A note on the alpha-level: Not all alphas can be justified

Casper Albers

University of Groningen, The Netherlands

Author Note

Casper Albers <https://orcid.org/0000-0002-9213-6743>

I have no conflict of interest to disclose.

Correspondence concerning this article should be addressed to Casper Albers, Heymans Institute for Psychological Research, University of Groningen, 9712TS Groningen, The Netherlands.

E-mail: c.j.albers@rug.nl

Submitted to *Meta-Psychology*. Participate in open peer review by commenting through hypothes.is directly on this preprint.

Abstract

In the past years, discussions about the ‘alpha-level’ in null hypothesis significance testing gained attention in psychological literature. Much has already been said in this debate. One point has, however, not been raised in the recent discussions: that sometimes it is illogical to fix alpha to a specific level, whether this is 0.005, 0.05 or some other level. In this note, I raise this point and offer recommendations.

In frequentist null-hypothesis significance testing (NHST), one has to specify a nominal level of significance, also known as the alpha-level. One rejects the null hypothesis if and only if the p -value falls below this alpha-level. As default, for decades most social scientists simply worked with $\alpha = .05$ without giving this choice much attention. This changed when Benjamin et al. (2018) argued that science would be better off if we set $\alpha = .005$ (as this would clearly lower the rate of false positive findings). In a response to this claim, Lakens et al. (2018) argued against using any default level – whether it is .05, .005 or some other value – and that one should justify one's alpha on an individual study basis.

In this note, I will stress the point that in psychological research often not all alpha-values are appropriate choices. As a consequence, in many published studies where the authors claimed to have worked with $\alpha = 0.05$, they actually worked with an alpha just below 0.05, thus performed a test more conservative than indicated.

This happens when one works with discrete data, which obviously is common practice: Likert-scales, recording whether or not a participant passed a test or has a certain clinical condition or, more general, satisfies a certain condition or not, are all discrete scales. Furthermore, often even continuous concepts such as participants' age are often recorded discretely (by rounding off the age in years).

The point stressed in this note is not new, but neither is the exchange of arguments over the alpha-level. As that debate has returned to the spotlights, it's important to have all arguments in this discussion on the table again.

In the following I will first provide some more background about null hypothesis significance testing. The issue will first be demonstrated through an intentionally simplified hypothetical example followed by an realistic example from psychological literature, and practical recommendations.

The definition of the alpha level

There are two main schools of frequentist hypothesis testing: the Fisherian one and the one following the ideas of Neyman and Pearson. Although they have many similarities, there are also differences between both approaches. In social scientific practice, often a mix of both approaches is used.

In the approach by Ronald Fisher only a null hypothesis is specified (e.g., $H_0: \mu = 0$) but no alternative hypothesis. In the Neyman-Pearson (NP) theory (Neyman & Pearson, 1933), both a null hypothesis (e.g., $H_0: \mu = 0$) and an alternative hypothesis

(e.g., $H_1: \mu \neq 0$) are specified. Rejecting H_0 , and then concluding that the population mean is non-zero, typically is a NP approach. Many psychologists focus their whole research on attempting to reject the null hypothesis and finding significant effects, thus seemingly they are operating under the NP dogma.

Another distinction between both schools is that the NP approach is a behavioural, decision-theoretic approach. The NP approach is to set up two competing hypotheses, collect and analyse data, and force a decision: pick either H_0 or H_1 . The Fisherian approach is more evidential in nature: with the p -value some form of ‘evidence’ is quantified¹. Under the Fisherian approach, one doesn’t *have* to take any follow up actions. Concluding, based on $p = .03$, that the current result is somewhat unlikely, yet far from impossible, under null effects is justified. The same holds for concluding, based on $p = .78$, that there is no clear evidence against H_0 without being forced to accept H_0 as some universal truth. This aspect of the Fisherian approach lies close to the practice of many social scientists, indicating that scientific practice contains a mix of both approaches.

In NP theory, one has to specify the level α , which is defined as the probability of rejecting the null in favour of the alternative when H_0 actually is true: $\alpha = P(\text{rejecting } H_0 \mid H_0)$. Under NP theory, a likelihood ratio-based test statistic is developed, one which is in most situations the best possible test (for linear models it is uniform most powerful and unbiased; Lehmann, 1997, Ch. 3). The p -value is the probability that, under H_0 being true, the current value of the test statistic, or a more extreme value, is observed. If $p < \alpha$ then H_0 is rejected, else it is not rejected.

This alpha-level serves to control the error rate, in this case the rate of Type I errors, or false positives, which certainly has very useful properties (cf. Mayo, 1996) in long run error control (but less so in an individual study).

The alpha-level is useful for the NP approach of making dichotomous decisions: p is either below α or it isn’t. From a Fisherian point of view, one could do without the alpha-level and just interpret p -values (whether they are $p = .04$ or $p = .06$) as they are. In modern frequentist NHST it is standard to only claim a significant effect whenever $p < \alpha$, thus applying this element of NP-theory.

¹ Bayesians, and some other researchers, will argue that “the probability of finding the current result or something more extreme, under the null hypothesis” isn’t evidence; and that one should aim for “the probability that the null hypothesis is true”. This is a separate discussion, outside the scope of this paper. Here, I restrict this point by putting the term evidence in parentheses.

Thus, modern NHST often is a blend between the NP and Fisherian school of inference. In empirical research, elements from both schools are often combined. An literature review of 12 psychological journals in the years 1990-2002 (Hubbard, 2004) confirms this view. Various authors (Hubbard & Bayarri, 2003; Hubbard, 2004; Sprenger, 2016; Gigerenzer, 2004, 2018) discuss this mixed approach in further detail, and discuss the confusion and practical considerations this causes.

For more details on the comparison of these two schools, see Berger (2003). For a full mathematical background of the NP-approach, see Lehmann (1997; Chapter 3).

Illustrative example

Consider the following textbook situation, which is kept intentionally simple for illustrative purposes. A coin is flipped $n = 12$ times, and the number of tails is counted and denoted by S . The probability that the coin lands on tails is denoted by p . The null hypothesis states that the coin is fair, thus $p = \frac{1}{2}$, whereas the alternative in this example states that $p < \frac{1}{2}$. (For simplicity I employ a one-sided alternative but the example also ‘works’ for two-sided alternatives.) The probability flipping tails s times can be calculated using the binomial distribution and the p -value is calculated based on these probabilities. Table 1 lists values of these probabilities for small values of s .

Suppose one wishes to employ the ‘conventional default’ $\alpha = 0.05$. When observing $s = 2$ tails, one has $p = 0.019$ and would thus reject H_0 . When observing $s = 3$ tails, one has $p = 0.073$ and one would not reject H_0 . For many readers, this will sound as solid statistical practice to perform a test at the $\alpha = .05$ level.

There is, however, a catch: the alpha-level of this procedure is not .05. One has now made the decision rule to reject H_0 only if $S \leq 2$ and thus not to reject H_0 if $S > 2$. Therefore, the alpha-level is now set at 0.019, not at 0.05! Note, from Table 1, that the same problem occurs when one wants to perform the test at $\alpha = .005$ or $\alpha = .01$ or some other round number; the problem noted here lies not in the choice of .05 but in the procedure.

This problem is, obviously, due to the discrete nature of the data: one can observe 2 or 3 tails, but no values in between. The discrete nature of the data causes that not all values between 0 and 1 can be found as p -value.

This observation, obviously, is not new. Neyman and Pearson suggested the following randomized approach which, from a mathematical point of view, is optimal.

First, the value k for which $P(S \leq k) < \alpha$, but $P(S \leq k + 1) > \alpha$ needs to be specified. Rejecting H_0 only if $S \leq k$ leads to a too small alpha-level, whereas always rejecting H_0 also if $S = k + 1$ leads to a too high alpha-level. The Neyman-Pearson solution is to compute a probability γ and, when observing $S = k + 1$, to reject H_0 with probability γ (and, thus, not reject with probability $1 - \gamma$). The value γ can be computed according to the formula

$$\gamma = \frac{\alpha - P_{H_0}(S \leq k)}{P_{H_0}(S = k + 1)}.$$

Then, based on the observed outcome s , rejecting H_0 or not is decided based upon the following three situations:

1. If $s \leq k$, reject H_0 ;
2. If $s = k + 1$, reject H_0 with probability γ ;
3. If $s > k + 1$, do not reject H_0 .

This way, the alpha-level is set at exactly the desired level.

In our example, $k = 2$ and $\gamma = (.05 - .019)/.053 = .572$. Thus, if we observe 2 or fewer tails we certainly reject H_0 , and if we observe 3 tails, we reject with probability 57%. If we observe more than 3 tails, we do not reject. This way, the probability of rejecting H_0 is .05, so the alpha-level is set at the desired level.

Although this approach ‘works’ mathematically, most researchers would not feel comfortable with letting a random number generator decide whether or not their study yields a significant result. (This would also introduce a new way to p -hack your study; as it would be impossible to check, in absence of preregistration, whether the generated number is truly the result of a random process).

In practice, one will simply perform the rule as outlined before: reject if and only if the p -value falls below some level, in our example .05. There is nothing wrong with this approach, as long as one realises that the actual alpha-level lies below this value: one performs a test at $\alpha = .019$. In our example with a very small sample size, this is very apparent, as with small samples the probability of a specific observation can be considerable. However, it is not the case that this issue is merely a nuisance but of no practical consequence for more realistic sample sizes. In Figure 1, I outlined how the actual alpha level behaves for various sample sizes, in case of one-sided testing in the context of a binomial distribution. As can be seen, there is an interesting double sawtooth pattern, but it is not the case that the real alpha level will be practically close to .05 soon: out of the 51 values with n between 100 and 150, 21

(41%) have an actual alpha level below .04; which is more than one fifth off the desired level.

It is important to note that the issue discussed here has *nothing* to do with the specific value of .05. To illustrate this, Figure 2 displays the same situation as Figure 1, but now with α set to the value .005 as was advised by Benjamin et al. (2018). Now, 30 (59%) out of 51 values are at least one fifth too low.

Realistic example

The illustrative example was chosen intentionally simple for demonstration purposes. Comparing a certain proportion to a pre-prespecified proportion is perhaps not a very common practice in psychological research. Working with discrete data, however, is common. The use of Likert-scale based questionnaires is abundant in psychological research. In essence, as an individual's sum score clearly is discrete, one works with discrete data as well. This implies that not all values (in a certain range) are possible outcomes of the test statistic, which is explained in detail in Brown & Heathers (2017). To demonstrate this effect in 'real data', consider the following example.

In Greer & Morris (1975), the authors study psychological attributes of women who develop breast cancer. A group of 160 clinical patients was studied, $n_1 = 69$ of which were diagnosed with breast cancer and the remaining $n_2 = 91$ with a benign breast disease. Questionnaires were administered in both groups, including the Eysenck Personality Inventory (EPI; Eysenck & Eysenck, 1964), and one of the measures is the 'Neuroticism scale', with a range 0 – 24. In total, the breast cancer group scored 734 points on the neuroticism scale (mean 11.29, sd 5.48), and the control group scored 967 points (mean 10.74, sd 2.04). The research questions focusses on whether the mean test scores differs between both groups. Note that even though the study is almost half a century old, the approach with Likert scales and samples of this size are still very common.

Student's t -test provided $p = .511$. (Note: Even though the t -test is developed for continuous data, it is commonly applied to discrete situations such as these, and the t -test is proven to be sufficiently robust for such applications, cf. Norman, 2010).

The smallest possible change is one point more or less for one of the two groups. Table 2 displays the four p -values one obtains when the sum scores of one of

the groups is decreased or increased by 1². The smallest possible change in the data thus leads to a change of about one percentage point in the p -value. Although this sounds like a small difference, it can make the difference between a significant 0.045 and a non-significant 0.055.

The effect is much smaller than in our previous example. This is, given the sample size of 160 and the 24-point range of the scale, not surprising: the more measurements you have, the smaller the influence of rounding off to integers. One of the scales of the EPI is the ‘Lie scale’ which consists of only 9 items. Scores between both groups are very similar ($p = .994$), but when the sum score of the cancer group decreases by only one, the p -value changes by .039 to $p = .955$. A change of .039 can not be considered small anymore, when whole procedure is focused on finding values below .05. Simply having $n > 150$ thus does not safeguard against this issue.

Discussion and recommendations

The issue discussed in this note occurs whenever the distribution of the test statistic is discrete, i.e. can only attain a finite number of values. In these situations, it is often undesired to set alpha exactly to some round number like .05 or .005. The only way to use these alpha-levels is by having the outcome of the experiment – significant or not – decided, in part, purely on chance. In practice, very few people are willing to do that. As a consequence, their actual alpha-level is actually somewhat lower than the .05 level commonly taken. Thus, the rate of false positives (Type I errors) will be smaller than one would assume when ignoring this discretisation issue.

In simple situations, as in the illustrative example this randomisation mechanism can be avoided by justifying another choice for α a priori. In our example, one could easily justify that $\alpha = .019$ is a better choice than $\alpha = .05$, as this choice removes the need for randomisation (see Table 1) and does not lead to a higher false negative rate.

Once n has been decided, it can easily be computed which alpha-values do not require the randomisation of the second step in the NP-process, for instance by simply computing all theoretically possible p -values. From these values, one can be selected. This can for instance be the largest value below .05 or .005 (see the solid lines in

² Note: When the score of one of the group’s individuals changes by ± 1 , the variance of that group will also change slightly, which has an effect on the t - and p -values. As I don’t have the data themselves but only the summary statistics, it’s impossible to assess how the variance will change. For demonstration purposes, I worked with the same group variances for all computations.

Figure 1 and 2), but also the smallest values above .05 or .005 (the dotted lines), or any other value one can justify.

This approach also works in slightly more complicated situations. In Box 1, sample *R* code is provided to do this for the situation when comparing two proportions (e.g., the proportion of students passing after following a conventional course with the proportion passing after following a ‘flipped classroom’ course). For $n_1 = n_2 = 50$, the closest alpha-value below .05 is .0495, and the closest above .05 is .0537. Thus, even with a moderate sample size of 100, the difference with .05 is small but non-negligible.

This approach, however, is not practical in all situations. Take, for instance, a two-group *t*-test on questionnaire sum scores, such as in the example from Greer & Morris (1975). The statistic for the *t*-test is based on five ‘ingredients’: sample sizes n_1, n_2 (known a priori), group means $\bar{y}_1, \bar{y}_2$ and pooled variance s_p^2 . Theoretically, it still is possible to work out *all* possible outcomes (cf. Heathers, Anaya, Van der Zee and Brown, 2018) and select the alpha-level closest to the desired level, but this clearly is too cumbersome for practical purposes. The best one can do here is to set α to some desired level (e.g., .05) and acknowledge that the actual alpha-level can be slightly different.

The simplest procedure is to select a α , e.g. 0.05, and state that you will work with the highest possible α that is no higher than 0.05. From a meta-analytic and open science point of view, colleagues have stressed the importance of explicitly stating the *p*-value, rather than a dichotomous statement on whether *p* lies above or below the significance-threshold. This message coincides with Fisherian rather than NP thinking. This paper provides yet another argument for this: after all, if the significance-threshold isn’t some fixed value (e.g., exactly .05) but an interval (e.g., (.0495, .0537) as in the example above) the context is different. This especially holds for small(ish) samples but, as shown in this paper, even for samples with $n > 150$, the effect still can be too large to ignore.

Author Contributions

CJA is the sole author of this manuscript. All contributions are his.

Data Availability

No empirical data is presented in this manuscript.

References

- Benjamin, D. J., Berger, J. O., Johannesson, M., ..., Johnson, V.E. (2018). Redefine statistical significance. *Nature Human Behaviour*, 2, 6-10, doi: 10.1038/s41562-017-0189-z
- Berger, J. O. (2003). Could Fisher, Jeffreys and Neyman have agreed on testing? *Statistical Science*, 18, 1-32, doi: 10.1214/ss/1056397485
- Brown, N. J. L., Heathers, J. A. J. (2017). The GRIM test: A simple technique detects numerous anomalies in the reporting of results in psychology. *Social Psychology and Personality Science*, 8, 363-369, doi: 10.1177/1948550616673876
- Eysenck, H. J., Eysenck, S. B. G. (1964). *Manual of the Eysenck Personality Index*. London: University of London Press
- Gigerenzer, G. (2004). Mindless statistics. *The Journal of Socio-Economics*, 33, 587-606, doi: 10.1016/j.socec.2004.09.033
- Gigerenzer, G. (2018). Statistical Rituals: The replication delusion and how we got there. *Advances in methods of practices in psychological science*, 1, 198-218, doi: 10.1177/2515245918771329
- Greer, S., Morris, T. (1975). Psychological attributes of women who develop breast cancer: A controlled study. *Journal of Psychosomatic Research* 19, 147-153, doi: 10.1016/0022-3999(75)90062-8
- Heathers J. A., Anaya J., Van der Zee T., Brown N. J. L. (2018). Recovering data from summary statistics: Sample Parameter Reconstruction via Iterative TEchniques (SPRITE). *PeerJ Preprints* 6:e26968v1. doi: 10.7287/peerj.preprints.26968v1
- Hubbard, R. (2004). Alphabet soup: blurring the distinction between p 's and α 's in psychological research. *Theory & Psychology*, 14, 295-327, doi: 10.1177/0959354304043638

- Hubbard, R., Bayarri, M. J. (2003). Confusion over measures of evidence (p 's) versus errors (α 's) in classical statistical testing. *The American Statistician*, 57, 171-177, doi: 10.1198/0003130031856
- Lakens, D., Adolphi, F.G., Albers, C.J., ..., Zwaan, R.A. (2018). Justify your alpha. *Nature Human Behaviour*, 2, 168-171, doi: 0.1038/s41562-018-0311-x
- Lehmann, E.L. (1997). *Testing statistical hypotheses* (2nd edition), Ney York: Springer
- Mayo, D. G. (1996). *Error and the growth of experimental knowledge*. Chicago: University of Chicago Press.
- Neyman, J., Pearson, E.S. (1933). On the problem of the most efficient tests of statistical hypotheses. *Philosophical Transactions of the Royal Society London, Series A*, 231, 289-337, doi: 10.1098/rsta.1933.0009
- Norman, G. (2010). Likert scales, levels of measurement and the “laws” of statistics. *Advances in Health Science Education*, 15, 625-632, doi: 10.1007/s10459-010-9222-y
- Sprenger, J. (2016). Bayesianism vs. frequentism in statistical inference. in Hájek, A., Hitchcock, C. (eds.): *Handbook of the Philosophy of Probability*, 382-405. Oxford: Oxford University Press

Table 1. Probabilities and cumulative probabilities for the running Binomial(12, $\frac{1}{2}$) example

s	0	1	2	3	4
$P_{H0}(S = s)$	< .001	.003	.016	.054	.121
$P_{H0}(S \leq s)$	.001	.003	.019	.073	.194

Table 2. Computation of p -values for the Greer & Morris (1975)-example with the sum scores per group as reported (first row), and slightly altered (other groups). The *R* code for the computations is provided in the appendix.

sum1	sum2	p-value
734	967	.511
733	967	.523
735	967	.499
734	966	.502
734	968	.519

Box 1. *R* code for finding the possible alpha-level closest to a desired alpha-level for a standard two-proportion comparison. For $n_1 = n_2 = 50$ and a desired level of .05, the closest possible alpha-level that doesn't require randomisation is .0495.

```
# Sample sizes and the desired alpha-level are known a priori
n1 <- 50
n2 <- 50
desiredalpha <- .05

# start from an unrealistic value, and improve it each time it's possible
closestalpha <- 1
for(s1 in 0:n1){
  for(s2 in 0:n2){
    p <- prop.test(x = c(s1, s2), n = c(n1, n2),
                  alternative = "two.sided")$p.value
    if(!is.na(p)){
      if(abs(p - desiredalpha) < abs(closestalpha - desiredalpha)){
        closestalpha <- p
      }
    }
  }
}
```

Figure 1. Actual alpha level when one performs the non-randomized approach to reject H_0 whenever $p < .05$ (solid line), or whenever k lies one observation higher (dotted line).

Figure 2. As Figure 1, but now with $\alpha = .005$ instead of $\alpha = .05$.

Appendix R code for the computations of the Greer & Morris-example

```
# Computations for the Greer & Morris (1975) example

# Neuroticism scale
# These scores are provided by, or deducible from, the paper itself (Table
2)
sum1 <- 734
n1 <- 65
sum2 <- 967
n2 <- 90

# The variances can be computed from the SD's (up to rounding off errors)
var1 <- 5.48^2
var2 <- 4.82^2

# Standard (equal. variances) t-test. Result coincides with that in Table 2.
t <- (sum1/n1 - sum2/n2) / (sqrt(((n1-1)*var1+(n2-1)*var2)/(n1+n2-2))) *
sqrt(1/n1+1/n2))
2*pt(-abs(t),n1+n2-2)

# t-test if cancer group scored one point lower
# (Assuming variances unchanged; see footnote)
t <- ((sum1-1)/n1 - sum2/n2) / (sqrt(((n1-1)*var1+(n2-1)*var2)/(n1+n2-2))) *
sqrt(1/n1+1/n2))
2*pt(-abs(t),n1+n2-2)

# t-test if cancer group scored one point higher
# (Assuming variances unchanged; see footnote)
t <- ((sum1+1)/n1 - sum2/n2) / (sqrt(((n1-1)*var1+(n2-1)*var2)/(n1+n2-2))) *
sqrt(1/n1+1/n2))
2*pt(-abs(t),n1+n2-2)

# t-test if benign group scored one point lower
# (Assuming variances unchanged; see footnote)
t <- (sum1/n1 - (sum2-1)/n2) / (sqrt(((n1-1)*var1+(n2-1)*var2)/(n1+n2-2))) *
sqrt(1/n1+1/n2))
2*pt(-abs(t),n1+n2-2)

# t-test if benign group scored one point higher
# (Assuming variances unchanged; see footnote)
t <- (sum1/n1 - (sum2+1)/n2) / (sqrt(((n1-1)*var1+(n2-1)*var2)/(n1+n2-2))) *
sqrt(1/n1+1/n2))
2*pt(-abs(t),n1+n2-2)

# "Lie scale"
# Same approach.
sum1 <- 288
n1 <- 65
sum2 <- 399
n2 <- 90
var1 <- 2.04^2
var2 <- 1.85^2

# Standard t-test on actual data
t <- (sum1/n1 - sum2/n2) / (sqrt(((n1-1)*var1+(n2-1)*var2)/(n1+n2-2))) *
sqrt(1/n1+1/n2))
2*pt(-abs(t),n1+n2-2)
# t-test if cancer group scored 1 point lower
t <- ((sum1-1)/n1 - sum2/n2) / (sqrt(((n1-1)*var1+(n2-1)*var2)/(n1+n2-2))) *
sqrt(1/n1+1/n2))
2*pt(-abs(t),n1+n2-2)
```