Running head: PUTTING THE P

In press: *The American Statistician*

Putting the P Value in its Place

Joachim I. Krueger

Brown University

Patrick R. Heck

Geisinger Health System

Authors' Note

Joachim I. Krueger is a professor of psychology at Brown University. Patrick R. Heck is a postdoctoral research fellow at Geisinger Health System. Some of the ideas described in this article were presented at a symposium on p values convened by the *American Statistical Association* in Bethesda, MD, October 2017.

Correspondence:

Joachim I. Krueger or Patrick R. Heck
Department of Cognitive, Linguistic & Psychological Sciences
Brown University
190 Thayer St.
Providence, RI 02912
Phone: (401) 863-2503
E-mail: Joachim@brown.edu, Pheck1000@gmail.com
Home page: <https://vivo.brown.edu/display/jkrueger>,
<https://www.patrickrheck.com>

Keywords: statistical significance testing, null hypotheses, p values, Bayes' Theorem, inference

Abstract

As the debate over best statistical practices continues in academic journals, conferences, and the blogosphere, working researchers (e.g., psychologists) need to figure out how much time and effort to invest in attending to experts' arguments, how to design their next project, and how to craft a sustainable long-term strategy for data analysis and inference. The present special issue of the *American Statistician* promises help. In this article, we offer a modest proposal for a continued and informed use of the conventional p value without the pitfalls of statistical rituals. Other statistical indices should complement reporting, and extra-statistical (e.g., theoretical) judgments ought to be made with care and clarity.

Dr. Livingstone, I presume!

~ Stanley to Livingstone, after rejecting the hypothesis that Livingstone was not Livingstone

Walk barefoot on the beach; it gets the toxins out! ~ A well-meaning neighbor

The most striking pair of facts about significance testing is its immense popularity among researchers, and the intensity of the critical opposition. Since modern sampling techniques and frequentist statistics became *de rigueur* in many fields of scientific study over half a century ago, critics have questioned the wisdom of significance testing in general and the reliance on p values in particular (Gigerenzer, 2004; Stigler, 1999). Significance tests and p values can be found in many sciences, and thus the context and the nature of the debate vary. In this article, we focus on research typical in experimental psychological science. We discuss statistical inference in the canonical case of the treatment-plus-control experiment designed to test a single, novel, and causal hypothesis with a limited (and often small) sample.

In the prototypical experimental scenario, a researcher who, after having done formal and informal (i.e., intuitive) theoretical work, comes to think that a particular stimulus or treatment might have an effect on participants' experience and behavior. To give a simple example, assume that the researcher thinks that taking a break from creative production increases the eventual total output. In the control condition, respondents generate as many diverse answers to a prompt such as "If all schools were abolished, what would you do to become educated?" as they can without taking a break. In the experimental condition, respondents take a short break during which they engage in a cognitively consuming task before resuming their creative efforts (Ostrowski, 2017; see also Gilhooly, Georgiou, & Devery, 2012). The statistical question is whether there is a significant difference in creative production between these two conditions.

Many experiments in psychological science take this general form. The properties of the data (e.g., their potential skew, the scale of measurement) help the researcher select an

appropriate significance test (e.g., by choosing between parametric and nonparametric tests), compute a test statistic, and ascertain its p value, that is, the probability of the observed test statistic – or any value more extreme than the observed one – assuming that the null hypothesis of no difference is true. Following convention, the researcher rejects the null hypothesis as being a poor fit with the data if $p < .05$. The cautious researcher infers that there is some (perhaps limited or preliminary) support for the ordinal and imprecise hypothesis that the treatment has an effect; in our example, taking a break from generating diverse responses increases the eventual total number of responses. In our experience, many researchers choose their words carefully when communicating what they have learned from a significant result. Few insist that they have proven their hypothesis, and many note that the data provide some evidence against the null hypothesis.

Yet, much of the criticism leveled at significance testing, and specifically its most common variant of *Null Hypothesis Significance Testing*, NHST, focuses on the limits of many researchers' knowledge of its logic, its proper use, and the meaning of its concepts (Goodman, 2008; Greenland, Senn, Rothman, Carlin, Poole, Goodman, & Altman, 2016). This sort of criticism does not challenge the validity of NHST directly, but it invites the audience to consider alternatives. We prefer a clear distinction between the properties of a method and its (flawed) reception by its practitioners. Yet, while focusing on the former, we hope to educate practicing researchers on the latter.

In the first section of the remainder of this short article, we discuss what we consider the primary contested issue regarding the p values produced by significance testing, which is their ability to support reverse inferences about the truth status of the tested hypothesis given the observed data. In the second section, we ask whether significance testing should be replaced *tout*

court with Bayesian methods. In the third section, we return to significance testing and discuss its implications for different types of decision error and how these errors are viewed. In the fourth and final section, we explore the psychological and sociological context of the current statistical debate and its implications for the researcher in the lab.

1.1 The predictive power of p

Chief among the concerns about researchers' ignorance is that they mistake the p value for the probability of the hypothesis given the data, $p(H|D)$. An informal – but serviceable – interpretation of the p value is that it is the probability of the data assuming that the (null) hypothesis is true, $p(D|H)$ (Wasserstein & Lazar, 2016). An answer to the question of whether the p value is subtly different from $p(D|H)$ lies beyond this article (Greenland et al., 2016). Many statisticians accept the idea of a close correspondence, while others insist that the differences are deep. Yet, most statisticians agree that to assume without checking that $p(D|H) = p(H|D)$ is to commit a grave reasoning error, namely the fallacy of reverse inference (Krueger, 2017). To claim that the null hypothesis is false with a probability of .95 because an experiment yielded $p = .05$ is to make an irresponsible statement.

The probability of the hypothesis given the data, $p(H|D)$, is given by Bayes' theorem, which states that $p(H|D) = \frac{p(H) \times p(D|H)}{p(H) \times p(D|H) + p(\sim H) \times p(D|\sim H)}$. The theorem shows that the p value (or $p(D|H)$) predicts $p(H|D)$ but that the prior probability of the hypothesis, $p(H)$, also matters, as does the probability of the data under *any* hypothesis, $p(D)$, which is shown in decomposed form in the denominator. The equality of $p(H|D) = p(D|H)$ may occur, but only for the special case of $p(H) = p(D)$ (Dawes, 1988). Otherwise, the relationship between $p(H|D)$ and $p(D|H)$ is imperfect. Using simulation experiments, we mapped the range of possibilities for this relationship (see Krueger & Heck, 2017, for details). We can summarize the two main findings

thus: First, there is a wide range of positive correlations between the two conditional probabilities; under no circumstance did we find a negative or zero correlation. Second, these correlations are higher inasmuch as the simulated conditions reflect the typical context of empirical work. Specifically, it is typically not the case that $p(H)$ and $p(D|H)$ are conditionally independent over experiments. If an experiment is very risky, that is, if the probability of the null hypothesis, $p(H)$, is very high, then it is also very unlikely to find a low p value. The second-order probability of finding a low value for $p(D|H)$ is low when $p(H)$ is high. As a result, positive correlations between $p(D|H)$ and $p(H|D)$ can be high, even in excess of .7. In other words, under realistic conditions of the typical experimental research environment, p values serve as useful heuristic cues for inferences about the posterior probability of the tested hypothesis, $p(H|D)$.

Our analysis supports the following conclusion: Whereas it is good to warn against equating the p value with the posterior of the tested hypothesis, it is unwise to suggest that no inference can be made. The absence of a one-to-one correspondence is not the absence of *any* correspondence. Indeed, the application of Bayes' theorem to the logic of NHST reveals that statistical inference is an inference under uncertainty and not a strict computation or logical implication. Hence, we regard the p value as a useful heuristic cue for estimating the tested hypothesis's posterior probability. In Krueger and Heck (2017), we describe estimation techniques that help researchers identify expected values of $p(H|D)$ and their associated ranges. We encourage researchers to use p values for first-pass, heuristic inferences.

1.2 The Bayesian alternative

The foregoing argument is pragmatic rather than purist. It does not make a clean separation between Bayesian and frequentist approaches. Instead, we treat the observed p value as input for the estimation of the posterior probability of the tested hypothesis. Absent additional

or contravening assumptions, a researcher might favor (if not ‘accept’) a hypothesis if it is more likely than turning up heads in a coin flip ($p(H|D) > .5$). Similar pragmatic integrations of Bayesian and frequentist operations have been proposed elsewhere (Cohen, 1994; Krueger, 2001; Nickerson, 2000; Trafimow, 2003). Much of the current debate, however, is characterized by committed advocacy, which asks researchers to endorse to a particular set of principles. The desire for a full conversion of researchers to a particular school of thought is understandable, given that the advocates of these schools take pains to work strictly within their own set of assumptions, and given that mixing approaches can beget confusion and contraction (Gigerenzer, 2004; Gigerenzer & Marewski, 2015). Yet, in this particular case – estimating $p(H|D)$ from $p(D|H)$ using Bayes’ theorem – we see little epistemological danger. Instead, we see an opportunity for researchers to better understand the properties of the p value and its implications for the hypothesis of interest. In the spirit of this special issue on statistical inference, we advise researchers to familiarize themselves with the basic assumptions underlying contemporary schools of statistical thought and to treat available methods as tools in a box, to be used judiciously.

Consider the central assumption about what is perceived to vary and what is perceived to be a fixed condition. Frequentists treat hypotheses as fixed parameters; it may only be the null hypothesis, but it could also be a non-null or substantive hypothesis, or it could be several hypotheses. Hence, analysis produces probabilities of data – which may vary due to sampling – and which are conditioned on the parameters (i.e., hypotheses). Bayesians, in contrast, treat the data as the conditions given by observation, while treating hypotheses as random variables, or distributed. Hence, Bayesians are interested in posterior probability distributions. If they select two discrete hypotheses, H and $\sim H$, from these distributions, they can compute odds ratios

$p(H|D) / p(\sim H|D)$ or ratios of posterior over the prior odds (Ortega & Navarrete, 2017). If they consider an entire distribution of hypotheses (i.e., where the hypothesis is a random variable), an infinite number of odds ratios awaits contemplation.

We suspect that many researchers find the Bayesian theory (theory = a way of seeing) of hypotheses as forming a distribution of possibilities counterintuitive – and not only because they have been trained in frequentist statistics. It seems more natural to ask if a *particular* idea or hypothesis is true or false than it is to consider a set of observations (the data) and represent knowledge and belief as a distribution over an infinite set of hypothesis. The latter does not allow a person to express the strength of belief in any particular idea but only to express how much stronger or weaker the belief is relative to some other particular idea. In contrast, the former (hypothesis-conditioned) way of thinking is aligned with any epistemological framework that begins with core theoretical assumptions, generates testable hypotheses, and ends with experimental data and inferences.

We are inclined – as we foreshadowed in section 1 – to recommend significance testing in small-sample experimentation and amend it with Bayesian inferences to estimate $p(H|D)$. The p value, or $p(D|H)$, is a useful heuristic to infer $p(H|D)$. The ratio of $p(D|H) / p(D|\sim H)$ may be a stronger heuristic, but becomes relevant only if researchers articulate a specific alternative to the null hypothesis. The ratio $p(D|H) / p(D|\sim H)$ is a likelihood ratio, LR, when it is computed with the values of the density function. The use of probabilities (areas under the density curve) and use of likelihoods (the value on the y axis of the density function) makes little difference. Our simulations corroborate what a look at a unimodal distribution reveals, namely that both probabilities and likelihoods become smaller as one moves into the tail of a unimodal distribution. For the normal distribution, for example, the correlation between the log-

transformed likelihoods and log-transformed probabilities is nearly perfect (Krueger & Heck, 2017).

Now consider the relationship between the p value (or $p(D|H)$) and the LR. Suppose the researcher has specified a discrete alternative to the null hypothesis such that $H \Rightarrow \delta = 0$ (i.e., there is no effect between control and treatment on some outcome measure) and $\sim H \Rightarrow \delta = .5$ (there is a standardized effect of .5 between control and treatment). When more than one experiment has been performed and some variation in the empirical effect size, d , has been observed, the hydraulic relationship between $p(D|H)$ and $p(D|\sim H)$ is evident. The p value now perfectly predicts the LR. If, however, as in our example from experimental psychology, the researcher cannot assign a specific effect size to the alternative hypothesis, and if the alternative (i.e., $\delta \neq 0$) is variable, then $p(D|\sim H)$ varies, while the observed data and $p(D|H)$ remain constant. In this scenario, the correlation between p and the LR is not defined. Finally, if there are several experiments, perhaps gleaned from meta-analysis or replication attempts, both $p(D|H)$ and $p(D|\sim H)$ are variable, and so is the LR. Here, a positive association between p and LR remains (Krueger & Heck, in revision). In other words, the p value continues to play a useful heuristic role in predicting the LR and $p(H|D)$.

Our analysis suggests the following conclusion: Researchers may wish to supplement p values with LRs if they can present a rationale for particular alternative hypotheses. In such cases, which remain rare in experimental psychology, they can estimate the posterior probabilities of both hypotheses with greater precision. Either way, it is unwise to discard the p value entirely as it is generally unwise to ignore the building blocks of composite scores such as ratios or discrepancies (Krueger, Heck, & Asendorpf, 2017).

1.3 *Errare humanum est, sed quem errorem?*

One version of the concern that p values seduce the unwary to infer the posterior probability of the tested hypothesis manifests in warnings about inflated Type I errors (Colquhoun, 2015; Ioannides, 2005; but see Stroebe, 2016). Type I errors are false positives, FP, in decision-theoretic terms (Swets, Dawes, & Monahan, 2000). In the context of NHST, FPs refer to mistaken beliefs in something that does not exist. FP aversion directly follows from the worry about the reverse inference fallacy discussed in the first section. If researchers rush to the judgment that $p(H|D) < .5$ if $p(D|H) < .05$, they may log many results as discoveries even though $p(H|D)$ remains greater than .5. This is most likely to happen when the work is risky, that is, when $p(H)$ is very high. FP aversion has spawned various recommendations for error control, most of which amount to a call for larger samples. Yet, the primary result of increased sample size is a reduction of Type II errors (i.e., failures to reject false null hypotheses). Although larger samples can raise FP rates, they tend to lower FP *ratios* (when Hits [rejections of false null hypotheses] rise faster than FPs; Krueger & Heck, 2017).

An alternative recommendation is to lower the significance threshold. Lowering significance thresholds to $p < .005$, for example (Benjamin et al., 2017; or $p < .001$, Johnson, Payne, Wang, Asher, & Mandal, in press) will reduce the number of FP errors, but at the cost of increasing the number of Type II errors to an unknown degree (Fiedler, Kutzner, & Krueger, 2012). As significance testing becomes more conservative, FP ratios (the proportion of significant findings leading to false inferences) may even increase (Krueger & Heck, 2017). We recommend, in the spirit of Neyman and Pearson (1933ab) that researchers carefully think about the utility they assign to the two types of error before ritualistically endorsing a conventional scheme such as $\alpha = .05$ and $\beta = .2$ (Cohen, 1977; 1992; Erdfelder, Faul, & Buchner, 1996).

To put these concerns in perspective, consider Meehl's (1978) "strong use" of significance testing as a thought experiment (see also Antonakis, 2017). Meehl, being a Popperian falsificationist at the time, asked us to imagine that a specific non-null hypothesis (which we have referred to as $\sim H$) is subjected to the ordeal of the significance test. The meanings of the two errors are now reversed. A Type I error (FP) has become the rejection of a true substantive hypothesis, inviting a belief in a false null. In contrast, a Type II error (Miss) is now the failure to reject a false substantive hypothesis, inviting a false belief that the null is false. Consider the implications of these two suggested tactics of error control, increasing power and reducing significance thresholds.

The main consequence of increasing sample size – and thereby power – is the reduction of Type II errors. In the Meehl-Popper strong-inference scenario, a true null hypothesis is more likely to be detected. This prospect might subtly motivate the researcher to limit data collection. From a Popperian perspective, it makes good sense to demand high power. Alternatively, a lowering of the significance threshold to, say, .005, will make it harder for researchers to demonstrate invariances (i.e., the absence of an effect). When theory predicts a substantive effect, the conservative researcher may wish to *relax* the criterion of significance, whereas the liberal researcher would tighten it. Again, our recommendation is for researchers to consult relevant theory before deciding whether they can (or should) put a non-null hypothesis to the test and how strict this test should be (see also Lakens et al., 2018). If theory or compelling convention does not require unbalanced *a priori* error rates, we think it prudent to use the same probabilities for both types of error (e.g., $\alpha = \beta = .1$; see also Bredenkamp, 1980; as well as Neyman and Pearson themselves, 1933ab).

The discourse of error presupposes *decisions* regarding significance, and decisions need criteria. The conventional criterion for declaring significance is $p < .05$, and much has been made of the presumed rigidity of its application (Greenland, 2017; Wasserstein & Lazar, 2016). A radical response is to surrender all pretension of decision-making and to limit inference to the task of judgment by, for example, estimating a probable range of values for $p(H|D)$ given the obtained p value and auxiliary assumptions (Krueger & Heck, 2017).

As experimentalists, we are reluctant to relinquish dichotomous decision-making entirely. The first reason for this traditionalism is that for many questions humans ask of nature, the null (or any particular tested hypothesis) may in fact be true.¹ Nature presents some true-false dichotomies, and it would be unwise to use a decision-making scheme that ignores this. A profusion of pseudo-, para-, and alternative sciences thrive on false claims. They assert the presence of not-nothing when there is nothing (Krueger, Vogrincic-Haselbacher, & Evans, in preparation). The second epigraph of this article may serve as an illustration.

It is easy to see the truth of this claim (this claim itself has a good probability of being true) when considering categorical data. We rightly treat categorical questions to true-or-false responses. In contrast, much of the statistical debate (including this short article) plays out in the world of continuous variables. Here, the claim that the tested hypothesis cannot be literally true is an artifact of allowing each specific point on the scale to be no more than one among an infinity of points. Such a point has no probability, only a likelihood – and as a number, this likelihood is meaningful only when compared with (e.g., divided by) another number.

¹ The argument that any point hypothesis is false (e.g., Cohen, 1990; Gelman & Carlin, 2017) is a mathematical artifact of the assumption that there is an infinite number of hypotheses (Krueger, 2001). It would be more accurate to say that the probability of such an hypothesis to be true is indeterminate; that is, such an hypothesis would be neither true nor false irrespective of the evidence.

Returning to our example of the experiment on creativity, the researcher is tasked with assigning respondents randomly to conditions. Comparing the experimental with the control condition, we may wonder if any difference, however small, might eventually yield significance without being an FP. If, however, we compare two control conditions, both created at random and each without a treatment, we would know that $p(H) = 1$ and that any $p < .05$ could only be an FP. We now infer that if the null hypothesis can be true if we ensure that it is, it can also be true when we did not.

The second reason for allowing a judicious use of dichotomous inference is psychological. Categorization is built into perception (Bruner, Goodnow, & Austin, 1956). To perceive is to categorize. We cannot see Argos without also seeing a dog, and we cannot see Eumaios without seeing a man, a Greek, and a swineherd. Categorization affords induction. Recognized as a dog, Argos tells us much about his conspecifics (e.g., love of his master); if we ignored the inductive power of categorical perception, we'd be condemned to encounter each dog as a novel creature. Ditto for Eumaios. Following Bruner and other pioneers, Tajfel (1959; 1969) proposed accentuation theory to formalize the interplay of categorization and perception of stimuli falling on a graded scale. Accentuation theory predicts that if there is a categorical line drawn somewhere on the continuum, stimuli falling to the left (lower) of this line will be seen as having been shifted to the left, while stimuli falling to the right (higher) will be seen as having been shifted to the right. The theory also predicts that this effect is strongest near the line itself, which results in a perceptual narrowing of each category (Krueger & Clement, 1994).

Consider the implications of accentuation theory for the perception of p values. Values $< .05$ will be regarded as small, while values $> .05$ will be regarded as large, with little discrimination among values falling on the same side of the divide. Some may take this

perceptual sharpening as evidence of the dangers of dichotomous statistical inference (Gelman & Stern, 2006); others will view it as an unavoidable byproduct of categorization. We side with the latter group, thinking that binary statistical inference ultimately stands in the service of action. Neyman and Pearson (1933ab) explicitly subordinated statistical decision-making to action, and Fisher (1956) did so implicitly – at best. When a fertilizer has been tested, the *agronome* must decide whether to use it. We ask researchers to be mindful of accentuation theory and to consider if and how their work is connected to decisions about action. In this exposition of our findings, we can only mention, but not explore, the differences between inference and decision, and the complexities arising along the road from statistical inference to scientific inference to practical inference. Suffice it to note our agreement with Fisher (1959, p. 100) who wrote that “an important difference is that decisions are final, while the state of opinion derived from a test of significance is provisional, and capable, not only of confirmation, but of revision.”

1.4 Statistics as a social process

The future of statistical practice will not be decided by logical proof or empirical test. Try to imagine what a critical experiment might look like! Would the data be analyzed with significance tests, model-fitting techniques, or subjective Bayesian interpretation? How would you go about managing the two types of error in this experiment? This is a difficult question: if our methods are designed to test the predictions derived from our theories, how might the methods themselves and their parent theories be tested? We submit that the acceptance of a set of methods is largely a matter of social process, rhetoric, and *Zeitgeist*. There surely are innovations, reforms, and refinements of methods that can be strictly justified on logico-mathematical grounds, but such advances are usually made *within* the context of a paradigm or school of thought. A recent example is the introduction of the *p-rep* index, which was meant to replace the

conventional p value (Killeen, 2005). Only $p\text{-rep}$, but not p , was to reveal the probability of successful replication. It soon turned out that $p\text{-rep}$ was a mere log-linear transformation of p , and so the index went as fast as it had come (Wagenmakers, 2007).² Other innovations have found traction, however, such as the now widespread use of multi-level regression models (Austin, 2017).

When schools of thought whose most central thoughts are incommensurable (Feyerabend, 1976; Fleck, 1935; Kuhn, 1962) vie for the affections of bench scientists, rhetoric and social process become relevant. The result is a cross-purpose debate. In the current climate, Bayesians emphasize the coherence of their methods compared with the ease with which significance tests produce incoherent patterns (e.g., inferences violating the axiom of transitivity). Significance testers, in contrast, emphasize the vulnerability of Bayesian methods to belief bias (Revlin, Leirer, Yopp, & Yopp, 1980). False priors make bad estimates, unless the data are really big – in which case neither Bayesian nor frequentist methods are needed. Frequentists may turn some of the criticisms leveled at significance testing back on the Bayesians. The dangers of HARKing (Hypothesizing After the Results are Known; Kerr, 1998), for example, apply to both schools, as do the dangers of limiting sample size (Krueger & Heck, 2017). Other attempts to reform research practice do not affect the analytical methods themselves. Preregistration regimens are intended to protect researchers from one another and themselves (e.g., their implicit biases and p -hacking self-deceptions), but are mute on the logic of inductive inference.

When there are no incontestable criteria to settle methodological debates, social processes unfold that might lead to the eventual dominance of one school of thought. These processes are

² New indices with little or no new informational input occasionally appear in the literature. Bayarri, Benjamin, Berger, and Sellke (2016), for example, suggested that p values be replaced by upper-bound Bayes Factors as

$$\frac{1}{-ep \log p}.$$

Here, all the variation is in the p value.

‘soft’ in that they cannot leverage irrefutable proof or unquestionable authority. The limits of authority may be found at the level of the journal editor who by way of *Diktat* declares certain methods out of bounds. Reporting of *p-rep* was declared obligatory by the journal *Psychological Science*; its demise was less formal. The editors of *Basic and Applied Psychology* banned *p* values from their pages (Trafimow & Marks, 2015). Professional societies convene task forces to work out recommendations for statistical reporting, but, as part of the social process, the staffing of such task forces must not be partisan (Wasserstein & Lazar, 2016; Wilkinson & APA Task Force, 1999). As a result, the recommendations tend to be encouragements rather than pre- or proscriptions.

More forceful demands surface when groups of statisticians and researchers band together to issue calls for changes in practice. Recently, Benjamin et al. (2017) proposed a lowering of the conventional significance threshold to .005, conceding that this was a patch, and that elegant solutions were being prepared, but that there was no consensus yet. The ‘et al.’ of this group were 71 individuals, many of whom are quite distinguished. This is a large number. It dramatizes the claimed consensus that *something* needs to be done. The social psychology of this tactic is to disarm resistance by leveraging the heuristic of ‘social proof’ (Cialdini, 1984). If all these statisticians agree – the bench scientist is nudged to infer – then I should shrink my alpha. In a *réplique*, Lakens recruited 87 co-authors to argue against any privileged alpha level. Reasonable as their arguments are, the use of social proof wears thin when it begins to smack of mimicry and when the ratio of the number of written words to the number of authors becomes distressingly low (here: 22). There ought to be a minimum criterion value for that!

We conclude that the social forces that shape research practice are part and parcel of any evolving science. Instead of wishing them away, we'd do well to understand them, and hope they will move the practice of science forward.

1.5 Conclusion

In 2001, one of us (Krueger, 2001) predicted that NHST would outlast many of its challenges. This prediction was partly based on the method's intrinsic value, partly on the beside-the-pointness of some of the critiques, and partly on the fact that by that time, NHST had already shown itself to be resilient. These three reasons are not independent of one another, and there is no reason for complacency. Like others before us (Abelson, 1995; Wasserstein & Lazar, 2016; Wilkinson & the APA Task Force on Statistical Inference, 1999), we wish to impress on research scientists that the p value is a mere heuristic. It has predictive value, but it guarantees that some false inferences will be made. The p value cannot do all of the inductive work; no single method can. We join those who recommend researchers use a toolbox of statistical techniques, employ good judgment, and keep an eye on developments in statistical and data science.

Appendix

To guard against the impression that our appeal to good judgment is mere handwaving and buck-passing, we offer two examples from recent experience and vivid memory. The first example features a trolleyologist³ explaining the decision to test the significance of 44 successes under the null hypothesis of 25. "I would have done a significance test if I had 50 out of 50 successes," the researcher continued, "because I know the editor would have demanded it." But why perform a significance test to reject the chance hypothesis when the results are so clear that

³ A trolleyologist is someone doing empirical research on philosophical thought experiments such as Philippa Foot's trolley-footbridge dilemma (Foot, 1978/1967).

a naked-eye test reveals the effect. To say that a test would have been performed if the result had been 50 out of 50 is to endorse and perpetuate an empty ritual. We advise the use of good judgment to oppose such rituals, and look forward to seeing more articles in which striking results are reported using only descriptive statistics.

The second example involves a correlation between a trait measure of moral orientation and a specific moralistic action. This correlation turned out to be $+.08$. With over 900 degrees of freedom, this correlation was “highly significant.” When asked whether a correlation of $+.02$, if significant in a much larger sample, would be deemed satisfactory as well, the researcher said ‘yes.’ What might be done to de-absurdify this approach to data analysis? For this case, we recommend that a plausible substantive hypothesis be selected for strong significance testing in Meehl’s sense or for the estimation of a likelihood ratio. To wit, a correlation of $+.3$ may be installed to represent $\sim H$, that is, the baseline expectation of how well general traits predict specific behaviors (Dawes & Smith, 1985; Mischel, 1968).

References

- Abelson, R. P. (1995). *Statistics as principled argument*. Hillsdale, NJ: Erlbaum.
- Antonakis, J. (2017). On doing better science: From thrill of discovery to policy implications. *The Leadership Quarterly*, 28, 5-21. <http://dx.doi.org/10.1016/j.leaqua.2017.01.006>
- Austin, P. C. (2017). A tutorial on multi-level survival analysis: Methods, models and applications. *International Statistical Review*, 85, 185-203. doi: 10.1111/insr.12214
- Bayarri, M. J., Benjamin, J., Berger, J. O., & Sellke, T. M. (2016). Rejection odds and rejection ratios: A proposal for statistical practice in testing hypotheses. *Journal of Mathematical Psychology*, 72, 90-103. doi: [10.1016/j.jmp.2015.12.007](https://doi.org/10.1016/j.jmp.2015.12.007)
- Benjamin, D. J., Berger, J., Johannesson, M., Nosek, B. A., Wagenmakers, E.-J., Berk, R., ... Johnson, V. (2017, July 22). Redefine statistical significance. Retrieved from osf.io/preprints/psyarxiv/mky9j
- Bredenkamp, J. (1980) *Theorie und Planung psychologischer Experimente* [Theory and design of psychological experiments]. Darmstadt, Germany: Stemkopff
- Bruner, J. S., Goodnow, J. J., & Austin, G. A. (1956). *A study of thinking*. New York: Wiley.
- Cialdini, R. (1984). *Influence: The psychology of persuasion*. Boston, MA: Allyn & Bacon.
- Cohen, J. (1977). *Statistical power analysis for the behavioral sciences*. New York, NY: Academic Press.
- Cohen, J. (1990). Things I have learned (so far). *American Psychologist*, 45, 1304-1312. doi: 10.1037/0003-066X.45.12.1304
- Cohen, J. (1992). A power primer. *Psychological Bulletin*, 112, 155-159. doi: 10.1037/0033-2909.112.1.155
- Cohen, J. (1994). The earth is round ($p < .05$). *American Psychologist*, 49, 997-1003. <http://dx.doi.org/10.1037/0003-066X.49.12.997>

- Colquhoun, D. (2015). An investigation of the false discovery rate and the misinterpretation of p -values. *Royal Society Open Science*, 1, 140216. <http://dx.doi.org/10.1098/rsos.140216>
- Dawes, R. M. (1988). *Rational choice in an uncertain world*. San Diego, CA: Harcourt, Brace, Jovanovich.
- Dawes, R. M., & Smith, T. L. (1985). Attitude and opinion measurement. In G. Lindzey & E. Aronson (Eds.), *Handbook of social psychology* (3rd ed., Vol. 1, pp. 509–566). New York: Random House.
- Erdfelder, E., Faul, F., & Buchner, A. (1996). GPOWER: A general power analysis program. *Behavior Research Methods, Instruments & Computers*, 28, 1-11. doi: 10.3758/bf03203630
- Feyerabend, P. K. (1976). *Against method*. New York, NY: Humanities Press.
- Fisher, R. A. (1956). *Statistical methods and scientific inference*. Oxford, England: Hafner Publishing Co..
- Fiedler, K., Kutzner, F., & Krueger, J. I. (2012). The long way from error control to validity proper: Problems with a short-sighted false-positive debate. *Perspectives on Psychological Science*, 7, 661–669. <http://doi.org/10.1177/1745691612462587>
- Fleck, L. (1935). *Entstehung und Entwicklung einer wissenschaftlichen Tatsache. Einführung in die Lehre vom Denkstil und Denkkollektiv*. Benno Schwabe & Co.
- Foot, P. (1978/1967). *Virtues and vices*. Oxford, UK: Blackwell. Originally published in the *Oxford Review*, 5, 1967.
- Gelman, A., & Carlin, J. (2017). Some natural solutions to the p-value communication problem—and why they won’t work. *Journal of the American Statistical Association*, 112, 899-901, DOI: 10.1080/01621459.2017.1311263

- Gelman, A., & Stern, H. (2006). The difference between “significant” and “not significant” is not itself statistically significant. *The American Statistician*, 60, 328-331.
<http://dx.doi.org/10.1198/000313006X152649>
- Gigerenzer, G. (2004). Mindless statistics. *The Journal of Socio-Economics*, 33, 587-606.
 doi:10.1016/j.socec.2004.09.033
- Gigerenzer, G., and Marewski, J. (2015). Surrogate science: The idol of a universal method for scientific inference. *Journal of Management*, 41, 421–440.
 doi:10.1177/0149206314547522
- Gilhooly, K. J., Georgiou, G., & Devery, U. (2012). Incubation and creativity: Do something different. *Thinking & Reasoning*, 19, 137-149. doi: [10.1080/13546783.2012.749812](https://doi.org/10.1080/13546783.2012.749812)
- Goodman, S. (2008). A dirty dozen: Twelve *p*-value misconceptions. *Seminars in Hematology*, 45, 135-140. doi:10.1053/j.seminhematol.2008.04.003
- Greenland, S. (2017). The need for cognitive science in methodology. *American Journal of Epidemiology*, 186, 639-645. <https://doi.org/10.1093/aje/kwx259>
- Greenland, S., Senn, S. J., Rothman, K. J., Carlin, J. B., Poole, C., Goodman, S. N., & Altman, D. G. (2016). Statistical tests, *P* values, confidence intervals, and power: a guide to misinterpretations. *European Journal of Epidemiology*, 31, 337-350. doi: 10.1007/s10654-016-0149-3
- Ioannidis, J. P. A. (2005). Why most published research findings are false. *PLoS Medicine*, 2, e124. <http://dx.doi.org/10.1371/journal.pmed.0020124>
- Johnson, V. E., Payne, R. D., Wang, T., Asher, A., & Mandal, S. (in press). On the reproducibility of psychological science. *Journal of the American Statistical Association*.
<http://dx.doi.org/10.1080/01621459.2016.1240079>

- Kerr, N. L. (1998). HARKing: Hypothesizing After the Results are Known. *Personality and Social Psychology Review*, 2, 196-217. https://doi.org/10.1207/s15327957pspr0203_4
- Killeen, P. R. (2005). An alternative to null-hypothesis significance tests. *Psychological Science*, 16, 345-353. doi: 10.1111/j.0956-7976.2005.01538.x
- Krueger, J. I. (2001). Null hypothesis significance testing: On the survival of a flawed method. *American Psychologist*, 56, 16–26. <http://doi.org/10.1037//0003-066X.56.1.16>
- Krueger, J. I. (2017). Reverse inference. In S. O. Lilienfeld & I. D. Waldman (Eds.), *Psychological science under scrutiny: Recent challenges and proposed solutions* (pp. 110-124). New York, NY: Wiley.
- Krueger, J., & Clement, R. W. (1994). Memory-based judgments about multiple categories: A revision and extension of Tajfel's accentuation theory. *Journal of Personality and Social Psychology*, 67, 35-47. doi:10.1037/0022-3514.67.1.35
- Krueger, J. I., & Heck, P. R. (2017). The heuristic value of p in inductive statistical inference. *Frontiers in Psychology*
<http://journal.frontiersin.org/article/10.3389/fpsyg.2017.00908/abstract>
- Krueger, J. I., & Heck P. R. (in revision). Testing significance testing.
- Krueger, J. I., Heck, P. R., & Asendorpf, J. B. (2017). Self-enhancement: Conceptualization and Assessment. *Collabra: Psychology*, 3(1), 28. <http://doi.org/10.1525/collabra.91>
- Krueger, J. I., Vogrincic-Haselbacher, C., & Evans, A. M. (in preparation). We need a credible theory of gullibility. To appear in Forgas, J. P. & Baumeister, R. F. (2019). *Homo Credulus: The social psychology of gullibility [The 20th Sydney Symposium on Social Psychology]*. New York: Taylor & Francis.
- Kuhn, T. S. (1962). *The structure of scientific revolutions*. Chicago, IL: University of Chicago

Press.

Lakens, D., Adolphi, F. G., Albers, C. J., Anvari, F., Apps, M. A. J., Argamon, S. E., . . . Zwaan,

R. A. (2018). Justify your alpha. *Nature Human Behavior*, 2, 168-171, doi:

10.1038/s41562-018-0311-x

Lehmann, E. L. (1993). The Fisher, Neyman-Pearson theories of testing hypotheses: One theory

or two. *Journal of the American Statistical Association*, 88, 1242-1249. doi:

[10.1080/01621459.1993.10476404](https://doi.org/10.1080/01621459.1993.10476404)

Meehl, P. E. (1978). Theoretical risks and tabular asterisks: Sir Karl, Sir Ronald, and the slow

progress of soft psychology. *Journal of Consulting and Clinical Psychology*, 46, 806-834.

<http://dx.doi.org/10.1037/0022-006X.46.4.806>

Mischel, W. (1968). *Personality and Assessment*. Mahwah, NJ: Erlbaum.

Neyman, J., & Pearson, E. S. (1933a). On the problem of the most efficient tests of statistical

hypotheses. *Philosophical Transactions of the Royal Society A: Mathematical, Physical*

and Engineering Sciences, 231, 694-706. doi: 10.1098/rsta.1933.0009

Neyman, J., & Pearson, E. S. (1933b). The testing of statistical hypotheses in relation to

probabilities a priori. *Mathematical Proceedings of the Cambridge Philosophical Society*,

29, 492-510. doi: 10.1017/S030500410001152X

Nickerson, R. S. (2000). Null hypothesis significance testing: a review of an old and continuing

controversy. *Psychological Methods*, 5, 241-301. [http://doi.org/10.1037/1082-](http://doi.org/10.1037/1082-989X.5.2.241)

989X.5.2.241

Ortega, A., & Navarrete, G. (2017). Bayesian hypothesis testing: An alternative to Null

Hypothesis Significance Testing (NHST) in psychology. In J. P. Tejedor (ed.). *Bayesian*

inference. Open access. doi: 10.5772/intechopen.70230

- Ostrowski, B. (2017). *Vujà Dé: The effects of incubation on creativity*. Unpublished honors thesis in psychology. Providence, RI: Brown University.
- Revlin, R., Leirer, V. , Yopp, H., & Yopp, R. (1980). The belief-bias effect in formal reasoning: The influence of knowledge on logic. *Memory & Cognition*, 8, 584-592.
- Stigler, S. M. (1999). *Statistics on the table: The history of statistical concepts and methods*. Cambridge, MA: Harvard University Press.
- Stroebe, W. (2016). Are most published social psychological findings false? *Journal of Experimental Social Psychology*, 66, 134-144.
<http://dx.doi.org/10.1016/j.jesp.2015.09.017>
- Swets, J. A., Dawes, R. M., & Monahan, J. (2000). Psychological science can improve diagnostic decisions. *Psychological Science in the Public Interest*, 1, 1–26.
<http://doi.org/10.1111/1529-1006.001>
- Tajfel, H. (1959). Quantitative judgement in social perception. *British Journal of Psychology*, 50, 16-29.
- Tajfel, H. (1969). Cognitive aspects of prejudice. *Journal of Social Issues*, 25, 79-97.
- Trafimow, D. (2003). Hypothesis testing and theory evaluation at the boundaries: Surprising insights from Bayes's theorem. *Psychological Review*, 110, 526-535.
<http://dx.doi.org/10.1037/0033-295X.110.3.526>
- Trafimow, D., & Marks, M. (2015). Editorial. *Basic and Applied Social Psychology*, 37, 1–2.
<http://doi.org/10.1080/01973533.2015.1012991>
- Wagenmakers, E.-J. (2007). A practical solution to the pervasive problems of *p* values. *Psychonomic Bulletin & Review*, 14, 779-804. doi: 10.3758/BF03194105
- Wasserstein, R. L., & Lazar, N. A. (2016). The ASA’s statement on p-values: Context, process,

and purpose. *The American Statistician*, 70, 129-133. doi:

10.1080/00031305.2016.1154108; 21 online expert commentaries available at

<http://amstat.tandfonline.com/doi/suppl/10.1080/00031305.2016.1154108?scroll=top>

Wilkinson, L., & the Task Force on Statistical Inference. (1999). Statistical methods in

psychology journals: Guidelines and explanations. *American Psychologist*, 54, 594–604.

doi:10.1037/0003-066X.54.8.594