Noise in the process: an assessment of the evidential value of mediation effects in marketing journals

AARON CHARLTON

AMANDA K. MONTOYA

JOHN PRICE

JOSEPH HILGARD*

* Aaron Charlton is an Assistant Professor of Marketing at Illinois State University (email: abchar1@ilstu.edu). Amanda K. Montoya is an Assistant Professor of Psychology at University of California, Los Angeles (akmontoya@psych.ucla.edu). John Price is a doctoral student in Marketing at Vienna University of Economics and Business (john.price@wu.ac.at). Joseph Hilgard was an Assistant Professor of Psychology at Illinois State University (jhilgard@gmail.com). The authors thank Quentin André for feedback on an early draft. They also thank student research assistants who helped in data collection and coding: Mukhtar Aghazada, Sonia Contreras, Caesar Garduno, Nancy Gomez, Kathleen Lamarque-Navarrete, Carolyn Moor, Valerie Polad, and Sabia Sultana. Correspondence can be sent to Aaron Charlton.

Abstract

Mediation analysis plays a central role in marketing research due to its usefulness in helping to explain complex processes. Like other forms of inference, mediation analyses are susceptible to false positive results. This is particularly true when analytic decisions are based on the data, rather than a priori hypotheses. To assess the collective evidential value of mediation analyses in marketing, we used an approach first implemented by Götz and colleagues (2021) that (1) measures the relative proximity of confidence intervals to zero (RP) and (2) aggregates a related set of RP scores into a single distribution. For our analysis, we compared the RP distribution of top marketing journals (2018-20) to simulations of low power, adequate power, and null effects. We also compared the marketing journals to real-world data from *Journal of Personality and Social Psychology* (JPSP) (2018-20). We found that, in terms of evidential value, mediation analyses in marketing substantially deviated from simulations of adequate power and JPSP but were similar to simulations of low power and null effects. We propose study preregistration, corrections for multiple testing, and increased statistical power as solutions to increase evidence quality going forward.

Keywords: Meta-analysis, mediation analysis, *p*-hacking, *p*-curve

Marketing journals, particularly those focused on consumer psychology, have developed a strong appetite for mediation analyses. The hope is that mediation models, whether simple mediation, moderated mediation, parallel mediation, or other similar models, can reveal the underlying process by which a predictor of interest affects the outcome variable. The adoption of mediation analysis has also been aided by the availability of free, easy-to-use software tools (e.g. Hayes 2018). Like other tests of inference, mediation analyses can suffer from problems of evidential value such as inflated type I errors (false positive results) when researchers have the latitude to run many tests and only report those that obtain statistical significance (Götz et al. 2021; Simmons et al. 2011).

Simonsohn and colleagues (2014) demonstrated that through a tool called '*p*-curve' problems of evidential value can sometimes be detected in the distribution of a related set of reported *p*-values. Mediation analyses are usually reported without a *p*-value, so the same inference about collective evidential value cannot be made. However, Götz and colleagues (2021) implemented a method they called relative proximity (RP), a statistic which can provide a comparable measure of degree of significance for confidence intervals (CI), which are typically reported for mediation analysis. The RP score is calculated by dividing the CI's distance to zero by its width. The RP score is negative if the CI overlaps zero; positive otherwise. Positive RP scores that are very close to zero are comparable to *p*-values that are only slightly below the .05 threshold in that they give relatively poor evidence against the null hypothesis.

In the research that follows, we apply *p*-curve analytic thinking (Simonsohn et al. 2014) using the aforementioned RP score to assess the collective evidential value of mediation studies published in marketing journals. We extracted information about CIs for tests of mediation from 4 marketing journals over 3 years (2018-2020), calculating RP for the first mediation test

reported in each article. For Study 1, we simulated how RP scores ought to be distributed for low power, adequate power, and null effects. Study 2 compensates for some of the weaknesses of simulations by comparing the RP distribution from marketing journals to real-world data from a peer journal, *Journal of Personality and Social Psychology (JPSP)*, for the same three-year period.

Visual comparisons of RP distributions showed that the marketing journals matched both low power and null effects simulations fairly closely, but deviated substantially from simulations of adequate power. We also found that the RP distribution from each marketing journal (median positive RP range: .11 to .13) showed substantially less evidential value than *JPSP* (median positive RP = .21). This approach cannot be used to determine whether the problems with evidential value were caused by underpowered studies, false positive results, or some combination of both. To increase evidential value in future mediation analyses, we propose (1) study preregistration, (2) corrections for multiple testing, and (3) increased statistical power. We also offer a discussion of publication bias and its effects on research outcomes.

MEDIATION ANALYSIS IN MARKETING JOURNALS

As most commonly used in consumer behavior and marketing research, a mediation analysis is a regression-based statistical model used to assess whether a proposed mediator (M) could be part of a causal chain between X and Y , such that X affects M which then affects Y . Mediation analysis is appealing because understanding why or how an effect occurs can broaden our theoretical understanding of the effect itself, and provides greater opportunity to capitalize on this mechanism through other antecedent variables. As with other statistical tests, mediation

analyses only measure covariance and do not provide information about causality, which must be established through study design and theory. Moreover, even true positive results may still be subject to other errors such as response style bias (e.g. acquiescence, midpoint responding; Baumgartner and Steenkamp 2001), “meaningless mediation” where M is really a different measure of Y or a manipulation check (Lynch 2011; Zhao, Lynch and Chen 2010), demand effects, confounds, or other problems (MacKinnon et al. 2000).

While a few marketing scholars have ably argued for a more colorful palette of epistemological and methodological approaches to marketing research (Inman et al. 2018; Lynch et al. 2012; Pham 2013; Peracchio, Luce and McGill 2014), evidence of mediation continues to play a prominent role in many leading marketing journals (Pieters 2017). This is true not only of experimental consumer psychology, but also of other adjacent sub-fields of marketing. This powerful push for process evidence has even spilled over into consumer culture theory, where researchers increasingly use qualitative methods to understand consumer processes (e.g. “hegemonic memory making” Brunk, Giesler, and Hartmann 2017; Giesler and Thompson 2016). Indeed, with the few exceptions of modeling and quantitative review papers, one is hard-pressed to find empirical papers that lack mediation analyses. In our own review, we found that of 647 papers in the top four marketing journals in 2018 - 2020, 307 (47%) contained at least one test of mediation.

Currently, bootstrap confidence intervals for the indirect effect, as proposed by Shrout and Bolger (2002), are the preferred method of inference with the most common tool of choice being Hayes’ (2018) PROCESS macro. Bootstrapping poses many advantages over Baron and Kenny’s (1986) approach, in particular focusing on the non-parametric estimation of the sampling distribution of the indirect effect. Bootstrapping methods outperform methods that

assume a normal sampling distribution by having higher statistical power and more appropriate type I error rates (Biesanz, Falk and Savalei 2010; Fritz and MacKinnon 2007; MacKinnon, Lockwood and Williams 2004). A notable downside to bootstrap confidence intervals is that they do not provide a p -value, which could be helpful for evaluating evidential value. As mentioned, however, the recent development of the RP statistic by Götz and colleagues (2021) allows a similar evaluation of mediation analyses.

Reasons to Question Evidence Quality of Mediation Analyses in Marketing Journals

When researchers have the freedom to run tests on a dataset until obtaining publishable results, type I errors are likely to occur (Ioannidis 2005; Simmons et al. 2011). This is because the number of tests they can run is indeterminate and the cumulative type I error rate substantially increases with each subsequent test after the first. For example, deciding whether to control for gender post hoc can inflate the expected type I error rate from 5% to 11.7% (Simmons et al. 2011). As with other statistical tests, researchers carrying out mediation analyses are incentivized to find results that are statistically significant. Peer reviewed journals tend to prefer to publish significant results over nonsignificant results (Mlinarić et al. 2017), presumably because they suggest novel discoveries or effective interventions. Even in cases where researchers have good intentions and subsequent tests can be logically justified, type I errors are likely to occur because researchers are susceptible to bias and are likely to apply motivated reasoning to obtain a result that benefits them (Lilienfeld 2017; Simmons et al. 2021).

There are at least two reasons to expect that mediation analyses in the field of marketing could suffer from evidence quality issues: (1) some amount of undisclosed data exploration

appears to be common in confirmatory mediation analyses and (2) direct replications have been largely unsuccessful (for all studies; not just mediation studies).

First, in marketing research it may be a common practice to adapt the mediation model *post hoc* to the data during analysis. This could be problematic given that mediation analyses are tools designed to test preplanned hypotheses and corrections for multiple testing and post hoc testing are rarely applied. Pham and Oh (2021), for example, implied that mediation models are too complex to specify a priori, and argued that “most consumer research should probably be regarded as exploratory” (p. 168), rather than confirmatory as it is typically reported. Janiszewski and van Osselaer (2021) provided survey data showing that only 25% of consumer research studies are purely deductive, meaning that exploration and post hoc testing play a role in the remaining 75% of studies. Additionally, some methods papers have called for systematic multiple testing in order to identify the “true mediator” (Fiedler, Harris and Schott 2018, p. 96; Fiedler, Schott and Meiser 2011) or to find the true direction of causality (Wiedermann and von Eye 2015). Misrepresentation of data exploration as confirmatory research and formal calls for systematic multiple testing may lead to inflated error rates and a blurring of the lines regarding what is an acceptable way to test a hypothesis.

Second, consider replications for marketing studies more generally. Although mainly only successful conceptual replications of marketing studies have been published (Lynch et al. 2015), it is noteworthy that a series of ten high-powered, preregistered, direct replications of *Journal of Marketing Research* and *Journal of Consumer Research* studies carried out by the research team of Simmons and Nelson (2019) resulted in only one unambiguously successful replication (Nelson 2021). Similarly, only two of ten marketing studies replicated in another mass replication project focused on psychological scarcity (O’Donnell et al. 2021). We also

identified six additional preregistered, direct replications of marketing studies (Harms et al. 2018; O'Donnell and Nelson 2015; Tunca and Yanar 2020; Verschuere et al. 2018; Ziano et al. 2020; Ziano and Tunca forthcoming). All six failed to replicate the original findings, making a total of three out of 26 (11.5%) direct replications of marketing studies that have unambiguously succeeded. While the majority of these replication failures are not mediation analyses, there may be spillover effects, such that we would not be surprised to see similar problems in mediation studies.

It is plausible that mediation studies suffer more from evidence quality issues than other studies for a variety of reasons. First, mediation analyses in marketing are rarely subjected to exact replication either within the article or by other researchers. Second, thousands of model configurations are usually possible because the researcher has tremendous flexibility in how they choose and arrange a set of mediators, moderators and covariates into a testable model. Third, there is no p -value or similar metric reported that can be used to detect issues with evidence quality (Simonsohn et al. 2014) and prevent their publication.

Given the lack of a clear distinction between exploratory and confirmatory research in marketing (Janiszewski and van Osselaer 2021), difficulty in directly replicating marketing findings, and challenges that relate more specifically to mediation analysis, a closer examination of the evidence of mediation in published marketing studies is warranted.

Approach to Assessing the Collective Evidential Value of Mediation Tests

Problems of evidential value can sometimes be detected by meta-analysis of p -values from a related set of studies (e.g. Simonsohn et al. 2014; Bartoš and Schimmack forthcoming).

The logic behind this approach is that true effects most commonly yield very low p -values (less than .01), whereas in the case of publication bias, p -hacking, or similar problems p -values of significant results tend to be higher and closer to the .05 cutoff. One accessible and commonly-used tool to test the collective evidential value of a set of p -values is p -curve (Simonsohn et al. 2014). P -curve compares the distribution of reported statistically-significant p -values to p -value distributions of both null effects and true effects that are low powered (33% power). A uniform or left-skewed p -curve indicates a lack of evidential value. P -curve has been used to identify problems with evidential value of entire literature streams. Examples include elderly priming (Lakens 2017) and power posing (Simmons and Simonsohn 2017).

In line with the recommended inferential approaches for mediation (Hayes 2018), the majority of mediation tests we evaluated from marketing journals were reported with CIs and not p -values. This is because mediation analyses differ from traditional statistical tests in that the CI indicating whether the effect is significant is usually calculated using a resampling method called bootstrapping, which does not yield a p -value (Efron and Tibshirani 1993). The CI can be used to determine statistical significance: For example, if a 95% CI excludes zero the coefficient is deemed statistically significant at a 0.05 level. However, it is not possible to glean whether, for example, a 99% CI would include zero based on the reporting of a 95% CI. This leaves a gap between what could be known with a p -value compared to a bootstrap CI, because readers cannot know whether an effect would be significant at any level other than what is reported. Because p -values have been helpful in detecting bias in submitted and published research, it is unfortunate that this information is not available for mediation analyses. This lack of an indicator of significance level could in itself lead to more problematic studies being published because it is more difficult to quickly judge a study to gauge how likely it is to have been p -hacked.

Fortunately, the reported CI of the indirect effect can be converted to a metric that allows us to do something similar to a *p*-value-based meta-analysis. Götz and colleagues (2021) introduced the “relative proximity” of the CI to zero (RP) based on the idea that if CIs are systematically too close to zero it is indicative of low evidence quality similar to what is signaled by *p*-values that are systematically too close to *p* = .05. This test statistic divides the distance to zero of an indirect effect’s CI by the range of the CI. A positive RP indicates a CI that excludes zero and is thus statistically significant, whereas a negative RP indicates a CI that overlaps with zero and is not statistically significant. As CIs move further from zero and/or become narrower due to greater precision, RPs grow larger, a phenomenon comparable to the increasing *t* or *z* statistics as statistical power improves. Götz and colleagues (2021) used the RP statistic to conduct a quantitative meta-analysis of the field of psychology, during which they discovered the presence of questionable research practices.

$$RP = \frac{\min(|CI_{upper}|, |CI_{lower}|)}{CI_{upper} - CI_{lower}} \text{ RP is negative if CI overlaps 0; positive otherwise.}$$

Even a true effect can result in an RP score that is very close to zero due to random chance. Because the chance of this happening is very low, however, we can aggregate a set of RP scores together to see if there is a pattern of RP scores that are too close to zero, which would indicate a lack of evidential value. To provide context, the RP distribution from reported results must be compared to another distribution, which can be either real or simulated data. In our case, we chose to compare the marketing RP distribution to both simulated data and real data from a related field with more transparent research protocols.

DATA COLLECTION, SIMULATION, AND ANALYSIS

Our first step in comparing RP distributions was to harvest reported CIs from marketing journals and convert them to RP scores. The data were compiled from four highly cited marketing journals and one psychology journal for comparison during a three-year period (2018-20). Next, we simulated mediation tests that were adequately powered (80-95% power), low-powered (33-79% power) and null effects (~5% type I error rate). We then compared the distribution of RP scores from marketing mediation tests to the simulated distributions. Lastly, to compensate for some of the shortcomings of simulations, we also compared RP distributions for each marketing journal to a closely related psychology journal. Greater details about our methodological approach and analyses follow.

Harvesting Test Statistics From Mediation Studies in Published in Marketing Journals

To test the collective evidential value of mediation analyses in marketing journals, we conducted a systematic review of reported test statistics. We chose the following journals for our analysis: *Journal of Marketing* (JM), *Journal of Marketing Research* (JMR), *Journal of Consumer Research* (JCR), and *Journal of Consumer Psychology* (JCP). We selected these journals because they have high impact factors and represent typical outlets for mediation studies in marketing. We analyzed three years of each of the preceding journals, which covered all articles which appeared in print between 2018 and 2020 for each journal. Our initial sample contained 647 publications. Using keywords such as *mediation*, *indirect*, *Hayes*, and *PROCESS*, a laboratory assistant and the third author manually reviewed each article in order to identify those containing at least one mediation model with accompanying CI. This initial screening of publications resulted in a sample of 308 publications. During the course of our analysis, we

excluded 13 publications which contained multilevel models (6), were repeated-measures designs (2), were retracted (2), were non-linear mediation models (1), or were in retrospect found not to contain any mediation model (2). An additional ten papers were unusable because they lacked a codable CI. This process resulted in a final dataset of 285 mediation analyses.

The articles that included mediation were then given to a team of six laboratory assistants who were trained to manually code mediation analyses, but naive to the hypothesis. We chose to apply manual harvesting of test statistics rather than mechanical harvesting because research articles tend to include many test statistics that are not central to the core hypotheses and therefore less informative regarding problems with evidence quality (Simonsohn 2015). Working independently, pairs of assistants were instructed to record the first reported hypothesised test of an indirect effect mentioned in the publication subject to two exceptions. If the first reported test was to rule out alternative explanations, the assistants were told to ignore these tests in favor of tests which were hypothesized. In the case of moderated mediation, the assistants were told to record the index of moderated mediation, if it was reported, even if it was not reported before any other indirect effects.

The assistants recorded the statement in the publication which contained the first reported test of an indirect effect, as well as the formal hypothesis on which the test was based. The assistants recorded point estimates of the indirect effect or the index of moderated mediation, as well as its standard error and the CIs around the point estimate. They also recorded details about the component paths in the mediation model, such as point estimates, effect sizes (group means and standard deviations), or inferential statistics (F -test or t -test), if they were reported. Finally, they recorded the confidence level of the analysis, e.g. 95%, the model number of the PROCESS macro used, and what kind of bootstrapping was performed, e.g. bias-corrected, and for how

many bootstrap samples. The laboratory assistants then met with the second and third authors of the paper to resolve any discrepancies. These were resolved through discussion such that there was unanimous agreement in order for the above data to be recorded in the final data set.

Next, we calculated the relative proximity score (RP; Götz et al. 2021) for each mediation analysis of interest by dividing the reported CI's distance to zero by its width. RP's are negative if the CI overlaps zero and positive otherwise. This was done so that the distribution of RP scores could later be compared with simulated distributions to see how well the reported results match adequately powered, low powered, and null results.

Study 1: Comparing Marketing Journals to Simulated Data

The purpose of Study 1 was to test how RP scores calculated from mediation tests reported in marketing journals compared to data simulated under a variety of circumstances. To this end, we simulated three different RP distributions (low power, adequate power, and null effects). The simulated distributions were separately plotted against the marketing journal distribution so that we could visually examine the similarity in form among the four distributions. In this way, we hoped to learn whether the reported data in aggregate most closely matched adequately-powered tests, low-powered tests, null effects, or a mix of these categories.

Data simulations. We ran a series of simulations to determine how the RP scores ought to be distributed for simple mediation analyses when adequately powered (80-95% power), low-powered (33-79% power) and with no power (null effects; 5% false positive rate). We used the 'simstudy::genCorData' R function (Goldfield and Wjciak-Jens 2019) to create each simulated dataset with pre-specified population-level correlations among the three variables (X ,

the independent variable; M , the mediator, and Y , the outcome variable). To conduct each test of the indirect effect, we used the `mediate` function from the ‘psych’ package for R (Revelle 2015) with default options ($\alpha = .05$; 5,000 bootstraps; percentile CIs). We verified that this tool gives results consistent with the widely used PROCESS macro (Hayes 2018). Just like with the test statistics harvested from marketing journals, RP scores were calculated by dividing the simulated CI’s distance to zero by its width. RP’s are negative if the CI overlaps zero and positive otherwise (Götz et al. 2021).

To simulate true effects, we ran simple mediation analyses on a series of simulated 3-variable datasets that each had a weak population-level indirect effect ($\beta=.04$) and a residual direct effect of zero (full model specifications: $r_{XM} = .20$, $r_{MY} = .20$, $r_{XY} = .04$; all means = 0, all SDs = 1). We settled on $\beta=.04$ because it allowed us to achieve desired power levels using sample sizes that are similar to those reported in marketing journals. This implies that reported indirect effects are also close to $\beta=.04$ on average, assuming they are mostly true effects with power in the range of 33% to 95%. We manipulated power by varying the sample size of each simulated data set while holding all other parameters constant. Sample sizes were drawn at random, uniformly, from a range based on Monte Carlo estimates (Schoemann 2017) for adequate power (80-95%; $N=238$ to 356) and for low power (33-79%; $N=124$ to 234). Sample size estimates used Schoemann's (2017) tool defaults (all standard deviations = 1.00, random seed = 1234, # of replications = 1,000, Monte Carlo draws per rep = 20,000). In cases where a range of sample sizes were estimated to arrive at the desired power level, the midpoint of the range was selected.

For null effects, we ran simple mediation analyses on a series of simulated 3-variable data sets in which all population correlations among variables was set to zero except the b -path (M - Y),

which was assigned to have weak ($r_{MY} = .1$), moderate ($r_{MY} = .3$), strong ($r_{MY} = .5$; Cohen 1988), or very strong correlations ($r_{MY} = .7$). We allowed the b -path to be significant because (1) a fully null model would have a type I error rate many times lower than the 5% target error rate, and (2) this model represents a common scenario in which an experimentally manipulated X variable has no true effect on downstream measured variables (M and Y). In this scenario, M and Y would be correlated (significant b -path) because they are obviously related measures, or simply due to response-style bias. Just like in the true effects simulations, all means = 0 and all SDs = 1 in the null effects models. Simulated sample size was varied based on 20th, 40th, 60th, and 80th percentile sample sizes reported in the mediation analyses we coded ($N = 138, 195, 250$, and 335 respectively). These 16 different combinations of M - Y correlation and sample size were used to simulate the heterogeneity in a real life null distribution and to assess whether there are meaningful differences among null distributions. All were combined for our key analysis when no meaningful differences emerged. The occurrence of type I error is fairly similar across sample sizes and correlations we tested except for the weak b -path effect size ($r = .1$) simulations, which showed very low type I error rates (see Figure 1 below).

Figure 1. Distributions of false positive RP scores in simulations of null indirect effects given various b -path effect sizes ($r = .10, .30, .50, .70$) and sample sizes ($N = 138, 195, 250, 335$). Type I error rates for each distribution are in bold.

Notes: All paths in the simple mediation model are set to zero in the population except the b -path. We modeled 95% CIs which have an implied type I error rate of .05.

We programmed all three sets of simulations to run until the highest positive point of each of their respective RP histograms was 1,000 units for a bin width of .05. This resulted in 7,252 tests of adequate power, 6,769 tests of lower power and 54,574 tests of null effects.

Constraining distributions by histogram height, rather than running a fixed number of

simulations was done out of convenience as it (1) eliminated any need to accurately predict the results of the simulations a priori, (2) afforded us the ability to downscale the simulations to match the height of the marketing journals distribution more cleanly and easily, and (3) helped us economize our server usage. In our final simulated distributions, mediation analyses reached statistical significance in 85.0% of tests (adequate power), 55.4% of tests (low power), and 4.4% of tests (null effects).

Comparative analysis of simulated and observed RP score distributions. To make a fair comparison of simulated data to the RP distribution from marketing journals, we scaled all simulated distributions down so that their highest positive peak would be 76 with a bin width of .05, equal to the peak of the marketing RP distribution with the same bin width. Next, each of the three distributions of simulated RP scores was paired with the distribution of RP scores from the marketing journals (see Figure 2). Nonsignificant results of both simulated and actual data were included because the absence of nonsignificant results in the reported data is itself somewhat diagnostic of multiple testing, selective reporting, and/or publication bias. This is consistent with Bartoš and Schimmack (forthcoming) and Götz and colleagues (2021) but deviates from Simonsohn and colleagues (2014). One reason for this deviation from *P*-curve (Simonsohn et al. 2014) is that our data sets are very large, much larger than a typical *p*-curve, giving us the ability to show more aspects of the distribution than would be practical with a 10- or 20-study *p*-curve.

Figure 2. A comparison of simulated relative proximity (RP) scores with observed RP scores calculated from mediation analyses reported in marketing journals (2018-2020)

Note: The portion of the null effects histogram higher than 80 mediation tests (between -0.5 and 0.0 RP) is not visible in this chart.

Visual comparison of reported and simulated results in Figure 2 show that in aggregate, the distribution of RP scores for mediation analyses reported in marketing journals is inconsistent with adequately-powered studies. Indeed, the RP scores from observed tests most closely resemble the simulated null distribution. The right tail of the observed tests is longer, however, indicating some statistical power, but less than the desired 80%. This is consistent with

the finding that too many CIs from mediation analyses just barely exclude zero in psychology journals (Götz et al. 2021). It is also clear there is some level of selection bias, whereby results with CIs including zero are not published at the same rate as would be expected under any of these distributions. Although we estimate the typical mediation analysis has poor statistical power, nearly all observed results were statistically significant. Indeed, the observed distribution of RP scores is highly suggestive of some form of selection bias: Even with good (80-95%) statistical power, some portion of tests should result in type II errors (5-20%), yielding negative RP scores. We observe that CIs which include zero are exceedingly rare in mediation analyses in marketing journals.

Study 2: Comparing Marketing Journals to Real-World Data from a Peer Journal

After comparing marketing journal RP distributions to simulated data, we next sought to compare the distributions to real-world data in a closely related discipline. For this comparison, we chose the *Journal of Personality and Social Psychology (JPSP)*, a premier outlet in the field of social psychology. If there are meaningful differences in RP scores between marketing and social psychology that could indicate a difference in false discovery rate. Such knowledge could potentially benefit the weaker of the two fields as it seeks to mimic research practices of the stronger field. We were not able to predict a priori how the marketing journals would compare against *JPSP*. On one hand, *JPSP* fared poorly in the first Open Science Collaboration's mass replication effort (2015). On the other hand, Götz and colleagues (2021) found the RP distribution of *JPSP* to be fairly consistent with other areas of psychology. Also, following low replication scores in 2015, it has become common for studies published in *JPSP* to be

preregistered. Consistent with Simmons and colleagues (2021) we found that preregistration in marketing journals is more rare than it is in *JPSP*.

Data collection method. Consistent with the coding of the marketing journals, we chose to code all *JPSP* articles from the years 2018 to 2020. The first author manually reviewed all 306 articles from that period in order to identify papers which included mediation analysis. Each article was mechanically searched for the terms, “mediat,” “hayes,” and “indirect effect.” If articles yielded zero results for these searches, the article was not coded. If the search came up positive, the first author then searched the paper to identify whether there was at least one mediation analysis with CIs. All articles with mediation analyses that included CIs were sent to student coders. Two student researchers independently, redundantly coded all of the mediation articles. Reconciliation of results between the two coders was conducted by the first and third author jointly. Consistent with the marketing journals coding project, students coded the first mediation analysis except in rare cases where the first test was not a hypothesis test, but was used to rule out confounds. In that case, students coded the hypothesis test. Student coders identified the following information, adapted from a *p*-curve disclosure table (Simonsohn et al. 2014): Hypothesis tested, statistical test (copy/pasted), study number, PROCESS model number, lower and upper CI, point estimate, standard error, %CI (95% in nearly every case), link to preregistration, sample size, and whether it was a multilevel or repeated measures model. Consistent with how we selected the marketing mediation analyses, all multilevel and repeated measures models (8 articles) were excluded from our analysis, leaving us with a total sample size of 123 *JPSP* articles ($123/306 = 40.2\%$ of all *JPSP* articles with codable mediation analyses).

In order to have as much consistency as possible between the two coding projects ((1) marketing journals and (2) *JPSP*), the third author worked on both. Two notable differences in

data collection are worth pointing out: (1) in the case of the marketing journals, we collected information we would need to calculate effect sizes and information about component paths but we did not do that for *JPSP*, and (2) we coded all mediation studies for marketing, regardless of whether they had CIs we could use for our meta-analysis but for *JPSP*, we only coded studies that had mediation tests with CIs. We found 12 articles in *JPSP* with mediation tests but no codable CIs.

Analysis. Once all the *JPSP* data were collected, we plotted a histogram of all RP scores from *JPSP* along with histograms for each of the four marketing journals that we coded (see Figure 3). In order to have some objective metrics to compare, we also calculated and plotted the median positive RP score for each journal. Analysis of RP scores from *JPSP* relative to RP scores from the marketing journals indicates a stark difference. Not only is the median positive RP score from *JPSP* (0.21) close to double that of the marketing journals (0.10, 0.11, 0.11, 0.13) but the shape of the *JPSP* distribution is also more consistent with the bell curve that we observed in simulations of adequately-powered true effects. Even in *JPSP*, however, there is evidence of a problem in the rather large spike between zero and .05. Although sample sizes of mediation studies in marketing journals were not significantly different from those found in *JPSP*, $t(288.42) = -1.021, p = .308$, we observed that a few large outliers in the marketing journals had masked some differences. Rather than removing outliers and reanalyzing, we simply note the differences in 25th, 50th and 75th percentiles for marketing journals: 153, 206.5, 304 and for *JPSP*: 151, 257, 382.5. The median *JPSP* sample size is 24.45% larger than that of the marketing journals.

Figure 3. Distribution of RP scores derived from mediation tests in marketing journals compared to those from a social psychology journal (*Journal of Personality and Social Psychology*) (2018-20)

Note: Black vertical lines indicate the median positive RP scores for each journal.

Our comparison of RP distributions indicates that mediation analyses in marketing collectively offer substantially less evidential value than mediation analyses in social psychology. Thus, it would seem wise to make a closer examination of the practices of social psychology and to attempt to mimic those that are likely to lead to lower type I error rates. For example, broader adoption of study preregistration could assist in addressing this problem. We caution, however, that the previously mentioned spike in the RP distribution between $RP = 0$ and $RP = 0.05$ from *JPSP* indicates a problem of evidence quality. It is unlikely that the majority of

the studies that fall in that range are adequately powered true results (though a few may be). This is consistent with the claim that RP scores in psychology, including social psychology, show evidence of bias (Götz et al. 2021). Our conclusion is that both *JPSP* and marketing journals could benefit from a more thorough embrace of practices that reduce false discovery. Only 14 out of 131 (10.7%) of mediation tests we coded from *JPSP* (2018-20) were preregistered. No studies in marketing journals included preregistration for mediation analyses.

DEALING WITH ISSUES OF EVIDENCE QUALITY

Disentangling the Causes of Evidence Quality Issues

Unfortunately, the data we have collected do not reveal *why* published statistics are inconsistent with adequately-powered studies of true effects. The problem of too many barely significant results has been explained in the past by underpowered studies (Cohen 1988), the file drawer problem (Rosenthal 1979), and *p*-hacking (Nelson et al. 2018). Our data appear to be consistent with any combination of these explanations.

One possibility is that researchers are studying true effects with insufficient statistical power, using samples too small to reliably detect those effects. Because the true effect size for any given study is always unknown, we cannot point to an individual study to say it is or is not underpowered. Similarly, when examining tests of mediation, we cannot know what a "typical" effect size is, and so we cannot say whether the body of literature is underpowered to detect that effect, but we can compare it to benchmarks from psychology. Götz et al. (2021) observed a median standardized indirect effect size of 0.08, which would require a sample size of about 130

to detect with 80% power. However, this is likely an overestimate due to publication bias (Lane and Dunlap 1978). Fritz and MacKinnon (2007) found that 75% of psychology studies did not have sample sizes sufficient to detect an indirect effect which involved one small effect (either *a*-path or *b*-path) at 80% power ($N = 385$). Using the same metric, 80.5% of marketing studies examined in this study are underpowered to detect an indirect effect with at least one small effect. In addition to being worse at detecting true effects, having underpowered studies increases the proportion of published results which are false positives (Ioannidis et al. 2011).

If studies are routinely underpowered, where are the nonsignificant results? Are they all languishing, unpublished, in a file drawer somewhere? It seems unlikely that researchers would consistently rely on a process in which they have multiple nonsignificant, unpublished results for every significant, published result. Such a process would be wasteful and frustrating; researchers would likely learn to power their studies appropriately (Nelson et al. 2018). Instead, it is possible that low-powered studies are nevertheless able to reach statistical significance through *p*-hacking. If the first test is not successful, researchers might run a second test, and likely more tests, trying to obtain significance. In the end, a significant *p*-value or 95% CI might be obtained, but with a substantially increased type I error rate. *P*-hacking is a plausible explanation for at least a portion of the problematic pattern we have observed in mediation analyses in marketing journals.

Improving the Evidence Quality of Mediation Analyses in Marketing

These different potential causes of poor evidence quality (low power, file drawer problem and *p*-hacking) can be handled in very similar ways and reforms that target one of these

problems can have secondary effects to help lessen the occurrence of the other problems. We propose three solutions to improve the evidence quality of confirmatory studies: (1) preregistration, (2) reducing errors from multiple testing, and (3) designing adequately-powered studies. We also offer a discussion of how publication bias can be mitigated.

Preregistration. Exploratory analyses are often mislabeled as confirmatory.

Preregistration creates documentary evidence that confirmatory analyses were indeed planned in advance of seeing the data. When studies are not preregistered, researchers may be tempted to make decisions based on the data, which can lead to *p*-hacking, even if unintentional. Many adjacent fields, particularly within psychology, have adopted the practice of preregistration to increase evidence quality and decrease *p*-hacking. Simmons and colleagues (2021) argued that preregistration promotes accurate reporting of confirmatory tests while still allowing for data exploration, as long as exploration is noted as such. Given more transparent reporting, astute readers would be able to make better inferences about the trustworthiness and replicability of research both prior to and after publication.

Corrections for multiple testing. Even when a study is preregistered and the preplanned analysis is adhered to, multiple testing can still be a source of error. This is because the cumulative error, or the chance of obtaining at least one false positive result increases with each subsequent test after the first. This phenomenon is known as the multiple comparisons, or multiple testing problem (Tukey 1991). Although corrections for multiple testing are not currently common in mediation analysis, corrections such as the Bonferroni adjustment (Bonferroni 1936) have been applied to other categories of statistical testing with great success in reducing false discovery rates. Tools such as PROCESS (Hayes 2018) could easily accommodate such adjustments by changing the level of the confidence intervals. For example, a researcher

planning to conduct 5 tests could use an $\alpha = 0.05/5 = 0.01$, and thus request a 99% CI from PROCESS or other similar tools. For further information about cumulative errors and adjustments, see Lakens (2020).

Ensuring Proper Statistical Power in Mediation Analyses. As we explained previously, low statistical power is one possible explanation for the RP distribution we observed in mediation analyses published in marketing journals. This implies that either the sample size or the effect size or both are too small to reliably detect true effects when they exist. Although sample sizes in marketing have increased in recent years with the availability of online samples, average effect sizes have actually decreased, leaving average power largely unchanged (Krefeld-Schwalb and Scheibehenne 2021). Mediation studies require larger sample sizes than do direct effects, and those that are more complex or that model subtle relationships require very large sample sizes. For example, a simple mediation that involves one small effect (a weak relationship between mediator and either predictor or outcome variable) requires a sample size of $N = 385$ to have 80% power (Fritz and MacKinnon 2007). Available tools can help determine the sample size necessary to obtain desired statistical power given information about what the strength of relationship among the various variables may be (Schoemann et al. 2017). Either increasing sample size, studying larger effect sizes, or both can increase statistical power.

Another area where power can be improved is through improved measurement and measurement modeling. Selecting measurement instruments with better reliability can improve statistical power to detect effects (Williams and Zimmerman 1989; Leon, Marzuk and Portera 1995). This can often mean using longer scales with more informative items, or selecting scales that have been vetted in previous research rather than designing your own scale for the specific study (Flake, Pek, Hehman 2017). Latent variable modeling can be used to account for

measurement error and improve power to detect indirect effects (Sim, Kim, Suh 2021). In many of the studies included in this meta-analysis either the outcome variable or the mediator could have easily been incorporated as a latent variable, and perhaps power would be improved.

Overcoming Publication Bias. Another potential explanation for the unfavorable RP distribution we've observed is publication bias. Publication bias is a problem on a number of levels in that it biases effect size estimates in meta-analysis (Frieze, Frankenbach 2020), reduces the number of true negatives in the published literature (Bakker, van Dijk, Wicherts 2012), and incentivizes researchers to "find" significant results because they cannot publish otherwise. Solutions to publication bias are not specific to mediation, but can have unique implications for mediation analysis specifically in marketing. Previous research on publication bias has focused on how the results of primary hypothesis tests can influence likelihood of publication, but the current research shows that even mediation analyses (which are typically secondary) also show strong evidence of publication bias.

Ultimately, to reduce publication bias reviewers and editors need to be open to publishing meaningful and interesting null results. This can be challenging, as null results may seem uninformative or uninteresting. This may be caused, in part, by the typical practice of null hypothesis significance testing, which does not allow one to claim support for the null hypothesis, only a failure to reject the null hypothesis. However, statistical methods such as equivalence testing help quantify the evidentiary value of failures to reject the null hypothesis (Lakens et al. 2014). Preregistration and registered reports can help to build a literature where the analysis plan is laid out before data collection so that *p*-hacking is less prevalent. Registered reports build on what's possible with preregistration by implementing peer review after the study is planned, but before data is collected. After peer review the manuscript is given an "in-principle

acceptance" which means that it will be published regardless of the results, so long as the research team is able to stick to their plan (Chambers 2013). This method prioritizes well-designed studies rather than significant results, since the decision to publish is made before data is collected. Evidence suggests that about 66% of registered reports are null results as compared to 4% in traditional publications (Scheel, Schijen and Lakens 2020). Registered reports have yet to be adopted by marketing journals (Montoya, Krenzer and Fossum 2021), and so one way to reduce publication bias in this field would be implementing this new publication mechanism at top journals in this field.

GENERAL DISCUSSION

Summary and conclusions

In aggregate, confidence intervals (CIs) reported for mediation analyses in marketing journals do not match a pattern of adequately-powered results in data simulations nor do they match up well with real-world data from an adjacent field with more transparent research practices. The bulk of reported CIs are too close to zero relative to their widths (low RP scores). We found that they more closely match both simulations of low power and simulations of null results. This is problematic because mediation analyses play a central role in helping us to understand the process by which predictor variables influence outcomes of interest. We were unable to determine whether the suboptimal CIs reported were caused by low statistical power, type I error, or some combination of the two. We note that both low power and false discovery are big problems that can be addressed in similar ways.

Our process of discovery involved comparing results reported in top marketing journals (*Journal of Marketing*, *Journal of Marketing Research*, *Journal of Consumer Research*, *Journal of Consumer Psychology*; 2018-2020) to both simulated data and reported data from *Journal of Personality and Social Psychology*. Our metric of interest, the RP score (Götz et al. 2021) was calculated by dividing the distance of the CI from the first mediation analysis in each article by its width. Simulations of adequate statistical power produced a distribution of RP scores that was substantially different from the distribution of marketing journal RP scores. With exception of a few outliers, marketing results more closely matched both low power and null effects. RP scores from marketing journals were also compared to those from *JPSP*. The median RP score in *JPSP* was about twice as high as the median RP score in the marketing articles. The comparison to *JPSP* was made for two reasons: (1) comparing the data of interest to real as well as simulated data is more thorough analytically, and (2) we were aware that *JPSP* has implemented many open science reforms that have not yet become the norm in marketing, so we wanted to get some idea of their effectiveness. Now knowing that *JPSP* outperformed the marketing journals in evidence quality, we propose adoption of similar practices and policies such as preregistration, adopting registered reports, and open data and materials. One caution, however, is that the troubling spike in RP scores just above zero and low number of preregistered studies in the *JPSP* distribution indicate that they also have room for improvement.

We've suggested study preregistration, multiple testing adjustments, and multiple methods of increasing power as solutions. We've also argued for changes that could decrease the influence of publication bias on our literature. Another issue to consider is how to handle non-confirmatory research. Conducting exploratory analyses and reporting them as confirmatory could be part of the problem we're seeing here. If true, that would indicate that there is a need or

at least an appetite for non-confirmatory, perhaps non-reductionist, approaches. Janiszewski and van Osselaer (2021) made a recent contribution in this area by showing how a combination of exploratory and confirmatory analyses could be reported. Additional, more flexible approaches to knowledge creation are needed.

Research implications

The problems of evidence quality that we have uncovered in published marketing studies should encourage authors, readers, reviewers and editors to exert caution in the future when evaluating process evidence. Moreover, the possible causes we have identified—low power, publication bias, and *p*-hacking (even if unintentional)—are not unique to mediation analyses. The same evidence quality issues could be present in other statistical tests, even if not detectable through meta-analysis. This research should help bring more emphasis to study preregistration, corrections for multiple testing, and methods of increasing statistical power to improve the reliability of future studies.

Our analysis also helps to explain the low direct replicability of marketing studies we described earlier. It could be argued that studies are not replicating due to differences in context between the original study and the replication (e.g. perhaps people respond differently to the language of an out-of-date individual differences scale) or due to problems of evidence quality in the original study. The differences in context argument implies that replicability cannot be improved and it nullifies claims of generalizability made in the original studies. The evidence quality argument gives hope that replicability (and therefore generalizability) is achievable by increasing the evidence quality of published studies. Simmons and Nelson's (2019) failure to

successfully replicate more than one out of 10 newly published marketing studies by matching the original experiments as exactly as possible weakens the differences in context argument. Our analysis further weakens this argument by directly showing substantial problems in collective evidence quality in published marketing studies. Once evidence quality is improved, we should expect to see higher direct replicability.

Study limitations

Simulation studies necessarily lack much of the complexity and diversity of a set of real-world studies. This is because (1) it is impossible to identify and quantify every force that influences data from actual studies, and (2) simulations that are too complex are uninterpretable. For example, while the empirical results from the marketing journals are a mix of simple mediation, multiple mediation, and moderated mediation, we only generated results from simple mediation models for the simulation. We selected parameters that we hoped would most closely mimic what is reported in marketing journals. It is likely that there are scenarios that we have overlooked. We don't believe we have overlooked any scenario that would diminish the findings we have reported. We offered a comparison to real-world data from *JPSP* to compensate for the limitations of simulations, but comparisons to real-world data also have limitations. For example, we assumed that the *JPSP* evidence quality looked better than that of the marketing journals because of open science practices such as preregistration, but other differences between fields such as studying obviously related constructs, non-experimental design, and data fabrication could also increase RP scores without achieving desired increases in evidence quality or usefulness. We judged social psychology to be one of the closest sub-disciplines to marketing but

in actuality, many *JPSP* studies are quite different from marketing studies, making a fully apples-to-apples comparison difficult.

It is important to note that even RP distributions that perfectly match or exceed what we would expect from highly powered true results can mask severe defects. The RP distribution only gives an indication of whether there is a substantial, obvious presence of false positive results or underpowered studies. The RP distribution gives no information about issues of causal order, confounds, lack of discriminant validity among modeled variables, and other issues that frequently arise and must be dealt with in inferential statistical tests such as mediation analyses.

REFERENCES

- Bakker, Marjan, Annette van Dijk and Jelte M. Wicherts (2012). "The Rules of the Game Called Psychological Science," *The Rules of the Game Called Psychological Science*, 7 (6), 543–54.
- Baron, Reuben M. and David A. Kenny (1986). "The Moderator–Mediator Variable Distinction in Social Psychological Research: Conceptual, Strategic, and Statistical Considerations," *Journal of Personality and Social Psychology*, 51 (6), 1173.
- Bartoš, František and Ulrich Schimmack (forthcoming). "Z-Curve 2.0," *Meta-Psychology*, forthcoming.
- Baumgartner, Hans and Jan-Benedict E. M. Steenkamp, (2001). "Response Styles in Marketing Research: A Cross-National Investigation," *Journal of Marketing Research*, 38 (2), 143–56.
- Biesanz, Jeremy C., Carl F. Falk and Victoria Savalei (2010). "Assessing mediational models: Testing and interval estimation for indirect effects." *Multivariate Behavioral Research*, 45, 661–701.
- Bonferroni, Carlo (1936). "Teoria Statistica delle Classi e Calcolo delle Probabilita," *Pubblicazioni del R Istituto Superiore di Scienze Economiche e Commerciali di Firenze*, 8, 3–62.
- Brunk, Katja H., Markus Giesler and Benjamin J. Hartmann (2017). "Creating a Consumable Past: How Memory Making Shapes Marketization," *Journal of Consumer Research*, 44 (6), 1325–42.
- Chambers, Christopher D. (2013). "Registered reports: A new publishing initiative at cortex." *Cortex*, 49, 609–10.
- Cohen, Jacob (1988), *Statistical Power Analysis For The Behavioral Sciences*, 2nd ed. Hillsdale,

NJ: Erlbaum.

Efron, Bradley and Robert J. Tibshirani (1993). “An introduction to the bootstrap.” Boca Raton, FL: Chapman & Hall.

Fiedler, Klaus, Chris Harris and Malte Schott (2018). “Unwarranted inferences from statistical mediation analyses—An analysis of articles published in 2015,” *Journal of Experimental Social Psychology*, 75, 95–102.

Fiedler, Klaus, Malte Schott and Thorsten Meiser (2011). “What Mediation Analysis Can (Not) Do,” *Journal of Experimental Social Psychology*, 47 (6), 1231–6.

Flake, Jessica K., Jolynn Pek and Eric Hehman (2017). “Construct Validation in Social and Personality Research: Current Practice and Recommendations,” *Social Psychological and Personality Science*, 8 (4), 370–8.

Friese, Malte and Julius Frankenbach (2020). “P-hacking and Publication Bias Interact to Distort Meta-analytic Effect Size Estimates,” *Psychological Methods*, 25 (4), 456–71.

Fritz, Matthew S. and David P. MacKinnon (2007). “Required sample size to detect the mediated effect,” *Psychological Science*, 18(3), 233–9.

Giesler, Markus and Craig J. Thompson (2016). “Process Theorization in Cultural Consumer Research,” *Journal of Consumer Research*, 43 (4), 497–508.

Götz, Martin, Ernest H. O’Boyle, Erink Gonzalez-Mulé, George C. Banks and Stella S.

Bollmann (2021). “The “Goldilocks” Zone: (Too) Many Confidence Intervals In Tests Of Mediation Just Exclude Zero,” *Psychological Bulletin*, 147 (1), 95–114.

Goldfield, Keith and Jacob Wujciak-Jens (2019). simstudy 0.1.14.

<https://CRAN.R-project.org/package=simstudy>

Harms, Christopher, Hanna A. Genau, Carolin Meschede and André Beauducel (2018), “Does It

Actually Feel Right? A Replication Attempt of the Rounded Price Effect,” *Royal Society Open Science*, 5 (4), 1–13.

Hayes, Andrew. F. (2018). *Introduction to mediation, moderation, and conditional process analysis* (2nd ed.). New York, NY: Guilford Press.

Ioannidis, John P. A., Robert Tarone and Joseph K. McLaughlin (2011). “The False-positive to False-negative Ratio in Epidemiologic Studies,” *Epidemiology*, 22 (4), 450–6.

Inman, J. Jeffery, Margaret C. Campbell, Amna Kirmani and Linda L. Price (2018). “Our Vision for the Journal of Consumer Research: It’s All about the Consumer,” *Journal of Consumer Research*, 44 (5), 955–9.

Janiszewski, Chris and Stijn M. J. van Osselaer (2021). “The Benefits of Candidly Reporting Consumer Research,” *Journal of Consumer Psychology*, 31 (4), 633–46.

Krefeld-Schwalb, Antonia and Benjamin Scheibehenne (2020) “Tighter Nets for Smaller Fishes: Mapping the Development of Statistical Practices in Consumer Research Between 2011 and 2018,” PsyArXiv. June 4. doi:10.31234/osf.io/xjgyk.

Lakens, Daniel (2017), “Professors Are Not Elderly: Evaluating the Evidential Value of Two Social Priming Effects Through P-curve Analyses,” PsyArXiv. January 15.
doi:10.31234/osf.io/3m5y9.

Lakens, Daniel (2020), What’s a family in family-wise error control? [blog post], retrieved from <http://daniellakens.blogspot.com/2020/03/whats-family-in-family-wise-error.html>.

Lakens, Daniël, Anne M. Scheel, Anne M. and Peder M. Isager (2014). “Equivalence Testing for Psychological Research: A Tutorial,” *Advances in Methods and Practices in Psychological Science*, 1 (2), 259–69.

Lane, David M. and William P. Dunlap (1978). “Estimating effect size: Bias resulting from the

- significance criterion in editorial decisions.” *British Journal of Mathematical and Statistical Psychology*, 31 (2), 107–12.
- Leon, Andrew C., Peter M. Marzuk and Laura Portera (1995). “More Reliable Outcome Measures Can Reduce Sample Size Requirements” *Arch Gen Psychiatry* 52 (10), 867–71.
- Lilienfeld, Scott O. (2017). “Psychology’s Replication Crisis and the Grant Culture: Righting The Ship,” *Perspectives on Psychological Science*, 12 (4), 660–4.
- Lynch, John G., Jr. (2011). “Meaningless Mediation,” *Society for Judgment and Decision Making Newsletter*, 30 (3), <http://sjdm.org/newsletters/11-oct.pdf>.
- Lynch Jr., John G., Joseph W. Alba, Aradhna Krishna, Vicki G. Morwitz and Zeynep Gürhan-Canli (2012). “Knowledge Creation In Consumer Research: Multiple Routes, Multiple Criteria,” *Journal of Consumer Psychology*, 22 (4), 473–85.
- Lynch Jr, John G., Eric T. Bradlow, Joel C. Huber and Donald R. Lehmann (2015). “Reflections on the Replication Corner: In Praise of Conceptual Replications,” *International Journal of Research in Marketing*, 32 (4), 333–42.
- MacKinnon, David P., Jennifer L. Krull and Chondra M. Lockwood (2000). “Equivalence of the Mediation, Confounding and Suppression Effect,” *Prevention Science*, 1 (4), 173–81.
- MacKinnon, David P., Chondra M. Lockwood and Jason Williams (2004). “Confidence limits for the indirect effect: Distribution of the product and resampling methods,” *Multivariate Behavioral Research*, 39, 99–128.
- Mlinarić, Ana, Martina Horvat and Vesna Šupak Smolčić (2017). “Dealing With the Positive Publication Bias: Why You Should Really Publish Your Negative Results,” *Biochemia Medica*, 27 (3), 447–52.
- Montoya, Amanda K., William L. D. Krenser and Jessica L. Fossum (2021) “Opening the Door

- to Registered Reports: Census of Journals Publishing Registered Reports (2013 - 2020),” *Collabra* 7 (1) , 24404.
- Nelson, Leif (2021), “The Value of Replications in the Behavioral and Consumer Sciences [video file],” *YouTube*, <https://www.youtube.com/watch?v=DflbIVB2MH0>
- Nelson, Leif D., Joseph Simmons and Uri Simonsohn (2018). “Psychology’s Renaissance,” *Annual Review of Psychology*, 69 (1), 511–34.
- O’Donnell, Michael, Amelia S. Dev, Stephen Antonoplis, Stephen M. Baum, Arianna H. Benedetti, N. Derek Brown, Belinda Carrillo et al. (2021), “Empirical Audit and Review and an Assessment of Evidentiary Value in Research on the Psychological Consequences of Scarcity,” *Proceedings of the National Academy of Sciences*, 118 (44), 1–3.
- O’Donnell, Michael and Leif D. Nelson (2015), “Wadhwa & Zhang 2015, Study 3 Replication,” See <https://osf.io/am32k/>.
- Open Science Collaboration (2015). “Estimating The Reproducibility Of Psychological Science,” *Science*, 349 (6251).
- Peracchio, Laura A., Mary F. Luce and Ann L. McGill (2014). “Building Bridges For An Interconnected Field Of Consumer Research,” *Journal of Consumer Research*, 40 (6), v-viii.
- Pham, Michel Tuan (2013). “The Seven Sins Of Consumer Psychology,” *Journal of Consumer Psychology*, 23 (4), 411–23.
- Pham, Michel Tuan, and Travis Tae Oh (2021). “Preregistration Is Neither Sufficient Nor Necessary For Good Science,” *Journal of Consumer Psychology*, 31 (1), 163–76.
- Pieters, Rik (2017). “Meaningful Mediation Analysis: Plausible Causal Inference And Informative Communication,” *Journal of Consumer Research*, 44 (3), 692–716.
- Revelle, William (2015). "Package ‘psych’." The comprehensive R archive network 337 (2015):

338.

- Rosenthal, R. (1979). "The File Drawer Problem And Tolerance For Null Results," *Psychological Bulletin*, 86 (3), 638–41.
- Scheel, Anne M., Mitchell R. M. J. Schijen and Daniël Lakens, (2020). "An excess of positive results: Comparing the standard psychology literature with registered reports." *Advances in Methods and Practices in Psychological Science* 4 (2), 1–12.
- Schoemann, Alexander M., Aaron J. Boulton and Stephen D. Short (2017). "Determining power and sample size for simple and complex mediation models." *Social Psychological and Personality Science*, 8 (4), 379–86.
- Shrout, Patrick E. and Niall Bolger (2002). "Mediation in experimental and nonexperimental studies: New Procedures and recommendations." *Psychological Methods*, 7, 422–45.
- Sim, Mikyung, Su-Young Kim and Younguk Suh (2021). "Sample Size Requirements for Simple and Complex Mediation Models," *Educational and Psychological Measurement*. 1–31.
- Simonsohn, Uri (2015), Falsely Reassuring: Analyses of ALL p-values [blog post], retrieved from <http://datacolada.org/41>.
- Simmons, Joseph P. and Leif D. Nelson (2019), Data Replicada [blog post], retrieved from <http://datacolada.org/81>.
- Simmons, Joseph P. and Uri Simonsohn (2017), "Power Posing: P-Curving the Evidence," *Psychological Science*, 28 (5), 687–93.
- Simmons, Joseph P., Leif D. Nelson and Uri Simonsohn (2011), "False-Positive Psychology: Undisclosed Flexibility in Data Collection and Analysis Allows Presenting Anything as Significant," *Psychological Science*, 22 (11), 1359–66.
- Simmons, Joseph P., Leif D. Nelson and Uri Simonsohn (2021), "Pre-registration: Why and

- How,” *Journal of Consumer Psychology*, 31 (1), 151–62.
- Simonsohn, Uri, Leif D. Nelson and Joseph P. Simmons (2014). “P-curve: A Key To The File-drawer,” *Journal Of Experimental Psychology: General*, 143 (2), 534–47.
- Sobel, Michael E. (1982). “Asymptotic Confidence Intervals for Indirect Effects in Structural Equation Models,” *Sociological Methodology*, 13, 290–312.
- Tukey, John W. (1991). “The Philosophy of Multiple Comparisons,” *Statistical Science*, 6 (1), 100–16.
- Tunca, Burak and Evrim Yanar (2020). “Women's Luxury Products as Signals to Other Women,” *Journal of Marketing Behavior*, 4 (2-4), 227–38.
- Verschuere, Bruno, Ewout H. Meijer, Ariane Jim, Katherine Hoogesteyn, Robin Orthey, Randy J. McCarthy, John J. Skowronski et al. (2018). “Registered Replication Report on Mazar, Amir and Ariely (2008),” *Advances in Methods and Practices in Psychological Science*, 1 (3), 299–317.
- Wiedermann, Wolfgang and Alexander von Eye (2015). “Direction Of Effects In Mediation Analysis,” *Psychological Methods*, 20 (2), 221–44.
- Williams, Richard H. and Donald W. Zimmerman (1989). “Statistical Power Analysis and Reliability of Measurement,” *The Journal of General Psychology*, 116 (4), 359–69.
- Zhao, Xinshu, John G. Lynch Jr. and Qimei Chen (2010). “Reconsidering Baron and Kenny: Myths and Truths about Mediation Analysis,” *Journal of Consumer Research*, 37 (2), 197–206.
- Ziano, Ignazio, Burak Tunca and Xu Wenting (forthcoming). “*Super Size Me: An Unsuccessful Preregistered Replication of the Effect of Product Size on Status Signaling*,” forthcoming at *Meta Psychology*.

Ziano, Ignazio, Jing Dan Yao, Yajing Gao and Gilad Feldman (2020). “Impact of Ownership on Liking and Value: Replications and Extensions of Three Ownership Effect Experiments,” *Journal of Experimental Social Psychology*, 89, 1–15.