Commentary:

Fallibility in science: Responding to errors in the work of oneself and others

D. V. M. Bishop

Department of Experimental Psychology
University of Oxford
OX1 3UD, UK
dorothy.bishop@psy.ox.ac.uk

Abstract

Fallibility in science cuts both ways: it poses dilemmas for the scientist who discovers errors in their own work, and for those who discover errors in the work of others. The ethical response to finding errors in one's own work is clear: they should be claimed and corrected promptly. Yet people are often reluctant to 'do the right thing' because of a perception this could lead to reputational damage. Open science practices can help avoid errors and also lead to recognition that mistakes are part of normal science. Indeed, a reputation for scientific integrity can be enhanced by admitting to errors. Conversely, failure to admit to errors can create internal conflict and damage a researcher's reputation in the longer term. I also consider the situation where errors are discovered in the work of others; in the case of honest errors, action must be taken to put things right, but this should be done in a collegial way that offers the researcher the opportunity to deal with the problem themselves. Difficulties arise if those who commit errors are unresponsive or reluctant to make changes, when there is disagreement about whether a dataset or analysis is problematic, or where deliberate manipulation of findings or outright fraud is suspected. I offer some guidelines about how to approach such cases. My key message is that errors are inevitable. In the long run, scientists will not be judged on whether or not they make mistakes, but on how they respond when those mistakes are detected.

Errors in your work: how to respond

Imagine the following scenario:
PhD student, David, has run a series of studies trying to find an impact of brain stimulation on language comprehension in stroke patients. After three studies with null findings, he has changed the design in various ways and is overjoyed when the fourth study gives a statistically significant effect. The paper is published in a prestigious high-impact journal, with David as first author and his eminent supervisor as last author.
The university press office promotes the study and it is featured on National Public Radio. Two weeks later, when preparing slides for a talk at the Society for Neuroscience, David finds the groups were miscoded, and in fact the sham treatment group obtained higher post-training scores.

When I use fictitious examples like this in seminars and ask the audience 'What should David do?' the usual response is that, of course, David should come clean, admit the error and ask for the paper to be retracted. But there is typically nervousness in the room. It is pointed out that that there are massive pressures on him not to do so: the general perception is that admission of error will mean that the reputation of both David and his supervisor will be in tatters, with David's prospects for a future career badly damaged.

Yet there are real-life examples of scientists admitting to honest errors that show that this doom-laden scenario is unrealistic. A recent study considered how reputation is affected by retraction, by comparing subsequent citations of earlier published papers for authors who had a paper retracted vs. a control group who had not (Azoulay, Bonatti, & Krieger, 2017). Retraction of a paper due to researcher misconduct led to a drop in subsequent citations of their earlier work, but there was a smaller effect when honest error was involved – with no evidence of reputational damage for junior researchers. In an interview study of 14 authors whose papers were retracted after they notified the journal of errors, Hosseini et al (2018) found that contrary to their expectations, self-retraction did not damage their reputation and in some cases improved it. This fits with more informal evidence suggesting that there can be reputational advantage from going public in correcting an error: you demonstrate you are someone who values scientific accuracy over your success in publishing (Retraction Watch, 2017). Nevertheless, there may be pressures from institutions or senior colleagues to hide errors, and journal editors are not always supportive. Hosseini et al noted: "*Many authors expected rapid, empathic and detailed responses from journal editors, but reported receiving short, unsympathetic and sometimes unpleasant ones instead.*" (p. 209).

I give some examples from online sources in Box 1. The thought of having to retract a paper can instil fear into the heart of scientists, who see it as equivalent to being named and shamed. There are currently few incentives for honesty, and keeping quiet about an error will often seem the easiest option. Recognising that the threat of bad consequences could act as a deterrent to honest admission of error, Retraction Watch instituted the 'Doing the

Right Thing' award, to 'honor those who clean up the scientific literature' (Oransky & Marcus, 2017).

Box 1

Examples of researchers who highlighted errors in their own work

Richard Mann, a postdoctoral researcher using statistical methods to study behavioural ecology, had published a paper on behaviour in prawns in PLOS Computational Biology with six co-authors. He shared the prawn dataset with a colleague who was looking for data to test out some ideas on numerical integration. On his blog, Mann (2013) described the moment when the colleague rang him to tell him of a fatal error in his analysis. As stated in the retraction notice: "Where each of 102 experiments should have been down-sampled to half the original size for computational efficiency, instead the number of experiments in the data set was repeatedly halved 102 times … results and conclusions were based on only one experimental study, rather than the 102 reported in the paper." The paper was retracted, the analysis redone giving similar findings, and Mann states that, although he had a terrible few months, he did not suffer any long-term stigma.

Pamela Ronald, a professor in plant pathology, became concerned when two of her postdocs could not replicate findings she had published in two high-profile papers on the basis of the immune response in rice. She notified the journal editors and then devoted the next 18 months to try and locate the source of the discrepancy. It turned out that the strains of microbes she had been using were mislabelled, and in 2013 the papers were retracted. The story was covered by Nature News (Gewin, 2015), who noted that this year Ronald published a paper correctly identifying the source of the immune response. She has changed her lab procedures so that three independent researchers now validate new experimental approaches.

Senior neuroscientist Russ Poldrack wrote computer code to classify a set of brain images into classes based on the task being performed. He had submitted a paper based on this analysis for publication, when a student collaborator told him that after obtaining far lower classification accuracy on the same dataset, he found an error in the code. Poldrack's (2013) response was to write a blogpost about this experience, encouraging everyone to share code, use better methods for checking code, and talk about their errors.

There are two further points to take from the David scenario. Awful and embarrassing as it is to admit to error, the alternative, hiding a known error, has to be worse. The person who does this is entering into a Faustian pact to reject science in favour of personal ambition. As data fraudster Diederik Stapel openly admitted, once you embark on this process, it is difficult to stop, but it creates considerable internal conflict (Stapel, 2014, pp. 128-131).

The second point is that although errors can never be eliminated, they can be reduced by adoption of open science practices. Even in situations where the raw data cannot be made completely open, usually because of confidentiality issues, it is often possible to deposit a version that has been modified to remove identifiable information, so others can reproduce

what was done (UK Data Service, undated). For sensitive data, a data-sharing agreement may be needed in addition to anonymization (Medical Research Council, 2017). Regardless of which level of security is required, there should be no barriers to researchers making their analysis code open, so that the analysis steps can be checked. The example from Russ Poldrack in Box 1 illustrates how easy it is even for an experienced scientist to make an error in coding that has serious consequences for results. People often worry that if they make code and data open, then errors will be found. There are two answers to that. First, if you know your code and data will be open, you are likely to check and double-check far more rigorously than if you know they will never be seen by anyone else. So open practices reduce the likelihood of error. But second, errors in analysis scripts are extremely common among scientists who have taught themselves to program (Merali, 2010), and so errors are likely to be present. This is really the whole point: we need to make code and data open so that the errors, which are very likely to be there, can be found. But to encourage people to do that, we must remove any stigma associated with detection of those errors. This is not condoning sloppy science: it is just accepting the reality that we are all fallible.
Of course, making analysis programs open does not guarantee they are free from errors: an example of widely-used neuroimaging software that was discovered to include a bug only after many years of use was reported by Eklund et al (2016a). The authors noted in a subsequent correction that they were not implying that all analyses using the software were erroneous, but rather that it was not possible to establish which were: "*Due to lamentable archiving and data-sharing practices, it is unlikely that problematic analyses can be redone*." (Eklund et al, 2016b). Quite simply, making code and data open does not prevent errors, but it does make it possible to detect them. And as amply documented elsewhere, it brings other benefits to researchers, in terms of improving their science as well as enhancing recognition of their work (Markowetz, 2015; McKiernan et al, 2016).

Errors in someone else's work: how to respond

The prior discussion of errors in one's own work should give clues about how to respond when you find errors in another's work. You would not want to be pilloried for an honest error, so don't pillory others for simple mistakes. In a comment on a blogpost on this topic, Anne Weil (2014) put it very well:

…my first prominent publication was a note tearing down someone else's work. That work had appeared in a major journal and caused quite a stir — but the apparent results were the product of a careless (not dishonest, just careless) mistake in the analysis.
The note pointing this out was not derogatory in tone, nor was it intended to shame, but was doubtless embarrassing to the authors.
Now that I am much older, a little wiser, and a little kinder (and a lot more employed, and thus less vulnerable to jerks) I would send the authors my analysis of their math first and give them the opportunity to correct.
And I hope that my colleagues would give me the same consideration if (when?) I make a stupid mistake.

Life, however, is not always so simple. The researcher whose error is remarked on may respond with anger, denial or silence. This is, of course, a normal human reaction, but it is not a sensible response if the error is unambiguous, as it can damage the author's reputation for integrity. In theory, it should be possible to resolve such issues via the journal that published the original article, but in practice, this process seldom proceeds smoothly. Allison, Brown, George and Kaiser (2016) described how their own attempts to correct substantial errors in others' work met with inaction or delaying tactics by authors and editors, and even demands for payment to publish a letter pointing out the errors. At the time of writing this commentary, it was possible to put the record right by adding a comment in PubMed Commons (Bastian, 2014). The comment was linked to the abstract of the original paper on PubMed and became part of the scientific record. Box 2, examples 1 and 2 illustrate how both authors and other researchers have used PubMed Commons to record a correction. However, despite its utility, PubMed Commons was not widely used by commentators, and was discontinued in February 2018, though the comments remain archived (National Center for Biotechnology Information, 2018).

Box 2

Examples of post-publication commentary on PubMed Commons

1. Author adding minor corrections
https://www.ncbi.nlm.nih.gov/pubmed/28436345
Jim van Os notes some numerical errors in a table.

2. Reviewer correcting an error
https://www.ncbi.nlm.nih.gov/pubmed/28461468
Pavel Nesmiyanov noted that β-endorphin, oxytocin, and dopamine were wrongly described as neuropeptides. Although authors did not respond on PubMed Commons, an erratum was published in the journal.

3. Reviewer critiquing methods
https://www.ncbi.nlm.nih.gov/pubmed/29153326
Franck Ramus criticises small sample size of paper on neurobiological correlates of dyslexia. Authors respond defending the small sample size and arguing their analyses were driven by a priori hypothesis derived from previous study.

4. Reviewer critiquing methods
https://www.ncbi.nlm.nih.gov/pubmed/28706072
Serge Ahmed suggests that a study of planning in ravens needs an additional control for learning of affective value of objects.

5. Reviewer noting over-hyped interpretation of results
https://www.ncbi.nlm.nih.gov/pubmed/28735725
Clive Bates notes that a study on association between vaping and smoking in adolescents has been widely interpreted in the media as showing causal link. Bates adds a link to a more detailed critique of the study.

6. Reviewer raising more serious concerns
https://www.ncbi.nlm.nih.gov/pubmed/17688420
David Nunan notes prior concerns about duplicate data in a paper on diet in congestive heart failure.

Errors in interpretation of data

Another scenario is when research results seem suspect because of concerns about methodology, rather than straightforward errors in calculation or scripting. For instance, a study may lack a control group, be underpowered, use an unreliable measure, or have a major confound. There may be strong suspicion that the author has engaged in p-hacking. These are not simple errors that can be corrected, but they affect the conclusions that can be drawn. All of these are situations where PubMed Commons provided a venue for raising the concerns, as illustrated in Box 2, examples 3-5. With the disappearance of PubMed Commons, there are limited options remain to those who want to engage in post-publication peer review, given that few journals have options for commenting. For researchers who do not have access to a blog, an alternative platform, PubPeer, is likely to become the method of choice for post-publication peer review: an important difference from PubMed Commons is that commentators can be anonymous. Probably because of this, PubPeer has been far more popular than PubMed Commons, but it is also noted for a harsh style of criticism, that can include accusations of malpractice (Dolgin, 2018). This is unfortunate, because it leads to the impression that post-publication peer review typically involves a personal attack. This can polarise debate and make many people reluctant to engage.

My recommendation is that when errors are found, the starting position should be that methodological weaknesses are due to ignorance rather than bad faith. Consider, for instance, p-hacking. The dangers of this practice were pointed out many years ago (de Groot, 2014), but it has been normative for decades in many branches of science, including psychology. Before he moved on to fraud, Stapel (2014) engaged in p-hacking, noting:

What I did wasn't whiter than white, but it wasn't completely black either. It was grey, and it was what everyone did. (p.102)

Even now that it has been prominently demonstrated that p-hacking is a major cause of false positive findings (Simmons, Nelson, & Simonsohn, 2011), many still do not recognise how seriously it can distort results (Nuzzo, 2014). Furthermore, it is likely that p-hacking is deemed acceptable, because it involved 'paltering', i.e. using a truthful statement (e.g., the p-value associated with this contrast is < .05) to mislead by failing to provide relevant contextual information (e.g., this was one of numerous comparisons and would not be statistically significant if correction was made for multiple contrasts) (Rogers, Zeckhauser, Gino, Norton, Schweitzer, 2017).

Scientists are particularly fallible when it comes to presenting the results of others in a literature review. The process of conducting a literature review is likely to be affected by confirmation bias – i.e. seeking and remembering evidence that supports one's position, and ignoring or forgetting evidence that does not (Nickerson, 1998). In this case, paltering is again implicated, because the misinformation in a biased review does not come from fabricated or distorted evidence, but simply from omission of key information. Rogers et al (2017) showed that people judge such omission as less dishonest than inclusion of untrue information. One way of counteracting bias in literature reviews is to require them to adopt the format of a systematic review, in which criteria are specified in advance for deciding which papers to include (Gough, Oliver, & Thomas, 2017; Wicherts, 2017).

Failure to replicate: an unreliable indicator of fallibility

I have focused so far on situations where there are either honest errors in the data or analysis, or methodological weaknesses that compromise conclusions that can be drawn. A much more complicated scenario is when there is difficulty in replicating a published result. Reproducibility has become a hot topic in psychology in recent years (Munafò et al., 2017), with failure to replicate findings in psychology being brought to the fore by an influential study published in Science (Open Science Collaboration, 2015). These developments coincided with growing awareness of p-hacking as an endemic problem for psychology (Simmons et al., 2011), making it easy to conclude that results that did not replicate were indicative of bad science. The key point to note is that while erroneous data, erroneous inferences and failure to control bias can lead to results that do not replicate, one cannot assume that failure to replicate is necessarily the result of any of these types of error. In psychology, where we are dealing with probabilistic phenomena, random noise is always a factor affecting results: our statistical methods are designed to guard against type I and type II errors, but there is an inevitable trade-off, and some statistically significant differences will be false positives, while some failures to find an effect will be false negatives (see Table 2). Replication is important precisely because our confidence in the robustness of a given finding cannot depend on a single study.

So, the question arises as to how researchers should respond when there is a failure to replicate prior work. Given the range of reasons for non-replication, it should not be assumed that a failure to replicate a result is evidence of poor science in the original study. Nevertheless, it is important to uncover reasons for discrepant findings. Ideally, both sets of researchers should work together to consider how to reconcile these. If the original researcher believes that contextual factors or researcher expertise are critical to obtaining the result, then it is up to them to specify more carefully the conditions under which the effect obtains, rather than simply putting forward hypothetical explanations for a null result. When there is a failure to replicate a finding, it is bad if the first response is to disparage the original researchers as incompetent, malign or fraudulent, but it is just as bad if those whose findings were not replicable dismiss the critics as having malevolent motives or lack of expertise. Again, the kudos will go to the researchers who show integrity in putting scientific truth before their own career ambitions.

Table 2
Possible reasons for failure to replicate a scientific result

---

- Initial result was a false positive due to chance variation (type I error)
- Replication study fails to detect true effect because of chance variation (type II error)
- Results are sensitive to contextual factors
- Method requires specific expertise that replicator lacks
- Initial results involved basic data entry, computational or statistical errors
- Initial results were obtained using questionable research practices such as p-hacking
- Researcher committed fraud

Deliberate omission, misrepresentation and fraud

I turn now to those unfortunate situations when it is hard to avoid concluding that a researcher is acting in bad faith. A particularly insidious kind of behaviour involves deliberate selective citation of the literature, or 'cherry-picking'. As with other methodological errors, no person should be pilloried for occasional bias in coverage of a review: indeed, it may be hard to tell if someone who conducts a biased review did so intentionally. Even if one strenuously attempts to avoid bias, use of search terms to identify papers on a topic may miss relevant articles if authors omit keywords unless they refer to positive findings. The question one has to ask is whether there is a persistent pattern of bias, with the author ignoring contrary evidence, even when it is readily available and drawn to their attention. Worse still are cases where results from cited studies are selectively or inaccurately portrayed, giving the impression of a large body of confirmatory work. Authors may claim that there are hundreds of papers supporting their position, yet when the cited sources are examined, it is found that many are only of weak quality or tangential relevance, or may even contain information that supports the opposite view. This is a standard ploy by those promoting pseudoscientific views (Grimes & Bishop, 2017) and needs to be robustly challenged. However, to do so effectively, it may be necessary to trawl through a huge amount of material to reveal the lack of substance in the claims, and meanwhile, amplified by confirmation bias and social media, the original article may have propagated a wildfire of misinformation that is hard to extinguish (Lewandowsky, Ecker, & Cook, 2017).

The next step after distortion of research findings is outright invention of fake data. It is generally assumed that this is rare, though it is by its very nature difficult to get accurate estimates of frequency. If a researcher suspects fraud by another scientist, they are placed in an uncomfortable position, and there is little formal guidance as to how to proceed. Table 2 shows the advice of Uri Simonsohn (2013), who used statistical methods to uncover the fraudulent work of two psychologists.

Table 2
Simonsohn's (2013) recommendations when fraud is suspected

- Replicate analyses across multiple studies before suspecting foul play by a
- given author.
- Compare suspected studies with similar ones by other authors.
- Extend analyses to raw data.
- Contact authors privately and transparently, and give them ample time to consider your concerns.
- Offer to discuss matters with a trusted statistically savvy advisor.
- Give the authors more time.
- If suspicions remain, convey them only to entities tasked with investigating such matters, and do so as discreetly as possible.

p. 1886

Uncovering fraud is extremely important work, but it is not for the faint-hearted. For a start, an accusation of fraud is serious business and requires rock-solid evidence, which can take hours of careful work to discover. Although one would hope that academic institutions would take seriously an accusation of fraud against a staff member, they can be slow to act; it is, of course, important that they consider the possibility that they are dealing with an unjustified attack by those with vested interests or fixed ideas. These do occur, but malign intent should not be the default assumption, unless there are several 'red flags' of the kind noted by Lewandowsky and Bishop (2016). Although there are some notable cases of good practice (e.g. Høj, 2013), there are also many historical instances where institutions closed ranks to protect an eminent researcher (Judson, 2004). This is short-sighted, as the ultimate reputational damage of being revealed to be supporting a fraudster is far worse than any bad publicity from early disclosure of a problem. But the scientist who is trying to put things right can find it to be a lonely and dispiriting process, as James Heathers (2017) documented on his blog. Furthermore, one can expect the fraudster to use every method possible to avoid discovery, because they have built a career on deceit. They are likely to be obstructive and may well attack back, accusing those who are raising questions of ulterior motives. Just like whistle-blowers in other areas of life, those who detect fraud tend to get little thanks from the community whose interests they serve.

General principles for responding to fallibility

Accusations of fraudulent science are thankfully rare, but the spotlight has started to shine increasingly on fallibility in psychology, with some hitherto well-established findings now looking less solid (e.g., O'Donnell et al, 2018). My general rule is that we should never use mockery or personal abuse against other scientists who make honest errors: such behaviour just reinforces people's unwillingness to be open about errors. Nor should we assume that failure to replicate a result is a sign of poor science in the original study: rather it is an indication that more work needs to be done to establish whether, and under what conditions, the result is robust. But a good researcher will not hesitate to note flaws in the scientific work of themselves and others. Criticism is the bedrock of the scientific method. It should not be personal: if one has to point to problems with someone's data, methods or conclusions this should be done without implying they are stupid or dishonest. This is important, because the alternative is that many people will avoid engaging in robust debate because of fears of interpersonal conflict; this is a recipe for scientific stasis. If wrong ideas or results are not challenged, we let down future generations who try to build on a research base that is not a solid foundation. Worse still, where the research findings have practical applications in clinical or policy areas, we may allow wrongheaded interventions or policies to damage the wellbeing of individuals or society. As open science becomes increasingly the norm, we will find that everyone is fallible. The reputation of a scientist will depend not on whether there are flaws in their research, but on how they respond when those flaws are noted.

Author contributions

This manuscript is the sole work of DVMB.

References

Allison, D. B., Brown, A. W., George, B. J., & Kaiser, K. A. (2016). Reproducibility: A tragedy of errors. *Nature, 530*(7588), 27-29. doi:10.1038/530027a

Azoulay, P., Bonatti, A., & Krieger, J. L. (2017). The career effects of scandal: Evidence from scientific retractions. *Research Policy*. doi:10.1016/j.respol.2017.07.003

Bastian, H. (2014). Editorial: A stronger post-publication culture Is needed for better science. *PLOS Medicine, 11*(12), e1001772. doi:10.1371/journal.pmed.1001772 de Groot, A. D. (2014). The meaning of “significance” for different types of research [translated and annotated by Eric-Jan Wagenmakers, Denny Borsboom, Josine Verhagen, Rogier Kievit, Marjan Bakker, Angelique Cramer, Dora Matzke, Don Mellenbergh, and Han L. J. van der Maas]. *Acta Psychologica, 148*, 188-194. doi:http://dx.doi.org/10.1016/j.actpsy.2014.02.001

Dolgin, E. (2018). PubMed Commons closes its doors to comments. *Nature*(News, 2nd February 2018).

Eklund, A., Nichols, T. E., & Knutsson, H. (2016a). Cluster failure: Why fMRI inferences for spatial extent have inflated false-positive rates. *Proceedings of the National Academy of Sciences of the United States of America, 113*(28), 7900-7905. doi:10.1073/pnas.1602413113

Eklund, A., Nichols, T. E., & Knutsson, H. (2016b. Correction: cluster failure: Why fMRI inferences for spatial extent have inflated false-positive rates (vol 113, pg 7900, 2016). *Proceedings of the National Academy of Sciences of the United States of America, 113*(33), E4929-E4929. doi:10.1073/pnas.1612033113

Gewin, V. (2015). Nature News: Rice researchers redress retraction (24 July 2015). *Nature*. doi:10.1038/nature.2015.18055

Gough, D., Oliver, S., & Thomas, J. (2017). *An Introduction to Systematic Reviews*. London: Sage Publications.

Grimes, D. R., & Bishop, D. (2017). Distinguishing polemic from commentary in science: Some guidelines illustrated with the case of Sage and Burgio, 2017. *Child Development, in press*. doi:10.1111/cdev.13013

Heathers, J. (2017). The buck stops nowhere: When research goes wrong, who’s responsible? Retrieved from https://medium.com/@jamesheathers/the-buck-stops-nowhere-8284a57c88c9

Høj, P. (2013). UQ investigates events leading to retraction: Statement from The University of Queensland President and Vice-Chancellor Professor Peter Høj Retrieved from https://www.uq.edu.au/news/article/2013/09/uq-investigates-events-leading-retraction

Hosseini, M., Hilhorst, M., de Beaufort, I., & Fanelli, D. (2018). Doing the right thing: A qualitative investigation of retractions due to unintentional error. *Science and Englneering Ethics, 24*(1), 189-206. doi:10.1007/s11948-017-9894-2

Judson, H. F. (2004). *The great betrayal: Fraud in science.* Orlando, Florida, USA: Harcourt Inc.

Lewandowsky, S., & Bishop, D. V. M. (2016). Don't let transparency damage science. *Nature, 529*(7587), 459-461

Lewandowsky, S., Ecker, U. K. H., & Cook, J. (2017). Beyond misinformation: Understanding and coping with the 'post truth' era. *Journal of Applied Research in Memory and Cognition, 6*, 353-369.

Markowetz, F. (2015). Five selfish reasons to work reproducibly. *Genome Biology, 16*. doi:10.1186/s13059-015-0850-7

McKiernan, E. C., Bourne, P. E., Brown, C. T., Buck, S., Kenall, A., Lin, J., . . . Yarkoni, T. (2016). How open science helps researchers succeed. *Elife, 5*. doi:10.7554/eLife.16800

Medical Research Council. (2017, 6th September). Using information about people in health research *MRC ethics series.* Retrieved from https://www.mrc.ac.uk/documents/pdf/using-information-about-people-in-health-research-2017/

Merali, Z. (2010). Computational science - error: why scientific programming does not compute. *Nature, 467*, 775-777. doi:10.1038/467775a

Munafò, M. R., Nosek, B. A., Bishop, D. V. M., Button, K. S., Chambers, C. D., Percie du Sert, N., . . . Ioannidis, J. P. A. (2017). A manifesto for reproducible science. *Nature Human Behavior, 1*(1: 0021). doi:10.1038/s41562-016-0021

National Center for Biotechnology Information. (2018). PubMed Commons to be discontinued. *NCBI Insights.* Retrieved from https://ncbiinsights.ncbi.nlm.nih.gov/2018/02/01/pubmed-commons-to-be-discontinued/

Nickerson, R. S. (1998). Confirmation bias: a ubiquitous phenomenon in many guises. *Review of General Psychology, 2*(2), 175-220.

Nuzzo, R. (2014). Scientific method: statistical errors. *Nature, 506*, 150-152. doi:10.1038/506150a

O'Donnell, M., Nelson, L. D., Ackermann, E., Aczel, B., Akhtar, A., Aldrovandi, S., . . . Zrubka, M. (2018). Registered Replication Report: Dijksterhuis and van Knippenberg (1998). *Perspectives on Psychological Science*. doi:10.1177/1745691618755704

Open Science Collaboration. (2015). Estimating the reproducibility of psychological science. *Nature, 349*(6251). doi:10.1126/science.aac4716

Oransky, I., & Marcus, A. (2017, 5th May). Introducing the Doing the Right Thing award, honoring those who clean up the scientific literature. Retrieved from https://www.statnews.com/2017/05/05/dirt-award-cleaning-scientific-literature/

Poldrack, R. (2013, 20th February 2013). Anatomy of a coding error Retrieved from http://www.russpoldrack.org/2013/02/anatomy-of-coding-error.html

Retraction Watch. (2017, 27th March). Authors who retract for honest error say they aren't penalized as a result. Retrieved from http://retractionwatch.com/2017/03/27/authors-retract-honest-error-say-arent-penalized-result/ - more-48973

Rogers, T., Zeckhauser, R., Gino, F., Norton, M. I., & Schweitzer, M. E. (2017). Artful paltering: The risks and rewards of using truthful statements to mislead others. *Journal of Personality and Social Psychology, 112*(3), 456-473. doi:10.1037/pspi0000081

Simmons, J. P., Nelson, L. D., & Simonsohn, U. (2011). False-positive psychology. *Psychological science, 22*(11), 1359-1366. doi:10.1177/0956797611417632

Simonsohn, U. (2013). Just post it: The lesson from two cases of fabricated data detected by statistics alone. *Psychological science, 24*(10), 1875-1888. doi:10.1177/0956797613480366

Stapel, D. (2014). *Faking Science: A True Story of Academic Fraud. Translated by Nicholas J. L. Brown*. Retrieved from https://errorstatistics.files.wordpress.com/2014/12/fakingscience-20141214.pdf

UK Data Service. (undated). Anonymisation. Retrieved from https://www.ukdataservice.ac.uk/manage-data/legal-ethical/anonymisation

Weil, A. (2014). Comment on blog by John Hutchinson: Co-rex-ions. Retrieved from https://whatsinjohnsfreezer.com/2014/05/10/co-rex-ions/ - comment-22328

Wicherts, J. (2017). The weak spots in contemporary science (and how to fix them). *Animals, 17*(12), 90. doi:10.3390/ani7120090

Acknowledgements

This commentary is based on a talk given on 7th July 2017 at a meeting on Reprodcible Science for Early Career Researchers. I thank David Mehler from the University of Cardiff for inviting me to present at the meeting and for proposing this topic. I am also grateful to Kendal Smith for constructive comments on a pre-print version of this paper.