

Too good to be true: Publication bias in two prominent studies from experimental psychology

Gregory Francis

Published online: 15 February 2012
© Psychonomic Society, Inc. 2012

Abstract Empirical replication has long been considered the final arbiter of phenomena in science, but replication is undermined when there is evidence for publication bias. Evidence for publication bias in a set of experiments can be found when the observed number of rejections of the null hypothesis exceeds the expected number of rejections. Application of this test reveals evidence of publication bias in two prominent investigations from experimental psychology that have purported to reveal evidence of extrasensory perception and to indicate severe limitations of the scientific method. The presence of publication bias suggests that those investigations cannot be taken as proper scientific studies of such phenomena, because critical data are not available to the field. Publication bias could partly be avoided if experimental psychologists started using Bayesian data analysis techniques.

Keywords Statistical inference · Repeated testing

Experimental psychologists are trained to use statistics to prevent faulty interpretations of their data. By formalizing the decision process, statistical analysis is supposed to remove the influence of the researcher's belief or desire. No researcher in experimental psychology would report an experiment that involved filtering out subjects who did not behave according to the researcher's expectations, because such actions would render the findings scientifically meaningless. Publication bias has a similar effect when replication across experiments is used to determine the evidence for experimental findings (Johnson & Yuan, 2007). When replication is the criterion by which scientific results are judged, a bias to publish positive

findings is essentially the same error as filtering out subjects who do not behave in a desired way. Even well-designed studies can be rendered scientifically useless if other studies are done poorly and publication bias contaminates the set.

Here, publication bias is investigated in two prominent sets of results from experimental psychology. These studies have attracted widespread attention in both academic and nonacademic reports, because they appear to challenge the established scientific understanding of the universe and the scientific method. As one example, Bem (2011) reported empirical evidence that humans can sense future events and use that information to judge the present, an ability that is usually referred to as *psi*. Convincing evidence for *psi* would necessitate major alterations in theories of psychology, biology, and physics. As another example, Schooler (2011) described the “decline effect,” in which early empirical investigations show the strong presence of a phenomenon, but later studies show weak or nonexistent effects. Schooler speculated that scientific studies might introduce something akin to the Heisenberg uncertainty principle, in which observations of effects modify the properties of what is being studied. If this speculation were true, it would imply a fundamental rethinking of causality and would put in question the ability of scientific investigations to reveal truths about the world. The present study reports new analyses of the data sets used to support these claims, leading to the deduction that their conclusions are unwarranted because both data sets suffer from publication bias. As a result, the studies at issue do not provide useful scientific information about the phenomena they have attempted to study.

Publication bias in Bem (2011)

The *psi* experiments reported by Bem (2011) have been criticized on both methodological and analytical grounds

G. Francis (✉)
Department of Psychological Sciences, Purdue University,
West Lafayette, IN, USA
e-mail: gfrancis@purdue.edu

(Wagenmakers, Wetzels, Borsboom & van der Maas, 2011). However, the methodological criticisms are partly speculative, because many experimental steps are not fully described in the published reports. The analytical criticisms are also only partially convincing. Although Wagenmakers et al. noted that the individual experiments in Bem (2011) do not meet the analytical criteria of a standard Bayesian analysis, Rouder and Morey (2011) used a Bayesian meta-analysis and found some evidence for the proposed ψ effect, although the latter authors emphasized that such evidence must be considered in the context of other, conflicting evidence.

Perhaps the most striking property of Bem (2011) is that nine out of 10 described investigations rejected the null hypothesis, thereby indicating evidence for ψ . For many scientists, replication of an effect across multiple experiments provides compelling evidence, but this interpretation is misguided, because it does not consider the statistical power of the experiments. If all of the experiments have high power (the probability of rejecting the null hypothesis when it is false), multiple experiments that reject the null hypothesis would indeed be strong evidence for an effect. However, if the experiments have low or moderate power, then even if the effect were real, one would expect to frequently not reject the null hypothesis.

Ioannidis and Trikalinos (2007) used this observation to develop a statistical test for publication bias. The basic idea is to measure the power of each experiment and to use those measures to predict how often one would expect the null hypothesis to be rejected. If the number of reported rejections is substantially different from what was expected, then the test has found evidence for some kind of publication bias. In essence, the test is a check on the internal consistency of the number of reported rejections, the reported effect sizes, and the power of the tests to detect those effect sizes.

Meta-analytic methods were applied to estimate the power of the experiments in Bem (2011), and the key statistical properties of those experiments are shown in Table 1. A pooled effect size was measured across the 10 experiments, to produce $g^* = .186$. This pooled effect size differs from the average effect size reported by Bem because here a correction has been applied for bias in the effect sizes, and the effect sizes were pooled by weighting the effect size value from each study with its inverse variance (Hedges & Olkin, 1985). This pooled effect size was then combined with the sample size of each experiment to produce an estimated power value (Champely, 2009; R Development Core Team, 2011) for a one-tailed test with $\alpha = .05$, which is the hypothesis test Bem used in the original analysis of the data. The penultimate column in Table 1 shows the estimated power of each experiment to detect the pooled effect size. The sum of the power values across the 10 experiments is 6.27, which is the expected number of rejections of the null

Table 1 Statistical properties of the Bem (2011) experiments on ψ

Experiment	Sample Size	Effect Size	Power From Pooled ES	Observed Power
1	100	.249	.578	.796
2	150	.194	.731	.765
3	97	.248	.567	.783
4	99	.202	.575	.639
5	100	.221	.578	.710
6a	150	.146	.731	.555
6b	150	.144	.731	.543
7	200	.092	.834	.365
8	100	.191	.578	.598
9	50	.412	.363	.890

A positive effect size is consistent with ψ .

hypothesis, given the pooled effect size and the designs of the experiments. The expected number of rejections of the null hypothesis is in stark contrast to the observed nine out of 10 rejections.

The probability of getting nine or more rejections for the 10 experiments reported by Bem (2011) was calculated with an exact test that computed the probability of every combination of nine or more rejections out of 10 experiments by multiplying the appropriate power or Type II error values. There are 11 ways to get nine or more rejections out of 10 experiments, and given the estimated powers of these experiments, the probability of getting a set of experiments with nine or more rejections by chance is .058, which is less than the .1 criterion frequently used for tests of publication bias (Begg & Mazumdar, 1994; Ioannidis & Trikalinos, 2007; Stern, Gavaghan, & Egger, 2000). This low probability suggests that the reported number of rejections of the null hypothesis is abnormally high, given the power of the experiments to detect the pooled effect size.

The use of a pooled effect size supposes a fixed common effect across the experiments, and this approach is consistent with previous interpretations of these experiments (Bem, 2011; Bem, Utts, & Johnson, 2011). It is worthwhile to consider the possibility that such pooling is not appropriate, and that each experiment has a different effect size. Such a calculation is frequently called *observed power*, and although it is a biased estimate of true power (Yuan & Maxwell, 2005), with the large sample sizes used in these experiments it should produce a good estimate of the true power, at least on average. These values are given in the last column of Table 1.

The sum of the observed power estimates across the 10 experiments is 6.64. The exact test reveals that the probability of getting nine or more rejections by chance from 10 experiments with these power values is .088. Again, the number of reported rejections of the null hypothesis

(evidence for ψ) in this set of experiments is higher than is to be expected for the properties of the tests and the magnitudes of the effect sizes being measured.

Publication bias in a set of studies showing a decline effect

Schooler (2011) was motivated to explore the decline effect because of its purported influence on reports of verbal overshadowing (Schooler & Engstler-Schooler, 1990). In verbal overshadowing, performance on a visual memory task is impaired after subjects provide a verbal description of the stimuli. The verbal overshadowing effect has a variable history, with some early experiments showing a strong effect, and other, later experiments showing no effect or an effect in the opposite direction. This weakening of an experimental effect has been labelled the *decline effect*, and it has been observed for other topics that depend on hypothesis testing, including studies of extrasensory perception. In addition to arguing that scientists should reconsider the fundamental nature of investigating the universe, Schooler (2011) suggested that the decline effect may be quite common but remains hidden because of publication bias. On the other hand, if publication bias were to be found to contribute to the findings in which a decline effect appears, then one must be skeptical about the conclusions drawn about the decline effect itself.

We applied the publication bias test to the set of published experiments identified in a meta-analysis of verbal overshadowing (Meissner & Brigham, 2001). Nine of those 18 experiments reported evidence of verbal overshadowing (rejected the null hypothesis in a direction consistent with the effect). The pooled effect size (twice the difference of the arcsine square-root proportions, as in Cohen, 1988) across all experiments ($h^* = 0.301$) and the sample sizes of each experiment were used to compute each study's power for showing evidence of verbal overshadowing with a two-tailed test. These values are shown in the penultimate column of Table 2. The sum of the power values across all of the published experiments is 4.65, which is the expected number of times that these studies would report evidence of verbal overshadowing. An exact test computed the probability of each of the 155,382 possible ways to have nine or more of these experiments report evidence for verbal overshadowing. The sum of these probabilities was .022, which is the chance probability that nine or more of these experiments would find evidence of verbal overshadowing. The conclusion is that a publication bias in these studies favored reporting evidence of verbal overshadowing. This appears to be true, even though only half of the published reports actually found evidence of the effect. A corollary of this analysis is that these experimental studies of verbal

Table 2 Statistical properties of the experiments on verbal overshadowing (Meissner & Brigham, 2001)

n_1	n_2	Effect Size	Power From Pooled ES	Power for Pooled Positive ESs
39	39	.453	.264	.377
33	34	.498	.233	.332
28	28	.644	.202	.286
44	44	.526	.292	.416
35	35	.692	.242	.344
23	27	-.675	.184	–
80	80	.102	.478	.654
60	60	.479	.378	.532
30	30	.175	.214	.303
30	30	.182	.214	.303
20	20	.339	.157	.218
40	40	.245	.270	.385
27	32	.320	.210	.297
97	68	.307	.478	.654
25	29	.659	.196	.276
30	30	-.403	.214	–
30	30	.000	.214	.303
30	30	.711	.214	.303

A positive effect size is consistent with verbal overshadowing

overshadowing were woefully underpowered. To determine whether the effect is real, investigators would need to run studies with larger sample sizes.

Because the studies of verbal overshadowing have tended to use relatively small sample sizes, it is not possible to estimate the power of each experiment with an observed power analysis. Thus, one possible criticism of the publication bias test is that it pooled together findings from experiments that actually investigated different effects. Indeed, the publication bias test is sensitive to heterogeneity of the effect sizes (Ioannidis & Trikalinos, 2007; Johnson & Yuan, 2007). The analysis above addressed this concern by using a selection of experiments that were identified by subject matter experts as attempted replications of the verbal overshadowing effect, but it could be that other experts would have made different choices, thereby leading to different outcomes of the publication bias test. If enough experiments with similar methods were available (e.g., experiments that used a particular set of instructions or experiments from one laboratory), it would be possible to run the publication bias test for subsets of experiment sets and then to pool them together to get an overall probability of the entire set.

Concerns about the heterogeneity of effect sizes are often not as worrisome as one might suspect. For example, a reviewer noted that two discrepant findings in the studies of verbal overshadowing showed a strong effect in the direction opposite the one that was typically reported. These

findings could be the result of normal variation, due to random sampling from a population with a small effect size (this is the interpretation for the analysis above), but these findings could alternatively be interpreted as investigations of an entirely different effect. If the latter interpretation were true, then these experiments should be removed from the meta-analysis. When this is done, the pooled effect size increases to $h^* = 0.373$ (the two negative experiments had fairly small sample sizes, so they did not strongly influence the pooled effect size). As the final column in Table 2 shows, the power of each experiment increases when considering this larger effect size.

However, the impact of the larger power values for the publication bias test is partly mitigated by the fact that one must now consider the probability of rejecting the null hypothesis nine times out of only 16 experiments. The sum of the power values for the new effect size is 5.98, and an exact test that considers the 26,333 ways that nine or more of these 16 experiments could reject the null hypothesis produces a result of .090, which is still below the .1 threshold. Thus, for this data set, even if one uses the outcomes of the experiments to determine whether an experiment is a replication, there is still evidence of publication bias. In general, though, using an experimental outcome to determine whether an experiment is an attempted replication is itself a type of bias and should be avoided. Ultimately, the choices about which experiments are replications of a basic effect should be defined by the methodological characteristics of the experiments or by a theoretical explanation of the phenomenon.

Conclusions

Just as no experimental psychologist would believe the findings of an experiment that was biased to report data from only subjects that showed a desired effect, so the presence of a publication bias means that the studies of psi and verbal overshadowing do not tell us anything scientifically useful about the phenomena, because critical data are not part of the results. The effects may be real, or they may be entirely due to bias. The set of studies are simply not scientific. Even worse, the publication bias test generally cannot identify which, if any, specific experimental results are valid, because it only tracks statistics across the entire set. Thus, although some researchers may report their experimental findings fully and properly, those experiments can be rendered scientifically useless by poor reporting practices from other researchers.

It might be feared that the publication bias test is so stringent that almost every set of studies would demonstrate a publication bias. In fact, the test used here is strongly inclined to not report a publication bias, because reported

effect sizes tend to overestimate reality (Ioannidis, 2008). Moreover, when the publication bias test indicates suppression of null or negative findings, the true effect size for the phenomena being studied is probably smaller than what is estimated by the biased set of reports. Thus, many of the estimated power values that form the basis of the test are larger than they should be, which means that the expected number of rejections is overestimated. Once evidence for bias is found, it is likely that the bias is even more pronounced than the test indicates.

When evidence of a publication bias is presented, many people think of the file drawer problem, which refers to the idea that a researcher runs many different experiments but only publishes the ones that tend to show evidence for an effect. This kind of bias could be due to the deliberate intention of the author to mislead the field, or by an inability to get null findings approved by reviewers and editors. Such a problem surely exists, and the test described above can detect its influence.

A closely related bias, with a similar result, occurs when an experimenter measures many different variables but then selectively reports only the findings that reject the null hypothesis. Simmons, Nelson, and Simonsohn (2011) demonstrated how such an approach (in combination with some other tricks) can suggest evidence for truly outlandish effects. Moreover, seemingly innocuous decisions at many different levels of research can produce a publication bias. Given the high frequency of errors in reporting statistical findings (Bakker & Wicherts, 2011; Wicherts, Bakker, & Molenaar, 2011), researchers can introduce a publication bias by being very careful when the results are contrary to what was expected, but not double-checking results that agree with their beliefs or expectations. Likewise, data from a subject that happens to perform poorly (given the experimenter's hopes) might be thrown out if there is some external cause, such as noise in a neighboring room or a mistake in the computer program, but data from a subject that happens to perform well under the same circumstances might be kept (and even interpreted as evidence of the effect's strength despite distractions).

One form of publication bias is related to experiment stopping rules. When gathering experimental data, it is easy to compute statistical tests at various times and to explore whether the experiment is likely to work. Such checks (sometimes called *data peeking*) introduce a bias if the experimenter allows that information to influence whether to continue the experiment (Berger & Berry, 1988; Kruschke, 2010; Wagenmakers, 2007). In a similar way, when an experiment ends but the data do not quite reach statistical significance, it is common for researchers to add more subjects in order to get a definitive answer. This approach (although it seems like good science to gather more data) is inconsistent with the traditional frequentist

approach to hypothesis testing. Such approaches inflate the Type I error rate (Strube, 2006).

The publication bias test cannot distinguish between the myriad ways for bias to appear, but since it provides evidence that the studies of psi and verbal overshadowing contain bias, one need not propose radical characteristics of the universe (Bem, 2011) or limits to the scientific method (Schooler, 2011) in order to explain the properties of those studies. The simpler explanation is that, as a set, the studies either are not reporting all of the relevant information or have rejected the null hypothesis more frequently than they should have because they were run improperly. Either way, these studies are at best anecdotal, and as such they need no scientific explanation at all.

Perhaps the most striking characteristic of both the Bem (2011) study and the set of studies reporting verbal overshadowing is that they meet the current standards of experimental psychology. The implication is that it is the standards and practices of the field that are not operating properly. Clearly, these standards and practices need to be improved to ensure that the frequency of reporting the null hypothesis is consistent with the power of the experiments. Such improvements are long overdue. Sterling (1959) noted that 97.3% of statistical tests rejected the null hypothesis for the major scientific findings reported in four psychology journals. A follow-up analysis by Sterling, Rosenbaum, and Weinkam (1995) found a similar result, with a null hypothesis rejection rate of 95.56%. These high percentages suggest that the power of the experiments (if the alternative hypothesis is true) must generally be well above .9, even though power values in the range of .25 to .85 are more common in psychology (Hedges, 1984). Formalizing this discrepancy between the observed and expected proportions of rejections of the null hypothesis is the core of the publication bias test developed by Ioannidis and Trikalinos (2007) and used above.

When publication bias has been reported in other fields (e.g., by Munafò & Flint, 2010), there is often a call to create a registry of planned experiments and to require researchers to describe the outcomes of the experiments, regardless of the findings (Schooler, 2011). Such a project would be an expensive undertaking, would require complex paperwork for every experiment, and would be difficult to enforce for a field like experimental psychology, where many of the investigations are exploratory rather than planned. A much simpler partial solution to the problem would be Bayesian data analysis.

A Bayesian approach has two features that mitigate the appearance of publication bias. First, in addition to finding evidence in support of an alternative hypothesis, a Bayesian analysis can find evidence in support of a null hypothesis. Thus, an experiment that finds convincing evidence for a null hypothesis provides publishable scientific information

about a phenomenon in a way that a failure to reject the null hypothesis does not. Second, a Bayesian analysis is largely insensitive to the effects of data peeking and multiple tests (Berger & Berry, 1988; Kruschke, 2010; Wagenmakers, 2007), so some of the methodological approaches that inflate Type I error rates and introduce a publication bias would be rendered inert. Bayesian data analysis is only a partial solution because the file drawer problem might still result from the actions of researchers, reviewers, and editors who deliberately or unintentionally choose to suppress experimental findings that go against their hopes. Such remaining cases could be identified by the publication bias test described in this article, or by a Bayesian equivalent.

The findings of the publication bias test in experimental psychology demonstrate that the care that is so rigorously taken at the level of an experiment is sometimes not being exercised at the next level up, when research findings are being reported. Such publication bias is inappropriate for an objective science, and the field must improve its methods and reporting practices to avoid it.

References

- Bakker, M., & Wicherts, J. M. (2011). The (mis)reporting of statistical results in psychology journals. *Behavior Research Methods*, *43*, 666–678. doi:10.3758/s13428-011-0089-5
- Begg, C. B., & Mazumdar, M. (1994). Operating characteristics of a rank correlation test for publication bias. *Biometrics*, *50*, 1088–1101.
- Bem, D. J. (2011). Feeling the future: Experimental evidence for anomalous retroactive influences on cognition and affect. *Journal of Personality and Social Psychology*, *100*, 407–425. doi:10.1037/a0021524
- Bem, D. J., Utts, J., & Johnson, W. O. (2011). Must psychologists change the way they analyze their data? *Journal of Personality and Social Psychology*, *101*, 716–719. doi:10.1037/a0024777
- Berger, J., & Berry, D. (1988). The relevance of stopping rules in statistical inference (with discussion). In S. S. Gupta & J. Berger (Eds.), *Statistical decision theory and related topics 4* (Vol. 1, pp. 29–72). New York: Springer.
- Champely, S. (2009). pwr: Basic functions for power analysis (R package version 1.1.1). Available at <http://CRAN.R-project.org/package=pwr>
- Cohen, J. (1988). *Statistical power analysis for the behavioral sciences* (2nd ed.). Hillsdale: Erlbaum.
- Hedges, L. V. (1984). Estimation of effect size under nonrandom sampling: The effects of censoring studies yielding statistically insignificant mean differences. *Journal of Educational Statistics*, *9*, 61–85.
- Hedges, L. V., & Olkin, I. (1985). *Statistical methods for meta-analysis*. San Diego: Academic Press.
- Ioannidis, J. P. A. (2008). Why most discovered true associations are inflated. *Epidemiology*, *19*, 640–648.
- Ioannidis, J. P. A., & Trikalinos, T. A. (2007). An exploratory test for an excess of significant findings. *Clinical Trials*, *4*, 245–253.
- Johnson, V., & Yuan, Y. (2007). Comments on “An exploratory test for an excess of significant findings” by JPA Ioannidis and TA Trikalinos. *Clinical Trials*, *4*, 254–255.

- Kruschke, J. K. (2010). Bayesian data analysis. *Wiley Interdisciplinary Reviews: Cognitive Science*, 1, 658–676. doi:10.1002/wcs.72
- Meissner, C. A., & Brigham, J. C. (2001). A meta-analysis of the verbal overshadowing effect in face identification. *Applied Cognitive Psychology*, 15, 603–616.
- Munafò, M. R., & Flint, J. (2010). How reliable are scientific studies? *The British Journal of Psychiatry*, 197, 257–258. doi:10.1192/bjp.bp.109.069849
- R Development Core Team. (2011). R: A language and environment for statistical computing. Vienna, Austria: R Foundation for Statistical Computing. Available at www.R-project.org.
- Rouder, J. N., & Morey, R. D. (2011). A Bayes factor meta-analysis of Bem's ESP claim. *Psychonomic Bulletin & Review*, 18, 682–689. doi:10.3758/s13423-011-0088-7
- Schooler, J. (2011). Unpublished results hide the decline effect. *Nature*, 470, 437. doi:10.1038/470437a
- Schooler, J. W., & Engstler-Schooler, T. Y. (1990). Verbal overshadowing of visual memories: Some things are better left unsaid. *Cognitive Psychology*, 22, 36–71. doi:10.1016/0010-0285(90)90003-M
- Simmons, J. P., Nelson, L. D., & Simonsohn, U. (2011). False-positive psychology: Undisclosed flexibility in data collection and analysis allows presenting anything as significant. *Psychological Science*, 22, 1359–1366. doi:10.1177/0956797611417632
- Sterling, T. D. (1959). Publication decisions and the possible effects on inferences drawn from test of significance—or vice versa. *Journal of the American Statistical Association*, 54, 30–34.
- Sterling, T. D., Rosenbaum, W. L., & Weinkam, J. J. (1995). Publication decisions revisited: The effect of the outcome of statistical tests on the decision to publish and vice versa. *The American Statistician*, 49, 108–112.
- Sterne, J. A., Gavaghan, D., & Egger, M. (2000). Publication and related bias in meta-analysis: Power of statistical tests and prevalence in the literature. *Journal of Clinical Epidemiology*, 53, 1119–1129.
- Strube, M. J. (2006). SNOOP: A program for demonstrating the consequences of premature and repeated null hypothesis testing. *Behavior Research Methods*, 38, 24–27. doi:10.3758/BF03192746
- Wagenmakers, E.-J. (2007). A practical solution to the pervasive problems of *p* values. *Psychonomic Bulletin & Review*, 14, 779–804. doi:10.3758/BF03194105
- Wagenmakers, E.-J., Wetzels, R., Borsboom, D., & van der Maas, H. L. J. (2011). Why psychologists must change the way they analyze their data: The case of *psi*. Comment on Bem (2011). *Journal of Personality and Social Psychology*, 100, 426–432. doi:10.1037/a0022790
- Wicherts, J. M., Bakker, M., & Molenaar, D. (2011). Willingness to share research data is related to the strength of the evidence and the quality of reporting of statistical results. *PLoS One*, 6, e26828. doi:10.1371/journal.pone.0026828
- Yuan, K. H., & Maxwell, S. (2005). On the post hoc power in testing mean differences. *Journal of Educational and Behavioral Statistics*, 30, 141–167.