Practical Tools and Strategies for Researchers to Increase Replicability

Michèle B. Nuijten¹

1. Department of Methodology and Statistics, Tilburg School of Social and Behavioral Sciences, Tilburg University, Warandelaan 2, 5037 AB, Tilburg, The Netherlands. Email: m.b.nuijten@tilburguniversity.edu.

Acknowledgements

Parts of this paper are based on sections from my doctoral dissertation. I would like to thank Jelte Wicherts and Marcel van Assen for their helpful comments on earlier drafts of those sections.

Abstract

A growing body of evidence indicates that many published effects in multiple scientific fields may be overestimated or even false. These problems have gained a lot of attention in psychology, where researchers have even started to speak of a “replication crisis”. Luckily, many different solutions have already been proposed and implemented; some were inspired by practices in other scientific fields, and some might inspire other fields in return. Here, I would like to briefly review the problems in psychology and provide an overview of some practical tools and strategies that researchers can implement in their own workflow to increase replicability and the overall quality of science.

What this Paper Adds

- Researchers can implement many practical tools and strategies to improve replicability
- Strategies include improving statistical inference, preregistration, multi-lab collaboration, and sharing data.
- Different scientific fields should look at each other’s best practices

How solid is scientific research? This question is asked more and more, and there is growing concern that many published findings are overly optimistic (1). A vast amount of research has found evidence for problems such as overestimated effects, low replicability rates, and statistical errors in medical research (e.g., in cancer biology; 2), economics (e.g., 3), behavioral ecology (e.g., 4), genetic behavior research (e.g., 5), and other scientific fields. Even though these problems seem to be omnipresent, they have received a lot of attention in psychology in particular. Psychological researchers have even started to speak of a “crisis of confidence” or a “replication crisis” (6, 7).

Since the start of the replication crisis in psychology around 2011 (8), many solutions have been proposed and implemented. A large number of these solutions were inspired by what can be considered “standard practice” in other fields. Other solutions have been designed specifically for psychology but should perhaps become standard practice in other fields. I think the current state of psychology is a good example of how scientific fields can learn from each other to improve their research.

In this review, I will first sketch an overview of the replicability problems in psychology and then discuss several practical tools and strategies that psychologists can implement and have implemented to increase replicability and overall scientific quality. My hope is that not only psychologists, but also researchers from other fields will implement these tools and strategies in their own workflow, or consider them an inspiration to come up with solutions that are tailored to their specific field.

The Replication Crisis in Psychology

Psychology is in a crisis. The growing concern about psychology’s trustworthiness is fueled by the repeated findings that large numbers of published psychological findings could not be replicated in novel samples. For instance, the large-scale, collaborative Reproducibility Project: Psychology (RPP) investigated the replicability of 100 psychology studies (9). Two of the main findings in this project were that the percentage of statistically significant effects dropped from 97%

in the original studies to only 36% in the replications, and that the effects in the replications were only about half the size of those in the original studies. Furthermore, the large majority of studies in psychology find support for the tested hypothesis (10). However, this is incompatible with the generally low statistical power of studies in the psychological literature (11, 12).

One of the main causes for this excess of significant effects seems to be publication bias; the phenomenon that statistically significant findings have a higher probability of being published than nonsignificant findings (13). The evidence that the field of psychology is affected by publication bias is overwhelming. Studies found that manuscripts without significant results are both less likely to be submitted and less likely to be accepted for publication (e.g., 14, 15). Furthermore, published studies seem to have systematically larger effects than unpublished ones (16).

This *de facto* requirement to report statistically significant results in journal articles can lead to strategic behavior in data analysis (11, 17). Data analysis in psychology is very flexible: there are many possible statistical analyses to answer the same research question (18). It can be shown that strategic use of this flexibility will almost always result in at least one significant finding; one that is likely to be a false positive (11, 19). This becomes even more problematic, if only the analyses that “worked” are reported and presented as if they were planned from the start (20, 21). Survey results show that many psychologists admit to such strategic choices (17), and use of study registers and later disclosures by researchers provide direct evidence that indeed some of these practices are quite common (22, 23).

So, are psychological researchers “evil scientists”, who only care about publishing and not about finding the truth? I sincerely doubt that. I think that the underlying cause of the replication crisis is a combination of issues that are at play in many scientific fields: publication bias, publication pressure, human biases, and inadequate statistical training and experience. A lack of awareness of these problems may also contribute to the crisis. I am convinced that most scientists are interested in the truth and want to do the best research they can. However, I also think that it is not always

straightforward how you can improve research. In the next section I will therefore give an overview of practical tools and strategies to increase replicability and overall scientific quality. Some of these suggested solutions are specifically tailored to psychology, but all of them were derived from general principles such as solid inference, collaboration, and transparency, which can be translated to concrete solutions for basically any field.

Strategies to Increase Replicability

Over the years, there have been many suggested solutions for the problems related to the replicability crisis in psychology. Some solutions should be implemented top-down, by journal editors, funders, or policy makers. Some solutions have to do with improving education, and others are focused on correcting the existing literature. Here, I chose not to focus on these, but instead focus on concrete solutions that researchers can implement in their own workflow to improve their own research.

Statistics

Many of the current problems in psychological science seem to be related to the use and reporting of statistics. For instance, many published papers contain statistical errors. This makes papers irreproducible: it becomes unclear where the reported numbers came from. Many of these errors seem to be caused by simple typos (e.g., 24). New tools are being developed that can be used to check your manuscript for such statistical mistakes. For instance, the free tool *statcheck* (25; <http://statcheck.io>) can scan a paper, extract statistics reported in APA style, and check whether the reported *p*-values match the reported test statistic and degrees of freedom. Similarly, the tool *GRIM* (26; http://www.prepubmed.org/grim_test/) can be used to check if reported means are compatible with the sample size and number of items. A typo is easily made, and using such “spellcheckers” for statistics may help in avoiding many unnecessary errors in reporting statistics.

Some researchers argue that null-hypothesis significance testing is inherently flawed and that we should focus on other ways of statistical inference. One alternative is to focus on reporting effect sizes and their confidence intervals (27). This would take the focus away from the dichotomous decisions that are surrounding p -values, and shift it to a continuous judgement of evidence. Another suggestion is to leave the frequentist framework altogether and move to Bayesian statistics (28). For researchers with relatively little statistical training, this may seem intimidating. To that end, the new statistical program JASP (<http://jasp-stats.org>) is being developed, to offer user-friendly, “SPSS-like” software to facilitate the use of Bayesian statistics (29).

It has also been argued that the main problem is not the specific statistical framework per se, but rules of thumb that are applied uncritically, such as the strong focus on $p < .05$ to draw dichotomous conclusions about the presence of an effect (30). Switching to a new framework would not necessarily immediately solve this; it is easy to come up with new rules of thumb (e.g., Bayes factors should be larger than 3, confidence intervals should not include the null, etc.). The solution would be to promote better understanding of statistics. Psychologists usually had relatively little statistical training in their undergraduate and graduate programs, and this probably also holds for scientists in other fields. One way to is to follow online courses. These courses are often free, you are flexible in planning them, and you have the freedom to pick specific modules that fit your needs. For example, there is a free Coursera course to improve your statistical inference, which also has a strong focus on open science principles (<https://www.coursera.org/learn/statistical-inferences>, by Daniël Lakens).

Preregistration

One promising way to avoid both publication bias and strategic use of flexibility in data analysis is preregistration (see 31). Preregistering studies has been standard practice in clinical trials for a long time, but has only recently gained popularity in psychology. When preregistering a study, researchers specify all planned aspects of their study, such as the study design, sampling plan, and

analysis plan, before the data are collected. This can be done online (e.g., on the website of the Open Science Framework; <http://osf.io>), but it is also possible to submit a preregistration for peer review. Over 100 journals are now even offering the possibility to submit registered reports (for a full list of participating journals, see <https://cos.io/rr/>). If the research plan is considered theoretically and methodologically sound, the study can get an “in-principle acceptance”. If the study is executed and reported according to the preregistration, it will be published, regardless of the outcomes. This is a direct prevention of publication bias.

Preregistration also removes room for undisclosed flexibility in data analysis; any deviation from the original analysis plan is now clearly identified as an exploratory analysis, and can no longer be presented as having been planned from the start. In other words, preregistration creates a clear separation between confirmatory and exploratory findings (21). Note, however, that preregistration can be done in different levels of detail, and several preregistration formats and templates have been suggested (e.g., 32). Strategic use of flexibility in analyses can only truly be avoided if the preregistration contains a high level of detail, and this is currently often not the case (33). Future research is still needed to investigate formats and implementations of preregistration, but preregistration seems a promising way to avoid undisclosed flexibility in data analysis and publication bias.

Multi-Lab Collaborations

Psychological studies are often vastly underpowered (e.g., see 12). This is already bad news in itself, since this means that we often will not find true effects. But what is more, the overestimation of effects due to flexibility in data analysis and publication bias becomes worse when power is low (11, 34). In short: low power can lead to an excess of false positive findings in the literature.

The solution might seem simple: run studies with higher power. But in many cases, this is easier said than done. For instance, in developmental psychology it can be difficult to reach

sufficient power; not only is collecting large samples of minors challenging, children can also display large amounts of variability in their behavior, which makes it hard to collect consistent measurements (35).

In cases where it is challenging or even impossible to run a sufficiently powered study, a multi-lab collaboration can be a solution. Multi-lab collaborations are relatively common in genetics (e.g., 36), but they have only recently started gaining popularity in psychology. One great example is the Many Babies project, in which the aim was to “bring as many diverse labs together, conduct pre-registered replications of well-established phenomena and explore the impact of multiple factors, including differences in set-up, infant age and population, and so forth” (<https://osf.io/rpw6d/wiki/home/>; 35). Other examples are the series of replication studies in the Many Labs projects (e.g., see 37), and the Reproducibility Project: Psychology that I mentioned above (9).

How do you get in touch with other labs to join or start such a multi-lab collaboration? In psychology, researchers are developing online platforms to facilitate collaboration and share resources. One example is the Psychological Science Accelerator (<https://psysciacc.wordpress.com>). This is a network of psychological science labs from all over the world that will collect data for democratically selected studies. A related, but slightly different initiative is StudySwap (<https://osf.io/9aj5g/>), an online platform where researchers can offer or request resources to initiate multi-site collaborations. For instance, if your participant pool is limited, you can post a request on StudySwap for other researchers to run your experiment in their labs. In turn, you might be able to run their experiments in your lab.

Multi-lab collaborations are also a great way to conduct high quality replication studies. A replication study, as any study, is most informative when it has high power and is not affected by publication bias and related biases caused by researcher degrees of freedom (34). One way to ensure both is through Registered Replication Reports (RRR; 38). An RRR is very similar to a

preregistered report (see above), but adds an element of multi-lab collaboration to facilitate large-scale replication attempts. In an RRR, researchers first submit a preregistration of the research plan, which is then reviewed by peer reviewers and the author(s) of the original study. If the plan is approved, it is posted publicly so that other labs can follow the same protocol and conduct their own replication of that study. RRRs therefore share the same benefits as preregistrations (see above), and the collaborative aspect allows for larger samples and thus increases power. Moreover, the multi-lab format allows the study of whether effects are moderated by substantive or methodological factors (e.g., types of sample or specifics of the data collection).

Transparency

One of the most recurring suggestions to improve the quality of our science is to increase transparency and openness (39-41). This call for transparency encompasses almost all aspects of a study. First, there is a strong call for open data. Some of the greatest advantages of open data are the possibility to run secondary analyses to answer new questions, verify analyses of published work or examine the robustness of the original analyses, and compute specific effect sizes for meta-analyses (see also 42). Several journals are now requiring or rewarding open data, and these policies are strongly related to increased data availability (43, 44).

Ideally, when data are shared, there should be a quality check to ensure they are relevant, complete, and usable (44, 45). One way to also increase the likelihood that data remain available over time, is to publish them in online data repositories. An example of such a repository is the Open Science Framework (<http://osf.io>). One advantage of this platform is that it also facilitates sharing other aspects of a study, such as materials and analysis scripts. With the original materials, a study can be more easily replicated, and if analysis scripts are posted online, any flexibility in data analysis can be openly discussed. Another way to ensure quality of published data, is to publish them in a data journal, such as the peer-reviewed *Journal of Open Psychology Data* (42).

Overall, there seems to be growing enthusiasm towards increasing transparency in research, and there is increasing talk of an “Open Science Movement” (46). A long list of journals from a wide range of disciplines has now signed the Transparency and Openness Promotion (TOP) guidelines to indicate their support for open science (see <https://cos.io/our-services/top-guidelines/> for the full list). Furthermore, the Society for the Improvement of Psychological Science (SIPS) was formed specifically to “[bring] together scholars working to improve methods and practices in psychological science” (<http://improvingpsych.org/mission/>), and the past two SIPS meetings have already resulted in tangible steps forward (e.g., StudySwap, the preprint service PsyArXiv, and the implementation of the tool statcheck in the peer review process of the journal *Psychological Science*). Organizations like SIPS offer a platform for scholars of all levels of seniority to share their experiences and ideas, to come up with solutions to improve the field.

Discussion

Several large-scale problems are affecting the validity and reproducibility of scientific research. Especially in psychology, these problems have gained a lot of attention, and there has been an increasing number of suggestions how to solve them. Many of the suggested solutions are systemic, focused on top-down implementation, or on the training of students, but many other solutions are practical tools and strategies that researchers can immediately implement in their own workflow. Researchers can use online tools to double check reported statistics, follow online courses to improve their statistical inference, or use different statistical frameworks. They can also start to preregister their research plans, engage in multi-lab collaborations, and increase their level of transparency and openness. Many of these tools and strategies were inspired by standard practices in other fields, and I think these and other fields can in turn be inspired by the way psychology has implemented them.

Not every suggested solution is always feasible for every researcher in every project. For instance, data sharing can be difficult in research that involves sensitive questions, or research

involving minors, due to privacy restrictions. Similarly, some research may require participants from a certain local population, which makes the project unsuitable for a multi-lab collaboration.

However, in both cases there is still a long list of other feasible tools and strategies to improve replicability: for instance, researchers can still preregister the research plan, share the materials and analysis scripts, and improve statistical inference by reporting effect sizes and confidence intervals and add Bayesian analyses.

The large number of potential solutions to improve science may seem a bit overwhelming, but I think that every researcher can make a selection out of the suggested solutions and tailor them to their specific research projects. Even if you pick only one of the solutions above for one single research project, science will already be more solid than it was yesterday.

References

1. Ioannidis JPA. Why most discovered true associations are inflated. *Epidemiology*. 2008;19(5):640-8.
2. Nosek BA, Errington TM. Reproducibility in cancer biology: Making sense of replications. *eLife*. 2017;6:e23383.
3. Ioannidis JPA, Stanley TD, Doucouliagos H. The power of bias in economics research. *The Economic Journal*. 2017;127(605).
4. Jennions MD, Moller AP. A survey of the statistical power of research in behavioral ecology and animal behavior. *Behavioral Ecology*. 2003;14(3):438-45.
5. Ioannidis JPA. Calibration of credibility of agnostic genome-wide associations. *American Journal of Medical Genetics Part B-Neuropsychiatric Genetics*. 2008;147B(6):964-72.
6. Baker M. 1,500 scientists lift the lid on reproducibility. *Nature News*. 2016;533(7604):452.
7. Pashler H, Wagenmakers EJ. Editors' Introduction to the Special Section on Replicability in Psychological Science A Crisis of Confidence? *Perspectives on Psychological Science*. 2012;7(6):528-30.
8. Spellman BA. A short (personal) future history of revolution 2.0. SAGE Publications Sage CA: Los Angeles, CA; 2015.
9. Open Science Collaboration. Estimating the reproducibility of psychological science. *Science*. 2015;349(6251).
10. Fanelli D. "Positive" Results Increase Down the Hierarchy of the Sciences. *PLoS One*. 2010;5(3):e10068.
11. Bakker M, van Dijk A, Wicherts JM. The rules of the game called psychological science. *Perspectives on Psychological Science*. 2012;7(6):543-54.
12. Button KS, Ioannidis JPA, Mokrysz C, Nosek BA, Flint J, Robinson ESJ, et al. Power failure: why small sample size undermines the reliability of neuroscience. *Nature Reviews Neuroscience*. 2013;14(5):1-12.
13. Greenwald AG. Consequences of prejudice against the null hypothesis. *Psychological Bulletin*. 1975;82:1-20.
14. Cooper H, DeNeve K, Charlton K. Finding the missing science: The fate of studies submitted for review by a human subjects committee. *Psychological Methods*. 1997;2(4):447-52.
15. Franco A, Malhotra N, Simonovits G. Publication bias in the social sciences: Unlocking the file drawer. *Science*. 2014;345(6203):1502-5.
16. Polanin JR, Tanner-Smith EE, Hennessy EA. Estimating the Difference Between Published and Unpublished Effect Sizes: A Meta-Review. *Review of Educational Research*. 2015:1-30.
17. John LK, Loewenstein G, Prelec D. Measuring the prevalence of questionable research practices with incentives for truth-telling. *Psychological Science*. 2012;23:524-32.
18. Gelman A, Loken E. The statistical crisis in science data-dependent analysis - a "garden of forking paths" - explains why many statistically significant comparisons don't hold up. *American Scientist*. 2014;102(6):460.
19. Simmons JP, Nelson LD, Simonsohn U. False-positive psychology: Undisclosed flexibility in data collection and analysis allows presenting anything as significant. *Psychological Science*. 2011;22:1359-66.
20. Kerr NL. HARKing: Hypothesizing after the results are known. *Personality and Social Psychology Review*. 1998;2:196-217.
21. Wagenmakers EJ, Wetzels R, Borsboom D, Maas HJvd, Kievit RA. An agenda for purely confirmatory research. *Perspectives on Psychological Science*. 2012;7:632-8.
22. Franco A, Malhotra N, Simonovits G. Underreporting in Psychology Experiments: Evidence From a Study Registry. *Social Psychological and Personality Science*. 2016;7(1):8-12.

This paper is accepted for publication in *Developmental Medicine & Child Neurology*. Please cite as Nuijten, M. B. (in press). Practical tools and strategies for researchers to increase replicability. *Developmental Medicine & Child Neurology*.

23. LeBel EP, Borsboom D, Giner-Sorolla R, Hasselman F, Peters KR, Ratliff KA, et al. PsychDisclosure.org: Grassroots Support for Reforming Reporting Standards in Psychology. *Perspectives on Psychological Science*. 2013;8(4):424-32.
24. Nuijten MB, Hartgerink CHJ, Van Assen MALM, Epskamp S, Wicherts JM. The prevalence of statistical reporting errors in psychology (1985-2013). *Behavior Research Methods*. 2016;48(4):1205-26.
25. Epskamp S, Nuijten MB. statcheck: Extract statistics from articles and recompute p values. R package version 1.0.0. <http://CRAN.R-project.org/package=statcheck2014>.
26. Brown NJL, Heathers JAJ. The GRIM Test: A Simple Technique Detects Numerous Anomalies in the Reporting of Results in Psychology. *Social Psychological and Personality Science*. 2016;8(4):363-9.
27. Cumming G. Understanding the new statistics: Effect sizes, confidence intervals, and meta-analysis: Routledge; 2013.
28. Wagenmakers EJ, Marsman M, Jamil T, Ly A, Verhagen J, Love J, et al. Bayesian inference for psychology, part I: Theoretical advantages and practical ramifications. *Psychonomic Bulletin & Review*. 2017.
29. Wagenmakers EJ, Love J, Marsman M, Jamil T, Ly A, Verhagen J, et al. Bayesian Inference for Psychology. Part II: Example Applications with JASP. *Psychonomic Bulletin & Review*. 2017.
30. Gigerenzer G. Mindless statistics. *The Journal of Socio-Economics*. 2004;33:587-606.
31. Chambers CD, Munafò MR. Trust in science would be improved by study pre-registration. . 2013.
32. Van 't Veer AE, Giner-Sorolla R. Pre-registration in social psychology - A discussion and suggested template. *Journal of Experimental Social Psychology*. 2016;67:2-12.
33. Veldkamp CLS, Bakker M, van Assen MALM, Cromptvoets EAV, Ong HH, Soderberg CK, et al. Restriction of opportunistic use of researcher degrees of freedom in pre-registrations on the Open Science Framework. 2017.
34. Nuijten MB, Van Assen MALM, Veldkamp CLS, Wicherts JM. The replication paradox: combining studies can decrease accuracy of effect size estimates. *Review of General Psychology*. 2015;19(2):172-82.
35. Frank MC, Bergelson E, Bergmann C, Cristia A, Floccia C, Gervain J, et al. A collaborative approach to infant research: Promoting reproducibility, best practices, and theory-building. Preprint retrieved from <https://psyarxivcom/27b43/>. 2018.
36. Davies G, Armstrong N, Bis JC, Bressler J, Chouraki V, Giddaluru S, et al. Genetic contributions to variation in general cognitive function: a meta-analysis of genome-wide association studies in the CHARGE consortium (N=53 949). *Molecular Psychiatry*. 2015;20:183.
37. Klein RA, Ratliff KA, Vianello M, Adams J, Reginald B., Bahník Š, Bernstein MJ, et al. Investigating variation in replicability: A “Many Labs” Replication Project. *Social Psychology*. 2014;45(3):142-52.
38. Association for Psychological Science. Registered Replication Reports. n.d.
39. Nosek BA, Alter G, Banks GC, Borsboom D, Bowman SD, Breckler SJ, et al. Promoting an open research culture. *Science*. 2015;348(6242):1422-5.
40. Vazire S. Quality uncertainty erodes trust in science. *Collabra: Psychology*. 2017;3(1).
41. Nuijten MB. Share Analysis Plans and Results. *Nature*. 2017;551:559.
42. Wicherts JM. Science revolves around the data. *Journal of Open Psychology Data*. 2013;1(1):e1.
43. Nuijten MB, Borghuis J, Veldkamp CLS, Dominguez-Alvarez L, Van Assen MALM, Wicherts JM. Journal data sharing policies and statistical reporting inconsistencies in psychology. *Collabra: Psychology*. 2017;3(1):1-22.
44. Kidwell MC, Lazarevic LB, Baranski E, Hardwicke TE, Piechowski S, Falkenberg L-S, et al. Badges to Acknowledge Open Practices: A Simple, Low-Cost, Effective Method for Increasing Transparency. *PLoS biology*. 2016;14(5):e1002456.

This paper is accepted for publication in *Developmental Medicine & Child Neurology*. Please cite as Nuijten, M. B. (in press). Practical tools and strategies for researchers to increase replicability. *Developmental Medicine & Child Neurology*.

45. Wilkinson MD, Dumontier M, Aalbersberg IJ, Appleton G, Axton M, Baak A, et al. The FAIR Guiding Principles for scientific data management and stewardship. *Scientific data*. 2016;3:160018.
46. Spellman B, Gilbert EA, Corker KS. Open Science: What, Why, and How. *PsyArXiv*. 2017.