arXiv:2105.03082v1 [physics.data-an] 7 May 2021

Comment on "Reproducibility and Replication of Experimental Particle Physics Results"

Andrew Fowlie

Department of Physics and Institute of Theoretical Physics,

Nanjing Normal University, Nanjing, Jiangsu 210023, China[*]

I would like to thank Junk and Lyons [1] for beginning a discussion about replication in high-energy physics (HEP). Junk and Lyons ultimately argue that HEP learned its lessons the hard way through past failures and that other fields could learn from our procedures. They emphasize that experimental collaborations would risk their legacies were they to make a type-1 error in a search for new physics and outline the vigilance taken to avoid one, such as data blinding and a strict 5σ threshold.

The discussion, however, ignores an elephant in the room: there are regularly anomalies in searches for new physics that result in substantial scientific activity but don't replicate with more data. For example, in 2015 ATLAS and CMS showed evidence for a new particle with a mass of about 750 GeV that decayed into two photons [2]. Whilst the statistical significance was never greater than 5σ [3, 4], the results motivated about 500 publications about the new particle, and countless special seminars and talks [5]. The effect did not replicate when the experimental teams analyzed a larger dataset about six months later [6, 7]. Although this was a particularly egregious example, experimental anomalies that garner considerable interest before vanishing are annual events [8].

We are motivated to attempt to control the type-1 error rate because type-1 errors damage our credibility and lead to us squandering our time and resources on spurious effects. Whilst these non-replications aren't strictly type-1 errors as the statistical significance didn't reach the 5σ threshold and no discoveries were announced, we incur similar damaging consequences, so they cannot be ignored. I shall refer to these errors — substantial scientific activity including publicly doubting the null and speculating about new effects when the null was in fact true — as type-$1'$ errors. Whilst type-1 errors appear to be under control in HEP, type-$1'$ errors are rampant. In the following sections, I discuss these errors in the context statistical practices at the LHC.

[*] Andrew.J.Fowlie@NJNU.edu.cn

I. EVIDENCE AND ERROR RATES

Searches for new physics at the LHC are performed by comparing a p-value, p, against a pre-specified threshold, α. There are two common interpretations of this procedure [9]:

1. **Error theoretic** (Neyman and Pearson [10]): By rejecting the null if $p < \alpha$, we ensure a long-run type-1 error rate of α. The threshold α specified the desired type-1 error rate and the p-value was a means to achieving it.

2. **Evidential** (Fisher [11]): The p-value is a measure of evidence of against the null hypothesis. The threshold α specified a desired level of evidence.

Even among adherents of p-values, the latter interpretation is considered unwarranted [12], and it is almost never accompanied by a theoretical framework or justification, or a discussion of the desired and actual properties of p as a measure of evidence.

Unfortunately, Junk and Lyons repeatedly implicitly switch from one to the other. Indeed, the authors interpret p as a measure of evidence and α as a threshold in evidence, e.g., justifying 5σ by "extraordinary claims require extraordinary evidence" and stating that "$[3\sigma]$ or greater constitutes 'evidence'." We know, however, that interpreted as a measure of evidence, p is incoherent [13, 14] and usually overstates the evidence against the null [15–17]. For example, there exists a famous bound [15, 18] implying that under mild assumptions $p = 0.05$ corresponds to at least about 30% posterior probability of the null. This was in fact the primary criticism in Benjamin *et al.* [19]. Consequently, one factor in the prevalence of type-1′ errors may be that

1. Physicists interpret p-values as evidence (as do Junk and Lyons)

2. Based on p-values, physicists overestimate the evidence for new effects

3. Substantial scientific activity on what turn out to be spurious effects

Unfortunately, p-values simply can't give researchers (including Junk and Lyons) what they want — a measure of evidence — leading to wishful and misleading interpretations of p as evidence [20]. This cannot be overcome by better statistical training; it is in inherent deficiency of p-values and no amount of education about them will imbue them with a coherent evidential meaning.

II. CONTROLLING ERRORS

Controlling error rates depends critically on knowing the data collection and analysis plan — the intentions of the researchers and what statistical tests would be performed under what circumstances — and adjusting the p-value to reflect that. There are, however, an extraordinary number of tests performed by ATLAS, CMS and LHCb at the LHC and elsewhere. This already makes it challenging to interpret a p-value at all and undoubtedly contributes to the prevalence of type-1′ errors.

Junk and Lyons rightly celebrate the trend in HEP to publicly release datasets and tools for analyzing them. This, however, raises the specter of data dredging. Massive public datasets [21] combined with recent developments in machine learning [22] could enable dredging at an unprecedented scale. We must think about what precautions we need to prevent misleading inferences being drawn in the future; e.g., pre-registration of planned analyses as a requisite to accessing otherwise open data. Other more radical proposals, to the problems here and elsewhere, include moving away from an error theoretic approach or any approach based on p-values.

III. FINAL WORDS ON "Response to Andrew Fowlie's Comments"

I would like to thank Junk and Lyons [23] for their detailed response to my comments. In the interests of brevity, I respond to only a few of the points (labeled A – H in Junk and Lyons [23]). First, I acknowledge that Junk and Lyons [1] aren't mistakenly equating p with the posterior of the null (C). My previous comment is at fault if it implied otherwise. Similarly, by "evidence" I assumed that the authors meant nothing more than observations that should change our opinion [24, 25]. The fact that p doesn't equal the posterior of the null is trivial and not that interesting. The fact that p is typically much less than the posterior and that there exist theorems demonstrating it across broad classes of priors isn't trivial and shouldn't be taken lightly (B).

Second, I find it an over-simplification to say that experimental particle physics usually only considers an error-theoretic interpretation of p-values (A). In reality, as in Junk and Lyons [1], our interpretation of p is an "anonymous hybrid" [9] of evidential and error theoretic. Why else would we respond at all to anomalies below 5σ if we didn't consider them evidence for new effects?

Lastly, there are two elements to consider regarding our responses to anomalies. First, what is the appropriate response from our community to a particular strength of evidence for a new effect? and second, are we misled by p about that strength of evidence, and would we respond differently if we recognized that p typically overstated the evidence or if we were able to easily incorporate the prior plausibility of the new effect?

REFERENCES

[1] Thomas R. Junk and Louis Lyons, "Reproducibility and Replication of Experimental Particle Physics Results," Harvard Data Science Review **2**, 4 (2020), arXiv:2009.06864 [physics.data-an].

[2] CERN, "ATLAS and CMS physics results from Run 2," (2015).

[3] Vardan Khachatryan *et al.* (CMS), "Search for Resonant Production of High-Mass Photon Pairs in Proton-Proton Collisions at $\sqrt{s} = 8$ and 13 TeV," Phys. Rev. Lett. **117**, 051802 (2016), arXiv:1606.04093 [hep-ex].

[4] Morad Aaboud *et al.* (ATLAS), "Search for resonances in diphoton events at $\sqrt{s} = 13$ TeV with the ATLAS detector," JHEP **09**, 001 (2016), arXiv:1606.03833 [hep-ex].

[5] Robert Garisto, "Editorial: Theorists React to the CERN 750 GeV Diphoton Data," Phys. Rev. Lett. **116**, 150001 (2016).

[6] Morad Aaboud *et al.* (ATLAS), "Search for new phenomena in high-mass diphoton final states using 37 fb^{-1} of proton–proton collisions collected at $\sqrt{s} = 13$ TeV with the ATLAS detector," Phys. Lett. B **775**, 105–125 (2017), arXiv:1707.04147 [hep-ex].

[7] Vardan Khachatryan *et al.* (CMS), "Search for high-mass diphoton resonances in proton–proton collisions at 13 TeV and combination with 8 TeV search," Phys. Lett. B **767**, 147–170 (2017), arXiv:1609.02507 [hep-ex].

[8] Dan Garisto, "The Era of Anomalies," Physics **13**, 79 (2020).

[9] Raymond Hubbard and M. J Bayarri, "Confusion over measures of evidence (p's) versus errors (α's) in classical statistical testing," Am. Stat. **57**, 171–178 (2003).

[10] J. Neyman and E. S. Pearson, "On the problem of the most efficient tests of statistical hypotheses," Philos. Trans. Roy. Soc. London Ser. A **231**, 289–337 (1933).

[11] R. A. Fisher, *Statistical Methods for Research Workers* (Oliver & Boyd, 1925).

[12] Daniël Lakens, "The practical alternative to the p value is the correctly used p value," Perspect.

Psychol. Sci. (2021).

[13] Mark J. Schervish, "*P* Values: What They are and What They are Not," Am. Stat. **50**, 203–206 (1996).

[14] Eric-Jan Wagenmakers, "A practical solution to the pervasive problems of *p* values," Psychon. Bull. Rev. **14**, 779 – 804 (2007).

[15] Thomas Sellke, M. J Bayarri, and J. O Berger, "Calibration of *p* values for testing precise null hypotheses," Am. Stat. **55**, 62–71 (2001).

[16] J. O Berger and Thomas Sellke, "Testing a Point Null Hypothesis: The Irreconcilability of *P* Values and Evidence," J. Am. Stat. Assoc. **82**, 112–122 (1987).

[17] J. O Berger and Mohan Delampady, "Testing precise hypotheses," Statist. Sci. **2**, 317–335 (1987).

[18] V. G. Vovk, "A logic of probability, with application to the foundations of statistics," J. Royal Stat. Soc. **B55**, 317–341 (1993).

[19] D. J. Benjamin, J. O Berger, M. Johannesson, B. A. Nosek, E. Wagenmakers, *et al.*, "Redefine statistical significance," Nat. Hum. Behav. **2**, 6–10 (2017).

[20] Jacob Cohen, "The earth is round ($p < .05$)," Am. Psychol. **49**, 997 – 1003 (1994).

[21] CERN, "CERN announces new open data policy in support of open science," (2020).

[22] Gregor Kasieczka *et al.*, "The LHC Olympics 2020: A Community Challenge for Anomaly Detection in High Energy Physics," (2021), arXiv:2101.08320 [hep-ph].

[23] Thomas Junk and Louis Lyons, "Response to Andrew Fowlie's Comments," Harvard Data Science Review (2021), 10.1162/99608f92.ebd73e8c.

[24] Thomas Kelly, "Evidence," in *The Stanford Encyclopedia of Philosophy*, edited by Edward N. Zalta (Metaphysics Research Lab, Stanford University, 2016) winter 2016 ed.

[25] Richard D. Morey, Jan-Willem Romeijn, and Jeffrey N. Rouder, "The philosophy of Bayes factors and the quantification of statistical evidence," J. Math. Psychol. **72**, 6–18 (2016).