



Pre-registration in social psychology—A discussion and suggested template



Anna Elisabeth van 't Veer^{a,b,*}, Roger Giner-Sorolla^c

^a Centre for Child and Family Studies, Leiden University, Leiden, The Netherlands

^b Department of Social Psychology, Tilburg Institute for Behavioral Economics Research, Tilburg University, Tilburg, The Netherlands

^c University of Kent, Canterbury, Kent, CT2 7NP, United Kingdom

HIGHLIGHTS

- We suggest elements for pre-registration in social psychology.
- We offer initial guidelines to facilitate the process of pre-registration.
- We provide a brief history of pre-registration in medicine and psychology.
- Two models of pre-registration are outlined: reviewed and unreviewed pre-registration.
- Benefits and drawbacks of both models are discussed.

ARTICLE INFO

Keywords:

Pre-registration
Research methods
Solid science
Reviewed pre-registration (RPR)
Unreviewed pre-registration (UPR)

ABSTRACT

Pre-registration of studies before they are conducted has recently become more feasible for researchers, and is encouraged by an increasing number of journals. However, because the practice of pre-registration is relatively new to psychological science, specific guidelines for the content of registrations are still in a formative stage. After giving a brief history of pre-registration in medical and psychological research, we outline two different models that can be applied—reviewed and unreviewed pre-registration—and discuss the advantages of each model to science as a whole and to the individual scientist, as well as some of their drawbacks and limitations. Finally, we present and justify a proposed standard template that can facilitate pre-registration. Researchers can use the template before and during the editorial process to meet article requirements and enhance the robustness of their scholarly efforts.

© 2016 Elsevier Inc. All rights reserved.

1. Introduction

In pre-registration, researchers describe their hypotheses, methods, and analyses before a piece of research is conducted, in a way that can be externally verified. Recently, a growing interest in transparency, reproducibility, and reducing publication bias has led scientists and journals to become more interested in the pre-registration of research. At the same time, pre-registration has been greatly facilitated by online tools that allow for public timestamping of plans and confirmatory predictions. This process can benefit both scientists and science; for example, when a researcher describes ahead of time which of several possible data analyses will be used, the resulting inferential statistics become more clearly interpretable, and the credibility of the claim increases. In this paper we discuss the advantages and disadvantages of pre-

registration. We arrive at some initial suggestions for how our own field of experimental social psychology, and other related areas, can implement this practice, and we differentiate two pre-registration models—*reviewed* and *unreviewed*—for doing so. Finally, we propose a flexible template for pre-registrations in social psychological research, for the benefit of creators as well as evaluators of pre-registered research.

Many aspects of pre-registration are still being worked out. To understand how and why research pre-registration has evolved, it is useful to know its general history. This history has mostly taken place in medical research.

2. Pre-registration in medical research

Pre-registration began, not as a check on the outcomes of research, but rather to help the research get done in the first place. Starting in the 1960's, limited registries of clinical trials in medicine were made available in several countries, to help recruit patients with the appropriate diagnosis (Dickerson & Rennie, 2003). Requirements to disclose the

* Corresponding author at: Centre for Child and Family Studies, PO 9555, 2300 RB, Leiden, The Netherlands.

E-mail addresses: a.e.van.t.veer@fsw.leidenuniv.nl, vantveer@gmail.com (A.E. van 't Veer), R.S.Giner-Sorolla@kent.ac.uk (R. Giner-Sorolla).

results of the eventual study were few. However, from the 1980s onward, investigations showed evidence of publication bias. That is, trials that yielded significant rather than nonsignificant (or “null”) results were substantially more likely to be published at all (Easterbrook, Berlin, Gopalan & Matthews, 1991; Simes, 1986) or in a timely manner (Ioannidis, 1998; Stern & Simes, 1997). Demonstrations of publication bias in specific medical literatures (e.g., Melander, Ahlqvist-Rastad, Meijer & Beermann, 2003; Turner, Matthews, Linardatos, Tell & Rosenthal, 2008), and of low replication rates of published medical research in registered clinical trials (e.g., Begley & Ellis, 2012; Mullane & Williams, 2013; Prinz, Schlange & Asadullah, 2011), led to calls for greater openness in registration.

The development of the Internet has allowed governmental and professional bodies to create accessible, centralized clinical trial registries. However, official oversight of their relation to scientific reporting did not begin until the mid-2000s. For example, in 2007, a new law in the United States required submission of results of trials involving FDA-approved treatments (Food and Drug Administration Amendments Act of 2007), and the World Medical Association's Declaration of Helsinki (2008) supported the principle that all results, regardless of outcome, should be made available. Efforts to improve the openness of registries have continued; the latest European regulation (Clinical trials-Regulation EU No 536/2014) requires reporting of results for all registered trials, as does a rule proposed recently in the US (Clinical Trials Registration and Results Submission, 2014). These recent developments seem to contribute to less selective reporting of medical research; preliminary evidence shows that the percentage of positive published results in one area of research dropped from 57% to 8% concurrent with the requirement to pre-register at clinicaltrials.gov (Kaplan & Irvin, 2015). However, a recent project comparing the specifics of pre-registered clinical trials in medicine to their published versions has found most articles to still contain some form of “outcome switching,” or failure to fully report the pre-specified analytic plan (Mahtani, February 5, 2016).

3. Pre-registration in psychological research

As in medical research, some psychologists and neuroscientists propose more pre-registration to resolve worries about the representativeness of research reports in the published literature (e.g., Wagenmakers, Wetzels, Borsboom, van der Maas, & Kievit, 2012). An open letter to the Guardian newspaper in June 2013 signed by 80 academics in psychology and neuroscience called for journals to adopt pre-registration as an option (Chambers & Munafò, 2013). Reflecting this development, psychology and neuroscience journals have recently shown increased willingness to adopt “registered reports” as a submission category (e.g., Cortex, Perspectives on Psychological Science), to designate a special issue for articles featuring pre-registered research (e.g., Journal of Experimental Social Psychology, Social Psychology), to implement a system of badges designating pre-registered research (see Eich, 2014; “Badges to Acknowledge Open Practices,” 2013), or, even more boldly, to dedicate a new journal in social psychology to such research (i.e., Comprehensive Results in Social Psychology, see “Challenging traditions in research reporting,” 2014; Jonas & Cesario, 2015). Online platforms for pre-registration include the Open Science Framework (OSF), which has recently offered a thousand prizes of \$1000 each to research teams in a pre-registration challenge (<https://cos.io/prereg/>), and the AsPredicted platform (<https://aspredicted.org/>). Additionally, pre-registration has been a requirement for most of the organized replication initiatives in psychology (e.g., Open Science Collaboration, 2012; Klein et al., 2014).

4. Two models of pre-registration and their uses

Two types of pre-registration are beginning to be used in psychology and related fields. The first type requires that studies undergo peer

review on the basis of their theoretical grounds and methods before data are collected. We refer to this model as *reviewed pre-registration* (RPR), which has also been called a “Registered Report” (Chambers, Feredoes, Muthukumaraswamy, & Etchells, 2014; Nosek & Lakens, 2014). This type of research is conducted with the expectation that, if the plan is carefully followed, the report will be published regardless of the outcome. By approving the registration, the peer review process grants *In Principle Acceptance* (IPA). During submission of the pre-registration, reviewers' suggested amendments to the planned study can still be incorporated before the study is run. Ideally, cooperation occurs between reviewers and researchers, to ensure that the most suited method for the research question is used. This type of pre-registration has been adopted, for example, by Cortex and *Comprehensive Results in Social Psychology* (for a continually updated list of journals see <https://osf.io/8mpji/wiki/home/>).

The second type of pre-registration, which we refer to as *unreviewed pre-registration* (UPR), does not involve reviewers before the data is collected. Authors write out and time-stamp their full plan before conducting the study in order to be able to refer back to it later. This self-registration allows authors to conduct research more or less as usual. Unreviewed pre-registration thus leads to a review process very similar to the standard model, but with the reassurance that the authors' reports of method and analytic procedures have been specified *a priori*.

We recognize that research papers can incorporate multiple forms of registration and non-registration. Some recent journal editorials, for example, have expressed a willingness to encourage authors to follow up non-registered findings that fall short of robustness with a registered replication (Giner-Sorolla, 2016; Vazire, 2015; see also Bostyn & Roets, 2016 for an example of a paper combining unregistered and registered studies). Authors themselves can take the initiative to follow up unregistered exploratory research with registered confirmatory research following either model. It is also possible to start with an unreviewed pre-registered study and extend the research with a reviewed registration, so that an initial proof of concept is followed by an extension that benefits from peer review and in principle acceptance. Therefore, these two models should not be seen as mutually exclusive. Rather, each contributes to different priorities in the research cycle.

5. Benefits to science

Can these developments benefit our science on the whole? Although any definitive conclusion on the basis of a few years' experience is premature, some positive outcomes can reasonably be expected.

5.1. Prioritizing theory and method

First of all, pre-registering studies puts emphasis on developing sound theory and methods—the very elements specified in the pre-registration—rather than on results. Positively valuing strong theory and methods, rather than merely accepting results that meet a certain standard of statistical consistency, has been suggested as a way for the field of psychology to become more confident in both positive and negative results when conducting and publishing research (LeBel & Peters, 2011; Murayama, Pekrun & Fiedler, 2013). We further suggest that re-emphasizing theory and methods, and moving away from the superficial appearance of results as the main criterion for judging research, is a common thread that runs through all other benefits that pre-registration holds for our science. For example, it is not enough simply to point to a series of significant study results at $p < .05$, without considering the full space of analytic decisions that were possible within the studies' theoretical constraints (Wasserstein & Lazar, *in press*), and pre-registration makes this full space more transparent.

From this viewpoint, pre-registration is particularly useful for studies that fall within a certain range on a spectrum of theoretical specification. At one extreme of this spectrum, we see studies that test

hypotheses derived from strong, pure tests of one or more theories. Such studies specify an outcome that would be disconfirming, without relying on unstated auxiliary assumptions (cf. Meehl, 1967, 1990). Strong tests like these would ideally have no need for pre-registration of hypotheses, because predictions would follow logically from the theory. However, it can be argued that even if the theory is crystallized, the methods used to test it could still benefit from clear *a priori* specification. At the other extreme, studies that start without any theory have no need to pre-register hypotheses either: any interpretation is by definition *post hoc*. In psychology, the typical study tends to fall in between these two extremes, and it is exactly this middle ground where pre-registration is beneficial, due to its ability to clarify main and auxiliary hypotheses, and specify solid methods.

5.2. Distinguishing confirmatory from exploratory research

Another reason to adopt pre-registration is to more clearly distinguish between exploratory and confirmatory tests. Ideally, research begins with an exploratory phase in which hypotheses and methods are tested without much prior evidence. It then follows through to a confirmatory phase in which already-observed hypotheses and methods are replicated to ensure the validity of initial findings. After this, these two approaches continue to interweave as research progresses (Tukey, 1980). However, in social psychology, it has not always been clear whether research described as confirmatory has indeed been specified *a priori* (Kerr, 1998). Many theories in psychology allow for multiple predictions (e.g., cognitive dissonance, theory of planned behavior), while many studies leave room for multiple interpretations of phenomena, allowing for misidentification of random patterns as meaningful (Gelman & Loken, 2014). The ambiguity surrounding exploratory research being presented as confirmatory may be due to perceived incentives for telling a clear and clean story in which hypotheses fit the findings (Giner-Sorolla, 2012; for evidence of this practice in organizational research, see Bosco, Aguinis, Field, Pierce, & Dalton, 2015). Although not mentioned by Kerr (1998), pre-registration presents itself as an appropriate solution for HARKing (hypothesizing after the results are known), because it limits the ability to covertly alter hypotheses and analyses. Although these might change in the course of research (and are allowed to!), the change is open for all to see. Pre-registration can thereby protect against the pitfalls of confusing exploratory and confirmatory phases,¹ such as drawing overly firm conclusions from a single exploratory study.

5.3. Reducing publication bias

Because pre-registration shifts emphasis in research from perfect results towards theory and method, and especially when reviewed pre-registration grants in principle acceptance regardless of eventual results, we can expect pre-registration to reduce a particular type of selective publication bias. It is true that the term “selective” can be applied to a great number of processes; there is selectivity in the topics chosen for investigation, in the methods that are used, and in the non-publication of entire lines of research that did not yield interpretable results. However, one particular kind of publication bias has been central in recent debates on science: the kind that happens when, for a given hypothesis, studies that do not yield a significant result in favor of the preferred (or any) conclusion are conducted but never published, while similar studies with positive outcomes are published (e.g., Begg & Berlin, 1988; Fanelli, 2010; Ferguson & Heene, 2012).

The actual impact of this practice on the field is a topic of debate. While some do not necessarily see it as a problem for interpreting past directional findings (e.g., Fabrigar & Wegner, 2016; Murayama et al. 2013), others criticize publication bias as undermining the real

and perceived integrity of findings (Giner-Sorolla, 2016; Vazire, 2015). Regardless of one's view on publication bias, pre-registration is of potential interest to those who want to present their results with some kind of reassurance that a full report of a given line of research is represented. Therefore, just as direct replication of research by independent labs has been promoted as a way to reduce the effects of publication bias (LeBel & Peters, 2011), pre-registration might eventually play a role here as well.

5.4. Reducing reporting bias

Bias in reporting analyses within a single study is another practice that has come under scrutiny (Carp, 2012; Dwan et al., 2008; Simmons, Nelson, & Simonsohn, 2011). For instance, because choices of statistical analysis in psychology are often subjective, there is a temptation to choose, out of many possible analyses, the one that gives the most consistent or significant results, and to dismiss as “exploratory” or “flawed” those elements of a study that fail to achieve the desired effect (LeBel & Peters, 2011). Concern about reporting bias has led to several recent practices proposed to enhance transparency in psychology. For instance, psychologists have been urged to disclose their full study design and to report analyses in a more complete way; specific protocols for doing so have been proposed (LeBel et al., 2013; Simmons et al., 2011) and various journals have adopted guidelines encouraging greater disclosure (e.g., Psychological Science, Eich, 2014; Personality and Social Psychology Bulletin, following Funder, et al., 2014; Journal of Experimental Social Psychology, Giner-Sorolla, 2016). In line with these steps, a pre-registration can serve as a verifiable means of full disclosure. And unlike standard disclosure statements, which cover only measures, manipulations, data collection, and exclusion rules, pre-registration requires the statement of all methods and data analytic strategies *a priori*.

We believe, then, that use of pre-registration does help create a more robust and credible science by strengthening emphasis on theory and methods; by increasing confidence that research reported as confirmatory is just that; and by increasing complete reporting of research lines and research studies. All of these benefits can also increase the efficient working of individual scientists if adopted on a large scale; for example, if further research uses power analyses based on a literature that includes null findings as well as positive findings, studies can be planned more realistically, reducing Type II error (false negative findings). Furthermore, there are also reasons to see direct benefits to individual researchers arising from pre-registration.

6. Why pre-registration is beneficial for individual researchers

6.1. Added review and input

At an early stage, the prospect of re-registration encourages individuals to thoroughly consider all the steps they will take in the research process. At the most basic level, writing out a plan is likely to entail more careful reasoning, especially if done knowing that the plan will be seen by an unknown audience, as research on the effects of accountability suggests (e.g. Lerner & Tetlock, 1999). Having a pre-registration also encourages all members of a research team to scrutinize the specifics of a plan before it is posted in their name. Keeping the registration plan in mind while the research is done will also ensure that each deviation from the plan has a good justification, and make miscommunication between team members less likely.

Planning for the pre-registration to be reviewed (in RPR) adds an extra layer of scrutiny to the project, so that any flaws can be corrected early on. In our own experience with RPR, suggestions by reviewers reflected a valuable collaborative effort to ensure the best possible test of the hypotheses (e.g., van 't Veer, Gallucci, Stel, & van Beest, 2015). Currently, many manuscripts are rejected by journals because of method flaws in studies. Manuscripts then often spend a long half-life

¹ In fact, it has been argued that inferential statistics should only apply to confirmatory analyses in which clear *a priori* criteria are applied (de Groot, 1956/2014).

bouncing from journal to journal, until they either add new and better evidence, or they find a combination of editor and reviewers willing to overlook the flaws (Nosek & Bar-Anan, 2012). The reviewed pre-registration model makes better use of reviewers' and editors' critical efforts, which are currently applied too late to do anything but improve future studies. In the large scale and long term, it might in fact reduce pressure on the reviewing system. Journals would no longer be haunted as much by the resubmitted ghosts of methodologically inadequate manuscripts, if the authors had the chance to do it right the first time.

6.2. Skill and chance

In its focus on validating theory and methods over results, the practice of reviewed pre-registration might especially benefit students and early career researchers. Often, doctoral students feel that they have relatively little time to achieve results, and are doubly dependent on chance: once from uncertainty about whether a plausible hypothesis actually reflects reality, and again from the vagaries of inferential statistics, which may yield an uncertain result even if the underlying idea is true (Hung, O'Neill, Bauer & Kohne, 1997). If a carefully planned study is deemed worthy of publication by the reviewers no matter how the results turn out, then this can be a way to show one's strength in theorizing, conceiving, and implementing quality research without being dependent on the results coming out a certain way. Importantly, and especially for people at a delicate career stage, pre-registration encourages a shift in incentive from quantity to quality, and from the content of results to the process by which they were produced.

6.3. Faster dissemination

With more studies being registered online, the chance of finding out that someone else has been working on a given topic are higher. A researcher working on this topic would be able to inform new research with the existing findings. For instance, rather than getting information only from personal communication (e.g., at conferences, seminars), it would become easier to find out whether a specific method had led to a dead end or whether there were unforeseen indications in the data that need to be followed up. To the extent that pre-registrations are openly available online, it is easier to see how a research project would more quickly add to existing knowledge cumulatively. In many cases, knowing that other labs are working on the same topic can lead researchers to join forces and share resources. Alternatively, researchers may also see a benefit in keeping the pre-registration private until publication is assured, to avoid idea theft. We discuss this possibility in more detail in the section on drawbacks.

6.4. Help with specific research types

Pre-registration can help researchers, as individuals and teams, to carry out specific types of research endeavor in which it is useful to agree on procedures ahead of time. One such type is *adversarial collaboration* between scholars with opposing views, sometimes proposed in psychology and related fields (e.g., Kahneman, 2003; Nier & Campbell, 2013; Rakow, Thompson, Ball, & Markovits, 2014), although published empirical examples are as yet few (e.g., Bateman, Kahneman, Munro, Starmer & Sugden, 2005; Matzke et al., 2015; Mellers, Hertwig & Kahneman, 2001). Mutually reviewed and agreed upon pre-registration is a logical structure for such collaborations; both sides can be satisfied that the other side is following an acceptable protocol. Registration, in fact, is also suited to other models of semi- or non-adversarial collaboration, such as when different theorists compete to develop and test interventions (e.g., Lai et al., 2014, in which procedures to reduce implicit prejudice were tested against each other), or when groups of researchers at different sites agree to follow and disseminate protocols for cooperative work (e.g., the Many Labs projects including Ebersole et al., 2016-in this issue). In other instances, such

as when a single research lab tests opposing predictions derived from different theories, pre-registration can likewise enhance confidence in the outcomes, which in turn has great value in informing further efforts.

Beside its benefits for collaboration, pre-registration has been seen as an essential procedure when replicating research across labs. Pre-registration plays a lead role in quality standards recently proposed for conducting close replications (e.g., Brandt et al., 2014), and as mentioned previously, has been an integral part of several replication initiatives. Rather than deriving an original methodology from a theory, close replication takes a published study as an *a priori* model, with the aim to confirm or disconfirm the underlying idea. Especially when the study's original authors are involved as collaborators or reviewers of the plan, it is useful to specify ahead of time exactly how the previous research will be replicated in a new context. This is especially true when considering that the replication may need to deviate from the previous study's literal procedures in order to create similar psychological states, given changes in context, culture, materials, or time since the original research was done.

7. Drawbacks of pre-registration

Pre-registration may also have several drawbacks and limitations. Some of these have been discussed already (e.g., in a recent editorial at AIMS Neuroscience, Chambers et al., 2014). Here, we articulate commonly voiced concerns, distinguishing between reviewed and unreviewed pre-registration.

7.1. More work?

An often-heard critique of pre-registration is that it requires more work for both the authors and reviewers. What are the likely sources of added effort? For authors, an unreviewed pre-registration requires them to format, review and time-stamp a plan that may already be present (e.g., in the design of the procedure or in an application for ethical approval). In many cases a more detailed analysis plan than usual will have to be written down, which will take considerably more effort.

There is also some added effort in the review process for both the reviewed (RPR) and unreviewed (UPR) models. For a reviewed pre-registration of a single study there is an additional review round before there are any data. Moreover, editors or reviewers should feel a need to check the submitted article against the pre-registration(s). Thus, pre-registration shifts some of the load from one phase of review to another, while increasing other work requirements. Specifically, we think of a standard review process as evaluating (a) rationale for research (b) methods of research (c) data analysis and (d) conclusions from the data. For RPR, the pre-registration review includes tasks a–c, while the review of the final publication needs to check if a–c were implemented correctly by the plan, review the validity of any exploratory analyses added, and evaluate whether the conclusions are justified by the data (task d). Extra work in both models, then, comes from the reassessment and checks of steps a–c.

It is true that it is easier to submit a pre-registration under the RPR model than to actually run and write up research. We can expect that a lowered barrier to submission will increase the number of submissions, in turn increasing the resources needed for initial review. At the same time, the review process as a whole can also become shorter under RPR. Revisions become simpler when the methods and analyses have gone through review beforehand, because additional studies and analyses are less likely to be required during the revision. For example, if reviewers want to see an additional condition they can suggest this before the data are collected instead of afterwards. Also, in a field with many opportunities for reviewed pre-registration, a given paper may also pass through fewer journals before it finds acceptance, as its methods would be improved from the start (see also Chambers et al., 2014). Another hidden savings, perhaps, would come from eliminating the effort that editors and reviewers spend in trying to figure out

whether authors might be using selective analysis to cover up less than perfect results.

If multiple sequential studies are reported in one manuscript, any work involved in reviewed pre-registration will multiply considerably, both in total hours and in the lengthening of the whole process. Reviewers, editors and authors would have to engage in multiple rounds of comment, revision and assessment, as each new round of the research would have to be peer-reviewed and approved in principle. Thus, reviewed pre-registration seems most effective when applied to a single-study paper, to a series of studies that do not depend crucially on each other's outcomes, or to a plan for multiple studies that includes alternate plans for later studies depending on earlier outcomes. Unreviewed pre-registration, by contrast, is particularly suited for a report written at the end of a series of studies that incorporated procedural changes and extensions one at a time, based on previous outcomes in the sequence. Registration may have been present from the beginning, or only for later studies. The important point is to make clear across and within studies which research elements are confirmatory and which are exploratory.

7.2. Too restrictive?

Another often-heard critique of pre-registration is that it leaves no room for exploration. However, we emphasize that authors *should* report exploratory analyses and *post hoc* interpretations, as long as these outcomes are clearly labeled as such. Exploration, after all, is what motivates the scientific endeavor and drives progress. A valuable part of exploration often comes about only after confirmatory evidence is seen. For example, an unforeseen factor may be the best explanation for why an effect did not generalize to another setting, or why an effect emerges in one measure but not another. To allow flexibility together with transparency, we strongly recommend that pre-registered plans should be allowed to include exploratory variables for which no clear predictions are made. Manuscripts, too, should not be penalized for reporting exploratory analyses, as long as they are clearly separate from the confirmatory ones. Where possible, exploratory findings can be retested with a new pre-registered study to have greater confidence in the reliability of the finding. In fact, from the limited experience our field has with pre-registration right now, it has become clear that small changes often have to be made after the pre-registration. These changes often will not undermine the validity of the registration, and can be discussed and agreed upon with the editor.

7.3. A null literature?

A near-certain consequence of the reviewed pre-registration model is that journals will more often report null results, as a necessary outcome of the goal to reduce publication bias. Some might worry that uninteresting null results will take up valuable journal space, but we think these worries are based on a misunderstanding of the purpose of scientific publishing. Scientific journals are not like popular magazines, newspapers or websites. What is of interest to the public or the non-specialist academic (counterintuitive findings, “sexy” topics, brief reports) is not necessarily what should be of interest to the academic specialist (theoretically grounded research, thorough reporting, the possibility that ideas are disconfirmed). Specifically, when a study is based on interesting theoretical predictions and proven methods, but yields disappointing results, the study can stimulate further refinement of theory and methods.

Null results are a prime example of information that is interesting to a specialist researcher rather than the general public. Someone engaged in research in a given field will certainly want to know which effects and methods are reliable and which are dead-ends. So, making public the specific attempts that were made to find an effect can eventually save the scientific community time and effort. Currently, the main ways of knowing about failed effects in psychology are through casual

conversation, conference presentations and blog entries. There is little incentive to share one's null results or to subject them to peer review. If researchers doing similar studies become aware of limitations sooner, they can adjust experiments to take them into account, or refrain from following blind alleys that others have found; all leading to a faster accumulation of accurate knowledge.

Additionally, the possibility that null results may be published should lead evaluators of the research—starting with the researchers themselves—to emphasize strong, tested and reliable methods, which we have identified above as a potential benefit of pre-registration. With scratch-built and untested manipulations and measures, null results are uninformative, because they can be blamed on shaky methods (Ferguson & Heene, 2012; LeBel & Peters, 2011). Likewise, null results with low statistical power to detect a reasonable effect size are also uninformative, due to the high likelihood of Type II error. A well-designed study or series of studies should allow for positive identification of “null” results, whether defined as literally nonsignificant, or as falling outside the range of some effect size of minimal interest.

7.4. Idea theft?

Another concern is that research ideas have a higher risk of being stolen if they are shared with reviewers before they are implemented. For unreviewed pre-registrations, this need not be a problem. The registration itself does not have to be made public until the article is submitted, so there is no more risk in exposing the final paper to reviewers than there is in a normal article. At the time of writing, the Open Science Framework is offering a flexible time limit of up to four years for registrations to remain private. And even after this period, registrations can be retracted.

For reviewed pre-registrations, however, it is possible for unscrupulous reviewers to conduct similar studies and quickly try to publish these elsewhere. But, because the reviewers are known to the editor, this move would be extremely obvious—the equivalent of a burglar leaving jam-covered fingerprints at the scene of a crime. Also, such concerns are not novel to pre-registration; they have always been present in grant review and journal publishing. At least with pre-registered plans, if the case is investigated, there can be no mistake about who registered what idea and when. We think anxieties over stolen ideas are particularly persistent because such investigations are rare, except perhaps in the court of public opinion. Indeed, our field finds it difficult to get to a point where any given theft is provable to all, and lacks clear procedures for disciplining people who steal ideas. Pre-registration by itself may thus not be the culprit.

8. Limitations of pre-registration

We move now from drawbacks, or the potential negative outcomes of pre-registration, to limitations, or problems that pre-registration is powerless to thwart.

8.1. Flexibility

An obvious limitation to pre-registered studies is the possibility that for any given analysis, all the parameters are difficult to pre-specify completely, so that authors may still knowingly or unknowingly build undisclosed flexibility into the analyses plan. Reviewed pre-registration, in contrast to unreviewed pre-registration, takes care of this problem to a large extent by allowing omissions in the registration to be pointed out, but this guarantee is only as good as the eyes of the reviewers. Although the template accompanying this paper is intended to reduce such practices, it is not so ambitious as to prescribe a program of analyses and reporting for each possible statistical situation (compare to the JARS questionnaire that regulates submissions to the APA Journal Archives of Scientific Psychology, see APA Publications and Communications Board Working Group on Journal Article Reporting Standards, 2008). However,

it is too soon to tell if this kind of flexibility will just find another way of expressing itself under pre-registration.

8.2. Fraud

Pre-registration does nothing at all to stop outright dishonesty—that is, when researchers make no attempt to imagine that they are still doing the right thing. Ways to intentionally cheat a pre-registration system are readily imaginable: multiple private unreviewed pre-registrations can be made, each with a different hypothesis, without disclosing this fact; dates can be misrepresented in order to falsely pre-register a study that was already run; the number of studies run can be misrepresented; and so on. Pre-registration sites can take some steps against the most egregious tricks by making registrations partially open. The unreviewed model of pre-registration, however, finds itself more limited in fighting publication bias, because the fact that a study is being done need not be disclosed until the manuscript is submitted. Thus, studies that yielded inconclusive or inconvenient results could potentially be hidden away in private pre-registrations without reporting them to anyone.

In general, the creation of multiple pre-registration sites (OSF, AsPredicted.org) is good, because it allows users to choose the model that best fits their needs. However, with more of these sites, a comprehensive check on public pre-registrations must make the rounds of all of them. Certainly for the field, relying on a small number of widely used, central sites would be preferable to archiving pre-registrations on a multitude of university sites or in uncheckable cloud data. The policy enacted by the Open Science Framework as of June 8, 2015, to put an expiration period on private registrations—after which they become public—is one way to improve trust in the unreviewed model while balancing this with the need to keep studies under wraps and avoid idea theft.

Regardless, if a researcher intends to commit fraud, there is little that any of the good research practices can do to either prevent or conclusively identify this, and pre-registration is no exception. In fact, we wonder why a fraudulent researcher would fiddle with the details of pre-registration, instead of simply forging the kind of data they would like to have. Nonetheless, it bears repeating that any falsifying of a pre-registration's status as an accurately date-stamped record of a single *a priori* protocol is outright scientific fraud, as much as tampering with a data file or fabricating responses.

8.3. Type of research

Sometimes, concerns are raised that pre-registration is fine for lab or experimental studies that collect new data, but does not cover the special needs of other types of research in psychology. Given the focus of this journal, we have written our template with experimental and correlational research in mind, conducted study-by-study in a lab, online or in a field setting. When extending beyond such paradigms, possibilities for pre-registration may be limited.

In evaluating pre-registration opportunities, we should keep in mind that pre-registration does not mean authors are expected to only report confirmatory analyses. In a large-investment longitudinal study, or one done on a hard-to-recruit population, researchers often optimize data collection by including as many measures as possible. Pre-registration does not mean that predictions are required for all these measures. In fact, researchers with massively multivariate data sets might especially benefit from pre-registration, because it will pre-empt any skepticism that confronts their truly *a priori* predictions.

Still, there are some types of research where pre-registration is not likely to be useful, requiring different assumptions than our template covers. Qualitative research has its own, different, and quite sophisticated ways of managing the dialogue between researchers' ideas and findings (Forrester, 2010). A completely exploratory study, one that

explicitly starts with few set ideas about the phenomenon and no set plan of data analysis, also will show little benefit from pre-registration until the research reaches a confirmatory phase. Studies analyzing complex patterns of observational, physiological, neurological or simulation data can all benefit from *a priori* specification of hypotheses and design, but a pre-registration template for those methods would likely be more specific in detail than we can cover here. Finally, projects analyzing existing data (such as archival research or meta-analysis) can in principle use pre-registration, as long as the earlier period of hypothesizing that forms the basis of the registration is clearly separated from the subsequent period of investigation and discovery. The difficulty of verifying that this practice has been followed may, in the eyes of some, reduce the value of pre-registration for secondary data analysis. On top of this, as Gelman and Loken (2014) point out, researchers who are continually in contact with pre-existing data may find it hard to draw such a precise line between exploration and confirmation, which limits the possibilities of pre-registration even further.

9. The elements of a pre-registration

As noted above, a concrete goal of pre-registration is to accurately describe hypotheses, methods and analyses before a study is conducted. Below we describe several elements that can be incorporated in a pre-registration, with a focus on utility for the field of experimental social psychology. These elements will likely see change as experience with pre-registration grows. As noted, it is also likely that the precise specifications will vary depending on the needs of specific fields and methodologies.

We present these elements as a template (see Appendix A) whose latest version can be downloaded online (<https://osf.io/k5wns/>), filled out, and time stamped via an online pre-registration form on osf.io once all collaborators agree on its details. In cases where a pre-registration platform provides its own structure, the template can be used to provide greater specification within each section of that structure. We emphasize that requiring an *a priori* plan should not prohibit researchers from pointing out *post hoc* deviations or subtleties—or even from saying *a priori* that there is no *a priori* plan—if circumstances dictate this. As long as the researchers can justify and explain these deviations to the evaluators of the registration, the registration itself is still of value.

9.1. Hypotheses

In the first section, confirmatory hypotheses are described in terms of predictions that connect the methods and outcomes of the proposed study to the theories and ideas underlying it. It is essential that predictions specify expected relationships between two or more variables. Ideally, they would specify a direction of relationship, and sometimes a more complex pattern if three or more variables are involved (e.g., describe an expected pattern of interactions and simple effects if two variables interact to predict a third). It is also advisable to number multiple hypotheses so that the analysis section can refer back to them. Rationales or theoretical frameworks for why a certain hypothesis is tested are helpful, but not necessary, and can be added to the template optionally. However, when predicting more than one outcome for a single test based on different ideas, the link between each outcome and its underlying idea should be made clear (for example, “if a top-down process is involved we would predict A, but if bottom-up, we predict B”).

For experimental methods, the template also requires explicit consideration of one kind of prediction that is often left out of the “hypotheses” section of an article. If an experiment manipulates a variable, some kind of positive control is often needed to demonstrate the success of this manipulation (i.e., a “manipulation check”). Specifically, manipulation checks test the effect of the manipulation on a measure representing the conceptual variable being manipulated. This practice should not be

confused with the useful but limited practice of including mere comprehension checks (for example, in a film manipulating perceived injustice in which Fred shouts at Barney, the participant's understanding that Fred and not Barney was the shouter is important to validate measures judging Fred, but does not establish that injustice is actually perceived).

If check variables exist, the template requires making predictions for them, and also asks for further explanation if manipulations without a check variable are included. We think that this step can lead researchers to take manipulation checks more seriously. Even if explicit checks are not advisable because they would influence participant awareness, it is still possible to run, for example, independent pilot tests with only the manipulation check as a dependent variable. Especially for reviewed pre-registration, establishing the validity of manipulations in this way is vital to interpreting null results should they arise.

Also, this step requires researchers to think about and express clearly what patterns in the manipulation check would support the validity of their method. This thinking is rarely expressed *a priori* in research articles in psychology. For example, if there are three progressively stronger levels of the manipulation, is it enough for the strongest to differ from the weakest, or should all comparisons be significant? If two variables are manipulated orthogonally, would it be a threat to validity if one manipulation had a significant effect on the other one's check variable, or is it acceptable that the "right" manipulation merely has a larger effect on its own check variable than on the other? Answering these questions explicitly ahead of time is one way pre-registration encourages a more careful approach to research.

9.2. Method

This section is similar in structure to a published method section—intentionally so, to save time later on in the process. The template starts with a description of the design, which informs the planned sample. Likewise, a description of the sample will lead authors to consider exclusion criteria. As a final step, the procedure is described, including materials. The information in this section should be detailed enough for reviewers to make an informed judgment on whether the hypotheses can successfully be tested with these methods.

9.2.1. Design

Here researchers describe the backbone of the experiment, outlining the independent variables with all their levels, whether they are within- or between-participants, the relationship between them (e.g., orthogonal, nested), as well as all dependent variables, and any third variables acting as covariates or moderators.

9.2.2. Planned sample

This part of the template describes the participant sample, giving sample size (and justification for it), an *a priori* description of the method used to recruit participants, and the stopping rule for collecting data. These descriptions directly address concerns that published results might be based on the undisclosed practice of collecting data in waves until a significant result is reached (Simmons et al., 2011). If this practice is not accounted for in statistical analyses (see Lakens, 2014) it inflates the overall alpha level of the test by taking advantage of multiple opportunities to stop data collection while looking for a desired result.

Most simply, a stopping rule can be based on a set number of participants. When resources are limited or specific lab rules (e.g., to run for a full week) do not allow a specific sample size to be given, a minimal sample size still ought to be given, and accompanied by the termination rule that will be applied (e.g., "All student participants who sign up for the study from the start until the end of the Spring term, minimum 80, maximum 160"). If the minimum number is not reached in a given

time, *post hoc* extensions of data collection, with target numbers specified, may be necessary and justifiable.

The data collection plan should also be informed by participant exclusion rules (see next section). If it can be verified, at some point after data collection, that some participants need to be excluded, the plan should specify whether additional participants will be recruited to make up the numbers, or whether the analysis will proceed with reduced numbers. Where possible, completing the planned numbers is preferable, to maintain pre-determined levels of statistical power and to ensure equivalent cell sizes in a categorical design.

Considerations of statistical power are useful in determining numbers of participants. In some cases information that comes to light during this process may lead to the realization that a different design would be more suitable, or that resources to adequately test the hypothesis do not exist. Where the registration is based on a known effect (e.g., when conducting a replication) power analysis can be based on the best estimate of that effect, or on a more conservative estimate if the original may be biased (Perugini, Gallucci, & Costantini, 2014). For previously unstudied effects, it is difficult to set exact guidelines, and the choice of an effect size may be arbitrary. In this case, the researcher can find comparable studies and decide on a range of effect sizes based on the power to detect an effect, or on what the smallest effect size of theoretical interest would be. Rather than setting forth a hard standard for power or sample size, we advocate being explicit about the reasoning that went into determining it, including assumptions about the effect size.

We also recommend saying where and how the data will be collected. This can give context for explaining later, unforeseeable circumstances that justify *post hoc* changes (e.g., "We started to collect data from passengers on a train but the conductor threw us off, so instead we collected data from people in a public library.").

9.2.3. Exclusion criteria

Here data exclusions are specified. Exclusion criteria can be on the participant, stimulus or trial level, and on the basis of missing, erroneous, or overly consistent responses. Examples include failed comprehension checks, demographic exclusions (e.g., analyzing only those who do not identify as group X in a study of prejudice against that group), outlier criteria, overly fast or slow reaction times, or ceiling/floor effects. With a greater number of exclusions anticipated in the pre-registration, there will be less need for exclusions to be determined *post hoc*.

Another optional element is to set fail-safe levels of exclusion at which the study needs to be stopped, altered, and restarted. For example, one might specify that if after running 20 participants, five or more of them do not show understanding of the instructions for their condition, the instructions need to be re-written and the study re-started. Of course, it is impossible to predict all such circumstances; sometimes this kind of circumstance has to be reported *post hoc*. Still, thinking through the decision process beforehand improves the value of the registration, especially if procedures are untested, or the context gives doubt about how many participants will yield valid data.

9.2.4. Procedure

As in a published manuscript, the details here should allow others to replicate the study, by describing all manipulations, measures, materials, and procedures, including the order of presentation, method of randomization, and "blindness" of experimenters and participants to condition (e.g., single or double blind). Tasks or measures reproducing previously published work do not have to be explained in full, but can be referenced, with any deviations from the published methods noted.

9.3. Analysis plan

9.3.1. Confirmatory analyses

Pre-registration asks that quantitative analyses be specified beforehand via an analysis plan. This procedure ensures that assumptions about analysis-wise alpha in null hypothesis significance testing are met, and not inflated by hidden flexibility in the methods and scope of the analysis. Having a plan is equally important, if not more so, for alternatives to null-hypothesis testing such as Bayesian analysis (cf. Wagenmakers, 2007), because these methods require prior assumptions about the effect sizes of null and alternative hypotheses (e.g., different point estimates, different functions). As with power analysis, while there is no clear consensus on a single method for deriving assumptions, it is important to explain the rationale beforehand to avoid doubt about whether the method chosen was influenced by its *post hoc* results.

The methods of quantitative data analysis are too diverse to cover comprehensively in our template. The important thing is that the key analytic decisions are based on hypotheses or method considerations, made ahead of time, carried out, and reported, while allowing for exploratory analyses to investigate unexpected aspects of the data. Minimally, the pre-registration should describe the analysis that will be carried out to test each numbered prediction from the hypotheses section, including: the key variables and how they are calculated from the original data; the statistical technique; and each variable's role in the technique (e.g., IV, DV, moderator, mediator). Anticipated covariates and their rationale (e.g., reducing variance in the DV, excluding a confound in the IV) should also be described here, reducing concerns about the use of covariates *post hoc* purely to achieve significant results (Simmons et al., 2011). If multiple simultaneous inferences are made, a method of correction for multiple comparisons can be described if appropriate (e.g., Bonferroni correction). Any analyses that are not described in this section, while completely permissible, should go under the heading of "exploratory" in the final paper.

Although it may be acceptable to state the analysis in general and obvious terms (e.g., "We will compare the mean memory task score across the three experimental and control conditions using one-way ANOVA and Tukey *post hoc* comparison tests"), a better practice, especially for complex analyses, would be to describe the analysis technically, so that it can be replicated by another person working with the same statistical software. At a high level of accuracy, but at the cost of additional effort, the registration can include a keyed list of variables and actual syntax for the planned analysis.

9.3.2. Contingencies and assumptions

The following considerations are optional in the template. An analysis plan can increase the *a priori* coverage of its procedures with more thorough plans in case the data violate statistical assumptions. Some common decisions that can be specified ahead of time, though by no means an exhaustive list, include:

1. A method for handling missing data (e.g., pairwise or listwise deletion, imputation, interpolation).
2. Criteria for scale reliability, and procedures to correct unacceptable levels of it (e.g., iteratively removing items with a low total correlation; treating items separately or via MANOVA).
3. Criteria for data transformations, such as departures from normality. This includes ceiling or floor effects, and procedures to correct this (e.g., using nonparametric tests, bootstrapping, transformation), and other transformations depending on the type of measures used (e.g., method of filtering out measurement noise in psychophysiological measures).
4. Criteria for problematic levels of heterogeneity in variance (e.g., Levene's test in ANOVA; sphericity testing in repeated measures ANOVA), and correction procedures (e.g., Games–Howell contrasts, Greenhouse–Geisser epsilon correction).

5. Criteria for overly high correlation between constructs in multivariate analysis (e.g., raw correlation levels, variance inflation factor) and steps to correct for it (e.g., aggregating highly related variables).
6. Criteria for identifying and handling outliers (e.g., in terms of interquartile range; Leys, Ley, Klein, Bernard & Licata, 2013; ESD procedure, Rosner, 1975; robust multivariate outlier criteria, Rousseeuw & van Zomeren, 1990; see also Bakker & Wicherts, 2014 for further critique of the usual method of excluding univariate outliers based on Z-scores, and solutions).

Many of these decisions (except for reliability) are rarely explicitly described in the psychological literature, *unless* they turn up problems with the assumptions of data analysis. As a result, it is hard to tell if these assumptions are always quietly checked in a principled way, or if they are instead resorted to mainly in an attempt to coax data into significance at $p < .05$, as Simmons et al. (2011) pointed out for transformation and outlier removal. Including them in the plan ahead of time can remove such doubts.

10. Conclusion

We conclude with some specific notes for using pre-registration, aimed individually at the creators and evaluators of research. For the *creators*—academics and students who carry out and report original research—we believe that pre-registration has enough benefits to encourage its regular use in research, without necessarily requiring it. Strong emphasis on sound theory and a clear divide between confirmatory and exploratory research can facilitate a shift towards solid science. As mentioned above, it seems that pre-registration can help individual researchers to realize well-thought-out studies and publicly gain acknowledgement for taking these steps. Additionally, valuable input from reviewers can be added early on, and knowledge about others who are working along similar lines and who may have valuable insights can be acquired faster. Further incentives to pre-register may come into sight as more journal editors see pre-registered research as indicative of more robust science, and even explicitly promote it as a way to confront doubts about publication and reporting bias in a manuscript (Giner-Sorolla, 2016; Vazire, 2015).

For *evaluators* of research—journal editors and reviewers—pre-registration requires a number of shifts in standards. Perhaps the most important shift is realizing that the outcome of a pre-registered series of studies testing a true hypothesis is not going to look perfect. Because of the little-appreciated "dance of the p-values" (that is, the variability in significance of demonstrations of a true effect; Cumming, 2014), not all individual study results are guaranteed to turn out significant, even though the overall picture gives strong support to the hypothesis. Evaluators of pre-registered studies need to keep this in mind; perfect-looking results across multiple studies are unlikely when the customary freedoms of selective reporting and analysis are constrained.

The other main shift in standards, as we have mentioned, is moving from a results-focused mindset to a methods-focused mindset, especially when evaluating reviewed pre-registrations, where methods but not results are available for inspection. It was our impression, prior to the current wave of methodological discussions, that if authors presented a significant result, then any flaws in methodology would only stop publication if they could have spuriously produced the result, not if they acted to suppress it. Going forward, the more that methods are evaluated independently of results, the more reviewers will need to be assured that the study is effective enough in its manipulations and measures so that they can trust even null results as informative.

In spite of our focus on academic publishing, registration of research need not be confined to that domain. Public granting agencies, for example, may see the research they fund as deserving dissemination no matter what the results, and no matter whether journals cooperate or not. In this

instance, pre-registration of studies and open reporting of findings might become a way to guarantee the return on investment in research and to ensure the accuracy of conclusions. Even without external funding, academics already spend much effort writing proposals to satisfy institutional review boards (IRBs) with detailed descriptions of methods and hypotheses. Those institutions would not need to ask much more in order to convert these efforts into actual pre-registrations, again helping the organization that hosts and facilitates the research to ensure that the outcomes of approved research are analyzed appropriately.

In conclusion, many authors, granting agencies, and journal editors in psychology and neuroscience are taking note of problems with publication bias and reproducibility, and are considering pre-registration of research as part of the solution. In the current paper, we have outlined several aspects of what pre-registration entails—its history, potential consequences both good and bad, and application. In sharing these thoughts, we aim to further the discussion and the use of pre-registration. It is our hope that the field as a whole will find ways to overcome its drawbacks and reap the benefits of this practice. We add that the current suggestions are not meant to be taken as set in stone. On the contrary, it is our genuine wish that more and more experience

with pre-registration will lead the field to fine-tune practices beyond these suggestions. We acknowledge that pre-registration is not always an option in some types of research, such as highly exploratory or qualitative research. When pre-registration is an option, however, we suggest that the benefits outweigh the costs both to individual scientists and science as a whole—especially when evaluators pick up the challenge and change their own standards, away from requiring an unnatural perfection in results and towards rewarding stronger theory and methods.

Acknowledgments

The authors would like to thank Christopher D. Chambers for comments on an earlier draft of this paper; also Brian A. Nosek and David Mellor for discussions on the pre-registration template; and finally, we greatly appreciate the engaged editorial comments of Laurie Rudman, Barbara Spellman, and the three signed reviewers. A.E. van 't Veer was funded by a European Research Council grant (ERC AdG 669249) to Marian J. Bakermans-Kranenburg.

Appendix A. Pre-registration in social psychology: a suggested template

Available on osf.io as an online pre-registration form that includes time stamping.

Section	Essential elements	Recommended elements
A. Hypotheses	<ol style="list-style-type: none"> Describe the (numbered) hypotheses in terms of relationships between your variables. For interaction effects, describe the expected shape of the interactions. If you are manipulating a variable, make predictions for successful check variables or explain why no manipulation check is included. 	<ol style="list-style-type: none"> A figure or table may be helpful to describe complex interactions. For original research, add rationales or theoretical frameworks for why a certain hypothesis is tested. If multiple predictions can be made for the same IV-DV combination, describe what outcome would be predicted by which theory.
B. Method	<p>Design</p> <p>List, based on your hypotheses from section A:</p> <ol style="list-style-type: none"> Independent variables and all their levels <ol style="list-style-type: none"> whether they are within- or between-participants; the relationship between them (e.g., orthogonal, nested). Dependent variables. Third variables acting as covariates or moderators. 	
Planned sample	<ol style="list-style-type: none"> If applicable, describe pre-selection rules. Indicate where, from whom and how the data will be collected. Justify planned sample size. Describe data collection termination rule. 	
Exclusion criteria	<ol style="list-style-type: none"> Describe anticipated data exclusion criteria. Some examples of exclusion criteria are: <ol style="list-style-type: none"> missing, erroneous, or overly consistent responses; failing check-tests or suspicion probes; demographic exclusions; data-based outlier criteria; method-based outlier criteria (e.g. too short or long response times). 	<ol style="list-style-type: none"> Set fail-safe levels of exclusion at which the whole study needs to be stopped, altered, and restarted.
Procedure	<ol style="list-style-type: none"> Describe all manipulations, measures, materials and procedures including the order of presentation and the method of randomization and blinding (e.g., single or double blind), as in a published Methods section. 	
C. Analysis plan		
Confirmatory analyses	<p>Describe the analyses that will test each main prediction from the hypotheses section. For each one, include:</p> <ol style="list-style-type: none"> the relevant variables and how they are calculated; the statistical technique; each variable's role in the technique (e.g., IV, DV, moderator, mediator, covariate); rationale for each covariate to be used, if any; if using techniques other than null hypothesis testing (for example, Bayesian statistics), describe your criteria and inputs towards making an evidential conclusion, including prior values or distributions. 	<p>Specify contingencies and assumptions, such as:</p> <ol style="list-style-type: none"> method of correction for multiple tests; the method of missing data handling (e.g., pairwise or listwise deletion, imputation, interpolation); reliability criteria for item inclusion in scale; anticipated data transformations; assumptions of analyses, and plans for alternative/corrected analyses if each assumption is violated.

References

APA Publications and Communications Board Working Group on Journal Article Reporting Standards (2008). *Reporting standards for research in psychology: hy do we need them? What might they be?* *American Psychologist*, 63(9), 839–851.

Badges to Acknowledge Open Practices (2013, February 19). Retrieved from <https://osf.io/tvyxz/>

Bakker, M., & Wicherts, J. M. (2014). Outlier removal, sum scores, and the inflation of the type I error rate in independent samples t tests: The power of alternatives and recommendations. *Psychological Methods*, 19(3), 409–427. <http://dx.doi.org/10.1037/met0000014>.

- Bateman, I., Kahneman, D., Munro, A., Starmer, C., & Sugden, R. (2005). Testing competing models of loss aversion: An adversarial collaboration. *Journal of Public Economics*, 89(8), 1561–1580.
- Begg, C. B., & Berlin, J. A. (1988). Publication bias: A problem in interpreting medical data. *Journal of the Royal Statistical Society. Series A (Statistics in Society)*, 151(3), 419–463. <http://dx.doi.org/10.2307/2982993>.
- Begley, C. G., & Ellis, L. M. (2012). Drug development: Raise standards for preclinical cancer research. *Nature*, 483(7391), 531–533. <http://dx.doi.org/10.1038/483531a>.
- Bosco, F. A., Aguinis, H., Field, J. G., Pierce, C. A., & Dalton, D. R. (2015). HARKing's threat to organizational research: Evidence from primary and meta-analytic sources. *Personnel Psychology*. <http://dx.doi.org/10.1111/peps.12111>.
- Bostyn, D. H., & Roets, A. (2016). The morality of action: The asymmetry between judgments of praise and blame in the action-omission effect. *Journal of Experimental Social Psychology*, 63, 19–25. <http://dx.doi.org/10.1016/j.jesp.2015.11.005>.
- Brandt, M. J., Ijzerman, H., Dijksterhuis, A., Farach, F. J., Geller, J., Giner-Sorolla, R., ... van 't Veer, A. E. (2014). The replication recipe: What makes for a convincing replication? *Journal of Experimental Social Psychology*, 50, 217–224. <http://dx.doi.org/10.1016/j.jesp.2013.10.005>.
- Carp, J. (2012). The secret lives of experiments: Methods reporting in the fMRI literature. *NeuroImage*, 63(1), 289–300. <http://dx.doi.org/10.1016/j.neuroimage.2012.07.004>.
- Challenging traditions in research reporting (2014, July). Retrieved from <http://newsroom.taylorandfrancisgroup.com/news/press-release/comprehensive-results-social-psychology#VXa3YIKGNg>
- Chambers, C., & Munafò, M. (2013, June 5). Trust in science would be improved by study pre-registration. Retrieved from <http://www.theguardian.com/science/blog/2013/jun/05/trust-in-science-study-pre-registration>
- Chambers, C. D., Feredoes, E., Muthukumaraswamy, S. D., & Etchells, P. J. (2014). Instead of "playing the game" it is time to change the rules: Registered reports at AIMS Neuroscience and beyond. *AIMS Neuroscience*, 1(1), 4–17. <http://dx.doi.org/10.3934/Neuroscience2014.1.4>.
- Clinical Trials Registration and Results Submission, 79 Fed. Reg. 225 (2014, November 21).
- Clinical Trials - Regulation EU No 536/2014 (2014). Retrieved from http://ec.europa.eu/health/human-use/clinical-trials/regulation/index_en.htm
- Cumming, G. (2014). The new statistics: Why and how. *Psychological Science*, 25(1), 7–29. <http://dx.doi.org/10.1177/0956797613504966>.
- de Groot, A. D. (1956/2014). The meaning of "significance" for different types of research. Translated and annotated by Eric-Jan Wagenmakers, Denny Borsboom, Josine Verhagen, Rogier Kievit, Marjan Bakker, Angelique Cramer, Dora Matzke, Don Mellenbergh, and Han L. J. van der Maas. *Acta Psychologica*, 148, 188–194.
- Dickerson, K., & Rennie, D. (2003). Registering clinical trials. *Journal of the American Medical Association*, 290(4), 516–523. <http://dx.doi.org/10.1001/jama.290.4.516>.
- Dwan, K., Altman, D. G., Arnaiz, J. A., Bloom, J., Chan, A. W., Cronin, E., ... Williamson, P. R. (2008). Systematic review of the empirical evidence of study publication bias and outcome reporting bias. *PLoS One*, 3(8), 1–31. <http://dx.doi.org/10.1371/journal.pone.0003081>.
- Easterbrook, P. J., Berlin, J. A., Gopalan, R., & Matthews, D. R. (1991). Publication bias in clinical research. *The Lancet*, 337(8746), 867–872. [http://dx.doi.org/10.1016/0140-6736\(91\)90201-Y](http://dx.doi.org/10.1016/0140-6736(91)90201-Y).
- Ebersole, C. R., Atherton, O. E., Belanger, A. L., Skulborstad, H. M., Allen, J. M., Banks, J. B., ... Nosek, B. A. (2016). Many labs 3: Evaluating participant pool quality across the academic semester via replication. *Journal of Experimental Social Psychology*, 67, 68–82 (in this issue).
- Eich, E. (2014). Business not as usual. *Psychological Science*, 25(1), 3–6. <http://dx.doi.org/10.1177/0956797613512465>.
- Fabrigar, L. R., & Wegener, D. T. (2016). Conceptualizing and evaluating the replication of research results. *Journal of Experimental Social Psychology*, 66, 68–80.
- Fanelli, D. (2010). Do pressures to publish increase scientists' bias? An empirical support from US states data. *PLoS One*, 5(4), e10271. <http://dx.doi.org/10.1371/journal.pone.0010271>.
- Ferguson, C. J., & Heene, M. (2012). A vast graveyard of undead theories: Publication bias and psychological science's aversion to the null. *Perspectives on Psychological Science*, 7(6), 555–561. <http://dx.doi.org/10.1177/1745691612459059>.
- Food and Drug Administration Amendments Act of (2007). 121 U.S. Stat., 823.
- Forrester, M. A. (Ed.). (2010). *Doing qualitative research in psychology: A practical guide*. London, England: Sage.
- Funder, D. C., Levine, J. M., Mackie, D. M., Morf, C. C., Sansone, C., Vazire, S., & West, S. G. (2014). Improving the dependability of research in personality and social psychology: Recommendations for research and educational practice. *Personality and Social Psychology Review*, 18, 3–12. <http://dx.doi.org/10.1177/1088868313507536>.
- Gelman, A., & Loken, E. (2014). The statistical crisis in science. *American Scientist*, 102(6), 460–465. <http://dx.doi.org/10.1511/2014.111.460>.
- Giner-Sorolla, R. (2016). Approaching a fair deal for significance and other concerns. *Journal of Experimental Social Psychology*, 65, 1–6. <http://dx.doi.org/10.1016/j.jesp.2016.01.010>.
- Giner-Sorolla, R. (2012). Science or art? How aesthetic standards grease the way through the publication bottleneck but undermine science. *Perspectives on Psychological Science*, 7(6), 562–571. <http://dx.doi.org/10.1177/1745691612457576>.
- Hung, H. M. J., O'Neill, R. T., Bauer, P., & Kohne, K. (1997). The behavior of the p-value when the alternative hypothesis is true. *Biometrics*, 53(1), 11–22. <http://dx.doi.org/10.2307/2533093>.
- Ioannidis, J. P. A. (1998). Effect of the statistical significance of results on the time to completion and publication of randomized efficacy trials. *JAMA: The Journal of the American Medical Association*, 279(4), 281–286. <http://dx.doi.org/10.1001/jama.279.4.281>.
- Jonas, K. J., & Cesario, J. (2015). How can preregistration contribute to research in our field? *Comprehensive Results in Social Psychology*, 1–7. <http://dx.doi.org/10.1080/23743603.2015.1070611>.
- Kahneman, D. (2003). Experiences of collaborative research. *The American Psychologist*, 58(9), 723–730. <http://dx.doi.org/10.1037/0003-066X.58.9.723>.
- Kaplan, R. M., & Irvin, V. L. (2015). Likelihood of null effects of large NHLBI clinical trials has increased over time. *PLoS One*, 10(8), e0132382. <http://dx.doi.org/10.1371/journal.pone.0132382>.
- Kerr, N. L. (1998). HARKing: Hypothesizing after the results are known. *Personality and Social Psychology Review*, 2(3), 196–217. http://dx.doi.org/10.1207/s15327957pspr0203_4.
- Klein, R. A., Ratliff, K. A., Vianello, M., Adams, R. B., Bahník, Š., Bernstein, M. J., ... Nosek, B. A. (2014). Investigating variation in replicability: A "many labs" replication project. *Social Psychology Quarterly*, 45(3), 142–152. <http://dx.doi.org/10.1027/1864-9335/a000178>.
- Lai, C. K., Marini, M., Lehr, S. A., Cerruti, C., Shin, J. E. L., Joy-Gaba, J. A., ... Frazier, R. S. (2014). Reducing implicit racial preferences: I. A comparative investigation of 17 interventions. *Journal of Experimental Psychology: General*, 143(4), 1765–1785. <http://dx.doi.org/10.1037/a0036260>.
- Lakens, D. (2014). Performing high-powered studies efficiently with sequential analyses. *European Journal of Social Psychology*, 44, 701–710. <http://dx.doi.org/10.1002/ejsp.2023>.
- LeBel, E. P., Borsboom, D., Giner-Sorolla, R., Hasselman, F., Peters, K. R., Ratliff, K. A., & Smith, C. T. (2013). PsychDisclosure.org: Grassroots support for reforming reporting standards in psychology. *Perspectives on Psychological Science*, 8(4), 424–432. <http://dx.doi.org/10.1177/1745691613491437>.
- LeBel, E. P., & Peters, K. R. (2011). Fearing the future of empirical psychology: Bem's (2011) evidence of psi as a case study of deficiencies in modal research practice. *Review of General Psychology*, 15(4), 371–379. <http://dx.doi.org/10.1037/a0025172>.
- Lerner, J. S., & Tetlock, P. E. (1999). Accounting for the effects of accountability. *Psychological Bulletin*, 125(2), 255–275. <http://dx.doi.org/10.1037/0033-2909.125.2.255>.
- Leys, C., Ley, C., Klein, O., Bernard, P., & Licata, L. (2013). Detecting outliers: Do not use standard deviation around the mean, use absolute deviation around the median. *Journal of Experimental Social Psychology*, 49(4), 764–766. <http://dx.doi.org/10.1016/j.jesp.2013.03.013>.
- Mahtani, K. (2016, February 5). How often are outcomes switched in clinical trials? And why does it matter? Retrieved from <http://compare-trials.org/blog/are-your-results-unusual-or-how-often-are-outcomes-switched/>
- Matzke, D., Nieuwenhuis, S., van Rijn, H., Slagter, H. A., van der Molen, M. W., & Wagenmakers, E.-J. (2015). The effect of horizontal eye movements on free recall: A preregistered adversarial collaboration. *Journal of Experimental Psychology: General*, 144(1), e1–e15. <http://dx.doi.org/10.1037/xge0000038>.
- Meehl, P. E. (1967). Theory-testing in psychology and physics: A methodological paradox. *Philosophy of Science*, 34(2), 103–115.
- Meehl, P. E. (1990). Why summaries of research on psychological theories are often uninterpretable. *Psychological Reports*, 66(1), 195–244. <http://dx.doi.org/10.2466/pr0.1990.66.1.195>.
- Melander, H., Ahlqvist-Rastad, J., Meijer, G., & Beermann, B. (2003). Evidence b(i)ased medicine—Selective reporting from studies sponsored by pharmaceutical industry: Review of studies in new drug applications. *BMJ*, 326(7400), 1171–1173. <http://dx.doi.org/10.1136/bmj.326.7400.1171>.
- Mellers, B., Hertwig, R., & Kahneman, D. (2001). Do frequency representations eliminate conjunction effects? An exercise in adversarial collaboration. *Psychological Science*, 12(4), 269–275.
- Mullane, K., & Williams, M. (2013). Alzheimer's therapeutics: Continued clinical failures question the validity of the amyloid hypothesis — But what lies beyond? *Biochemical Pharmacology*, 85(3), 289–305. <http://dx.doi.org/10.1016/j.bcp.2012.11.014>.
- Murayama, K., Pekrun, R., & Fiedler, K. (2013). Research practices that can prevent an inflation of false-positive rates. *Personality and Social Psychology Review*, 18(2), 107–118. <http://dx.doi.org/10.1177/1088868313496330>.
- Nier, J. A., & Campbell, S. D. (2013). Two outsiders' view on feminism and evolutionary psychology: An opportune time for adversarial collaboration. *Sex Roles*, 69(9–10), 503–506. <http://dx.doi.org/10.1007/s11199-012-0154-2>.
- Nosek, B. A., & Bar-Anan, Y. (2012). Scientific utopia: I. Opening scientific communication. *Psychological Inquiry*, 23(3), 217–243. <http://dx.doi.org/10.1080/1047840X.2012.692215>.
- Nosek, B. A., & Lakens, D. (2014). Registered reports: A method to increase the credibility of published results. *Social Psychology Quarterly*, 45(3), 137–141. <http://dx.doi.org/10.1027/1864-9335/a000192>.
- Open Science Collaboration. (2012). An open, large-scale, collaborative effort to estimate the reproducibility of psychological science. *Perspectives on Psychological Science*, 7(6), 657–660. <http://dx.doi.org/10.1177/1745691612462588>.
- Perugini, M., Gallucci, M., & Costantini, G. (2014). Safeguard power as a protection against imprecise power estimates. *Perspectives on Psychological Science*, 9(3), 319–332. <http://dx.doi.org/10.1177/1745691614528519>.
- Prinz, F., Schlange, T., & Asadullah, K. (2011). Believe it or not: How much can we rely on published data on potential drug targets? *Nature Reviews. Drug Discovery*, 10(9), 712. <http://dx.doi.org/10.1038/nrd3439-c1>.
- Rakow, T., Thompson, V., Ball, L., & Markovits, H. (2014). Rationale and guidelines for empirical adversarial collaboration: A thinking & reasoning initiative. *Thinking & Reasoning*, 21(2), 167–175. <http://dx.doi.org/10.1080/13546783.2015.975405>.
- Rosner, B. (1975). On the detection of many outliers. *Technometrics*, 17(2), 221–227. <http://dx.doi.org/10.2307/1268354>.
- Rousseeuw, P. J., & van Zomeren, B. C. (1990). Unmasking multivariate outliers and leverage points. *Journal of the American Statistical Association*, 85(411), 633–639. <http://dx.doi.org/10.1080/01621459.1990.10474920>.

- Simes, R. J. (1986). Publication bias: The case for an international registry of clinical trials. *Journal of Clinical Oncology*, 4(10), 1529–1541.
- 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. <http://dx.doi.org/10.1177/0956797611417632>.
- Stern, J. M., & Simes, R. J. (1997). Publication bias: Evidence of delayed publication in a cohort study of clinical research projects. *BMJ*, 315(7109), 640–645. <http://dx.doi.org/10.1136/bmj.315.7109.640>.
- Tukey, J. W. (1980). We need both exploratory and confirmatory. *The American Statistician*, 34(1), 23–25. <http://dx.doi.org/10.1080/00031305.1980.10482706>.
- Turner, E. H., Matthews, A. M., Linardatos, B. S., Tell, R. A., & Rosenthal, R. (2008). Selective publication of antidepressant trials and its influence on apparent efficacy. *The New England Journal of Medicine*, 358(3), 252–260.
- van 't Veer, A. E., Gallucci, M., Stel, M., & van Beest, I. (2015). Unconscious deception detection measured by finger skin temperature and indirect veracity judgments—Results of a registered report. *Frontiers in Psychology*, 6(672). <http://dx.doi.org/10.3389/fpsyg.2015.00672>.
- Vazire, S. (2015). Editorial. *Social Psychological and Personality Science*, 7(1), 3–7. <http://dx.doi.org/10.1177/1948550615603955>.
- Wagenmakers, E. -J. (2007). A practical solution to the pervasive problems of *p* values. *Psychonomic Bulletin & Review*, 14(5), 779–804. <http://dx.doi.org/10.3758/BF03194105>.
- Wagenmakers, E. -J., Wetzels, R., Borsboom, D., van der Maas, H. L. J., & Kievit, R. A. (2012). An agenda for purely confirmatory research. *Perspectives on Psychological Science*, 7(6), 632–638. <http://dx.doi.org/10.1177/1745691612463078>.
- Wasserstein, R. L., & Lazar, N. A. (2016). The ASA's statement on *p*-values: Context, process, and purpose. *The American Statistician*. <http://dx.doi.org/10.1080/00031305.2016.1154108> (in press).
- World Medical Association (2008). WMA Declaration of Helsinki – Ethical principles for medical research involving human subjects. Retrieved July 13, 2015 from <http://www.wma.net/en/30publications/10policies/b3>