The inner workings of Registered Reports

Zoltan Dienes

University of Sussex

Correspondence: zoltan.dienes@gmail.com

A first draft for the *Cambridge Handbook of Research Methods and Statistics for the Social and Behavioral Sciences, Volume 2*, edited by Austin Lee Nichols & John E. Edlund.

Abstract

Registered Reports provide one way to address shortcomings in the current way we manage research, right from the design of studies to their publication. The format requires pre-specifying why a design may crucially test a theory, what auxiliary assumptions are required for the experiment to be such a test, what outcome neutral tests are required in turn to test those assumptions, what specific crucial tests will therefore be used test the theory (of the many tests that could be used), and why those tests could provide evidence for no effect of interest given the proposed numbers of trials and participants. Reviewers and authors are then constructively involved in optimizing the experiment before it is run. The agreement between reviewers and authors, as adjudicated by the editors, then defines in advance the proposed method and analytic protocol, virtually guaranteeing acceptance of the paper, no matter what position, if any, the results support. Here we go through what problems the format solves, and why it must therefore be approached in a way that is little understood. Common pitfalls are discussed. The paper thus provides an argument for how to approach Registered Reports for readers, authors, and editors of the format.

Science involves discovering objective relations between theories, predictions and evidence (Popper, 1972). However, we are humans, and in assessing the relations between theory and evidence, we are subject to biases (Chavalarias & Ioannidis, 2010; Kerr, 1998; Ritchie, 2020). Thus, it behoves us to assess the relations under conditions that minimize bias. A way of addressing bias is to make judgments that are subject to public critical scrutiny and argument; and to make judgments without knowing if those judgments would lead to the results that would support a theory or not (Wagenmakers, Wetzels, Borsboom, et al., 2012). Hence the development of the Registered Report (Chambers, 2013, 2019), in which the motivation, method and analytic protocol are settled by critical discussion before the data are collected.

Initially a submitted manuscript, providing introduction, method and planned analysis sections, is, after editorial triage, sent to review. Reviewers and authors then work together to revise the theory, derivation of predictions, method and analytical protocol to optimize all aspects of the experiment as best they can. Once the editor judges that revisions have produced a workable study, Stage 1 In Principle Acceptance (IPA) is granted. The authors run the study. The data are analysed according to the agreed protocol; in addition, the authors can introduce new analyses in a different sub-section of the results section. The paper is finally accepted at Stage 2 if the authors followed the procedure they said they would, reported the analyses they said they would, planned quality control checks are passed, and the discussion is appropriate.

In this paper the logic of what the Registered Report achieves is discussed in order to motivate how a Registered Report should be approached. Common mistakes are addressed and guidance is provided for authors, reviewers and editors. The full blown Registered Report is discussed. A weaker but still valuable process is pre-registration. Pre-registration is not an article type; it is something any author can do for any paper. In pre-registration, the author writes up a plan in as much detail as is seen fit and it is posted in a public repository like OSF. Then the study is conducted and submitted as normal to a journal. Reviewers may disagree with the planned analysis; but at least they know what was planned and what was not, and that can inform judgment. The pre-registration misses out on many of the benefits of a Registered Report, which we now discuss – while gaining, for better or worse, in flexibility (the pre-registration is completely under the authors control). For example, when mere pre-registration is used rather than a Registered Report, the extent to which authors write clear preregistrations and then follow what they did is not enforced (Ikeda et al., 2019).

I first present a view of the workings of science as seen at the scale of a single study or a small sequence of studies. The possible fault lines are discussed and why the Registered Report helps address the problems is indicated. Then I go through each aspect of this account to indicate how a Registered Report can therefore be approached. The paper aims to provide a framework for understanding the nature of Registered Reports; for simple practical guidance see PCI RR (2023)

A view of science

On a certain level of resolution, science, at least an essential core of it, consists of testing a substantial theory, by deriving predictions of phenomena, using well-tested and otherwise simple assumptions; data then provide evidence for or against the phenomena (Figure 1). A substantial theory is a claim that could be found wrong by the study, and from which various predictions may follow. The theory functions as an explanation of those predictions. For example, the claim that “extrinsic motivation impairs intrinsic motivation because of a shift in attributions about the causes of behaviour” is a substantial theory (Deci & Ryan, 2000). Various interesting phenomena can be derived from it as predictions. For example, “rewarding children for eating their greens by giving them desert reduces the desirability of the greens” (cf. Newman & Taylor, 1992) is a phenomenon predicted by the theory. A phenomenon is an empirical generalization (Haig 2014); or what Popper (1959) calls a low level hypothesis claiming what the results are. “Extroverts drink more coffee than introverts in the morning” is a phenomenon. A substantial theory of extroversion may predict it. The

theory will do so on the basis of assumptions. For example, if the theory of extroversion refers to arousal, an assumption is that caffeine is relevant to the sort of arousal postulated by the theory. In fact, the study can only test the prediction by making assumptions; for example, by making the assumption that the extroversion questionnaire measured extroversion (Flake et al., 2022). Ideally, such assumptions have been well tested or else are simple. If the assumptions can be taken as safe, then if the evidence counts against the phenomenon, the substantial theory also takes a hit. Conversely, if the evidence supports the phenomenon, the substantial theory is corroborated. These relations are objective in that scientists can be wrong about their nature; that is, we need to *find out* what predictions follow from a theory in a simple way with well tested assumptions, and *find out* what the data say about the predictions (Popper, 1972).

If we can be wrong about the actual theoretical and evidential relationships, these relations should be assessed under conditions where bias in their assessment is minimized. When the relations are assessed in the light of the data, several biases are possible. First, when data are seen first before deriving predictions from a theory, complex assumptions may be introduced to make the theory fit the phenomena (Figure 1 B); the data are then no longer evidential regarding the theory. Gelman and Loken (2014) present examples in which papers have “predicted” complex effects (which were significant) but the stated theory in fact makes simpler predictions (which were not significant). Kerr (1998) also describes several variants of this practice. Second, one may give up on postulating any substantial theory, and simply declare whatever phenomenon is found as a discovery (Figure 1 C) (Kerr, 1998). In the absence of theory, the temptation is to want to find phenomena of low prior plausibility (Fiedler, 2017) for which a conventionally low evidential threshold may be inappropriate (Benjamin, Berger, Johannesson, et al., 2017). Further, without coherent explanatory research programs, such research can veer away from science to mere phenomena chasing (Klein, 2014). Some pursuit of phenomena themselves, i.e. natural history, is a crucial part of science (e.g. Haig, 2014; for an ingenious systematic way of exploring phenomena see Baribault, Donkin, Little, et al., 2018). But pursuing phenomena should not replace theory testing, as it is only in the context of a theoretical research program that science develops (Alger, 2019; Deutsch, 2011; Guest & Martin, 2020; Lakatos, 1970). Finally, when data are seen before planning analyses, p-hacking (or B-hacking for Bayes factors) may make support for the predictions sufficiently probable regardless of the data so that the data, as analysed, are no longer evidential regarding phenomena (Figure 1D) (Chuard, Vrtilek, Head, & Jennions, 2019; Dienes, 2023; Mayo, 2018; Simmons, Nelson, & Simonsohn, 2011).

A single paper need not cover all these aspects of science. A paper is valuable if it helps in any aspect of: developing a substantial theory (e.g. in exploring data, McIntosh, 2017; Yanai & Lercher, 2020); or in determining what assumptions are safe in making predictions (establishing phenomena); or in showing how predictions may be severely tested (developing methodology). Registered Reports are an article type that can be used when all these aspects happen to come together (and sometimes in other cases). Specifically, in a Registered Report, a paper is accepted before data are collected based on: 1) a substantial theory being tested; 2) assumptions connecting theory to predictions being safe; and 3) analytic flexibility being tied down while ensuring sensitive results. In this way, the Registered Report seeks to reduce the biases just reviewed that potentially corrupt the scientific process. We now go through the three corresponding desiderata.

Figure 1 Biases that may arise from seeing data first

A) An ideal view of science as objective relations between theory, predictions, and data. From a substantial theory predictions of phenomena can be derived, using auxiliary assumptions. Data provide evidence for or against the predictions, and therefore count for or against the theory. B) Seeing data first may tempt one to corrupt the perceived relation between theory and predictions. The predictions stated may no longer be those that follow from the theory with simple and safe assumptions. C) Seeing data first may tempt one to give up on using theory, because one can guarantee finding some phenomenon or other, if one pokes around long enough. D) Seeing data first may tempt one to corrupt the perceived relation of data to predictions by hacking. Data may seem to support predictions when in fact the data provide no evidence for them.

A substantial theory

A substantial theory functions as an explanation of the phenomena it predicts. It is also a claim that could be shown wrong by the results of the study. Exactly what is needed here varies by journal. *Peer Community In Registered Reports* (PCI RR, 2023), as well as *Cortex* (Cortex, 2020), *the European Journal of Neuroscience* (EJN, 2020), and *Royal Society Open Science* (RSOS, 2020), indicate only that submissions will be evaluated with respect to “the scientific validity of the research question(s).” *Nature Human Behaviour* (NHB, 2020), on the other hand, state as the criterion “the importance of the research question(s) and relevance for a broad, multidisciplinary audience.” While these criteria do not mention a substantial theory as such, in my experience as both editor and reviewer of Registered Reports, the lack of a clear substantial theory has been grounds for revision at Stage 1. In principle, a Registered Report may be used to test simply the existence of a phenomenon, or to estimate its size. The theory in this case may be that the phenomenon exists. However, if the phenomenon is interesting because it bears on theory, at some point a substantial theory will need to be postulated to explain the phenomenon; and if one is postulated at Stage 1, the formulation of the theory and its relation to predicted phenomena may benefit from critical discussion before the results are known. In this situation, an appropriate strategy is to conjecture the most general bold claim that could still be tested and shown wrong by the specific study. For an example of a Registered Report whose sole aim was to estimate a practically important quantity in an unbiased way (related to the accuracy of measuring the presence of SARS-CoV-2), see Khan and Cheung (2020). In this case, a substantial theory was not explicitly stated, though there was relevant theory at play in setting up the parameter to be estimated, and there was a key theory at stake: That the size of the parameter estimated rendered the measurement relevant to its practical function.

Safe assumptions connecting theory to predictions: Outcome neutral tests

All tests of theories involve assumptions, claims that must be true for the test to be a test of the substantial theory (Popper, 1963). For example, one must assume that the independent variable manipulates what it is claimed to, and the dependent variable measures what it is claimed to. The truth of these assumptions is independent of the truth of the substantial theory. Thus, tests of these assumptions are called outcome neutral tests. An example of an outcome neutral test is a positive control, such as a manipulation check; for example, a check that a manipulation of mood changes mood ratings. More generally, a positive control is a control which the manipulation should affect if the procedure works as it should. An example of a positive control would be showing one’s manipulation affects something it is known to affect; for example, showing that a repetitive Transcranial Magnetic Stimulation protocol (where an alternating magnetic field is applied to the scalp to disrupt underlying cortex) affects a function it is known to affect, in order to prove the protocol is being employed correctly. An example of a positive control that is not a manipulation check as such is showing that one’s measuring instrument detects a relevant difference where one is known to exist; for example, showing that an indirect measure of pain is sensitive to a relevant noxious stimulus. Outcome neutral tests include checks of implementation fidelity or intervention fidelity (Carroll, Patterson, Wood et al., 2007; French, Diekemper, Irwin, et al., 2015), for example, for the delivery of a therapy. Were there intervention manuals; and were therapists shown to know their contents?

According to the *PCI RR* (2023) guide to Registered Reports, one should consider whether there are “sufficient outcome-neutral conditions (e.g. absence of floor or ceiling effects; positive controls; other quality checks) for ensuring that the obtained results are able to test the stated

hypotheses or answer the stated research question(s).” A Registered Report likely includes outcome neutral tests. For example, if a mood manipulation is involved, test that the mood was changed appropriately, independent of the substantial theory tested. Imagine that a substantial theory predicts that a negative rather than positive mood speeds up learning of a certain sort. A researcher might say “I don’t need a separate manipulation check; I know if the mood changed if the learning was sped up; that is the only reason the learning could change.” Now there is a get-out clause if the prediction fails. If learning does not speed up, instead of the substantial theory being disconfirmed, the researcher can say “the prediction failed because mood did not change.” The researcher is then free to try again with a new mood manipulation, and then again, until learning changes - all the while maintaining confidence in the substantial theory, despite successive predictive failures. By using appropriate outcome neutral tests, this scenario is avoided.

In addition to assuming the independent variable manipulated what it claims to, the dependent variable is also assumed to measure what it claims to, and this may need to be established (Flake & Fried, 2019). Outcome neutral tests can also be used to rule out the problem of post hoc disregarding of data from particular experimenters because their failure to get the effect shows they are unskilled (e.g. Baumeister, 2016); if the outcome neutral tests are passed, the data from that experimenter should be included.

If predictions are not confirmed, key assumptions should be safe, so that the substantial theory takes the blame. Consider the substantial theory that “belief in free will induces one to overcome automatic habits and hence behave prosocially” (e.g. Schooler, 2010). Given certain assumptions, it makes the prediction that “After reading statements aimed at reducing a belief in free will rather than control statements, people will give more milligrams of hot sauce in to someone who doesn’t like it” (cf. Baumeister, Masicampo, & Dewart, 2009). An assumption is that the intervention (reading statements about free will) actually changes free will beliefs. If we fail to confirm predictions, we could just as plausibly reject this auxiliary assumption as reject the substantial theory. So there needs to be an outcome neutral test testing that belief in free will changes as a result of the intervention (for example, by using a belief in free will questionnaire).

In a Registered Report, some outcome neutral tests are specified as ones that must be passed for the test to test the substantial theory (e.g. belief in free will was changed). These can be distinguished from checks that are useful but not essential (e.g. did participants take equal amount of time to read intervention and control statements? To deal with this, if the time difference is within stated bounds, time might be added as a covariate). An advantage of Registered Reports is that reviewers and the editor can help identify and shape the outcome neutral tests that genuinely test key assumptions.

Outcome neutral tests may be useful even for studies not testing a substantial theory per se but e.g. estimating an important parameter. There will still be auxiliary assumptions for the estimation to be useful, and these assumptions should be carefully considered to determine if any need be tested by outcome neutral tests for the estimation to be regarded relevant. For example, if estimating the degree of racism, or anything socially sensitive, a researcher could check for indications of honesty in a pilot by comparing the standard test to one in which the person is sometimes instructed to roll a die and report the opposite response for binary answers. (This procedure means the participant is aware that the researcher does not know whether that participant’s answer is in fact true.)

If your manuscript does not contain one or more outcome neutral tests, consider what tests might be needed. Include checks of study or data quality broadly construed, such as absence of floor or ceiling effects. Explicitly list the outcome neutral tests that are crucial, that is, which must be passed for the results to count for or against the theory; and those that are just desirable (and which for example might just change the key analyses). Some quality control checks may simply be a basis for excluding participants (e.g. checks on whether the participant was concentrating). According to many journal guidelines, guaranteed publication at Stage 2 may depend on the crucial outcome neutral tests being passed. However, I am unaware of any Stage 2 submissions that were

not published for this reason; even a failed outcome neutral test can be useful in showing a standard procedure does not perform as it should.

Analytic flexibility being tied down while ensuring sensitive results

There are always a range of equally legitimate analyses to address a research question (Carp, 2012; Orben, & Przybylski, 2020; Simonsohn, Simmons, & Nelson, 2019; Silberzahn, Uhlmann, Martin, et al. 2018; Steegen, Tuerlinckx, Gelman, & Vanpaemel, 2016). Steegen et al. call this space of analyses the multiverse. There are always various ways of coding the variables, of excluding outliers, of dealing with assumptions, of including covariates, of modelling structure, and so on. The factorial combination of all these different possible analytic methods is the multiverse (the complete set of different analyses given the stated different possible variants of each stage of analysis). The multiverse will typically allow a range of conclusions. For example, in Figure 2 A, Bayes factors are used which allow the conclusion good enough evidence for H_0 , good enough evidence for H_1 , and no evidence to speak of. If the data speak relatively clearly to a question, one conclusion may predominate; in Figure 2A, it is that there is evidence for H_0 . However, if the researcher has the data in front of them while making analytic decisions, they can take a stroll through the garden of forking paths (Gelman & Loken, 2014; that is, make coding and analysis choices in the light of data), and post hoc justify each analytic choice given the way it pushes towards a desired result. There is always a compelling reason for any choice: It is tougher one way/tougher another way/more conservative/more inclusive/more widely used/more recently used, and so on. Thus, a researcher, while thinking they are making decisions for scientific reasons, may be fooling themselves: They may be making decisions because it supports the conclusion they personally want. As illustrated in Figure 2 B, they may therefore end up in a small region of the multiverse which gives an uncharacteristic conclusion (see Steegen et al. for a real example).

One solution is to conduct a multiverse analysis, that is, perform all analyses implied by any combination of coding and analytic choices that have been identified. This can be a good solution, but it is not practicable in general. In general, one cannot for every test conduct maybe 1000s of analyses and then pore over how to interpret them. Further, there is not one objective list of possible analyses for any given effect of interest; there is a multiverse of these multiverses. Further, there will be a multiverse of ways for interpreting each of these first order multiverses. Instead, there is a short cut (Dienes, 2016, p. 86). If in advance one randomly picked one point in a multiverse, the result probably yields the most common outcome, and therefore probably reflects the objective evidential relation of the data to the prediction. This can be achieved by specifying the analytic protocol in advance of the data, while making sure there is no analytic flexibility left.

Removing all analytic flexibility takes more specification of the analytic protocol than researchers have typically been used to. So in specifying your analyses, ask yourself: Have I left myself any wriggle room to bias the results in light of the data? If so, you need to tighten up to make sure where you end up in the multiverse is independent of where your prejudices would want you to be. Analytic decisions can still be conditional on properties of the data (Devezer et al., 2020), but in pre-specified ways so that they are not conditional on the outcome.

Consider how all variables will be coded, how assumptions will be checked, what you will do about assumption violation, on what precise grounds individuals or trials will be removed, how structure will be modelled. Typically, this cannot be simply guessed for a new paradigm, or for a paradigm you have not used yourself. It is therefore often useful to run a small pilot. The pilot need not be the full design of the study. But making sure you understand how the sort of data you are dealing with in advance of specifying full analyses can be very helpful.

One might decide that not all analytic decisions can be decided in advance due to the complexity of the data. A useful tool in this case may be blind analyses (e.g. Dutilh, Vandekerckhove, Ly, Matzke, et al. 2017; MacCoun & Perlmutter, 2015). The analyst is presented with several data sets. One data set is real; the others have had one or more key variables shuffled or

otherwise transformed. The analyst does not know which set is the real set. The analyst can then analyze all data sets until they are satisfied that the same analysis is appropriate for all data sets, and that the analysis deals with assumptions and unexpected features adequately. If the blinding has been successful, the analyst makes decisions without knowing how those decisions produce different conclusions. *PCI RR* and various journals allow blinding procedure as part of the analytical protocol for a Registered Report. How blinding is to be done, and when it will be broken, needs to be clearly specified.

A)

B)

C)

Figure 2 Picking a point in the multiverse

A) Shows a multiverse of different possible analyses. A majority of analyses favour H_0 . B) illustrates that if one took a stroll through the garden of forking paths, i.e. checking which way to go in the light of how much a theory is supported, one may end up with an outcome that few analyses show. C) illustrates that choosing the location of the multiverse in advance probably yields most common conclusion: Objective evidential relation between data and hypotheses are probably respected.

Obtaining sensitive results. The extent to which the analysis allows conclusions contrary to predictions should be determined. For example, if the theory predicts a difference, to what extent does the analysis allow the conclusion that there is no effect? This is crucial to determine in order that the study can count against the theory. Three (non-exhaustive) ways of answering this question are: by use of power; inference by intervals (including equivalence testing); and Bayes factors (see Dienes, 2021a, for all three). Each method is considered in turn.

Power. Part of tying down flexibility is specifying a stopping rule for frequentist statistics, that is, stating the precise conditions under which one will stop collecting data. The simplest method is to calculate the number of participants to achieve a relevant power. According to *PCI RR* (2023), “Studies involving Neyman-Pearson inference should include a priori statistical power analysis, with estimated effect sizes justified with reference to the existing literature or theory. Since publication bias overinflates published estimates of effect size, power analysis should be based on the lowest available or meaningful estimate of the effect size.” Note that specific power and significance requirements will vary by journal.

In testing a theory, the effect size used for calculating power should be the one that defines the minimal plausible effect of interest for that theory or practical context. In order to control the long run probability of missing an effect of interest (Type II errors), power must be calculated for the smallest effect one does not want to miss out on detecting. Otherwise, a non-significant result does not in itself count against a theory predicting an effect. Dienes (2021a) presents heuristics for determining a theoretically-relevant minimally interesting effect size. One heuristic is calibration. Calibration is the process of converting from one scale about which it is hard to judge what is of interest, to another scale for which, as it happens, we can judge. For example, a dependent variable for measuring distress in an animal may be how distracted it is, e.g. the proportion of time it wanders close to stimuli it would typically avoid (Braithwaite, 2010). What proportion of time is of minimal interest? This can’t be answered in the abstract, nor with standardized effect sizes. Instead, the question is, how much distress do we care about? That question may have a reasonable answer. For example, given a scale of painful stimuli from “non-painful squeeze”, “slight pinch”, “slight bruising of skin”,.... “broken skin” ... the level of distress caused by a “slight bruising of skin” may be relevant for welfare considerations. From a plot of proportion of time distracted against the scale of painful stimuli, the proportion corresponding to “slight pinch” can be read off. Note this process involves finding out answers to real questions; data may have to be collected. Notice also that if calibration were not used, and the minimal interesting effect was a standardized effect size plucked from the air (“Cohen’s $d = 0.5$ ” which is just a signal to noise ratio dependent on e.g. how many trials are used), the in-practice minimal effect may be one corresponding to considerable animal suffering. See Dienes (2021b) for other examples and heuristics.

Royal Society Open Science has no set power requirements; the authors can define the criteria between themselves, the reviewers and the editor, or have no specified criteria if hypothesis testing is not involved. In cases where it may be hard to gather data, for example in neuropsychology, the use of a Registered Report can still be valuable, despite the sensitivity of the data not being guaranteed. (However, *Royal Society Open Science* requires that outcome neutral tests must be passed. If an outcome neutral test is passed only by showing no difference between two conditions, and if Neyman Pearson statistics are being used, then the test must be sufficiently powered for the outcome neutral test to satisfy consumers of the paper.)

Inference by intervals. Another method for allowing the conclusion that there is no relevant effect is to define a null interval, i.e. the interval of effects less than a minimally interesting one, m.

Then if a confidence or credibility interval lies within the interval $[-m, +m]$, there is support for the null interval hypothesis (Greenwald, 1975; in the frequentist case, this is known as equivalence testing, Westlake, 1972). The same principles for finding minimally interesting effect sizes apply for this method, as for power; see Dienes (2021a) (and also Dienes, 2021b, for further examples in specifically a Bayesian context). Having determined the relevant null interval, in planning the study size, use simulations to determine the number of participants needed to achieve an X% chance (e.g. 90%) of the 95% CI being within the null interval if the null interval hypothesis were true - to make sure one has resources to obtain evidence against the theory if it were false. A stopping rule can be running until the 95% CI (or 90% CI etc.) is either completely inside or completely outside the null interval (cf. Dienes, 2016; Kruschke, 2018, for discussion). For advice on writing up inference by intervals, see Dienes (2021b).

Bayes factors. It may be easier not to determine a minimally interesting effect size, but a roughly expected effect size, or range of effect sizes. The minimally interesting effect size is typically the most arbitrary aspect of the range of effect sizes predicted by a theory, and can often be approximated as 0 if big data are not involved. The roughly expected effect size, or its plausible range, is what is required for a Bayes factor. A Bayes factor compares the ability of a model of H1 to predict data compared to a model of H0 (Jeffreys, 1939; Wagenmakers, Verhagen, Ly, Matzke et al., 2017). The Bayes factor is then the evidence for that model of H1 relative to H0 provided by the data. In order for the Bayes factor to be relevant to testing a theory, the model of H1 must be an adequate model of the predictions of that theory. That is, the effect size used to scale the predictions of the model of H1 must be the rough effect size predicted by the theory. Dienes (2019, 2021a,b) provide heuristics for determining what rough effect size a theory predicts. This may involve collecting some pilot data, and not necessarily for the full design. For example, if the aim of the study is to determine whether a manipulation removes or impairs an effect of caffeine, the pilot could estimate simply the caffeine effect (call that c); the manipulation can then be predicted to produce an effect no bigger than this, that is, an effect plausibly in the range $[0, c]$. When using Bayes factors, report how one modelled H1 (see Dienes, 2021b, for advice on this and also on writing up Bayes factors). H0 is typically modelled as a point (or spike of probability) at zero slope (or zero difference between means etc). However, H0 could be something else, and so the model of H0 (cf Palfi & Dienes, 2019) should be explicitly stated too.

Guidelines vary according to journal concerning whether a certain size Bayes factor needs to be reached. *Cortex* asks for a Bayes factor of more than 6 (or less than $1/6$). *Nature Human Behaviour* uses a cut off of 10 rather than 6. Most journals use 3. *Royal Society Open Science* has no set thresholds. One can use as a stopping rule: Collect data until the Bayes factor is more than 6 (or whatever the cut off is) or less than a $1/6$ (Rouder, 2014; Rouder & Haaf, 2019). If a decision threshold of, for example, 3, is used, the stopping rule could use a larger threshold, for example 6, to ensure robustness of conclusions. In planning the study, estimate roughly how many participants may be needed to achieve the desired cut off; see Schönbrodt, Wagenmakers, Zehetleitner, and Perugini (2017) and Schönbrodt and Wagenmakers (2018) for using simulations to estimate the sample size that might be required. Make sure the simulations use the same model of H1 as you will use. A very quick way of calculating an approximate likely number of participants, which involves no simulations, is provided by Palfi and Dienes (2019b, version 3, Table 1). According to *PCI RR* (2023), “If the stopping rule is dependent on the Bayes factor, authors should indicate a maximum feasible sample size after which sampling will stop, regardless of the Bayes factor. In this case, reviewers should judge whether the maximum possible sample size is sufficiently large that an inconclusive Bayes factor at that sample size would nevertheless provide an important message for the field.”

Whichever method one uses, obtaining support for no effect is always relative to the size of effect that is relevant to detect – and that is a matter of scientific context. For example, in the case of minimally interesting effect sizes, the minimal acceptable level of distress may vary according to context (farming versus toxicology). In the case of Bayes factors, if the theory tested is that compound XY will relieve pain by binding (to an unknown degree) to receptor A, then if the

existing analgesic that binds optimally to receptor A reduces pain by 5 units on a 0-9 Likert scale, H1, the predicted degree of pain relief produced by XY, can be modelled as a uniform [0, 5 Likert units]. On the other hand, if the theory were that XY binds to an unknown degree to receptor B, then if the existing analgesic that binds optimally to receptor B reduces pain by 2 units on a 0-9 Likert scale, H1 can be modelled as a uniform distribution [0, 2 units]. The different theories make different predictions, and that changes how easy it is to obtain evidence for no effect. There is no such thing as evidence for no effect in the absence of theory. So make sure minimal interesting effects (for power or inference by intervals) or expected effect sizes (for Bayes factors) are justified for reasons relevant to the theory being tested. Only by using theoretically relevant effects can the theory being put up to test be severely tested (Dienes, 2023).

Lining up predictions and tests

Consider a planned study, submitted to a journal, measuring Stroop interference (a reaction time difference) at four points in time (four blocks of trials, within-participants), and testing two groups, anxious versus non-anxious. The stated prediction is “Interference will decrease over blocks more slowly for anxious than non-anxious participants.” The authors will analyze the data with a 2 (group, anxious vs non-anxious) X 4 (time, block1 vs block2 vs block3 vs block4) ANOVA. The number of participants is determined by a power analysis: “The power to detect Stroop interference being above zero is 0.90 with $N = 30$ for $\alpha = .02$.” But that $N=30$ has been chosen by a power analysis on a different test (testing the mere existence of interference) than the test of their prediction (differential change in interference over time). So the paper is sent back at triage. It comes back with a new calculation: “The power for the two –way interaction group x time ($df = 3$) is 0.90 with $N = 150$ for $\alpha = .02$.” Now the sample size has been determined by reference to the group X time interaction, and the prediction is indeed about how the change over time is different for the different groups. The paper is still sent back at triage for revising this precise point. Why?

The prediction is about something decreasing; that is, it is about one thing moving as a whole. It is a one degree of freedom prediction. The suggested test has three degrees of freedom; that is, it is a test of a complex pattern in which three independent things can move in three different directions at once. So the test is not precisely a test of the prediction. The prediction refers to the extent to which interference decreases over blocks. This can be measured by a linear contrast. Let $B1$ be the interference in block one, $B2$ the interference in block two, and so on. Then let $L = (-\frac{3}{4}) \times B1 + (-\frac{1}{4}) \times B2 + (\frac{1}{4}) \times B3 + (\frac{3}{4}) \times B4$. This measures the linear change in interference over blocks. Note the positive coefficients sum to 1 and the negative ones to -1; the expected size of L in raw units is just a bit smaller than the expected difference between the two extreme conditions. Thus, the heuristics discussed in Dienes (2019, 2021a,b for the difference between conditions may be useful for such contrasts. (Useful rule of thumb: In calculating a contrast, make positive coefficients sum to 1, negative to -1, then it may be easy to think about expected effect sizes for the contrast in raw units.)

We can now phrase the test of prediction: $L_{\text{anxious}} - L_{\text{non-anxious}} > 0$, where $L_{\text{non-anxious}}$ is the linear contrast L for the non-anxious group, and L_{anxious} is L for the anxious group. This test is just a t-test. Power should be determined for this test.

Consider what test precisely tests your prediction. If your prediction says something will be greater or lesser than something else, it is one degree of freedom, and so the test should be likewise. This is partly a matter of tying down analytic flexibility: If you allow a multi-degree of freedom test to be the test of a one-degree of freedom prediction you are allowing a range of outcomes to support your theory. Consider also that you will have more power for the same number of participants for a one degree of freedom test than a multi-degree of freedom test. Further, any multi-degree of freedom result will need to be unpacked into single degree of freedom effects. What does the theory really predict? Does it really not make one-degree of freedom predictions?

You should line up all your predictions, both for outcome neutral tests and for testing the substantial theory. For each prediction determine a precise test that tests exactly that prediction. For each of those tests you need to show you have sufficient sensitivity (i.e. a sufficient samples size for adequate power, or probability of achieving conclusions with inference by intervals or by Bayes factors).

According to *PCI RR* (2023) guidelines “For full Stage 1 manuscripts ... describing quantitative research that involves the testing of hypotheses, inclusion of a study design template is a mandatory requirement; for all other quantitative and qualitative research modes it is strongly encouraged but not required....

Question	Hypothesis	Sampling plan	Analysis Plan	Rationale for deciding the sensitivity of the test for confirming or disconfirming the hypothesis	Interpretation given different outcomes	Theory that could be shown wrong by the outcomes

Guidance Notes

- **Question:** articulate each research question being addressed in one sentence.
- **Hypothesis:** where applicable, a prediction arising from the research question, stated in terms of specific variables rather than concepts. Where the testability of one or more hypotheses depends on the verification of auxiliary assumptions (such as positive controls, tests of intervention fidelity, manipulation checks, or any other quality checks), any tests of such assumptions should be listed as hypotheses. Stage 1 proposals that do not seek to test hypotheses can ignore or delete this column.
- **Sampling plan:** For proposals using inferential statistics, the details of the statistical sampling plan for the specific hypothesis (e.g. power analysis, Bayes Factor Design Analysis, ROPE etc). For proposals that do not use inferential statistics, include a description and justification of the sample size.
- **Analysis plan:** For hypothesis-driven studies, the specific test(s) that will confirm or disconfirm the hypothesis. For non-hypothesis-driven studies, the test(s) that will answer the research question.
- **Rationale for deciding the sensitivity of the test for confirming or disconfirming the hypothesis:** For hypothesis-driven studies that employ inferential statistics, an explanation of how the authors determined a relevant effect size for statistical power analysis, equivalence testing, Bayes factors, or other approach.
- **Interpretation given different outcomes:** A prospective interpretation of different potential outcomes, making clear which outcomes would confirm or disconfirm the hypothesis.
- **Theory that could be shown wrong by the outcomes:** Where the proposal is testing a theory, make clear what theory could be shown to be wrong, incomplete, or otherwise inadequate by the outcomes of the research.”

Conclusion

In preparing a Registered Report, ask yourself these questions: 1) What is the most general theory that could be disconfirmed? (Be clear about your substantial theory.) 2) What is the prediction, is it

one -degree of freedom? (It need not be, but often in my experience it is, even when that fact is not at first recognized by the authors.) 3) How does the prediction follow from theory? (What are the assumptions?) 4) Do these assumptions need to be tested? (Determine crucial outcome neutral tests.) 5) Is there sensitivity to establish everything that needs to be established? Line up a statistical test for each prediction and each assumption. Make sure you have resources to run an adequate sample size for each prediction and each assumption. A study design template lining up each prediction and assumption with an appropriate test is likely to be useful.

Registered Reports as a concept go back some decades and were first used in parapsychology in the 1970s and then clinical trials in the 2000s (Chambers & Tzavella, 2020). However, in none of these cases were papers accepted before results were in; and at least in the latter case, adherence to the principles has not been rigorous (Goldacre et al., 2018). The first attempt to enforce their rigorous use was in the explicitly labelled article type Registered Reports in *Cortex* spearheaded by Chris Chambers (2013), and the independently conceived Registered Replication Reports of *Perspectives on Psychological Science* (Simons et al., 2014). People soon appreciated the arguments: Currently, over 300 journals use the Registered Reports format as either a regular article type or a format for special issues (COS, 2023). (As this figure increases constantly, by the time you read this it will be false.) *Comprehensive Results in Social Psychology* set up in 2016 as a journal that deals only in Registered Reports. Registered Reports can now be submitted for any scientific discipline, in for example *Royal Society Open Science*, or *PCI RR*.

Registered Reports have been useful in providing apparently unbiased evidence for whether well-known phenomena predicted by a theory actually exist. The worry arises that with too much analytic flexibility and selective publishing whole literatures might arise around non-existent effects, or over-estimated ones; meta-analyses of those literatures will be distorted and thus not answer the question of what effect there is (e.g. Friesse & Frankenbach, 2020). For example, in the first Registered Replication Report, Alogna et al. (2014) investigated the “verbal overshadowing effect”. Originally, Schooler and Engstler-Schooler (1990) found that verbally summarizing a just witnessed event harms later visual identification, a finding which triggered a small literature. Francis (2012) suggested that given the average observed power for these studies was low, the effect may be purely a product of publication bias. In two studies, involving 31 and 22 labs respectively, with a minimum of 50 and 30 participants per lab, respectively, Alogna et al found evidence for the effect, with a precise estimate of its size.

Similarly, a special issue of *Comprehensive Results in Social Psychology* investigated various claimed effects of power posing (Cesario, Jonas, & Carney, 2017). Carney et al. (2010) suggested that holding an expansive versus contractive posture for two minutes would change one’s psychological and hormonal states towards those that co-occur with having a position of power. The theory triggered many subsequent studies; an analysis of the pattern of p-values (p-curve analysis) suggested to Simmons & Simonsohn (2015) that selective reporting of studies of an actually non-existent phenomenon may be at work. This is of course a problem that Registered Reports can address. Seven registered reports were submitted to the special issue addressing different aspects of the theory. No study found evidence for any of the hormonal and behavioural predictions of the theory. In a meta-analysis over all seven studies of the dependent variable common to them all, Gronau et al. (2017) concluded there was evidence for people subjectively feeling (or rating themselves as feeling) more powerful after expansive rather than contractive poses. In conclusion, the Registered Reports helped indicate the effect was considerably more limited than originally claimed, but that a core phenomenon remained (perhaps created by demand characteristics? Cf. Dienes et al., 2022).

One argument for Registered Reports is that it reduces the bias that can arise in every stage of the process of testing a theory when decisions about setting up the test can be made in the light of the data obtained from the test. It is an empirical matter whether bias is actually reduced by Registered Reports as they are in practice conducted. There is no guarantee that these biases do often arise just because data are seen first, and no guarantee that Registered Reports will reduce bias (Devezer et al., 2020). So far the evidence supports the claim that Registered Reports reduce biases

that arise in conventional publishing. In a pre-registered study, Scheel, Schijen, and Lakens (2019) found that published Registered Reports obtained results supporting predictions 44% of the time, whereas that rate was 96% in the corresponding standard literature. Similarly, Allen and Mehler (2019) estimated that only 39.5% of predictions in Registered Reports were supported, compared to 80-95% in the standard literature. Further, there is higher reproducibility of the main results from the original data than for regular articles (Obels et al. 2019). Methodological rigour and quality of Registered Reports rated higher than non-Registered Reports papers by judges given sections to read, blind to the status of the article (Soderberg et al., 2021). There is not a single answer to the question of how much Registered Reports reduce bias; the answer depends on context, and on how Registered Reports are actually practised. How much the different biases postulated in this paper are each addressed is unknown. The meta-scientific exploration of Registered Reports should be ongoing. But we already have evidence that the pursuit of Registered Reports is worthwhile.

Peer Community In Registered Reports (PCI RR, <https://rr.peercommunityin.org>) has introduced some publishing innovations to deal with limitations of previous Registered Reports. One difficulty people have faced is the long lead in from submitting the Stage 1 to being able to start collecting data, as there will be normal length of time back and forth between reviewers and the author before the Stage 1 receives In Principle Acceptance. PCI RR has a “scheduled review”, in which in the first instance the authors submit a “snapshot” or one-page summary of the proposed research, and they promise to deliver a Stage 1 manuscript a fixed period of time later, e.g. in 6 weeks (https://rr.peercommunityin.org/help/guide_for_authors#h_61998243643551613309672490). The editor (or “recommender” as they are called in PCI RR) then recruits reviewers who guarantee a review within a few days of the fixed date of Stage 1 submission. This dramatically cuts down the initial review period to only a few days. PCI RR also has programmatic reports, where a single Stage 1 can set out several studies, each of which can be submitted as a separate Stage 2 paper (https://rr.peercommunityin.org/help/guide_for_authors#h_52492857233251613309610581), thereby also reducing the total time involved in getting out the final Stage 2 papers. PCI RR allows for the fact that some of the benefits of Registered Reports can be achieved for existing data sets. One just needs to make explicit the degree of bias control achieved, and note it is less to various extents depending on the situation. For example, the existing data may have been inaccessible by the authors; or, by contrast, the authors could in principle have accessed the data but claim they have not (thus a lower degree of bias control needs to be noted). PCI RR explicitly marks six different levels of bias control (https://rr.peercommunityin.org/help/guide_for_authors#h_95790490510491613309490336).

One argument against Registered Reports might be that the timing of a prediction relative to data is metaphysically irrelevant. Imagine there are 10 different possible differences that could be obtained. Ten people pick one difference as the key one they predict, based only on a gut reaction they have. Andy predicts difference 1, Bonnie predicts difference 2, ...Jenny predicts difference 10. Data are collected and there is evidence for difference 2, and there is evidence against all other differences. What are we to make of this? First, from the point of view of advancing science, it is irrelevant that Bonnie predicted difference 2. That is, predicting in advance is not in itself important. Now change the scenario so that 10 different theories potentially explaining a phenomenon, not necessarily mutually exclusive, each predict one of the differences. Now it is very relevant that theory 2 predicts difference 2. We have learned a lot; theory 2 is supported and the rest of the theories disconfirmed. What is important for science is the objective relation between theory, predictions and data. The timing of the different components is not metaphysically important. But notice that the (effectively time-independent) objective relations between theory, predictions and data was the view of science this paper based its arguments on for Registered Reports. In sum, we as humans have to judge the relations between the components, and that may be one reason why using data to support a theory that came later has a bad name: Introducing data first motivates biased judgments. Thus, the argument from bias motivates the use of Registered Reports, regardless of any metaphysical status accorded timing per se.

When one considers the components that go into testing a theory, and tries to make each component as rigorous as possible to reform the scientific mess we are in, and puts the pieces together, the result is a tighter integration than some of us may have been used to, where, for example, theory penetrates even the statistical tests conducted. Go through this process perhaps even only once, and you might find you do not look upon science quite the same again.

Acknowledgements. Thanks to Chris Chambers, Rob McIntosh, and John Edlund for many helpful suggestions.

REFERENCES

Alger, B. E. (2019). *Defense of the Scientific Hypothesis: From Reproducibility Crisis to Big Data*. New York: Oxford University Press.

Allen, C., & Mehler, D. M. A. (2019). Open science challenges, benefits and tips in early career and beyond. *PLOS Biology*, 17(5), e3000246.

Alogna, V. K., Attaya, M. K., Aucoin, P., Bahník, S., Birch, S., Birt, A. R., ... Zwaan, R. A. (2014). Registered replication report: Schooler & Engstler-Schooler (1990). *Perspectives on Psychological Science*, 9, 556–578.

Baribault, B., Donkin, C., Little, D. R., Trueblood, J., Oravecz, Z., van Ravenzwaaij, D., White, C. N., de Boeck, P., & Vandekerckhove, J. (2018). Metastudies for robust tests of theory. *Proceedings of the National Academy of Sciences*, 115, 2607-2612.

Baumeister, R. F. (2016). Charting the future of social psychology on stormy seas: Winners, losers, and recommendations. *Journal of Experimental Social Psychology*, 66, 153-158

Baumeister, R. F., Masicampo, E. J., & Dwall, C. N. (2009). Prosocial benefits of feeling free: disbelief in free will increases aggression and reduces helpfulness. *Personality and Social Psychology Bulletin*, 35, 260-268.

Benjamin, D. J., Berger, J. O., Johannesson, M., et al. (2017). Redefine statistical significance. *Nature Human Behaviour*, doi: 10.1038/s41562-017-0189-z

Braithwaite, V. (2010). *Do Fish Feel Pain?* Oxford: Oxford University Press.

Carp, J. (2012). On the plurality of (methodological) worlds: estimating the analytic flexibility of fMRI experiments. *Frontiers in Neuroscience*, <http://dx.doi.org/10.3389/fnins.2012.00149>

Carney, D. R., Cuddy, A. J. C., & Yap, A. J. (2010). Power posing: Brief nonverbal displays affect neuroendocrine levels and risk tolerance. *Psychological Science*, 21, 1363–1368.

Carney, D. R., Cuddy, A. J. C., & Yap, A. J. (2015). Review and summary of research on the embodied effects of expansive (vs. contractive) nonverbal displays. *Psychological Science*, 26, 657–663.

Carroll, C., Patterson, M., Wood, S. Booth, A., Rick, J., & Balain, S. (2007). A conceptual framework for implementation fidelity. *Implementation Science*, 2, 40. <https://doi.org/10.1186/1748-5908-2-40>

Cesario, J., Jonas, K. J. & Carney, D. R. (2017) CRSP special issue on power poses: what was the point and what did we learn? *Comprehensive Results in Social Psychology*, 2:1, 1-5, DOI: 10.1080/23743603.2017.1309876

Chambers, C. D. (2013). Registered reports: a new publishing initiative at Cortex. *Cortex*, 49, 609-610.

Chambers, C. D. (2019). What's next for registered reports? *Nature*, 573 (7773), 187-189.

Chambers, C. D., & Tzavella, L. (2020). Registered Reports: Past, Present and Future. <https://doi.org/10.31222/osf.io/43298>

Chavalarias, D., & Ioannidis, J. P. A. (2010). Science mapping analysis characterizes 235 biases in biomedical research. *Journal of Clinical Epidemiology*, 63, 1205-1215.

Chuard, P. J. C., Vrtílek, M., Head, M. L., & Jennions, M. D. (2019) Evidence that nonsignificant results are sometimes preferred: Reverse P-hacking or selective reporting? *PLoS Biol* 17(1): e3000127. <https://doi.org/10.1371/journal.pbio.3000127>

Cortex (2020).

http://cdn.elsevier.com/promis_misc/PROMIS%20pub_idt_CORTEX%20Guidelines_RR_29_04_2013.pdf Downloaded 20 April 2020.

COS (2023). <https://www.cos.io/initiatives/registered-reports> Viewed 16 March 2023.

Deci, E. L., & Ryan, R. M. (2000). The “what” and “why” of goal pursuits: human needs and the self-determination of behavior. *Psychological Inquiry*, 11, 227–268.

Deutsch, D. (2011). *The beginning of infinity: Explanations that transform the world*. London: Penguin.

Devezer, B., Navarro, D. J., Vandekerckhove, J., Buzbas, E. O. (2020). The case for formal methodology in scientific reform. bioRxiv 2020.04.26.048306; doi: <https://doi.org/10.1101/2020.04.26.048306>

Dienes, Z. (2016). How Bayes factors change scientific practice. *Journal of Mathematical Psychology*, 72, 78-89.

Dienes, Z. (2019). How do I know what my theory predicts? *Advances in Methods and Practices in Psychological Science*, 2, 364-377. <https://doi.org/10.1177/2515245919876960>

Dienes, Z. (2021a). Obtaining evidence for no effect. *Collabra: Psychology*, 7 (1): 28202. <https://doi.org/10.1525/collabra.28202>

Dienes, Z. (2021b). How to use and report Bayesian hypothesis tests. *Psychology of Consciousness: Theory, Research, and Practice*, 8, 9–26

Dienes, Z. (2023). Testing theories with Bayes factors. In Austin Lee Nichols & John E. Edlund (Eds), *Cambridge Handbook of Research Methods and Statistics for the Social and Behavioral Sciences, volume 1: Building a program of research*, pp 494-512. Cambridge University Press.

Dienes, Z., Palfi, B., & Lush, P. (2022). Controlling phenomenology by being unaware of intentions. In Weisberg, J. (Ed.), *Qualitative Consciousness: Themes from the Philosophy of David Rosenthal*. (pp 229-242). Cambridge University Press

Dutilh, G., Vandekerckhove, J., Ly, A., Matzke, D., et al. (2017). A test of the diffusion model explanation for the worst performance rule using preregistration and blinding. *Attention, Perception, & Psychophysics*, 79, 713–725. doi.org/10.3758/s13414-017-1304-y

EJN (2020). https://wol-prod-cdn.literatumonline.com/pb-assets/assets/14609568/Registered_Report_author_guide.pdf Downloaded 30 July 2020

Fiedler, K. (2017). What Constitutes Strong Psychological Science? The (Neglected) Role of Diagnosticity and A Priori Theorizing. *Perspectives on Psychological Science*, 12, 46-61.

Flake, J. K., & Fried, E. I. (2019). Measurement Schmeasurement: Questionable Measurement Practices and How to Avoid Them. <https://doi.org/10.31234/osf.io/hs7wm>

Flake, J. K., Davidson, I. J., Wong, O., & Pek, J. (2022). Construct validity and the validity of replication studies: A systematic review. *American Psychologist*, 77(4), 576 -588.

Francis, G. (2012). Too good to be true: Publication bias in two prominent studies from experimental psychology. *Psychonomic Bulletin & Review*, 19, 151–156.

French, C. T., Diekemper, R. L., Irwin, R. S., and on behalf of the CHEST Expert Cough Panel (2015). Assessment of Intervention Fidelity and Recommendations for Researchers Conducting Studies on the Diagnosis and Treatment of Chronic Cough in the Adult: CHEST Guideline and Expert Panel Report. *Chest*, 148, 32–54.

Friese, M., & Frankenbach, J. (2020). p-Hacking and publication bias interact to distort meta-analytic effect size estimates. *Psychological Methods*, 25(4), 456–471.

Gelman, A., & Loken, E. (2014). The statistical crisis in science: data-dependent analysis--a 'garden of forking paths'--explains why many statistically significant comparisons don't hold up. *American Scientist*, 102, 460+. <https://link.gale.com/apps/doc/A389260653/AONE> Accessed 21 Apr. 2020.

Goldacre, B, DeVito, N. J., Heneghan, C., Irving, F., Bacon, S., Fleminger, J., & Curtis, H. (2018). Compliance with requirement to report results on the EU Clinical Trials Register: cohort study and web resource. *BMJ* 362, k3218.

Greenwald, A. G. (1975). Consequences of prejudice against the null hypothesis. *Psychological Bulletin*, 82, 1–20. DOI: 10.1037/h0076157

Gronau, Q. F., van Erp, S., Heck, D. A., Cesario, J., Jonas, K. J., & Wagenmakers, E. J. (2017). Bayesian model-averaged meta-analysis of the power – pose effect with informed and default priors: The case of felt power. *Comprehensive Results in Social Psychology*, 2, 123–138

Guest, O., & Martin, A. E. (2020, February 19). How computational modeling can force theory building in psychological science. <https://doi.org/10.31234/osf.io/rybh9>

Haig, B. D. (2014). *Investigating the psychological world: Scientific method in the behavioural sciences*. MIT Press: London.

Ikeda, A., Xu, H., Fuji, N., Zhu, S., & Yamada, Y. (2019). Questionable research practices following pre-registration. *心理学評論*, 62(3), 281-295.

Jeffreys, H. (1939). *The theory of probability*. Oxford, England: Oxford University Press.

Khan, K. A. & Cheung, P. (2020). Presence of mismatches between diagnostic PCR assays and coronavirus SARS-CoV-2 genome. *Royal Society Open Science*, 7200636
<http://doi.org/10.1098/rsos.200636>

Kerr, N. L. (1998). HARKing: Hypothesizing after the results are known. *Personality & Social Psychology Review*, 2, 196-217.

Klein, S. B. (2014). What can recent replication failures tell us about the theoretical commitments of psychology? *Theory & Psychology*, 24 (3), 326-338. DOI: 10.1177/0959354314529616

Kruschke, J. K. (2018). Rejecting or Accepting Parameter Values in Bayesian Estimation. *Advances in Methods and Practices in Psychological Science*, 1, 270-280.

Lakatos, I. (1970). Falsification and the methodology of scientific research programmes. In Lakatos, I., & Musgrave, A. (Eds) (1970) *Criticism and the growth of knowledge*. Cambridge University Press.

MacCoun, R., & Perlmutter, S. (2015). Hide results to seek the truth. *Nature*, 526, 187–189.

Mayo, D. G. (2018). *Statistical inference as severe testing: How to get beyond the statistics wars*. Cambridge: Cambridge University press.

McIntosh, R. D. (2017). Exploratory reports: A new article type for *Cortex*. *Cortex*, 96, A1-A4.

Newman, J., & Taylor, A. (1992). Effect of a Means-End Contingency on Young Children's Food Preferences. *Journal Of Experimental Child Psychology*, 64, 200-216.

NHB (2020). <https://www.nature.com/naturebehav/registeredreports> Downloaded 20 April 2020.

Obels, P., Lakens, D., Coles, N. A., Gottfried, J., & Green, S. A. (2020). Analysis of open data and computational reproducibility in registered reports in psychology. *Advances in Methods and Practices in Psychological Science*, 3(2), 229-237.

Orben, A., & Przybylski, A. K. (2020). Teenage sleep and technology engagement across the week. *PeerJ* 8:e8427 <https://doi.org/10.7717/peerj.8427>

Palfi, B., & Dienes, Z. (2019a). When and how to calculate the Bayes factor with an interval null hypothesis. <https://doi.org/10.31234/osf.io/9chmw>

Palfi, B., & Dienes, Z. (2019b). The role of Bayes factors in testing interactions.
<https://doi.org/10.31234/osf.io/qjrg4>

Palfi, B., & Dienes, Z. (2020). Why Bayesian “evidence for H1” in one condition and “evidence for H0” in another does not mean Bayesian evidence for a difference between conditions. *Advances in Methods and Practices in Psychological Science*, 3, 300–308.

PCI RR (2023) https://rr.peercommunityin.org/help/guide_for_authors

- Popper, K. R. (1959). *The logic of scientific discovery*. London: Hutchinson.
- Popper, K. R. (1963). *Conjectures and Refutations: The growth of scientific knowledge*. London: Routledge.
- Popper, K. R. (1972). *Objective Knowledge: An Evolutionary Approach*. Oxford: Oxford University Press.
- Rouder, J. N. (2014). Optional stopping: No problem for Bayesians. *Psychonomic Bulletin & Review*, 21, 301–308.
- Rouder, J. N., & Haaf, J. M. (2019). Optional Stopping and the Interpretation of Bayes Factor. <https://osf.io/uv456/>
- RSOS (2020). <https://royalsocietypublishing.org/rsos/registered-reports> Downloaded 20 April 2020.
- Schooler, J. (2010). What Science Tells Us About Free Will. In Baumister R.F, Mele A.R, Vohs K.D (Eds), *Free Will and Consciousness: How Might They Work?* (pp 191-218). Oxford: Oxford University Press.
- Schooler, J. W., Engstler-Schooler, T. Y. (1990). Verbal overshadowing of visual memories: Some things are better left unsaid. *Cognitive Psychology*, 22, 36–71.
- Scheel, A. M., Schijen, M., & Lakens, D. (2020). An excess of positive results: Comparing the standard Psychology literature with Registered Reports. [https://osf.io/p6e9c\(2020\)doi:10.31234/osf.io/p6e9c](https://osf.io/p6e9c(2020)doi:10.31234/osf.io/p6e9c)
- Schönbrodt, F. D., & Wagenmakers, E. J. (2018). Bayes factor design analysis: Planning for compelling evidence. *Psychonomic Bulletin & Review*, 25, 128–142.
- Schönbrodt, F. D., Wagenmakers, E.-J., Zehetleitner, M., & Perugini, M. (2017). Sequential hypothesis testing with Bayes factors: Efficiently testing mean differences. *Psychological Methods*, 22, 322–339.
- Silberzahn, R., Uhlmann, E. L., Martin, P., Anselmi, P. et al. (2018). Many Analysts, One Data Set: Making Transparent How Variations in Analytic Choices Affect Results. *Advances in Methods and Practices in Psychological Science*, 1, 337-356.
- Simmons, J., Nelson, L., and Simonsohn, U. (2011). False-positive psychology: Undisclosed flexibility in data collection and analysis allow presenting anything as significant. *Psychological Science* 22, 1359–1366.
- Simons, D. J., Holcombe, A. O., & Spellman, B. A. (2014). An Introduction to Registered Replication Reports at Perspectives on Psychological Science. *Perspectives on Psychological Science*, 9, 552-555.
- Simmons, J., & Simonsohn, U. (2015, May). Power posing: Reassessing the evidence behind the most popular TED talk. Retrieved from <http://datacolada.org/2015/05/08/37-power-posing-reassessing-the-evidence-behind-the-most-popular-ted-talk/>

Simonsohn, U., Simmons, J. P. & Nelson, L. D. (2019). Specification Curve: Descriptive and Inferential Statistics on All Reasonable Specifications. Available at SSRN: <https://ssrn.com/abstract=2694998> or <http://dx.doi.org/10.2139/ssrn.2694998>

Soderberg, C. K., Errington, T. M., Schiavone, S. R., Bottesini, J., Thorn, F. S., Vazire, S., ... & Nosek, B. A. (2021). Initial evidence of research quality of registered reports compared with the standard publishing model. *Nature Human Behaviour*, 5(8), 990-997.

Steege, S., Tuerlinckx, F., Gelman, A., & Vanpaemel, W. (2016). Increasing Transparency Through a Multiverse Analysis. *Perspectives on Psychological Science*, 11, 702-712.

Wagenmakers, E.-J., Verhagen, A. J., Ly, A., Matzke, D., ... Morey, R. D. (2017). The need for Bayesian hypothesis testing in psychological science. In S. Lilienfeld & I. Waldman, (Eds.), *Psychological science under scrutiny: Recent challenges and proposed solutions* (pp. 123-138). New York: John Wiley and Sons.

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, 627-633.

Westlake, W. J. (1972). Use of confidence intervals in analysis of comparative bioavailability trials. *Journal of Pharmaceutical Sciences*, 61, 1340-1341. DOI: 10.1002/jps.2600610845

Yanai, I., & Lercher, M. (2020). A hypothesis is a liability. *Genome Biology*, 21, 231 <https://doi.org/10.1186/s13059-020-02133-w>