To Have Practical Relevance, Scientific Claims Need to be Accurate

Rolf A. Zwaan
Erasmus University Rotterdam

Anita Eerland
Utrecht University

Corresponding author:

Rolf A. Zwaan
Department of Psychology
Erasmus University Rotterdam
3000 DR ROTTERDAM
Netherlands
rolfzwaan@gmail.com
Twitter: @RolfZwaan

Bullet Points

- Replication and preregistration are crucial if research is meant to have an impact on society.
- Transparency is crucial if research is meant to have an impact on society.
- From prediction to implementation in 12 steps.
- How producers, evaluators, communicators, and consumers of scientific information can contribute toward making our science more accurate.

Abstract

To have practical relevance, scientific claims should be accurate. Everyone needs scientific claims, and practices derived from them to be accurate and reliable: researchers, reviewers, editors, policy makers, funding agencies, journalists, and the general public. Nevertheless, there is evidence that questions the reliability and accuracy of findings in the social and behavioral sciences. In this article, we identify several prevalent threats to accuracy. We will then discuss ways in which these threats can be removed, or at least minimized. Finally, we will discuss how different agents in the system, researchers, reviewers, editors, policy makers, funding agencies, journalists, and the general public, can contribute to making science more accurate and thus more practically relevant.

We start with a simple notion: to have practical relevance, scientific claims should be accurate and reliable. Everyone needs scientific claims, and practices derived from them to be accurate and reliable: researchers, reviewers, editors, policy makers, journalists, and the general public. However, several threats to accuracy exist in the scientific system. We will discuss several of these threats in detail. Then we will present ways to eliminate or at least seriously reduce them. Next, we will sketch how, in a science that prioritizes accuracy and reliability, producers, evaluators, and consumers of scientific evidence can contribute to making science more accurate and thus more practically relevant.

Threats to accuracy

Every research methods class teaches us that there are two general types of errors scientists might make: Type I and Type II errors. Type I errors, *false positives*, occur when we falsely conclude that an effect exists. Type II errors, or *false negatives*, occur when we falsely conclude that no effect exists. Several threats to accuracy have become ingrained in the way we conduct science. They primarily lead to false positives but also to false negatives. Here we discuss some of the major ones.

P-hacking

P-hacking is a set of common practices by which researchers, often inadvertently, increase the likelihood of false positives. One such practice is *optional stopping*: looking at the results while data collection is in progress and making decisions on whether to stop or continue based on this rather than based on some previously determined sample size (Simmons, Nelson, & Simonsohn, 2011). Data snooping while they are being collected is not problematic if it is decided in advance that this will be done and the alpha-level is adjusted accordingly. There exist procedures to do this within a frequentist statistical framework (e.g., Frick, 1998; Lakens, 2014) and within a Bayesian statistical framework (Schönbrodt, Wagenmakers, Zehetleitner, & Perugini, 2017). In fact, such procedures can be efficient, given that data collection can be terminated if a decision about the presence/absence of an effect can be made before the total number of subjects is run.

Other common forms of p-hacking include strategically omitting dependent variables or conditions, and rounding down p-values (John, Loewenstein, & Prelec, 2012). Researchers only report the dependent measure(s) and/or condition(s) for which they obtained a significant effect (possible with optional stopping).

HARKing

HARKing is Hypothesizing After the Results are Known (Kerr, 1998).

It is a common procedure in psychological research. If the data do not turn out as predicted, an explanation for the observed pattern is generated and this post-hoc explanation is transported back in time to a moment supposedly before the data were collected and is now treated as a prediction. This reversal of chronological order has the negative consequence, as Kerr (p. 205) notes, that it can translate false positives into theory. Given that HARKing appears commonplace and has in the past been recommended by prominent researchers as a proper procedure, it is a serious threat to scientific accuracy. One of the reasons that HARKing was encouraged and is used (John et al., 2012) is that it is consistent with our need for coherent stories (see the next section). However, as Kerr (1998) is right to point out, the constraints under which a storyteller operates are different than the ones under which a scientist is expected to operate.

Publication bias

Publication bias manifests itself in an overrepresentation of positive findings in the literature relatively to what would be expected statistically, given the power of these studies. Power in psychology has been estimated to be .35 in psychology experiments, while the proportion of significant results is far higher, namely .95 (Bakker, van Dijk, & Wicherts, 2012). Publication bias is a systemic problem that has been known for decades (e.g., Greenwald, 1975). Part of the explanation lies in our general preference for coherent stories with a resolution, as has been widely documented by researchers in the field of text and discourse comprehension (e.g., Schank & Abelson, 1977; Graesser, Singer, & Trabasso, 1994; van Dijk & Kintsch, 1983). This preference likely also impacts our evaluation of scientific products. We prefer a series of events like "We predicted X and set out to investigate. After many trials and tribulations, we found X" over one such as: "We predicted X and set out to investigate. After many trials and tribulations, we failed to reject the Null Hypothesis." It is like reading the entire *Lord of the Rings* trilogy only to find out, after slogging through three volumes, that Frodo in the end failed to throw the One Ring into Mount Doom! Yet this is the kind of story we will often have to tell if we take our research seriously.

The preference for cognitive closure is shared by authors, reviewers, editors, and consumers of science alike. As a result, papers with coherent stories, and thus positive findings (sometimes obtained via HARKing) are more likely to get submitted, accepted, and published than papers with negative findings. Moreover, replications of published findings

that present null results are much less likely to get published. The result of publication bias is that the literature gives a systematic overestimation of effect sizes and may contain many false positives, rendering meta-analyses problematic as one cannot make a silk purse out of a sow's ear.

Low power

Many psychological experiments, especially those with between-subjects designs, are underpowered. Low power has been a long-standing issue in psychology, observed at least since Cohen (1962). Low power leads to false negatives: researchers concluding that there was no support for a hypothesis, even though they had not collected sufficient data to warrant such a conclusion. Because of publication bias, false negatives, along with true negatives, are often relegated to the file drawer. This unfortunate fact renders false negatives irreversible (Fiedler, Kutzner, & Krueger, 2012). As a result, low power makes our science less accurate than it could be. Fiedler and colleagues are right to note that the issue of false negatives goes beyond false negatives in a statistical sense and also includes false negatives in a theoretical sense: prematurely committing to a particular theoretical account and failing to see alternative explanations.

Perverse incentives

Kerr (1998) notes that the earlier mentioned advice to HARK did not occur before the late 1980s and early 1990s. He attributes this to the increase of publication pressure. And indeed, given that researchers are rewarded (even financially) for the number of publications that they have and the numbers of citations that these publications receive, there exists pressure to cut corners. Ideally, every study or set of studies that is performed should lead to a publication. Given publication bias, however, this means that every study or set of studies should have positive results. HARKing is one avenue for achieving this.

In addition to an increase of pressure to publish in scientific journals, there also is an increased pressure for researchers to get their work out in the media. University communication departments' are charged with writing press releases. Universities need to show their societal relevance. Journalists are highly motivated to tell stories that will give them many readers. And the public likes to read juicy stories or get simple answers to life's questions. These systemic pressures force scientists to cut corners and encourage universities and journalists to exaggerate scientific effects and their implications. In short, perverse incentives permeate the system.

Proposed solutions

How can we remedy this situation and make our science more accurate? Several solutions to this problem have been proposed, some many decades ago. It appears that now they are finally being given serious consideration.

Replication

A reliable science cannot be based on a set of findings from procedures that were used only once and led to publications if $p < .05$. Rather, what matters is not having a single finding but having a procedure that reliably produces certain effects (Dunlap, 1926).

It is often argued that direct replication is not useful because contextual factors vary across experiments. If true, this would inevitably hamstring applications of research in the real world. After all, if a behavior cannot even be expected to reproduce in a lab very similar to the one that observed the original finding, how can it be expected to be reproduced outside the lab? More productive than frivolously raising contextual differences as a barrier to replicability is the proposal to have researchers specify the constraints on the generality of their effects (Simons et al., 2017). Not only does this have the benefit that theories become more explicit (Zwaan, Etz, Lucas, & Donnellan, in press), it also is more honest about applications of the procedure. It may furthermore promote the search for more robust effects.

Direct replication is necessary but not sufficient. Once an effect is robust across labs, which means that a protocol exists that reliably produces a certain finding (Dunlap, 1926), and the constraints on generality are stated, tests of the effect can be extended in incremental steps, e.g., to other populations, to other environments, to other stimuli, to other outcome measures.

Preregistration

Kerr (1998) already observed that replication is an antidote, albeit not effective in all cases, against HARKing. An explanation for research findings obtained via HARKing can become an a priori hypothesis that can be tested in a replication. It is not completely effective because the hypothesis that motivated the original study has been replaced with the post-hoc hypothesis and is thus absent from the literature. The fact that this hypothesis, which seemed a priori plausible, was not supported (assuming the study was sufficiently powered) is relevant information and should therefore be part of the literature.

Four decades before Kerr, de Groot (1956/2014) stressed the importance of separating hypothesis testing from conducting exploratory analyses. Predictions are made

before the data are collected. Forming hypotheses when exploring data is a normal and important part of science. However, the key point is that hypotheses thus formed cannot be tested on the same data from which they have been derived, as this renders the statistics that are reported meaningless. Therefore, HARKing is an acceptable practice only if it is made explicit that the analysis is exploratory and not confirmatory in nature. The hypotheses obtained via exploration can then be tested in a confirmatory manner on a new data set (De Groot, 1969).

To fully separate the confirmatory from the exploratory phase, researchers have proposed using preregistered reports (Chambers, 2017, pp. 174-217). In such reports, researchers state explicitly all the details of the experiment: predictions, sampling plan, design, exclusion criteria, planned analyses. The plan is then reviewed. Among other things, reviewers assess the significance of the research question, the logic, rationale, and plausibility of the hypotheses, the soundness of the methodology, the feasibility of the research plan, whether the level of detail is sufficient to enable an independent replication. Crucially, they assess all of these things *before* any data are collected. As a result, their judgments are not colored by whether findings were significant or not, but purely by the merits of the theoretical conceptualization, the design of the study, and the proposed methods. A positive editorial decision will lead to an in-principle-acceptance of the manuscript and data collection can proceed. Final acceptance occurs when the authors have carried out the research plan as preregistered and when the preregistration is augmented with a results and a discussion section. At the same time, there is room for unregistered exploratory analyses, as long as these are identified as such.

Preregistration combats the threats to accuracy that we have reviewed: HARKing, p-hacking, low power, and publication bias. HARKing is prevented because the hypotheses have to be preregistered. P-hacking is prevented because the sample size and analysis must be preregistered. Low power is prevented because the sampling plan must be approved and power is one of the criteria that is considered. Finally, publication bias is prevented because the journal has committed to publication before the results are known.

Submission of preregistered reports is not always possible. For example, student theses sometimes have to be completed within a limited amount of time (e.g., six months). In such cases, one cannot wait with data collection for the preregistration to be approved by the journal. However, it is possible to write a preregistration and post it in a public place, for example on the Open Science Framework (<https://osf.io>) or on www.aspredicted.org. One

can then provide a link to the preregistration in the manuscript when it is submitted. It is important in such cases that reviewers compare the preregistration to the manuscript for consistency. This is not as airtight as submitting a preregistered report but it is perhaps the only logistically viable form of preregistration in some cases.

Preregistration is not a remedy against fraud. Preregistration can be gamed; for instance, by first collecting the data but pretending that they have not yet been collected. This would require tampering with the timestamps of the data files, which clearly amounts to fraud. Preregistration is designed as a means to keep the vast majority of researchers and editors who operate in good faith from motivated reasoning, something all scientists are likely to engage in (Nosek, Spies, & Motyl, 2012). As such, it is a major tool to make science more accurate.

Large-scale collaborations

Replication is the most convincing when it is accomplished across labs because this means there is a protocol that reliably produces a certain finding, rather than knowledge that might remain partly implicit (Dunlap, 1926), collaborative efforts among labs are required. These efforts are also important because studies in the social and behavioral sciences are often underpowered, especially in fields where data collection is difficult. Fortunately, recent developments will make it easier to acquire larger numbers of subjects for experiments, even with special populations.

One such development is *StudySwap* (<https://osf.io/view/StudySwap/>), a web-based interface facilitating collaboration among labs across the world. Researchers can post brief descriptions of resources (i.e., subject hours) that they either have available for others or need themselves for a particular project. Swapping of resources makes it easier to acquire a large number of subjects for experiments even when one's own participant pool is limited. This might be especially true for experiments that require special populations.

A related initiative that aims to address the need for large-scale collaborations is the *Psychological Science Accelerator* (<https://psysciacc.org>), a global network of laboratories working in psychological science. Researchers can propose a study; this proposal will be reviewed, and once approved data collection will be distributed across the laboratory network. Having multiple labs across the world being responsible for data collection increases the robustness and reliability of science.

Transparency

Open data and code. Most researchers consider data sharing to be a norm of science (Anderson, Martinson, & De Vries, 2007). Nevertheless, sharing data is not common in psychology (Wicherts, Borsboom, Kats, & Molenaar, 2006). Even as recent as 2012, only 38% of researchers who published in APA journals were willing to share their data (Vanpaemel, Vermorgen, Deriemaecker, & Storms, 2015). Making research data publicly available opens them up for scrutiny by others. The fact that other researchers can reanalyze shared data increases the accuracy of scientific results in two ways. First, when reanalysis shows the same outcome as the original analysis, we can be more confident in the outcome of a study. Second, when reanalysis yields a different outcome than the original study, errors might be detected and corrected. In both cases, the availability of open data increases the accuracy of our science..

Open data do not only enable reanalysis but also additional analyses that were not conducted or even conceived by the original authors. This includes (1) using the data for meta-analysis, and (2) using the data for other analyses to address the same research question as the original authors. In a time in which new methods for data analysis are being developed constantly and researchers do not always agree on what the best method is given a particular data set, it is extremely important that data are shared and are available for later analyses. This way, previous collected data will also benefit from new insights.

To make findings optimally reproducible, analysis code should be published along with the paper. This allows other researchers full insight into the analytical steps taken by the original researchers.

As with all measures proposed to promote transparency, some have addressed concerns that open data will make science more vulnerable (see for example, Lewandowsky & Bishop, 2016). It is furthermore important to ensure that public data are compliant with legislation, such as the General Data Protection Regulation. There are obviously also practical issues (data storage) and legal issues to consider.

Open materials. In addition to sharing data, researchers can share the materials they used in their study. In fields such as psycholinguistics and discourse comprehension, sharing materials has been common for decades, but this is not necessarily true for other areas. Sharing materials means that they will be open to scrutiny, allowing others to verify scientific findings, conduct replication studies more easily, and interpret the results in light of the materials used. All these activities will increase the accuracy and reliability of science.

Open peer review. In most cases, an article that is published by a scientific journal deviates from the manuscript that was originally submitted for publication due to the peer review process. This process is far from transparent, while transparency could provide useful insights into the decisions that lead to publication. Science would benefit from making reviews open upon publication. Open reviews will make reviewers more accountable for their writings, thereby promoting constructive feedback and improving the quality of reviews. Signed open reviews allow for the possibility to give reviewers credit for their work. Writing reviews is part of being an academic but researchers are often not credited for this work. The publication of reviews alongside research articles will make this easier. In fact, written reviews could become part of a researchers academic output.

Open access. An important prerequisite for science to influence policies is that policy makers and taxpayers in general have access to scientific findings. The reality is that most research articles are behind paywalls. However, recent developments will make it easier for researchers, policy makers, and the general public to gain access to scientific findings.

One such development is the increase in opportunities to open access publishing. Traditionally, research findings are published in scientific journals that can be accessed by anyone with a journal subscription. The introduction of the Internet allowed for an alternative way of disseminating research findings by making articles freely available online. This so called open access publishing started in the early 1990s after which the number of open access journals grew very rapidly. From the year 2000 the number of open access journals grew annually with 18% and the number of open access publications even with 30% (Laakso, Welling, Bukvova, Nyman, Björk, & Hedlund, 2011). This is far more than the annual growth of 3.5% of published articles in general (Ware & Mabe, 2009). This rapid growth of open access journals and articles is also displayed by the fact that the Directory of Open Access Journals (<https://doaj.org>) registered almost 5000 journals in 2009, and to date registered over 11,500 journals, and over 3,000,000 articles. The increased availability of freely accessible scientific findings allow for other researchers, policy makers, journalists, and the general public to be better informed.

A second development that will make the accessibility and dissemination of scientific findings easier is the availability in open access repositories for posting pre-prints. A preprint is the latest draft of a manuscript before it gets accepted by a peer-reviewed journal. *PsyArXiv* (<https://psyarxiv.com>) is a free preprint service for the psychological

sciences, including linguistics. Researchers can post any work they are the legal owner of on PsyArXiv. This might be the latest prepublication version of an article or the published manuscript itself, depending on the journal that accepted or published the paper (see the SHERPA/RoMEO website for more information to review the policies of journals, <http://www.sherpa.ac.uk/romeo/index.php>). Researchers might also post all their file-drawer work on preprint servers like PsyArXiv. This includes studies with null results that are hard to get published in scientific journals but can be extremely valuable. Research posted on PsyArXiv is freely available to everyone and will be indexed by research databases.

From lab to society in 12 steps

The discussion so far enables us to envision a path, laid on a bed of accurate and reliable science, that leads from the lab to society. We hasten to add that we are using an accuracy-driven rather than a pragmatic model. Obviously, the envisioned sequence of actions involves large groups of researchers and a large number of steps. It thus requires large investments in time and money that researchers and funding agencies might currently be uncomfortable with or be ill-equipped to handle. Nevertheless, we believe this is the only viable way forward if we value accuracy over expedience.

Here are the 12 steps from lab to society we envision in the ideal case. We should note that it is possible to accomplish some steps simultaneously. For example, steps 1-5 could be accomplished in one fell swoop using Psychological Science Accelerator.

- (1) Researcher derives prediction from theory.
- (2) Researcher designs and preregisters study.
- (3) Prediction is supported by the data.
- (4) Researcher specifies constraints on generality (COG) of effect.
- (5) Effect is replicated by other labs in preregistered direct replications.
- (6) Extensions of the experiment are conducted (to examine the COG).
- (7) The effect is extended.
- (8) A meta-analysis is conducted on all preregistered studies.
- (9) A preregistered field study is attempted that is within the current COG and informed by the meta-analysis.
- (10) The field study yields effects consistent with the hypothesis.
- (11) Further preregistered field studies are launched.

(12) The procedure is implemented.

Of course, it is possible that the preregistered prediction is not supported by the data. In that case, the researcher moves back to step 1. There are several other stages at which the researcher may need to return to an earlier stage. Figure 1 captures the dynamics of the entire process.

Figure 1. 12 steps from prediction to implementation.

Several robust effects from the cognitive psychology literature that have made an impact on society have undergone a process roughly similar to this. For example, the *testing effect* and the *spacing effect* have been leveraged into teaching tools (Sense, Behrens, Meijer, & van Rijn, 2016; Van Rijn, van Maanen, & van Woudenberg, 2009). As another example, signal-detection theory has led to insights regarding the relative diagnostic values of sequential versus simultaneous lineups in eyewitness testimony. Whereas sequential lineups were considered superior to simultaneous lineups, signal-detection analysis has

shown in lab studies that that the reverse is actually more likely to be the case (e.g., Carlson & Carlson, 2014; Dobolyi & Dodson, 2013; Mickes, Flowe, & Wixted, 2012). After this effect was shown to be replicable in the lab, a field study was conducted at a Houston police station, which yielded similar results to the lab results (Wixted, Mickes, Dunn, Clark, & Wells, 2016). This set of lab and field studies also established that eyewitness memory is not as unreliable as previously thought. Witnesses' confidence about their identification is a good predictor of their accuracy. These studies show a progression from (1) a strong psychological theory (signal detection) to (2) a reliable set of lab findings, to (3) a field study providing converging evidence with the lab results. Research on reading comprehension might be poised to have similar societal impact (Kendeou, McMaster, & Christ, 2016).

The 12-step procedure holds the promise of moving from a finding established in the lab to application in some practical domain in small incremental steps. It should be noted that COGs are not etched in stone. In some cases, a finding may be more generalizable than was perhaps assumed. For example, for some findings in the areas of attention, memory, and language it does not matter whether these experiments are conducted in the lab or online and even if subjects have recently participated in the same experiment or a slightly different version of it; effect sizes remain unaltered (Zwaan, Pecher, et al., in press), even though this was not necessarily expected. Given that the original studies in which these effects, such as the serial position effect, were obtained date from before the era of online experiments, extension to online environments was not considered at the time and the expectation might have been that findings would not generalize to these environments.

We envision this 12-step procedure to be maximally transparent and accessible, with open data and code, open materials, and open reviews.

How can agents in the scientific system contribute to a more accurate science?

While most of the 12 steps are performed by researchers, evaluators, non-researcher communicators, and consumers of scientific evidence can also contribute to a more accurate science. We list them here.

- (a) Researchers: use preregistration, conduct replications, collaborate (let other people replicate your studies), make your data, code, and materials available. Don't communicate findings to a broader audience until they have been replicated across labs.

- (b) Reviewers/editors: accept pre-registered studies, publish replications, require open data/materials, and allow for open reviews.
- (c) Policy makers: support open science practices and accept that it takes a large number of studies to get from the lab to the real world.
- (d) Funding agencies: value quality of research rather than quantity, value replications, require open access publications, open data, and open materials. Accept that it takes a large number of studies to get from the lab to society. This will likely mean longer funding schemes.
- (e) Journalists: accept and communicate that research can be open-ended. Require transparency.
- (f) The general public: accept that research can be open-ended.

Summary and Conclusion

We started out with the simple idea that in order to have practical relevance, science should be accurate. We then identified several threats to accuracy: p-hacking, HARKing, publication bias, low power, and perverse incentives. Next, we discussed ways to combat these threats: replication, preregistration, multi-lab collaborations, and transparency. Finally, we outlined a 12-step procedure that moves science in a reliable, accurate, and transparent manner from the lab to society. We finally showed how the various producers and consumers of scientific evidence can contribute to this process.

References

- Anderson, M.S., Martinson, B.C., & Vries, R. de (2007). Normative dissonance in science: Results from a national survey of U.S. scientists. *Journal of Empirical Research on Human Research Ethics*, 2, 3-14.
- Bakker, M., van Dijk, A., & Wicherts, J. M. (2012). The rules of the game called psychological science. *Perspectives on Psychological Science*, 7, 543-554. doi: 10.1177/1745691612459060
- Carlson, C.A., & Carlson, M.A. (2014). An evaluation of perpetrator distinctiveness, weapon presence, and lineup presentation using ROC analysis. *Journal of Applied Research in Memory and Cognition*, 3, 45-53.
- Chambers, C. (2017). *The seven deadly sins of psychology: A manifesto for reforming the culture of scientific practice*. Princeton, NJ: Princeton University Press.
- Cohen (1962). The statistical power of abnormal-social psychological research: a review. *Journal of Abnormal and Social Psychology*, 65, 145-153.
- De Groot, A.D. (1969). *Methodology: Foundations of inference and research in the Behavioral sciences*. The Hague: Mouton.
- De Groot, A.D., Wagenmakers, E.-J., Borsboom, D., Verhagen, J., Kievit, R., Bakker, M., ... van der Maas, H.L.J. (2014). The meaning of "significance" for different types of research. *Acta Psychologica*, 148, 188-194. doi: 10.1016/j.actpsy.2014.02.001
- Dobolyi, D.G., & Dodson, C.S. (2013). Eyewitness confidence in simultaneous and sequential lineups: A criterion shift account for sequential mistaken identification overconfidence. *Journal of Experimental Psychology - Applied*, 19, 345-357.
- Dunlap, K. (1926). The experimental methods of psychology. In: C. Murchison (Ed.) *Psychologies of 1925* (pp. 331-353). Worcester: Clark University Press.

- Fiedler, K., Kutzner, F., & Krueger, J.I. (2012). The long way from α -error control to validity proper: Problems with a short-sighted false-positive debate. *Perspectives on Psychological Science*, 7(6), 661-669. doi: 10.1177/1745691612462587
- Frick, R.W. (1998). A better stopping rule for conventional statistical tests. *Behavior Research Methods, Instruments, & Computers*, 30, 690-697.
- Graesser, A.C., Singer, M., & Trabasso, T. (1994). Constructing inferences during narrative text comprehension. *Psychological Review*, 101(3), 371-395.
- Greenwald, A.G. (1975). Significance, nonsignificance, and interpretation of an ESP experiment. *Journal of Experimental Social Psychology*, 11, 180-191.
- John, L.K., Loewenstein, G., & Prelec, D. (2012). Measuring the prevalence of questionable research practices with incentives for truth telling. *Psychological Science*, 23, 524-532.
- Kendeou, R., McMaster, K., & Christ, T.L. (2016). Reading comprehension: Core components and processes. *Policy Insights from the Behavioral and Brain Sciences*, 3, 62-69.
- Kerr, N.L. (1998). HARKing: Hypothesizing after the results are known. *Personality and Social Psychology Review*, 2, 196-217.
- Laakso, M., Welling, P., Bukvova, H., Nyman, L., Björk B-C., & Hedlund, T. (2011). The development of open access journal publishing from 1993 to 2009. *PLoS ONE*, 6(6): e20961. doi: 10.1371/journal.pone.0020961
- Lakens, D. (2014). Performing high-powered studies efficiently with sequential analyses. *European Journal of Social Psychology*, 44(7), 701-710.
- Lewandowsky, S., & Bishop, D. (2016). Research integrity: Don't let transparency damage science. *Nature*, 529, 459-461. doi: 10.1038/529459a

- Mickes L., Flowe, H.D., & Wixted, J.T. (2012) Receiver operating characteristic analysis of eyewitness memory: Comparing the diagnostic accuracy of simultaneous and sequential lineups. *Journal of Experimental Psychology - Applied*, 18, 361-376.
- Nosek, B.A., Spies, J.R., & Motyl, M. (2012). Scientific utopia: II. Restructuring incentives and practices to promote truth over publishability. *Perspectives on Psychological Science*, 7, 615-631.
- Schank, R.C., & Abelson, R.P. (1977). *Scripts, plans, goals, and understanding: An inquiry into human knowledge structures*. New Jersey: Erlbaum. ISBN 0-470-99033-3.
- Schönbrodt, F.D., Wagenmakers, E.J., Zehetleitner, M., & Perugini, M. (2017). Sequential hypothesis testing with bayes factors: Efficiently testing mean differences. *Psychological Methods*, 22(2), 322-339. doi: 10.1037/met0000061
- Sense, F., Behrens, F., Meijer, R.R., & van Rijn, H. (2016). An individual's rate of forgetting is stable over time but differs across materials. *Topics in Cognitive Science*, 8, 305-321. doi: 10.1111/tops.12183
- 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(11), 1359-1366. doi: 10.1177/0956797611417632
- Simons, D.J., Shoda, Y., & Lindsay, D.S. (2017). Constraints on generality (COG): A proposed addition to all empirical papers. *Perspectives on Psychological Science*, 12(6), 1123-1128. DOI: 10.1177/1745691617708630
- Van Dijk, T.A., & Kintsch, W. (1983). *Strategies of discourse comprehension*. New York: Academic Press.

- Vanpaemel, W., Vermorgen, M., Deriemaecker, L., & Storms, G. (2015). Are we wasting a good crisis? The availability of psychological research data after the storm. *Collabra* 1, 1–5. doi: 10.1525/collabra.13
- Van Rijn, H., van Maanen, L., & van Woudenberg, M. (2009). Passing the test: Improving learning gains by balancing spacing and testing effects. In A. Howes, D. Peebles, & R.P. Cooper (Eds.), *Proceedings of the 9th International Conference on Cognitive Modeling* (pp. 110–115). UK: Manchester.
- Ware, M., & Mabe, M. (2009). The stm report: An overview of scientific and scholarly journals publishing. International Association of Scientific, Technical and Medical Publishers. Available: http://www.stm-assoc.org/2009_10_13_MWC_STM_Report.pdf. Accessed 2018 May 28.
- Wicherts, J.M., Borsboom, D., Kats, J., & Molenaar, D. (2006). The poor availability of psychological research data for reanalysis. *American Psychologist*, 61, 726–728. doi: 10.1037/0003-066X.61.7.726
- Wixted, J.T., Mickes, L., Dunn, J.C., Clark, S.E., & Wells, W. (2016). Estimating the reliability of eyewitness identifications from police lineups. *Proceedings of the National Academy of Sciences*, 113, 304–309.
- Zwaan, R.A., Etz, A., Lucas, R.E., & Donnellan, M.B. (in press). Making replication mainstream. *Behavioral and Brain Sciences*. doi: 10.1017/S0140525X17001972
- Zwaan, R.A., Pecher, D., Paolacci, G. et al. (in press). Participant nonnaiveté and the reproducibility of cognitive psychology. *Psychonomic Bulletin and Review*. <https://doi.org/10.3758/s13423-017-1348-y>