Why p-values are not measures of evidence

Daniël Lakens

<https://orcid.org/0000-0002-0247-239X>

Eindhoven University of Technology

Cite as:

Lakens. D. (2022). Why p-values are not measures of evidence. Trends in Ecology and Evolution.

Correspondence can be addressed to Daniël Lakens, Human Technology Interaction Group, ATLAS 9.042, PO Box 513, 5600MB Eindhoven, The Netherlands. E-mail: D.Lakens@tue.nl. This work was funded by VIDI Grant 452-17-013 from the Netherlands Organisation for Scientific Research.

Muff, Nilsen, O'Hara, and Nater [1] propose to implement the recommendation "to regard P-values as what they are, namely, continuous measures of statistical evidence". This is a surprising recommendation, given that p-values are not valid measures of evidence [2]. If the null-hypothesis is true, p-values are uniformly distributed, and it is just as likely to observe a p-value of 0.001 as it is to observe a p-value of 0.999. Therefore, the interpretation of $p = 0.001$ as 'strong evidence' cannot be defended just because the probability to observe this p-value is very small. After all, if the null hypothesis is true, the probability of observing $p = 0.999$ is exactly as small.

Small p-values do not guide us in the direction of true effects because they are rarely observed when the null-hypothesis is true, but because they are *relatively* less likely to be observed when the null hypothesis is true than when the alternative hypothesis is true. Statisticians have explained that the concept of evidence is necessarily relative. As Goodman and Royall [3] write, "The p-value is not adequate for inference because the measurement of evidence requires at least three components: the observations, and two competing explanations for how they were produced."

Why p-values do not quantify the strength of evidence

The evidential value of a single p-value depends on the statistical power of the test. When the null hypothesis is true $\alpha\%$ of the observed p-values will fall below the alpha level, as p-values are uniformly distributed under the null-hypothesis. The horizontal grey line in Figure 1 illustrates the expected p-value distribution for a two-sided independent *t*-test if the null-hypothesis is true. As every p-value is equally likely, they cannot quantify the strength of evidence against the null hypothesis.

The black curve in Figure 1 illustrates the p-value distribution for a test with a statistical power of 50% for an alpha level of 5%. A p-value of 0.168 is more likely when there is a true effect that is examined in a statistical test with 50% power than when the null hypothesis is true (as illustrated by the black curve being above the grey line at $p = 0.168$), and thus evidence *for* an alternative hypothesis examined with 50% power, compared to the null hypothesis. If an effect is examined in a test with 99% power (the dotted line in Figure 1) a p-value of 0.168 is much more likely to be observed when the null-hypothesis is true than when a hypothesis is examined with 99% power. Thus, a p-value of

0.168 is evidence *against* an alternative hypothesis examined with 99% power, compared to the null hypothesis. The same p-value can thus correspond to different levels of evidence, which is an undesirable property for a measure. Figure 1 also illustrates that with 99% power even a ‘statistically significant’ p-value of 0.04 is evidence *for* the null-hypothesis. This fact, which is counterintuitive when first encountered, is known as Lindley’s paradox [4].

Figure 1: P-value distributions for a statistical power of 0% (grey line), 50% (black curve) and 99% (dotted black curve).

All p-values smaller than 1 correspond to evidence for some non-zero effect

Although it is not possible to relate a single p-value to a specific level of evidence, for any single analysis a lower p-value corresponds to more evidence against the null-hypothesis. If the alternative hypothesis is not specified, any p-value smaller than 1 should be treated as at least *some* evidence for *some* alternative hypothesis. It is therefore not correct to follow the recommendations of the authors in their Table 2 to interpret p-values above 0.1 (e.g., a p-value of 0.168) as “no evidence” for a relationship. This also contradicts their own arguments about the importance of evidence in meta-

analyses. Combining three studies with a p-value of 0.168 in a meta-analysis would lead one to reject the null hypothesis (see the forest plot in Figure 2). However, it is obviously undesirable to treat any $p < 1$ as unquantifiable evidence against the null hypothesis, as it makes it impossible to falsify the null-hypothesis. Yet, this is a price one must pay when attempting to use evidential language without specifying an alternative hypothesis [5].

Figure 2: Forest plot for a meta-analysis of three identical studies yielding $p = 0.168$.

Quantify evidence with measures of evidence

The recommendations by Muff and colleagues are an incoherent approach to statistical inferences, and should only be used if one wants to signal a misunderstanding of p-values. Coherent alternatives to quantify evidence exist, such as likelihoods [2] and Bayes factors [6]. Therefore, researchers should not follow the recommendation by Muff and colleagues to report $p = 0.08$ as ‘weak evidence’, $p = 0.03$ as ‘moderate evidence’, and $p = 0.168$ as ‘no evidence’.

References

- [1] S. Muff, E. B. Nilsen, R. B. O'Hara, and C. R. Nater, "Rewriting results sections in the language of evidence," *Trends Ecol. Evol.*, Nov. 2021, doi: 10.1016/j.tree.2021.10.009.
- [2] R. Royall, *Statistical Evidence: A Likelihood Paradigm*. London ; New York: Chapman and Hall/CRC, 1997.
- [3] S. N. Goodman and R. Royall, "Evidence and scientific research.," *Am. J. Public Health*, vol. 78, no. 12, pp. 1568–1574, 1988.
- [4] A. Spanos, "Who should be afraid of the Jeffreys-Lindley paradox?," *Philos. Sci.*, vol. 80, no. 1, pp. 73–93, 2013, doi: <https://doi.org/10.1086/668875>.
- [5] D. Lakens, A. M. Scheel, and P. M. Isager, "Equivalence testing for psychological research: A tutorial," *Adv. Methods Pract. Psychol. Sci.*, vol. 1, no. 2, pp. 259–269, 2018, doi: 10.1177/2515245918770963.
- [6] R. E. Kass and A. E. Raftery, "Bayes factors," *J. Am. Stat. Assoc.*, vol. 90, no. 430, pp. 773–795, 1995, doi: 10.1080/01621459.1995.10476572.

The code for Figures 1 and 2 is available from:

<https://gist.github.com/Lakens/553a0787873e8b7514aa55307e76de8c>