

# Fact or fiction: reducing the proportion and impact of false positives

D. Stahl and A. Pickles

Department of Biostatistics and Health Informatics, Institute of Psychiatry, Psychology and Neuroscience, King's College London, London, UK

## Invited Review

**Cite this article:** Stahl D, Pickles A (2018). Fact or fiction: reducing the proportion and impact of false positives. *Psychological Medicine* **48**, 1084–1091. <https://doi.org/10.1017/S003329171700294X>

Received: 7 April 2017

Revised: 1 September 2017

Accepted: 7 September 2017

First published online: 27 November 2017

### Key words:

Bayes' factor; cross-validation; external validation; false positives; meta-analyses; replication crisis; researcher degrees-of-freedom; statistical learning

### Author for correspondence:

A. Pickles, E-mail: [andrew.pickles@kcl.ac.uk](mailto:andrew.pickles@kcl.ac.uk)

## Abstract

False positive findings in science are inevitable, but are they particularly common in psychology and psychiatry? The evidence that we review suggests that while not restricted to our field, the problem is acute. We describe the concept of researcher 'degrees-of-freedom' to explain how many false-positive findings arise, and how the various strategies of registration, pre-specification, and reporting standards that are being adopted both reduce and make these visible. We review possible benefits and harms of proposed statistical solutions, from tougher requirements for significance, to Bayesian and machine learning approaches to analysis. Finally we consider the organisation and methods for replication and systematic review in psychology and psychiatry.

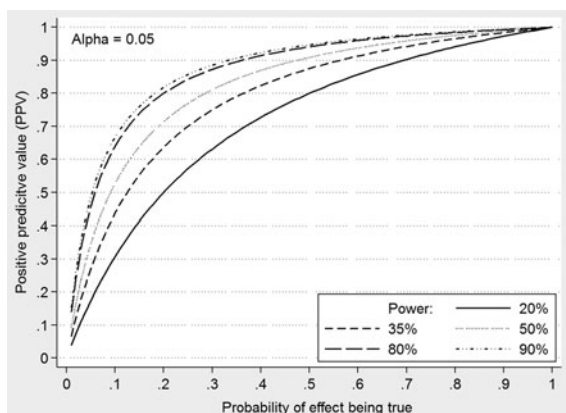
## Introduction

When regarded from the positivist perspective of the scientific method, psychology and neuroscience should be an accumulative, iterative, self-correcting endeavour. While mistakes are inevitable they will typically be identified by the onward march of the research process. However, like all human activity, psychology and neuroscience research is also a social process. It relies on many inter-personal and organisational processes and displays all the social phenomena we observe elsewhere, such as fashions, concern with status and reputation, group-identification, collective judgements, social norms, competitive and defensive actions (Edmonds *et al.* 2011). We are active participants in this complexity though we may sporadically rail against its norms and culture.

While the importance attached to scientific rigour is undoubtedly more than lip-service, the needs of careers, funders, scientific institutions and commercial stakeholders have priorities in addition to or other than the generation of reliable and unbiased findings. That published reports in psychology were over-represented by positive tests of hypotheses has long been known. Already in 1959 and again with colleagues in 1995 Sterling had reported that almost all statistical hypotheses in contemporary psychology journals were significant (Sterling, 1959; Sterling *et al.* 1995) and Fanelli (2010) could not show an improvement in this century, finding more than 90% of published hypotheses were described as positive in spite of typical studies being insufficiently powered for such 'success' (Bakker *et al.* 2012).

Pashler & Wagenmakers (2012) describe this implausible reality as a 'crisis of confidence' and examples of published reports of studies of premonition having achieved statistical significance and survived peer review seemingly particularly damning (Bem, 2011). The Open Science Foundations Reproducibility Project (<https://osf.io/ezcuj/>) attempted to provide a more objective assessment of the reliability of findings by attempting to replicate studies. Repeating almost 100 studies published in three high-ranking psychology journals, they found that less than half of the reported results could be replicated, the figure being somewhat higher for formal experimental studies than observational (Aarts *et al.* 2015). This and many other studies confirmed Ioannidis' claim in his seminal paper (Ioannidis, 2005) that more than half of claimed positive discoveries in medical and psychological studies are false positives. Figure 1 explains Ioannidis' reasoning behind his assumption, now strongly supported by empirical evidence, and why underpowered studies are a major problem for the replication crisis.

However, false positive results from underpowered studies ought to be detected and self-corrected by the 'vigorous and uncompromising scepticism' of the scientific method (Sagan, 1995). However, self-correction is not always swift and it may need time and effort to uproot belief in false positive reports. Pickles (2009) has described a common natural history for many 'findings'. It begins with a small study that finds a significant association and an editor persuaded of its interest. Many small studies follow with the great majority failing to replicate the finding. Everyone knows that small studies have low power so the lack of replication comes as no surprise. Few are published but preference is given to those that do find a significant association, apparently confirming the interesting finding. Often studies are of different



**Fig. 1.** Why more than 50% of published research findings are false. Ioannidis (2005) claimed that more than 50% of published 'significant' research findings in natural sciences are false. As others before him explained that the  $p$  value of a statistical test does not allow making probabilistic statements about the underlying validity of an effect. For that we need to make assumptions about the prior probability that an effect exists. The probability that a statistical test with a positive 'significant' result is a true effect is the positive predicted value (PPV). The PPV depends on our prior probability, the power of a test (the probability of a test to detect a specified effect with a given sample size) and the alpha error. Assume we are assessing 1000 risk factors for a mental health problem and 10% of the risk factors are real (but we do not know which ones). Our statistical tests have got 80% power to get a significant result if a specific effect exists and we usually specify an alpha error level of 0.05 or 5%. We would detect 80% of our 100 risk factors and thus obtain 80 true positive test results. 20 risk factors would not be detected (false negative results). The remaining 900 factors are not risk factors but we would expect that 5% or 45 would be falsely declared as such (false positives). We would therefore obtain  $80 + 45 = 125$  positive results but only 80 (64%) are true positives and about 1/3rd of the positive results are false positives, far more than our alpha error of 5% would suggest. However, the PPV also depends on the power (and therefore the sample size) and the alpha error. Meta-analyses have shown that typically studies in psychology and psychiatry often have little power: an estimate for animal studies with neurological diseases is 20% power (Tsilidis *et al.* 2013) and 35% for psychological studies (Bakker *et al.* 2012) although low power is not restricted to these research domains (Dumas-Mallet *et al.* 2017a). The figure shows the PPV (Y-axis) against the probability that a hypothesis is true for different powered studies (25%, 35%, 50%, 80%, and 90%) at an alpha error of 5%. Even without publication bias towards positive results the PPV of positive statistical results is far less than we may have expected. It underlines the importance of conducting and presenting a power analyses where feasible.

design and non-significant results are explained by small conceptual differences in the design. Eventually, a large study fails to replicate what is now an apparently well-established finding, and because of the study's size and the fact this failure is now seen as overturning a received wisdom, publication in a prominent journal follows. Such frequent diversions into cul-de-sacs, with research following these false trails, slow the real progress that we could be making. With research funding in psychology and mental health being so disproportionately low (Luengo-Fernandez *et al.* 2015), researchers in these areas, in particular, cannot afford such diversions. Moreover, corrupting incentives persist or are increasing. We struggle with replication of simple main effects, so the high-level of enthusiasm for stratified or personalised medicine in which effects are shown in some isolated sub-group serves to exacerbate the problem, and the expectation that researchers show 'impact' encourages clinical dissemination, commercialisation and incorporation into health guidelines prior to proper corroboration.

As the above suggests, we have moved from expressions of concern by a few lone voices, through a period of more widespread unease and some embarrassment, to a series of more

objective assessments and ad hoc proposals and initiatives. However, the problem is many faceted, complicated in technical, sociological and economic terms, and insidious. To change the way we do science requires a form of political movement, so perhaps it is unsurprising therefore that we now have manifestos, for example the Manifesto for Reproducible Science (Munafò *et al.* 2017) that promotes a set of these proposals. Interestingly, while arguing for their broad adoption by researchers, institutions, funders and journals, the authors make explicit that there is also the need for the iterative evaluation and improvement of such proposals. As in any other area of policy, what may sound sensible could nonetheless have unintended and possibly negative consequences that will need evaluation.

In this review, while we highlight a fuller understanding of the statistical and methodological issues and suggestions as to how to improve the unsatisfactory situation, we also underscore how statistics alone is not sufficient. Both as investigators and reviewers we need to reassess what we should be valuing, and how that influences what we do. Our review falls into two major parts. We begin by a description of researcher 'degrees-of-freedom' (Simmons *et al.* 2011), the recognised and unrecognised flexibility of the researcher to 'chase significance'. From this a whole series of practices – both methodological and organisational – have been proposed to make this flexibility visible and then, where required, to curtail it. We discuss methodological solutions to the problem including machine learning methods and explain how many current applications can still mislead.

### Fraud, false positives, and researcher degrees-of-freedom

While there are cases of deliberate fraud, where data have been made-up or intentionally incorrect analysis presented, these remain fortunately relatively rare. Nonetheless a recent meta-analyses show worryingly that about 2% of scientists admitted to having fabricated, falsified or modified data (Fanelli, 2009). But these represent the extreme end of a much larger problem where researchers may be turning a blind eye to what they know is not fully rigorous science or may be entirely unaware that what they have done results in misrepresentation.

We tend to think about the choices available to a researcher as being concerned with the study design and the instruments and schedule to be used in the measurement protocol. However, once the data are collected, the researcher still has a vast number of choices to be made before a report of findings is finalised. These choices imply alternatives and the extent to which the researcher can pick one or another, dependent upon how appealing are the results of such choices, and are referred to as degrees of freedom. Perhaps the most familiar are those associated with having multiple risk variables and multiple outcomes. A single  $p$  value is interpretable against the standard probability threshold of 0.05 if it is the only test being undertaken. With multiple risks and multiple outcomes we face the problem of multiple comparisons, for which Bonferroni proposed a well-known correction, tightening the required threshold the more comparisons the researcher makes (Bender & Lange, 2001).

However, before users of multiple-comparison adjustment methods feel too self-righteous, we should consider some of the other researcher degrees of freedom (Simmons *et al.* 2011). There may be multiple alternative definitions or constructions of even single risk factors or outcome disorders. Over how long/what follow-up period should these be measured? What transformations can be applied and what does the researcher do

with potential outliers? What participants should be included or excluded? What statistical test should be used? Will additional data be collected if a finding is not quite significant such that there is an implicit series of interim analyses? Will univariate pre-screening be performed to reduce the number of variables? Gelman & Loken, (2014) describe how these researcher degrees of freedom substantially increase the multiple testing problem but may not feel like fishing for significant results to the researcher. Exploiting these degrees-of-freedom can dramatically increase the actual false positive rate and rarely yields replicable results (Simmons *et al.* 2011; Harrell, 2015). Many researchers are not aware of the consequences of such subtle types of fishing and they need to be more aware of the negative impact of ignoring researcher's degrees of freedom. Researchers should list their use of degrees of freedoms to determine the potential for bias e.g. by using a checklist (Wicherts *et al.* 2016) and if possible correct for them in their analyses (Harrell, 2015). However, given the current state of affairs, a paper from 72 researchers from around the world (Benjamin *et al.* 2017) have proposed the blunt approach of reducing the standard alpha level for claiming significance from the gold standard of 0.05 to 0.005, with consequent implications for larger sample size to maintain power.

The null-hypothesis testing framework of Fisher, Neyman and Pearson works well if correctly applied, especially in double blinded placebo controlled randomised trial with a single outcome. However, it becomes more problematic with increasing numbers of often unplanned tests and choices, giving  $p$  values that are too small, confidence intervals too narrow and the final model from model selection too complex (Freedman, 1983; Harrell, 2015). Some have proposed Bayesian hypothesis testing instead, using the Bayes' factor -the ratio of the evidence for the alternative compared with the null (Rouder *et al.* 2009; Ly *et al.* 2016). Johnson (2013) showed that the significance level of 0.05 generally corresponds to a Bayes Factor of 5 or less, a range generally regarded as providing only modest evidence. Unsurprisingly therefore, given the use of a more demanding criterion, a Bayesian re-analysis of Bem's extrasensory perception (ESP) study by Wagenmaker and colleagues found little evidence for ESP (Wagenmakers *et al.* 2011) and Wetzels *et al.* (2011) reanalysis of 259 psychology articles found only 'anecdotal evidence' for 70% of the findings previously claimed as significant. Thus these arguments for a Bayesian approach as a solution appear to share much with the proposal for a stricter 0.005 criterion for nominal significance and may not escape the limitations of dichotomous thinking (Cumming, 2014).

Fixing upon an even lower alpha level may create new problems and could prevent innovation in science. Many sound research studies would become infeasible where funding envelopes and participant availability is fixed. Small but well-designed studies are already not expected to produce significant results at an alpha level of 0.05 and yet can still provide important insight for future research (see e.g. Parmar *et al.* (2016) for a discussion in the context of clinical trials). Making the practice of science more rigorous by reducing researcher degrees of freedom is therefore also necessary.

### Prespecification, registration, and oversight

An explicit consideration of these and similar questions have become routine for those working in pharmaceutical trials, where both the financial incentive and the potential harms to patients are large and obvious, and drug regulatory authorities

and independent trialists have been formalising procedures to defend against false positives.

In trials, the intention to undertake a study is made public by registration (e.g. at <http://www.isrctn.com> or <http://www.clinicaltrials.gov>), the details of the intended design and measures published in the Trial Protocol, and sometimes too the Statistical Analysis Plan, that has been first approved by the independent Data Monitoring Committee as a fair and complete specification of the intended analysis. These documents reduce the degrees-of-freedom of the researcher, laying the foundation for effective disclosure and oversight, and for a direct comparison of the study report against the originally intended study (<http://compare-trials.org/>). For the field of psychology and psychiatry as a whole, registration also helps improve the visibility of the evidence base for otherwise unpublished negative findings (Munafa & Black, 2017).

Any investigator first experiencing this process finds it challenging. Firstly, it soon becomes apparent just how many research questions had been previously poorly defined and how many decisions, some large and many small, had been left implicit, postponed or simply not made. Secondly, each of these decisions corresponds to giving up a researcher degree-of-freedom, and many find this quite painful. In particular, committing to one definition of risk, outcome and one analysis is seen as 'risking' missing important positive findings. Of course those 'missed' findings would have been much more likely to be a false positive than any positive finding from a pre-specified test of a well-founded research hypothesis.

It should also be understood that pre-specification is not intended to bind investigators come what may. Departures from the prespecification can, and often should be made. But when, how and the reasons why should be explicit in the study publication (Simmons *et al.* 2011) so that HARKing (Hypotheses after the results are known (Kerr *et al.* 1999), the practice of presenting results of posthoc or other explorative analyses as research hypotheses, can be prevented without denying the researcher the right to present explorative, novel findings.

### Reporting and reviewing

First published in 1996 the Consolidated Standards of Reporting Trials (CONSORT) has become a model for more than 350 reporting guidelines that now span laboratory measurement to epidemiology, and now helpfully collated by the Equator Network (<http://www.equator-network.org>). Reporting guidelines ease the writing of a protocol for pre-registrations and a report of a study drafted following these guidelines should allow an independent researcher to reproduce the published findings. The fact that such a report would exceed the word length of traditional journals is, with the use of online supplements, now of little consequence.

Pre-specification and reporting guidelines make the task of the reviewer much simpler – clarity of research question, appropriateness of design and the evidence of rigorous implementation should now be obvious. However, it is important to recognise, especially for editors, that the combination of pre-specification and reporting standards that press for full disclosure leave the authors exposed to unconstructive criticism and constrained in their ability to respond, for example to requests for unplanned revised analysis. Publishing reviewer's comments and responses may also be useful to encourage constructive criticism.

## Organisation of research

Description and awareness of researcher's degrees of freedom, while good practice, does not solve the problem of underpowered studies. University and college-based undergraduate and post-graduate psychology is hugely popular, and the common training vehicle is the small-scale structured experiment. Schaller (2016) suggests that the tested hypotheses are often little more than personal opinion, rarely being based on formal deduction from well-founded theory (e.g. Festinger & Hutte, 1954) and often with no or an over-optimistic consideration of likely effect-size, both increasing the chances of a false positive. Small enough for completion by a single student or scientist rehearsing basic study design and with greater emphasis on methodology to achieve high internal validity, little priority is given to external validity and generalisability. A simple way to improve the situation was suggested by Daniel Kahneman (cited in Yong, 2012) by turning undergraduate and graduate projects into replication studies.

In some areas of research increased scale can offer advantages over and above a reduction in false positives, such as where industrial scale brings the opportunity for industrial processing with improvements in standardisation, quality control and economy. But for studies involving, say, neuro-imaging, face-to-face interviewing or coding of observational data, as yet there seem few such economies. Thus, outside of genetics and a small number of epidemiological studies, truly large studies are rare in our field and more common has been the formation of consortia. A frequently adopted model is one where contributing groups agree a common set of core measures that will enable large sample analysis of a limited number of key research questions, while retaining scope for site specific independent extensions. Such consortia have the advantages of being potentially sustained by a mixture of core and site-specific funding, and can also provide a framework for intellectual collaboration, training and exchange. However, they can also incur significant costs with respect to much time-consuming discussion about process, agreements, standards, publication rights and so on. For scientists used to the rapid pursuit of flights of scientific imagination this can be anathema. However, incorporating a wider brief, including training and a sharing of scientific ideas, can make such organisations both productive and healthy environments for early career and experienced researchers alike. The BASIS (British Autism Sibling Infant Study) Consortium (<http://www.basisnetwork.org>) is one such which is more focussed on protocol sharing and making a core study sample available to a wider group of researchers who undertake largely independent research.

## Machine learning and cross-validation

Increasing study size or establishing consortium is not always feasible and current methodology needs to be optimised as well. Increasingly journals already encourage the use of new statistics by focusing on confidence intervals and effect sizes (Eich, 2014; Giofrè *et al.* 2017) or the use of Bayesian statistics (Dienes, 2011; Gallistel, 2015).

An interesting alternative statistical approach to reducing the problem of false positive results is assessing prediction accuracy of unseen cases as a measure of evidence is based on work by statisticians in the 1970s (Allen, 1974; Stone, 1974; Browne, 1975; Geisser, 1975; Harrell *et al.* 1996) and nowadays widely used in the field of machine learning. The contribution of machine learning to look at inference as a search through a space of possible

hypotheses to identify one or more supported by the data (Hunter, 2017) is an important one, and unlike the frequentist and Bayesian approaches, it allows the inclusion of data-preprocessing steps, such as variable transformation, variable selection or imputation of missing data within the model selection process (Kuhn & Johnson, 2013) and can thus account for some of the problems of researcher's degrees of freedom. The prediction capability of a model for independent, unseen data is different to 'classical' inferential statistics where efficient unbiased estimation of parameters is the goal. This means that the best predictive model may differ from the best explanatory model (Hoerl & Kennard, 1970; Sober, 2006; Hastie *et al.* 2009). The first minimises the prediction error of unseen data while the second minimises prediction error of the training dataset (Shmueli, 2010; Shmueli & Koppius, 2011).

Less well understood is that model selection and model assessment are two separate goals and cannot be assessed on the same unseen dataset (Varma & Simon, 2006; Hastie *et al.* 2009). This is often ignored in the machine learning community (Cawley & Talbot, 2010). If we want to perform model selection we would need to randomly divide the dataset into three parts: a training set to fit the models, a validation set to estimate prediction error, and a test set to assess the generalisation error of the final chosen model (Hastie *et al.* 2009). This three-way split-sample approach may be feasible for big-data but is usually not possible in medical research and is anyway inefficient and potentially unreliable (Steyerberg, 2009; Harrell, 2015). Better alternatives are nested cross-validation or bootstrap validation (see supplementary procedures). Being more complex and computationally demanding these are often avoided, simpler two-split cross-validation being used instead with users often unaware that this provides incomplete correction (Stone, 1974; Hastie *et al.* 2009; Harrell, 2015). Also, the fashion to maximise prediction by using Support Vector Machines or Deep Learning should also be questioned, as these offer poor interpretability as compared with regularised methods and simpler models may perform better in practical replication studies (Hand, 2006).

## External validation, replication, and triangulation

Cross-validation can deliver sound internal validation. For external validation to confirm the generalisability of findings replication with a new study sample of sufficient size to have good power for a plausible effect size is the gold standard. Journals need to accept replication studies and these should be linked to the original studies with the original paper. However, three highly influential journals refused to publish a (negative) replication of Bem's ESP study because of their policy not publishing straight replication studies (Yong, 2012; Gelman & O'Rourke, 2014). Even 2017, only 3% of 1151 psychology journals welcome scientists to submit replication studies (Martin & Clarke, 2017). Pre-registration of replication studies with a commitment to publish them irrespective of the outcome would help. Still better would be the encouragement of complementary studies to partner, sometimes in reciprocal agreements, enabling primary and replication studies to appear in the same or linked reports. The Dutch cohort study Generation-R and the UK Avon Longitudinal Study of Pregnancy and Childhood (ALSPAC), have examined the consistency of the effects of parental depression and anxiety during pregnancy (Van Batenburg-Eddes *et al.* 2013), finding maternal symptoms increased attentional problems in both studies, but neither convincingly excluding the possibility

of a confounder induced association. The ability to nominate a respected study as a replication partner within a research funding proposal is an important strength, emphasising the willingness of the primary study to put its findings under immediate independent testing. Outside of genetic studies, however, such partnering remains rare.

The term replication presupposes much about the equivalence of the primary and replication populations. For example, sometimes primary study samples for developing a classifier are selected as pure cases and non-cases, whereas the replication sample might include more every-day patients. Discrepancies can occur even when the causal process of interest is the same in the two samples, but the pattern of confounders is different. It is therefore recommended to quantify the predictive accuracy of a prediction model in different samples from the same or similar target populations or domains (Debray *et al.* 2015). This needs also to be considered for 'classical' statistical modelling. External validation studies may range from temporal (e.g. sample from the same hospital or primary care practice only later in time), to geographical (e.g. sample from different hospital, region, or even country), to validations across different medical settings (e.g. from secondary to primary care setting or vice versa) or different target populations or domains (e.g. from adults to children) with increasingly different study samples or case mix between development and validation samples. But is it exact/direct replication that is being sought, where as far as possible everything is held identical to the original study, or is it reproducibility (Drummond, 2009) – the replication of the concept, where some aspects of the study are allowed or designed to be different? The latter is more appealing to funders and editors in that it has novel elements, but the associated researcher degrees of freedom that comes with these must be made visible and controlled. Triangulation could be considered a special case of a reproducibility study, where the differences in the studies relate to the differences in the likely impact or measurement of confounders rather than the core variables (Lawlor *et al.* 2016).

### Meta-analysis/systematic review

Systematic review, the exhaustive collation of research findings, their structured evaluation against indicators of quality, and their formal quantitative combination has become the accepted basis for recommendations from organisations such as NICE (National Institute for Clinical Excellence <http://www.nice.org>) charged with formulating evidence based recommendations as to good clinical practice for the National Health Service in England. The best method available to establish consensus in multiple studies addressing the same issue is meta-analysis (Cumming, 2014). In therapeutic research systematic review has been led by the Cochrane collaboration that started in 1993. However at the end of 2014, of the more than 5000 Cochrane Reviews, only 2% (119) were within the areas of developmental, psychosocial and learning problems. It appears that many reviews fail to identify sufficient studies of eligible quality to justify a formal review panel. This cannot be considered satisfactory.

The small number of eligible studies also does not allow an adequate assessment of publication or other biases, such as selective analyses or outcome reporting towards significant studies (Ioannidis, 2008; Ioannidis & Karassa, 2010; Sterne *et al.* 2011). Based on their expected power, Tsilidis *et al.* (2013) examined animal studies of neurological disorders finding that almost twice as many than expected showed a significant result. The

discrepancy was larger in studies with small sample sizes. Even in randomised clinical trials for the efficacy of psychotherapies for major depressive disorder an excess of significant studies is evident (49% expected *v.* 58% reported, Flint *et al.* 2015). Although researchers are nowadays fully aware of the file drawer problem it seems still to exist. Finally, results are also sensitive to study inclusion criteria e.g. meta-analyses of moderating effect of serotonin transporter genotype on the association between stressful life events and depression resulted in different conclusions based on different (but sensible) inclusion criteria (Taylor & Munafò, 2016). Psychiatric studies may be particular prone to this problem because of the variety of outcomes measures of similar cognitive deficits and of types of control treatments or groups.

Thus even results from Cochrane-based meta-analyses need to be treated with care, and biases towards positive results assessed (Ioannidis & Trikalinos, 2007), perhaps using the online *p* value analyser (<http://www.p-curve.com>) of Simonsohn *et al.* (2014) and used by Taylor & Munafò (2016). Schuit *et al.* (2015) suggest a simple correction based on the number of studies, the number of significant results and the sample size of studies to keep the nominal alpha error at the pre-specified (i.e. 5%) level without loss of power. The use of network meta-analyses (Caldwell *et al.* 2005; Caldwell, 2014), which allows the comparison of a network of related studies addressing a common condition or outcome should also be encouraged.

We have already noted that there are instances where small scale studies remain essential (Parmar *et al.* (2016)). Small trials of good quality may also be desirable from the point of view of generalisation, with results of meta-analyses involving several smaller independent studies potentially being more reliable than performing one large study as long as all studies are available for a meta-analysis (Cappelleri *et al.* 1996; Contopoulos-Ioannidis *et al.* 2005).

### Conclusions

Statisticians are professional sceptics. That more than half of published findings are not reproducible came as no surprise to us. The environment in which scientists work, namely what we teach, the training and career opportunities, data and publication infrastructures, and individual and institutional incentives all require careful monitoring and, in many cases, revision. We would encourage modernisation of the curriculum for research methods to get away from presenting context-free statistical tests, to highlight instead the systematic and where possible pre-specified approaches to each step of the research process, and statistical modelling and analysis methods that more properly account for and describe uncertainty – the latter requiring further development work from statisticians. In spite of the claims of some, machine learning alone will not solve the problem, and badly done will make the problem worse. Well done, which is currently not that easy, it has a valuable place in our overall endeavour. Statistical learning, the study of properties of learning algorithm from the perspective of statistical theory can serve as a unifying of statistical modelling and machine learning (Hastie *et al.* 2009). Perhaps we might even need to change the way research excellence is defined, away from the number of high impact publications and high h-index towards a replication index which measures how often other scientists could replicate their results as Chamber & Sumner (2012) suggested in a newspaper article.

Of course, the problem does not stop in the scientific community. Dumas-Mallet *et al.* (2017b) reported that newspapers also selectively cover mainly exciting positive findings of lifestyle risk factors studies, often triggered by exaggerated press releases from Universities (Sumner *et al.* 2014). Here psychiatry performed poorly with only ~25% of studies being supported by later meta-analyses. Even more worryingly, newspapers rarely inform the public about the subsequent null finding. The crisis of replication is thus not only a problem of science but also of the media, as the autism and vaccination ‘debate’ illustrated only too well.

**Supplementary material.** The supplementary material for this article can be found at <https://doi.org/10.1017/S003329171700294X>.

**Acknowledgements.** Our work was in part supported by the National Institute for Health Research (NIHR) Biomedical Research Centre for Mental Health at the South London and Maudsley NHS Foundation Trust and Institute of Psychiatry, Psychology & Neuroscience, Kings College London. The views expressed in this publication are those of the author(s) and not necessarily those of the NHS, the National Institute for Health Research or the Department of Health.

## References

- Aarts AA, Anderson JE, Anderson CJ, Attridge PR, Attwood A, Axt J, Babel M, Bahnik S, Baranski E, Barnett-Cowan M, Bartmess E, Beer J, Bell R, Bentley H, Beyan L, Binion G, Borsboom D, Bosch A, Bosco FA, Bowman SD, Brandt MJ, Braswell E, Brohmer H, Brown BT, Brown K, Bruning J, Calhoun-Sauls A, Callahan SP, Chagnon E, Chandler J, Chartier CR, Cheung F, Christopherson CD, Cillessen L, Clay R, Cleary H, Cloud MD, Cohn M, Cohoon J, Columbus S, Cordes A, Costantini G, Alvarez LDC, Cremata E, Crusius J, DeCoster J, DeGaetano MA, Della Penna N, den Bezemer B, Deserno MK, Devitt O, Dewitte L, Dobolyi DG, Dodson GT, Donnellan MB, Donohue R, Dore RA, Dorrrough A, Dreber A, Dugas M, Dunn EW, Easey K, Eboigbe S, Eggleston C, Embley J, Epskamp S, Errington TM, Estel V, Farach FJ, Feather J, Fedor A, Fernandez-Castilla B, Fiedler S, Field JG, Fitneva SA, Flagan T, Forest AL, Forsell E, Foster JD, Frank MC, Frazier RS, Fuchs H, Gable P, Galak J, Galliani EM, Gampa A, Garcia S, Gazarian D, Gilbert E, Giner-Sorolla R, Glockner A, Goellner L, Goh JX, Goldberg R, Goodbourn PT, Gordon-McKeon S, Gorges B, Gorges J, Goss J, Graham J, Grange JA, Gray J, Hartgerink C, Hartshorne J, Hasselman F, Hayes T, Heikensten E, Henninger F, Hodsoll J, Holubar T, Hoogendoorn G, Humphries DJ, Hung COY, Immelman N, Irsik VC, Jahn G, Jakel F, Jekel M, Johannesson M, Johnson LG, Johnson DJ, Johnson KM, Johnston WJ, Jonas K, Joy-Gaba JA, Kappes HB, Kelso K, Kidwell MC, Kim SK, Kirkhart M, Kleinberg B, Knezevic G, Kolorz FM, Kossakowski JJ, Krause RW, Krijnen J, Kuhlmann T, Kunkels YK, Kyc MM, Lai CK, Laique A, Lakens D, Lane KA, Lassetter B, Lazarevic LB, LeBel EP, Lee KJ, Lee M, Lemm K, Levitan CA, Lewis M, Lin L, Lin S, Lippold M, Loureiro D, Luteijn I, Mackinnon S, Mainard HN, Marigold DC, Martin DP, Martinez T, Masicampo EJ, Maticotta J, Mathur M, May M, Mechin N, Mehta P, Meixner J, Melinger A, Miller JK, Miller M, Moore K, Moschl M, Motyl M, Muller SM, Munafò M, Neijenhuijs KI, Nervi T, Nicolas G, Nilsson G, Nosek BA, Nuijten MB, Olsson C, Osborne C, Ostkamp L, Pavel M, Penton-Voak IS, Perna O, Pernet C, Perugini M, Pipitone RN, Pitts M, Plessow F, Prenoveau JM, Rahal RM, Ratliff KA, Reinhard D, Renkewitz F, Ricker AA, Rigney A, Rivers AM, Roebke M, Rutchick AM, Ryan RS, Sahin O, Saide A, Sandstrom GM, Santos D, Saxe R, Schlegelmilch R, Schmidt K, Scholz S, Seibel L, Selterman DF, Shaki S, Simpson WB, Sinclair HC, Skorinok JLM, Slowik A, Snyder JS, Soderberg C, Sonnleitner C, Spencer N, Spies JR, Steegen S, Stieger S, Strohminger N, Sullivan GB, Talhelm T, Tapia M, te Dorsthorst A, Thomae M, Thomas SL, Tio P, Traets F, Tsang S, Tuerlinckx F, Turchan P, Valasek M, van 't Veer AE, Van Aert R, van Assen M, van Bork R, van de Ven M, van den Bergh D, van der Hulst M, van Dooren R, van Doorn J, van Renswoude DR, van Rijn H, Vanpaemel W, Echeverria AV, Vazquez M, Velez N, Vermue M, Verschoor M, Vianello M, Voracek M, Vuu G, Wagenmakers EJ, Weerdmeester J, Welsh A, Westgate EC, Wissink J, Wood M, Woods A, Wright E, Wu S, Zeelenberg M, Zuni K
- Open Science Collaboration** (2015) Estimating the reproducibility of psychological science. *Science* 349, 943–952.
- Allen DM** (1974) Relationship between variable selection and data augmentation and a method for prediction. *Technometrics* 16, 125–127.
- Bakker M, van Dijk A, Wicherts JM** (2012) The rules of the game called psychological science. *Perspectives on Psychological Science* 7, 543–554.
- Bem DJ** (2011) Feeling the future: experimental evidence for anomalous retroactive influences on cognition and affect. *Journal of Personality and Social Psychology* 100, 407–425.
- Bender R, Lange S** (2001) Adjusting for multiple testing – when and how? *Journal of Clinical Epidemiology* 54, 343–349.
- Benjamin DJ, Berger JO, Johannesson M, Nosek BA, Wagenmakers E-J, Berk R, Bollen KA, Brembs B, Brown L, Camerer C, Cesarini D, Chambers CD, Clyde M, Cook TD, De Boeck P, Dienes Z, Dreber A, Easwaran K, Efferson C, Fehr E, Fidler F, Field AP, Forster M, George EI, Gonzalez R, Goodman S, Green E, Green DP, Greenwald AG, Hadfield JD, Hedges LV, Held L, Hua Ho T, Hoijtink H, Hruschka DJ, Imai K, Imbens G, Ioannidis JPA, Jeon M, Jones JH, Kirchler M, Laibson D, List J, Little R, Lupia A, Machery E, Maxwell SE, McCarthy M, Moore DA, Morgan SL, Munafò M, Nakagawa S, Nyhan B, Parker TH, Pericchi L, Perugini M, Roulder J, Rousseau J, Savalei V, Schönbrodt FD, Sellke T, Sinclair B, Tingley D, Van Zandt T, Vazire S, Watts DJ, Winship C, Wolpert RL, Xie Y, Young C, Zinman J & Johnson VE** (2017) Redefine statistical significance. *Human Nature Behavior* 1, 1–5.
- Browne MW** (1975) Comparison of single sample and cross-validation methods for estimating mean squared error of prediction in multiple linear regression. *British Journal of Mathematical and Statistical Psychology* 28, 112–120.
- Caldwell DM** (2014) An overview of conducting systematic reviews with network meta-analysis. *Systematic Reviews* 3, 109–109.
- Caldwell DM, Ades AE, Higgins JPT** (2005) Simultaneous comparison of multiple treatments: combining direct and indirect evidence. *BMJ* 331, 897–900.
- Cappelleri JC, Ioannidis JPA, Schmid CH, deFerranti SD, Aubert M, Chalmers TC, Lau J** (1996) Large trials vs meta-analysis of smaller trials – how do their results compare? *Jama-Journal of the American Medical Association* 276, 1332–1338.
- Cawley GC, Talbot NLC** (2010) On over-fitting in model selection and subsequent selection bias in performance evaluation. *Journal of Machine Learning Research* 11, 2079–2107.
- Chamber C, Sumner P** (2012) Replication is the only solution to scientific fraud. *The Guardian* [online] (<https://www.psychologicalscience.org/observer/bayes-for-beginners-probability-and-likelihood>). Accessed 8 September 2017.
- Contopoulos-Ioannidis DG, Gilbody SM, Trikalinos TA, Churchill R, Wahlbeck K, Ioannidis JPA, Project, E.-P.** (2005) Comparison of large versus smaller randomized trials for mental health-related interventions. *American Journal of Psychiatry* 162, 578–584.
- Cumming G** (2014) The new statistics. *Psychological Science* 25, 7–29.
- Debray TPA, Vergouwe Y, Koffijberg H, Nieboer D, Steyerberg EW, Moons KGM** (2015) A new framework to enhance the interpretation of external validation studies of clinical prediction models. *Journal of Clinical Epidemiology* 68, 280–289.
- Dienes Z** (2011) Bayesian versus orthodox statistics: which side are you on? *Perspectives on Psychological Science* 6, 274–290.
- Drummond C** (2009) Replicability is not reproducibility: nor is it good science. In *Proceedings of the Twenty-Sixth International Conference on Machine Learning*. ICML: Montreal, Canada, p. 4.

- Dumas-Mallet E, Button KS, Boraud T, Gonon F, Munafo MR (2017a). Low statistical power in biomedical science: a review of three human research domains. *Royal Society Open Science* 4, 160254.
- Dumas-Mallet E, Smith A, Boraud T, Gonon F (2017b). Poor replication validity of biomedical association studies reported by newspapers. *Plos ONE* 12, 1–5.
- Edmonds B, Gilbert N, Ahrweiler P, Scharnhorst A (2011) Simulating the social processes of science. *Jasss-the Journal of Artificial Societies and Social Simulation* 14, 14.
- Eich E (2014) Business not as usual. *Psychological Science* 25, 3–6.
- Fanelli D (2009) How many scientists fabricate and falsify research? A systematic review and meta-analysis of survey data. *Plos ONE* 4, 1–11.
- Fanelli D (2010) “Positive” results increase down the hierarchy of the sciences. *Plos ONE* 5, 1–10.
- Festinger L, Hutte HA (1954) An experimental investigation of the effect of unstable interpersonal relations in a group. *Journal of Abnormal and Social Psychology* 49, 513–522.
- Flint J, Cuijpers P, Horder J, Koole SL, Munafo MR (2015) Is there an excess of significant findings in published studies of psychotherapy for depression? *Psychological Medicine* 45, 439–446.
- Freedman DA (1983) A note on screening regression equations. *American Statistician* 37, 152–155.
- Gallistel R (2015) Bayes for beginners: probability and likelihood. *Observer Magazine* 28(7) [online] (<https://www.psychologicalscience.org/observer/bayes-for-beginners-probability-and-likelihood>). Accessed 8 September 2017.
- Geisser S (1975) Predictive sample reuse method with applications. *Journal of the American Statistical Association* 70, 320–328.
- Gelman A, Loken E (2014) The statistical crisis in science. *American Scientist* 102, 460–465.
- Gelman A, O'Rourke K (2014) Discussion: difficulties in making inferences about scientific truth from distributions of published *p*-values. *“Biostatistics (Oxford, England)”* 15, 18–23.
- Giofrè D, Cumming G, Fresc L, Boedker I, Tressoldi P (2017) The influence of journal submission guidelines on authors' reporting of statistics and use of open research practices. *PLOS ONE* 12, e0175583.
- Hand DJ (2006) Classifier technology and the illusion of progress, *Statistical Science* 21, 1–14.
- Harrell F (2015) *Regression Modeling Strategies*. Springer: New York, USA.
- Harrell FE, Lee KL, Mark DB (1996) Multivariable prognostic models: issues in developing models, evaluating assumptions and adequacy, and measuring and reducing errors. *Statistics in Medicine* 15, 361–387.
- Hastie T, Tibshirani R, Friedman J (2009) *The Elements of Statistical Learning*. Springer-Verlag: New York.
- Hoerl AE, Kennard RW (1970) Ridge regression – biased estimation for non-orthogonal problems. *Technometrics* 12, 55–67.
- Hunter L (2017) An introduction to machine learning for statisticians. [online] ([http://compbio.ucdenver.edu/Hunter\\_lab/Hunter/ml-for-stat.txt](http://compbio.ucdenver.edu/Hunter_lab/Hunter/ml-for-stat.txt)). Accessed 8 September 2017.
- Ioannidis JPA (2005) Why most published research findings are false. *Plos Medicine* 2, 696–701.
- Ioannidis JPA (2008) Interpretation of tests of heterogeneity and bias in meta-analysis. *Journal of Evaluation in Clinical Practice* 14, 951–957.
- Ioannidis JPA, Karassa FB (2010) The need to consider the wider agenda in systematic reviews and meta-analyses: breadth, timing, and depth of the evidence. *British Medical Journal* 341, c4875.
- Ioannidis JPA, Trikalinos TA (2007) The appropriateness of asymmetry tests for publication bias in meta-analyses: a large survey. *CMAJ: Canadian Medical Association Journal* 176, 1091–1096.
- Johnson VE (2013) Uniformly most powerful Bayesian tests. *Annals of Statistics* 41, 1716–1741.
- Kerr NL, Niedermeier KE, Kaplan MF (1999) Bias in jurors vs bias in juries: new evidence from the SDS perspective. *Organizational Behavior and Human Decision Processes* 80, 70–86.
- Kuhn M, Johnson K (2013) *Applied Predictive Modelling*. Springer-Verlag: New York.
- Lawlor DA, Tilling K, Davey Smith G (2016) Triangulation in aetiological epidemiology. *International Journal of Epidemiology* 45, 1866–1886.
- Luengo-Fernandez R, Leal J, Gray A (2015) UK research spend in 2008 and 2012: comparing stroke, cancer, coronary heart disease and dementia. *British Medical Journal Open* 5, e006648.
- Ly A, Verhagen J, Wagenmakers E-J (2016) Harold Jeffreys's default Bayes factor hypothesis tests: explanation, extension, and application in psychology. *Journal of Mathematical Psychology* 72, 19–32.
- Martin GN, Clarke RM (2017) Are psychology journals anti-replication? A snapshot of editorial practices. *Frontiers in Psychology* 8, 1–6.
- Munafo MR, Black S (2017) Personality and smoking status: a longitudinal analysis (vol 9, pg 397, 2007). *Nicotine and Tobacco Research* 19, 129–129.
- Munafo MR, Nosek BA, Bishop DVM, Button KS, Chambers CD, Percie du Sert N, Simonsohn U, Wagenmakers E-J, Ware JJ, Ioannidis JPA (2017) A manifesto for reproducible science. *Nature Human Behaviour* 1, 0021.
- Parmar MKB, Sydes MR, Morris TP (2016) How do you design randomised trials for smaller populations? A framework. *BMC Medicine* 14, 183.
- Pashler H, Wagenmakers EJ (2012) Editors' introduction to the special section on replicability in psychological science: a crisis of confidence? *Perspectives on Psychological Science* 7, 528–530.
- Pickles A (2009) What clinicians need to know about statistical issues and methods. In *Rutter's Child and Adolescent Psychiatry* (ed. M. Rutter, D.V.M. Bishop, D. S. Pine, S. Scott, S. Stevenson, E. Taylor, A. Thapar), pp. 111–122. Wiley: Oxford.
- Rouder JN, Speckman PL, Sun D, Morey RD, Iverson G (2009) Bayesian t tests for accepting and rejecting the null hypothesis. *Psychonomic Bulletin and Review* 16, 225–237.
- Sagan C (1995) *The Demon-Haunted World: Science as A Candle in the Dark*. Random House: New York.
- Schaller M (2016) The empirical benefits of conceptual rigor: systematic articulation of conceptual hypotheses can reduce the risk of non-replicable results (and facilitate novel discoveries too). *Journal of Experimental Social Psychology* 66, 107–115.
- Schuit E, Roes KC, Mol BW, Kwee A, Moons KG, Groenwold RH (2015) Meta-analyses triggered by previous (false-) significant findings: problems and solutions. *Systematic Reviews* 4, 57.
- Shmueli G (2010) To explain or to predict? *Statistical Science* 25, 289–310.
- Shmueli G, Koppius OR (2011) Predictive analytics in information systems research. *Management Information Systems Quarterly* 35, 553–572.
- Simmons JP, Nelson LD, Simonsohn U (2011) False-positive psychology: undisclosed flexibility in data collection and analysis allows presenting anything as significant. *Psychological Science* 22, 1359–1366.
- Simonsohn U, Nelson LD, Simmons JP (2014) P-curve: a key to the file-drawer. *Journal of Experimental Psychology General* 143, 534–547.
- Sober E (2006) Parsimony. In *The Philosophy of Science: An Encyclopaedia* (ed. S. Sarkar, J. Pfeifer), pp. 531–538. Routledge: Oxford.
- Sterling TD (1959) Publication decisions and their possible effects on inferences drawn from tests of significance – or vice versa. *Journal of the American Statistical Association* 54, 30–34.
- Sterling TD, Rosenbaum WL, Weinkam JJ (1995) Publication decisions revisited – the effect of the outcome of statistical tests on the decision to publish and vice-versa. *American Statistician* 49, 108–112.
- Sterne JAC, Sutton AJ, Ioannidis JPA, Terrin N, Jones DR, Lau J, Carpenter J, Rücker G, Harbord RM, Schmid CH, Tetzlaff J, Deeks JJ, Peters J, Macaskill P, Schwarzer G, Duval S, Altman DG, Moher D, Higgins JPT (2011) Recommendations for examining and interpreting funnel plot asymmetry in meta-analyses of randomised controlled trials. *British Medical Journal* 343, D002.
- Steyerberg EW (2009) *Clinical Prediction Models: A Practical Approach to Development, Validation, and Updating*. Springer: New York.
- Stone M (1974) Cross-Validatory choice and assessment of statistical predictions. *Journal of the Royal Statistical Society Series B-Statistical Methodology* 36, 111–147.
- Sumner P, Vivian-Griffiths S, Boivin J, Williams A, Venetis CA, Davies A, Ogden J, Whelan L, Hughes B, Dalton B, Boy F, Chambers CD (2014) The association between exaggeration in health related science news and academic press releases: retrospective observational study. *British Medical Journal* 349, g7015.
- Taylor AE, Munafo MR (2016) Triangulating meta-analyses: the example of the serotonin transporter gene, stressful life events and major depression. *BMC Psychology* 4, 23.

- Tsilidis KK, Panagiotou OA, Sena ES, Aretouli E, Evangelou E, Howells DW, Salman RAS, Macleod MR, Ioannidis JPA** (2013) Evaluation of excess significance bias in animal studies of neurological diseases. *Plos Biology* **11**.
- Van Batenburg-Eddes T, Brion MJ, Henrichs J, Jaddoe VWV, Hofman A, Verhulst FC, Lawlor DA, Smith GD, Tiemeier H** (2013) Parental depressive and anxiety symptoms during pregnancy and attention problems in children: a cross-cohort consistency study. *Journal of Child Psychology and Psychiatry* **54**, 591–600.
- Varma S, Simon R** (2006) Bias in error estimation when using cross-validation for model selection. *BMC Bioinformatics* **7**, 91–91.
- Wagenmakers EJ, Wetzels R, Borsboom D, van der Maas HLJ** (2011) Why psychologists must change the way they analyze their data: the case of psi: comment on Bem (2011). *Journal of Personality and Social Psychology* **100**, 426–432.
- Wetzels R, Matzke D, Lee MD, Rouder JN, Iverson GJ, Wagenmakers E-J** (2011) Statistical evidence in experimental psychology. *Perspectives on Psychological Science* **6**, 291–298.
- Wicherts JM, Veldkamp CLS, Augusteijn HEM, Bakker M, van Aert RCM, van Assen MALM** (2016) Degrees of freedom in planning, running, analyzing, and reporting psychological studies: a checklist to avoid p-hacking. *Frontiers in Psychology* **7**, 1–12.
- Yong E** (2012) Bad copy. *Nature* **485**, 298–300.