. W., Floyd, R. G., Fagan, T. K., & Smithson, K. (2011). Journal article citation classics in school psychology: analysis of the most cited articles in five school psychology journals. *J Sch Psychol*, 49(6), 649-667. doi:10.1016/j.jsp.2011.10.001

- Regier, D. A., Farmer, M. E., Rae, D. S., Locke, B. Z., Keith, S. J., Judd, L. L., & Goodwin, F. K. (1990). Comorbidity of mental disorders with alcohol and other drug abuse. Results from the Epidemiologic Catchment Area (ECA) Study. *Jama*, 264(19), 2511-2518.
- Reising, K., Ttofi, M. M., Farrington, D. P., & Piquero, A. R. (2019). Depression and anxiety outcomes of offending trajectories: A systematic review of prospective longitudinal studies. *Journal of Criminal Justice*, 62, 3-15.
doi:<https://doi.org/10.1016/j.jcrimjus.2018.05.002>
- Richards, S. H., Anderson, L., Jenkinson, C. E., Whalley, B., Rees, K., Davies, P., . . . Taylor, R. S. (2017). Psychological interventions for coronary heart disease. *Cochrane Database Syst Rev*, 4(4), Cd002902. doi:10.1002/14651858.CD002902.pub4
- Simmons, J. P., Nelson, L. D., & Simonsohn, U. (2011). False-positive psychology: undisclosed flexibility in data collection and analysis allows presenting anything as significant. *Psychol Sci*, 22(11), 1359-1366. doi:10.1177/0956797611417632
- Stanley, T. D., Carter, E. C., & Doucouliagos, H. (2018). What meta-analyses reveal about the replicability of psychological research. *Psychol Bull*, 144(12), 1325-1346.
doi:10.1037/bul0000169
- Steenen, S., Tuerlinckx, F., Gelman, A., & Vanpaemel, W. (2016). Increasing Transparency Through a Multiverse Analysis. *Perspect Psychol Sci*, 11(5), 702-712.
doi:10.1177/1745691616658637
- Sterne, J. A. C., Sutton, A. J., Ioannidis, J. P. A., Terrin, N., Jones, D. R., Lau, J., . . . Higgins, J. P. T. (2011). Recommendations for examining and interpreting funnel plot asymmetry in meta-analyses of randomised controlled trials. *BMJ*, 343, d4002. doi:10.1136/bmj.d4002
- Swendsen, J. D., Merikangas, K. R., Canino, G. J., Kessler, R. C., Rubio-Stipec, M., & Angst, J. (1998). The comorbidity of alcoholism with anxiety and depressive disorders in four geographic communities. *Comprehensive Psychiatry*, 39(4), 176-184.
doi:[https://doi.org/10.1016/S0010-440X\(98\)90058-X](https://doi.org/10.1016/S0010-440X(98)90058-X)
- Szucs, D., & Ioannidis, J. P. A. (2020). Sample size evolution in neuroimaging research: An evaluation of highly-cited studies (1990-2012) and of latest practices (2017-2018) in high-impact journals. *Neuroimage*, 221, 117164.
doi:10.1016/j.neuroimage.2020.117164
- Tajika, A., Ogawa, Y., Takeshima, N., Hayasaka, Y., & Furukawa, T. A. (2015). Replication and contradiction of highly cited research papers in psychiatry: 10-year follow-up. *Br J Psychiatry*, 207(4), 357-362. doi:10.1192/bjp.bp.113.143701
- Veroniki, A. A., Jackson, D., Viechtbauer, W., Bender, R., Bowden, J., Knapp, G., . . . Salanti, G. (2016). Methods to estimate the between-study variance and its uncertainty in meta-analysis. *Res Synth Methods*, 7(1), 55-79. doi:10.1002/jrsm.1164
- Viechtbauer, W. (2007). Confidence intervals for the amount of heterogeneity in meta-analysis. *Stat Med*, 26(1), 37-52. doi:10.1002/sim.2514
- Wagenmakers, E.-J., Beek, T., Dijkhoff, L., Gronau, Q. F., Acosta, A., Adams, R. B., . . . Zwaan, R. A. (2016). Registered Replication Report: Strack, Martin, & Stepper (1988). *Perspectives on Psychological Science*, 11(6), 917-928. doi:10.1177/1745691616674458
- Wen, J., Ren, Y., Wang, L., Li, Y., Liu, Y., Zhou, M., . . . Tian, W. (2008). The reporting quality of meta-analyses improves: a random sampling study. *J Clin Epidemiol*, 61(8), 770-775.
doi:10.1016/j.jclinepi.2007.10.008

PRISMA 2009 Flow Diagram

Figure 1. PRISMA flow diagram of the study selection process

Table 1. Meta-analytic estimates and sensitivity analyses of ratio of odds ratios (RORs) for observational designs.

Comparison	N	ROR	95% CI	I^2 (95% CI)	τ^2
HC study vs summary estimate	40	1.42	1.09 to 1.87	98 (95 to 99)	0.55
Excluding HC study estimates recalculated from primary papers ^{*)}	37	1.47	1.08 to 1.99	98 (95 to 99)	0.64
RRs & HRs substituted for ORs ^{**)}	34	1.54	1.11 to 2.14	98 (95 to 99)	0.69
Excluding coined estimates	32	1.30	0.97 to 1.74	97 (92 to 99)	0.51
HC study is the earliest study	21	1.77	1.07 to 2.94	98 (95 to 99)	1.07
Meta-analytic comparison mentioned in HC study abstract	27	1.60	1.08 to 2.36	95 (86 to 98)	0.77
Exposure & outcome clinical/demographic/life events	25	1.16	0.94 to 1.42	91 (80 to 97)	0.17
Exposure or outcome non-clinical/ surrogate	15	2.22	1.11 to 4.45	99 (97 to 100)	1.55
Variance of meta-analysis < median variance of the cohort	20	1.18	0.92 to 1.52	98 (92 to 99)	0.26
Variance of meta-analysis > median variance of the cohort	20	1.91	1.14 to 3.21	87 (66 to 95)	0.93
HC study vs largest study	35	1.99	1.33 to 2.99	98 (97 to 99)	1.17
Excluding HC study estimates recalculated from primary papers ^{*)}	32	2.08	1.33 to 3.28	97 (94 to 99)	1.33
RRs & HRs substituted for ORs ^{**)}	30	2.27	1.41 to 3.66	97 (93 to 98)	1.38
Excluding coined estimates	28	1.62	1.06 to 2.48	98 (96 to 99)	1.04
HC study predates largest study	33	1.99	1.29 to 3.07	98 (96 to 99)	1.25
Meta-analytic comparison mentioned in HC study abstract	26	2.14	1.31 to 3.48	97 (93 to 99)	1.22
Exposure & outcome clinical/demographic/life events	22	1.49	0.97 to 2.3	98 (96 to 99)	0.81
Exposure or outcome non-clinical/ surrogate	13	3.31	1.5 to 7.32	89 (71 to 97)	1.69
Variance of meta-analysis < median variance of the cohort	17	1.71	1.22 to 2.41	98 (92 to 99)	0.40
Variance of meta-analysis > median variance of the cohort	18	2.51	1.17 to 5.41	91 (82 to 97)	2.12

Note. ES, effect size; HC, highly-cited; HR, hazard ratio;; RR, risk ratio; OR, odds ratio

*) Excluded HC studies: Caspi, 2003; Regier, 1990; Moffitt, 2002

**) The comparison excludes cases where RRs and HRs could not be converted and were considered equivalent to ORs

Table 2. Meta-analytic estimates and sensitivity analyses of ratio of odds ratios (RORs) for experimental designs.

Comparison	N	ROR	95% CI	I^2 (95% CI)	τ^2
HC study vs summary estimate	25	1.29	1.01 to 1.63	69 (44 to 85)	0.20
Excluding HC study estimates recalculated from primary papers ^{*)}	24	1.30	1.01 to 1.67	67 (42 to 84)	0.22
RRs & HRs substituted for ORs ^{**)}	24	1.34	1.06 to 1.70	66 (37 to 84)	0.19
Excluding coined estimates	17	1.21	0.97 to 1.51	47 (9 to 77)	0.09
HC study is the earliest study	10	1.32	1.07 to 1.62	0 (0 to 61)	0
Meta-analytic comparison mentioned in HC study abstract	21	1.33	1.02 to 1.75	74 (52 to 88)	0.26
HC study vs largest study	18	2.02	1.60 to 2.57	17 (0 to 61)	0.04
Excluding coined estimates	13	2.18	1.62 to 2.95	30 (0 to 74)	0.09
HC predates largest study	17	2.09	1.62 to 2.68	18 (0 to 62)	0.05
Meta-analytic comparison mentioned in HC abstract	15	2.07	1.58 to 2.72	32 (0 to 73)	0.09

Note. ES, effect size; HC, highly-cited; HR, hazard ratio; RR, risk ratio; OR, odds ratio

^{*)} Excluded HC studies: Berkman, 2003

^{**)} The comparison excludes cases where RRs and HRs could not be converted and were considered equivalent to ORs

Table 3. Meta-analytic primary analyses estimates and sensitivity analyses expressed as differences in standardized mean differences

Design	Comparison	N	dSMD	95% CI	τ^2
Observational	HC study vs summary estimate	40	0.19	0.04 to 0.34	0.55
	Meta-analyses reporting SMDs	15	0.45	0.04 to 0.86	1.66
	HC study vs largest study	35	0.38	0.15 to 0.61	1.17
	Meta-analyses reporting SMDs	14	0.73	0.28 to 1.18	1.91
Experimental	HC study vs summary estimate	25	0.14	0.007 to 0.27	0.20
	Meta-analyses reporting SMDs	13	0.24	0.05 to 0.44	0.25
	HC study vs largest study	18	0.39	0.26 to 0.52	0.04
	Meta-analyses reporting SMDs	9	0.41	0.17 to 0.65	0.18

Note. dSMD, difference in standardized mean differences; HC, highly-cited; SMD, standardized mean difference

Figure 2. Odds ratio in the a) highly-cited studies versus the corresponding summary estimate in the meta-analysis and b) highly-cited studies versus the largest study estimate in observational designs.

Diagonal lines represent equal effects between the highly-cited studies and the a) meta-analysis or b) largest studies. Not shown are two very large outliers (odds ratio of 883 and 49 in highly-cited study). For b) also not shown are five topics where the highly-cited study was the largest study

Figure 3. Odds ratio in the a) highly-cited studies versus the corresponding summary estimate in the meta-analysis and b) highly-cited studies versus the largest study estimate in experimental designs.

Diagonal lines represent equal effects between the highly-cited studies and the a) meta-analysis or (b) largest studies. For b) also not shown are seven topics where the highly-cited study was the largest study