SCALING THE SLIPPERY SLOPE:

How *p*-VALUES RUINED SCIENTIFIC PRACTICE IN PSYCHOLOGY, AND HOW WE CAN FIX IT IMMEDIATELY

Zachariah Basehore*

Draft version 3.1, January 23, 2017. Article currently under review at *Perspectives on Psychological Science*. Please cite as:

Basehore, Z. (2017). Scaling the slippery slope: How p-values ruined scientific practice in psychology, and how we can fix it immediately. Manuscript under review at *Perspectives on Psychological Science*. Retrieved from PsyArXiv, <http://osf.io/qhevg>

* Please address correspondence regarding this article to Zach Basehore, Department of Psychology, Bowling Green State University, Bowling Green, Ohio, 43403, USA; e-mail: zdbaseh@bgsu.edu

Department of Psychology, Bowling Green State University, OH, USA

Abstract

For over a half-century, psychologists, educators, and researchers have criticized the common misuses of statistics in the social sciences. Here, I summarize some of the various objections to the blind use of p -values and propose simple adjustments to 1) ameliorate the weaknesses inherent in current statistical practice, and 2) to paint a more complete picture of a study's results.

Keywords: research, methodology, statistics, p -values, NHST, best research practices, significance, Bayesian tests, Bayesian stats, JASP

As many students and professors in such disciplines as psychology, sociology, education, and biomedical sciences will readily attest, the typical goal of research is usually to find a correlation, ANOVA, or t test that yields a p -value of .05 or less. This exclusive focus on statistical significance has been called idolatrous (Gigerenzer & Marewski, 2015, in the article's title; Ziliak & McCloskey, 2008, in the preface), "mindless" (Gigerenzer, 2004, title), "embarrassing" (Gigerenzer, Krauss, & Vitouch, 2004, p. 6), and an "educational failure" (Hubbard & Armstrong, 2005, title) that yields a "cult"-like devotion among practitioners (Ziliak & McCloskey, 2008; 2009, titles)!

But these criticisms aren't new; rather, they stretch back over 50 years! For example, see Stevens' characterization of null hypothesis significance testing (NHST) procedures as "a meaningless ordeal of pedantic computations" (Stevens, 1960; p. 276). See also Bakan's remark later in the same decade that "the test of significance has been carrying too much of the burden of scientific inference" (Bakan, 1966, p. 423), or Lykken's observation that statistical significance in an NHST procedure is "perhaps the least important attribute of a good experiment; it is never a sufficient condition for claiming that a theory has been usefully corroborated, that a meaningful empirical fact has been established, or that an experimental report ought to be published" (Lykken, 1968; p. 151), or that significance testing constitutes "one of the worst things that ever happened in the history of psychology" (Meehl, 1978; p. 817), or that it is "mindless hypothesis testing in lieu of doing good research" (Luce, 1988; p. 582).

This strict focus on p -values below .05 in many fields ultimately leads to articles like the highly influential article "Why most published research findings are false" by Ioannidis (2005), who cites the common practice of "claiming conclusive research findings solely on the basis of a single study assessed by formal statistical significance, typically for a p -value less than 0.05" (p. 0696) as a major reason for the frequent failure to replicate findings in biomedical science. It is not unreasonable to assume that this same practice is also an important factor behind the failure of replication in

psychology studies, as summarized by the Open Science Collaboration (2015) [but see Gilbert, King, Pettigrew, & Wilson (2016) for a contravening interpretation. Also see Etz & Vandekerckhove (2016) for a Bayesian re-analysis of the results that showed that, although a majority of the original studies provided only weak or ambiguous evidence in favor of the alternative hypothesis, none of the replications strongly contradicted the original effects].

Readers may be tempted to dismiss such objections as the abstruse nit-picking of cantankerous old intellectuals. But anybody who has struggled to understand the logic of NHST should be quick to realize that NHST is “fatally flawed” (Cumming, 2014, p. 11). It is crucial for instructors and researchers alike to understand these criticisms from a historical and a philosophical perspective: the only way to achieve constructive change in our discipline is to stop teaching NHST as though it’s the only way to analyze research—and to stop behaving in the same manner ourselves!

So settle in, and enjoy the trip through a historical and philosophical minefield. Before we begin, I’d also like to note that progress has been made since Wilkinson and the APA Task Force report was published in 1999; Thompson (2004) characterized this progress as “inexorably, albeit glacially” present (p. 608). Hopefully, this manuscript will help this progress proceed a bit less glacially.

A full and detailed overview of the objections to NHST would be far beyond the scope of this manuscript, and would likely be prohibitively long and detailed. Few would want to read such an *oeuvre*, and even fewer would want to write it! Instead, the focus of this article will be to briefly summarize the major criticisms and to present an easily—and immediately—implementable strategy to improve statistical practice in the field of psychology.

The objections

Objection 1. The first objection to NHST, and (I think) the most damning, comes from Morgan (2003) and Gigerenzer, Krauss, and Vitouch (2004), among others. These authors present a compelling

case that NHST as it is typically used (and continues to be taught!) is a conceptually incoherent hodgepodge of two mutually incompatible philosophies.

Sir R. A. Fisher regarded the ‘.05 rule’ described above as a flexible guideline rather than a hard-and-fast line of demarcation; later in his life, Fisher decried the fixed and automatic use of this criterion (Gigerenzer, Krauss, & Vitouch, 2004). More fundamentally, Fisher saw the p -value as a characteristic of the experimental data; a reflection of the likelihood that a result *at least* as extreme as the observed data would have occurred due to sampling error, *given the truth of a null hypothesis with a known distribution*. Statisticians Jerzy Neyman and Egon Pearson, however, viewed the critical alpha level as a characteristic of the test itself, therefore arguing that researchers must set the alpha level before the experiment, and report only whether or not the p -value was smaller than the critical alpha level (Gigerenzer, Krauss, & Vitouch, 2004). By clearly stating these fundamental philosophical differences underlying the different approaches to significance testing, Gigerenzer, Krauss, and Vitouch (2004) make evident that the current NHST procedures give rise to irreconcilable issues.

One such issue is the debate over whether p should be reported as an exact value, like $p = .017$, or as having exceeded a certain pre-set threshold, as in $p < .05$. If you believe the former, you are adhering to Fisher’s perspective. If you argue the latter, you are following the Neyman-Pearson approach. Though exact p -values are frequently used in the modern literature, as recommended by Wilkinson & TSFI (1999; p. 599), there remain some who teach and/or practice the dichotomous classification of $p < .05$ or $p > .05$ (Head, Holman, Lanfear, Kahn, & Jennions, 2015). There are arguments to be made for each case, but Fisher, Neyman, and Pearson (who argued vitriolically over whose approach to statistical practice in research was better-justified) would likely be equally horrified by the Frankensteinian hybrid of today’s statistical practice!

Objection 2. Cowles and Davis (1982) argued that the .05 criterion did not originate with Fisher’s assertion that it is “...convenient to take this point as a limit in judging whether a deviation is

to be considered significant or not” (Fisher, 1925, p. 47, as cited in Cowles and Davis, p. 553). They presented the case that Fisher’s rule of 20-to-1 odds of finding a given result by chance is similar to the ‘probable error’ of 3.2 (30-to-1 odds) that researchers commonly used before Fisher popularized NHST. Cowles and Davis concluded that Fisher’s .05 rule is therefore justified.

But this criterion remains subjective. Who’s to say that a 5% risk of a Type I error (defined in current NHST practice as an incorrect rejection of the null hypothesis, often explained as a “false alarm”) is acceptable but a 10% risk is not? Since 30-to-1 odds (corresponding to a p -value of .033) would be a more stringent criterion than the current 20-to-1 ($p = .05$), should we return to that rule to determine what is publishable? Or should we instead loosen the criterion to 15-to-1 ($p = .066$), 10-to-1 ($p = .10$), or even 8-to-1 ($p = .125$)?

Such a debate is utterly misguided. The ‘acceptable level of risk’ often depends upon the question (Ziliak & McCloskey, 2008; 2009). My undergraduate statistics professor communicated the point this way: if I were to offer you an AIDS vaccine, stating that there was only a 5% chance that you’d contract the deadly disease from the vaccine, would you consent to receive the injection? The rhetorical point of this example is that few people from a low-risk group would volunteer to receive this vaccine; a Type I error rate of .001 would be *much* better than a Type I error rate of .05 in such an instance (i.e. a 0.1% chance of a backfiring AIDS vaccine is preferable to a 5% chance).

But researchers and journal editors alike almost invariably use the .05 significance level as a cutoff for determining importance. Ziliak and McCloskey (2008) lament this practice, calling it the “sizeless stare of statistical significance” (p. 17). Is a p -value of .049 as important as a p -value of .001? Are the two studies equally valuable? Do they indicate equally replicable effects?

Unfortunately, no standard answer to these questions is possible. The answers to these critical questions depend upon a variety of factors, including: the quality of each study’s experimental design,

the use of random sampling, the size of the samples in each study, the magnitude of the effect, and the topic being studied.

To illustrate this, imagine that the p -value of less than .001 came from a repeated-measures study comparing the height of plants at 3 PM on a Monday to the height of those same plants at 3 PM the following Monday. Despite the impressively small p -value, the conclusion that the plants grew taller over the course of that week is not exactly earth-shattering.

Objection 3. Selective publication, also known as the ‘file-drawer problem’ or ‘publication bias,’ has often been cited (e.g. Ioannidis, 2005; Nosek, Spies, & Motyl, 2012; Head et al., 2015) as a major biasing factor in how scientific knowledge is disseminated. This happens when a study is published when the p -value is .05 or lower, but studies on the same topic are not published when p is greater than .05. This practice could, and probably does, result in a single study with spuriously significant ‘findings,’ while numerous additional studies on the same topic are not published because they ‘didn’t work’—meaning that they did not achieve statistically significant results at an alpha level of .05. This practice results in the publication and dissemination of a *single study* showing an effect (a correlation, a difference between groups, an important risk factor, etc.), while 19 other studies on the same topic—studies that did *not* yield a ‘publishable’ effect—would be known to only a handful of researchers.

The ‘file-drawer problem’ thereby circumvents the self-correcting nature of science, at least in the short term (Nosek, Spies, & Motyl, 2012). If one study shows evidence for an effect but a different study shows no evidence, impartial observers have to judge which one they believe is correct. (By “correct,” I’m referring to the degree to which the research finding reflects the state of affairs in the overall population being studied). If both experiments are well-designed, there may be little or no objective basis to inform the observer’s judgment about which finding is “correct.”

This demonstrates the need for well-designed experiments to be published, regardless of the results. If we want to know how people think, it is every bit as important to know about the non-significant findings as it is to know about the significant ones! The practice of methodological pre-registration, practiced by a growing number of journals like *Perspectives on Psychological Science*, is a good step in the right direction, though this practice alone is inadequate to fully correct the ‘file-drawer problem’ (Nosek, Spies, & Motyl, 2012).

Objection 4. Another damning objection to current NHST procedures is that researchers always set up a null hypothesis of no difference or no correlation, and then specify an alternative hypothesis that, miraculously, nearly always aligns with the actual finding (Gigerenzer, Krauss, & Vitouch, 2004; Nosek, Spies, & Motyl, 2012; Head et al., 2015)!

The big problem here is a logical one. We’re testing whether or not an effect is large enough to preclude chance as a reasonable explanation. If the p -value is .05 or less, then the result is taken as evidence that the data cannot be explained by sampling error alone. Even if we agree that .05 is the best cutoff point for a decision rule—which, as we’ve seen, isn’t always the best policy—this only means that we have evidence that the null is false; the test provides no evidence in favor of a particular alternative hypothesis (Carver, 1978; Morgan, 2003)!

Anybody who has studied even rudimentary persuasion or rhetoric has surely heard of the ‘straw-man’ fallacy. The straw-man fallacy is when one sets up an argument that’s similar to an opponent’s argument (though different in some not-obvious but important way), then attacks that argument instead of the opponent’s *actual* argument. Say that two candidates for a political office are in a public debate. Candidate P advocates lessening the penalties for possession of illegal drugs. Candidate Q objects that Candidate P is “soft on crime” and would let convicted criminals run loose among the general populace. This simple example illustrates the straw-man fallacy in action.

In the case of significance-based research, we're all committing a version of this fallacy. If the p -value is smaller than .05, that is conventionally taken as evidence that sampling error (or 'chance') is not a likely explanation for the observed results. However, even when we show that chance is a poor explanation for the results, an infinite number of potential explanations remains! When we conduct research, we want to show that a particular explanation *is* the reason why an effect is present; what NHST procedures actually show is that chance is an unlikely explanation *if there is no effect*. But these are not the same: ruling one explanation out is different from supporting an alternative explanation!

Here's a simple, fictitious example to which I will refer throughout this essay: I'm trying to show that Medication X influences cognitive function—I'm not sure whether it will improve recall or impede it; I only suspect that Medication X will have an impact. My measurement of "cognitive function" is operationally defined by free recall of words from a list. A treatment group of 30 people takes Medication X daily, while a control group of 30 people remains on a waiting list and takes no medication. An independent-samples t test shows that after six weeks, people in the treatment group show a higher number of words recalled from a list, $t(58) = 2.384, p < .05$.

This analysis does not provide good evidence that Medication X is the *cause* of the better recall in the treatment group. The difference *may* be due to Medication X...but it may also be due to a placebo effect, or to unusually great effort by people in the treatment group (or to unusually low effort by people who are miffed at having been waitlisted), or to the behavior of the researchers toward members of the treatment group, or to a failure to properly assign participants randomly to condition, or to some combination of these and other factors.

After conducting the analysis described above, I only have evidence that sampling error is *not* a good explanation; I have no evidence of which explanation *is* good! The observed difference in means may just as easily have come from a placebo effect as from Medication X; no statistical test can tell me which of the explanations listed above is the best fit for the data, *because the experiment itself was not*

designed to answer that question! The p -value therefore fails to tell us what we want to know (Cohen, 1994). We want an explanation for why something happens. NHST does not give us that.

Objection 5. As others have noted (e.g. Lykken, 1968; Meehl, 1990; Gigerenzer, 2004; Wagenmakers, 2007), there is rarely, if ever, a difference (or a correlation) of *exactly* zero between groups of people (or between variables in a correlation) on some dimension. The idea behind this is that researchers are not selecting random, unrelated things to study—there is likely to be a reason to suspect that a relationship exists; otherwise, the study would not be very compelling. Since we can expect a sample to be even less likely to exhibit zero correlation or zero difference, a large sample is therefore likely to yield a statistically significant NHST result. In fact, it is a virtual mathematical certainty (see Wagenmakers, 2007, p. 781)! I will explain using an independent-samples t test.

Even a tiny mean difference and moderate variance in each group, combined with an enormous sample size, can result in $p < .05$ —the magic happens in Step 2. Here's how:

- Step 1) Pooled variance is calculated. If the samples are equal in size, the pooled variance is halfway between the variance of each group; otherwise, it is weighted by sample size such that the pooled variance is closer to the variance of the larger sample. Nothing special here.
- Step 2) Find standard error. This consists of two main steps:
 - 2a) divide the pooled variance by the sample size of the first group, divide the pooled variance by the sample size of the second group, and add those values, and
 - 2b) take the square root of that number. The resulting value is the standard error.
 - In Step 2a), pooled variance is divided by the huge sample size of each group. Dividing some number by a much bigger number results in a very small value (for example: $10/100 = 0.1$; whereas $10/100,000 = 0.0001$). Unless the mean difference is also in the thousands, a large sample size such as 2000 is certain to result in a

quotient of less than 1, because the sample size is in the denominator of step 2a. In Step 2b), note that the square root of any real number except 1 itself approaches, but never quite reaches, 1 (e.g. $\sqrt{0.64} = 0.80$). Taking the square root of a quotient between 0 and 1 will *always* result in a number that is still smaller than 1!

- Step 3) Subtract the mean of one group from the mean of the other group (it doesn't matter which is subtracted from which) to get the mean difference.
- Step 4) Divide the mean difference by the standard error. The result of this calculation gives the value of t . For any study with a sample of at least 6 and an alpha level of .05, a mean difference of 1.96 points or more and a standard error below 1 is guaranteed—*guaranteed*—to give a statistically significant value for t ! Dividing a positive number by any value smaller than 1 results in a quotient that is larger than the numerator. And, as we saw in Step 2b), standard error is virtually certain to be less than 1 if you have an enormous sample.

Any readers who have had difficulty following the above explanation can quickly and simply test this for themselves by entering the following numbers at <http://graphpad.com/quickcalcs/ttest1/> **Format=SD**: Group One mean = 50, Group Two mean = 49, Group One $SD = 28$, Group Two $SD = 35$, Group One $n = 10,000$, Group Two $n = 10,000$. Click “calculate now” and observe the resulting values of t and p . After this, enter the same numbers for the mean and SD of each group, but change the n of each group to 100, and click “calculate now” to observe the resulting values of t and p . While comparing these values, consider what has changed about the effect in this example—has the mean difference changed? Has the variance changed? If not, then why should the researcher's decision about whether there's a 'real' or 'important' effect change?

The sample size doesn't even have to be that extreme in order to observe this effect! Suppose that, in comparing the means of two different groups, I find a 2-point difference on a 100-point scale. If

each group has a standard deviation of 9 and an n of 15, I will not find a significant difference [$t = 1.7638$]. But the same mean difference and SD, with a sample size of 20 per group instead of 15, *would* yield a significant result [$t = 2.0548$]*—despite the fact that neither the means nor the variance changed at all!* Not to belabor the point, but it's worth making explicit: the exact same difference, with the exact same variance, can be either statistically significant at the conventional .05 level or not significant, *based only on the sample size* (see Rosnow & Rosenthal, 1989). If a given effect is not worthy of consideration when the sample has 15 people, why should the same data be worthy of consideration when there are 20 people in the sample?

Ziliak & McCloskey (2008; 2009) argue that NHST procedures measure sampling precision. Depending on the goal of a particular study, that may be useful. Oftentimes, however, researchers are (or should be) less concerned with the odds of falsely rejecting the null, and more concerned with the magnitude and direction of an effect. For this purpose, NHST procedures are not ideal, in part because of the excessive influence of sample size on the test. After all, the only real difference between a t test and Cohen's d is the fact that the t test features standard error—which includes both variance and sample size—in the denominator, while Cohen's d relies instead on the pooled SD in the denominator. Strip out the sample size from a t test, and you end up with a measure of effect size!

Meehl (1990) discussed the same principle in regard to correlations; a survey of 57,000 high school students in the U.S. state of Minnesota yielded statistical significance for a whopping 92% of relationships between the various “leisure activities” and “high school organizations” listed in the survey! Meehl also noted that 78% of those relationships would have been significant at an impressive alpha level of .00001 (p. 205)!

All the above is to say that any researcher can tip the scales in favor of obtaining a statistically significant (hence, publishable) result simply by collecting a sample numbering in the thousands. With such tools as online data collection or online access to archival survey data, it is easier than ever to

gain access to such large samples. Thereby, researchers can (intentionally or not) manipulate their p -values to get a more desirable result, *while simultaneously advertising their apparent statistical sophistication by appealing to the law of large numbers!* That entire disciplines filled with PhDs can maintain such Orwellian ‘doublethink’ for decades is a truly impressive feat! Perhaps Upton Sinclair was right to remark how difficult it is to get a man to understand something when his salary depends upon his not understanding it.

I will stop here, as I think these five objections are *each* sufficient to consign the still-common practice of “if p is less than .05, then publish” to the rubbish heap. As Meehl (1990) wrote: “If the reader is impelled to object at this point, ‘...you are practically saying that the whole tradition of testing substantive theories...by null hypothesis refutation is a mistake, despite R. A. Fisher and Co. in agronomy,’ that complaint does not disturb me because that is *exactly* what I am arguing” (p. 198; italics in original). This leaves us with two crucial questions: why did such a fatally flawed approach reach prominence in the first place, and what should we do about it?

Why did p -values rise to dominance?

It would be irresponsible of me to claim expertise in the history of psychological research, so I will keep this section brief. I have a conjecture, however, about which the reader will have to determine the quality of the argument.

Rucci and Tweney (1980) observed that the ANOVA was described in the psychological literature as early as 1923, and was used to analyze psychological data beginning in 1934. Use of the ANOVA increased especially quickly after World War II, and the proportion of published articles featuring the ANOVA and the t test overtook the proportion using the critical ratio as a statistical inferential technique shortly before 1950. By 1962, A.W. Melton, editor of the prestigious *Journal of Experimental Psychology*, insisted that a p -value of .05 was necessary for publication, and that a p -value of .01 was preferable, thereby enshrining p -values as the standard by which research is judged

(Gigerenzer, 2004). For better or for worse—mostly for worse—the entire field of psychology had hopped aboard the p -value train by this point, and many other social sciences followed suit.

I would like to submit the plausible conjecture that this occurred because then, as today, social scientists desired to make their chosen discipline respectable and scientific in the view of others. The economist F. A. Hayek (1942) observed that the great success of the physical and biological sciences led thinkers in other fields to imitate the methods and techniques of the scientific disciplines (p. 268). In much the same way that behaviorism was conceptualized and popularized in part as a reaction to allegations that psychology wasn't truly scientific (Watson, 1913; Roediger, 2004), I suspect that the motivation to adopt automatic inferential techniques like the non-negotiable cutoff of .05 stemmed, at least in part, from a similar desire to make the discipline more scientifically respectable (Levine, Weber, Hullet, Park, & Massi Lindsey, 2008).

After all, budding scientists are taught that a person's judgment is fallible, but a mathematical test is not. The adoption of automatic inferential techniques gave researchers an easy way to deflect criticism and to make the discipline seem more scientific: either your results reached the .05 cutoff or they didn't. See; no room for subjectivity here! Science at work!

But the key word above is “seem.” Making psychology *seem* more scientific is not making it more scientific *in actuality*. Instead of emphasizing tightly controlled manipulations, testing an effect several times before accepting it as “real,” and thinking critically about what the results mean about the effect under investigation—the very essence of science rests on these factors alone—psychologists instead substitute mechanized pseudo-statistical procedures! The inevitable result is the current ‘replication crisis’ in psychology, and the continued dismissal of the field as ‘not a *real* science’ by many physical scientists who see through this charade.

In this light, it becomes rather evident why these poor statistical practices rose to prominence, and persisted in the face of 50+ years worth of objections: researchers were grasping for a means by

which to defend the value of their work! When Fisher advocated the p -value in the 1920s and 30s as the outcome of the scientific process, and journal editors like Melton later codified it, people in a variety of disciplines latched onto this procedure as evidence that they are, in fact, doing science. That is understandable, and I can sympathize with the impulse.

But it does not change the reality that such researchers are actually replacing scientific thought with a mechanized interpretation of a mathematical procedure of limited import. Instead of thinking carefully about a study's results, and thoughtfully defending that interpretation, researchers instead slide down a slippery research slope by substituting the simple question "Is p less than .05?" If the answer is 'yes,' then we submit for publication. If not, we don't waste our time writing about it, because it is unlikely to be published anyway. But published or not, it still happened, as Meehl (1990) observed: "*A fact that occurs in the laboratory or exists in the clinic files is just as much a fact whether Jones elects to write it up for a refereed journal or to forget it*" (p. 214; italics in original).

As Gigerenzer (2004) noted, historic discoveries in science did not depend upon p -values. The discovery of the laws of thermodynamics did not require p -value-based statistical inference. Darwin's discovery of natural selection did not include a p -value, nor did Mendel's discovery of genetic heritability, Curie's discovery of X-rays, Crick and Watson's discovery of the double-helix structure of DNA, van Leeuwenhoek's discovery of bacteria, Goldberger's discovery that pellagra was induced by diet rather than germs, Pavlov's work on classical conditioning, or Piaget's theory of cognitive development in children. If discoveries of such enduring importance did not require mechanized inferential procedures like automatic p -value cutoffs, then why do these procedures persist as the near-universal standard of determining importance today?

Scaling the slippery slope

What should we do instead? Now that we've slid down the slippery slope of poor research practices, how can we clamber back up that slope? Though many excellent recommendations for better

statistical practices have been made already (e.g. Cumming & Finch, 2001; Morgan, 2003; Gigerenzer, Krauss, & Vitouch, 2004; Wagenmakers, 2007; Simmons, Nelson, & Simonsohn, 2011; Nosek, Spies, & Motyl, 2012; Cumming, 2014; Kruschke & Liddell, in press), the fact remains that most researchers have been trained in NHST procedures like correlations and t tests at both the undergraduate and the graduate levels; many also boast lengthy and distinguished careers in research and publication, usually based on calculating and interpreting p -values. They are therefore understandably loathe to abandon those practices and to learn new ones (Gigerenzer, 2004; Thompson, 2004). Gigerenzer, Krauss, and Vitouch (2004) shine a very revealing light on this psychological inertia:

“When Loftus (1993) became the editor of *Memory & Cognition*, he made it clear in his editorial that he did not want authors to submit papers in which p -, t -, or F -values are mindlessly being calculated and reported. Rather, he asked researchers to keep it simple and report figures with error bars, following the proverb that ‘a picture is worth more than a thousand p -values.’ We admire Loftus for having had the courage to take this step. Years after, one of us (GG) asked Loftus about the success of his crusade against thoughtless significance testing. Loftus bitterly complained that most researchers actually refused the opportunity to escape the ritual. Even when he asked in his editorial letter to get rid of dozens of p -values, the authors insisted on keeping them in. There is something deeply engrained in the minds of many researchers that makes them repeat the same action over and over again” (p. 14)

In defense of p -values. It strikes me as silly for a psychologist to ignore principles like psychological inertia, not to mention evidence such as that provided by the above story, when dispensing advice to other psychologists. So I believe it is more sensible to account for the weaknesses of p -values by acknowledging them, and then providing additional information, to give readers a more comprehensive view of the results in the manner of “meta-analytic thinking” as described by Cumming

and Finch (2001, p. 555) and Cumming (2014; p. 14). I would rather advise small but productive steps toward better statistical practice that can be implemented immediately, than to issue yet another call to completely overhaul our current procedures—only to be ignored.

Despite the calls of statistically sophisticated social scientists over the years to ban NHST, I will not join in such calls, as they are knee-jerk reactions to the poor statistical habits that remain common. NHST has its uses, chiefly to gain preliminary evidence for the existence of an effect. If we're trying to gather evidence that an effect exists, we want to rule out sampling error as an explanation—which is exactly what NHST does. Banning NHST denies the possibility that it could ever be useful; such a position is every bit as extreme as that of automatically using such a procedure. Furthermore, problems like Objection 3 are inherent to the publication process, and would therefore persist to some degree, regardless of which statistical methods are used (Head et al., 2015).

Therefore, despite the objections detailed in the previous pages, I will not advise readers to abandon the use of NHST procedures or p -values, in no small part because they would not heed that advice if I did! Though calls to abandon NHST in psychology have been made for over 50 years, a heavy reliance on NHST remains, the headline-grabbing p -value-ban by the journal *Basic and Applied Social Psychology* notwithstanding (Trafimow & Marks, 2015).

With this in mind, I would like to direct the reader back to the Medication X example described in Objection 4. What does the p -value of less than .05 indicate? It indicates that, *if* the null of zero effect of Medication X on memory is true, the probability is less than 5% that sampling error would explain a mean difference *at least* as large as that observed, given a sample size of 60. Therefore, we have some limited evidence that the null of no effect is *not* true. While this isn't what I'm trying to show (remember, I'm trying to show that Medication X affects memory), it's a start: I have shown that there is, in fact, a difference that is not easily explained by chance. But we only know that the p -value is less than .05, we don't know whether it's .00001 or .049!

This brings us to another bone of contention: nearly everyone would admit that, for a given sample size, there's little difference between a p -value of .047 and .049. So, why should the difference between .049 and .051 be enough to make a difference between publication and non-publication? When academic researchers perceive (whether rightly or wrongly) that publication is crucial to getting a desirable tenure-track job, to getting tenure, and to getting the professional respect of their peers (Nosek, Spies, & Motyl, 2012), the difference between $p = .049$ and $p = .051$ can literally change a person's future, even though a difference of the exact same magnitude between two values below the typical .05 cutoff is considered trivial! I therefore argue that instead of two categories—bigger than .05 and smaller than .05—we should instead use a stepwise classification. I further reiterate that p should be reported as an exact value, as Fisher argued, since it represents an exact probability estimate on a continuous distribution (Rosnow & Rosenthal, 1989).

A stepwise classification system. An existing example of a stepwise classification procedure for the result of a statistical test is Jeffreys' (1961) guideline for classifying a Bayes factor. Briefly put, Bayesian tests result in a Bayes Factor, denoted as either a BF_{10} or BF_{01} value.

A BF_{10} value above 1 indicates that Hypothesis 2 (usually described as the alternative hypothesis) is more likely than Hypothesis 1 (usually called the null). I refer to "Hypothesis 1" and "Hypothesis 2" rather than the "null" and "alternative" hypotheses here because Bayesian tests can evaluate competing explanations. The familiar null of "no difference" or "no effect" is not a necessary feature of a Bayesian test, though it *can* be one of the hypotheses if the researcher so chooses (Wagenmakers, 2007; Kruschke, 2013; Scheibehenne, Rieskamp, & Wagenmakers, 2013).

For example, a BF_{10} of 13.71 indicates that Hypothesis 2 is 13.71 times more likely to account for the observed data pattern than Hypothesis 1 (see Scheibehenne, Rieskamp, & Wagenmakers, 2013), while a BF_{10} of 0.10, for example, means that Hypothesis 2 is one-tenth as likely to account for the observed data pattern as Hypothesis 1 (and Hypothesis 1 is therefore 10 times as likely to account for

the data—making Hypothesis 1 the superior explanation). A BF_{10} that is equal to 1 indicates that the evidence does not support either hypothesis. For an example of the less-reported BF_{01} values: a BF_{01} of 5 indicates that Hypothesis 1 is 5 times more likely than Hypothesis 2, given the observed data.

Wetzels and Wagenmakers (2012; p. 1060) tweaked the wording of Jeffreys’ system, and I further substitute “Hypothesis 1” and “Hypothesis 2” for H_0 and H_1 respectively, as follows:

BF_{10}	Label
> 100	Decisive evidence for Hypothesis 2
$30 - 100$	Very strong evidence for Hypothesis 2
$10 - 30$	Strong evidence for Hypothesis 2
$3 - 10$	Substantial evidence for Hypothesis 2
$1 - 3$	Anecdotal evidence for Hypothesis 2
1	No evidence for either hypothesis
$1 - 1/3$	Anecdotal evidence for Hypothesis 1
$1/3 - 1/10$	Substantial evidence for Hypothesis 1
$1/10 - 1/30$	Strong evidence for Hypothesis 1
$1/30 - 1/100$	Very strong evidence for Hypothesis 1
$< 1/100$	Decisive evidence for Hypothesis 1

Though Bayes factors easily lend themselves to interpretation on a continuous scale independent of sample size, I see little reason not to use a similar approach to classify p -values. Many researchers already view p on a continuous scale (Wagenmakers, 2007, p. 780), as implied by the famous comment that “surely, God loves the .06 nearly as much as the .05” (Rosnow & Rosenthal, 1989, p. 1277).

There are arguments against the p -value as a measure of the strength of evidence, primarily based on the reality (as described in Objection 5) that sample size influences the p -value. However, it is true that, if the sample size is held constant, a smaller p -value does indicate a lower probability that

the results occurred by chance than a larger p -value does—thereby indicating that a p -value provides some evidence (albeit limited) of the strength of an effect.

Indeed, a careful reading of Open Science Collaboration (2015) reveals that a smaller p -value was, in fact, associated with a greater probability of a successful replication (p. aac4716-6). This gives us some experimental data to support the idea that a smaller p -value often does suggest a stronger effect and higher probability of replication than a larger p -value, given a reasonable sample size and the absence of intentional “ p -hacking.” Simulation data by Anderson (submitted manuscript) reveals that p , Cohen’s d , and equal-priors Bayes factor values are monotonically related to one another, thereby providing further evidence that p can indeed provide a general indication of the strength of the evidence against the null for a given sample size.

A new proposal. Having already noted several limitations of p -values, and contingent upon a sample size that is neither too small nor excessively large for the comparison under investigation, I therefore offer the following proposal to classify p -values in a stepwise manner:

Classification of p	
$p < .001$	Overwhelming evidence against chance
$.001 \leq p < .01$	Very strong evidence against chance
$.01 \leq p < .05$	Strong evidence against chance
$.05 \leq p < .10$	Limited (or “some”) evidence against chance
$.10 \leq p < .20$	Weak evidence against chance
*Classification system is conditional upon a sample size between 25 and 250	

In defense of a stepwise classification system. Some statistically inclined readers will surely object to interpreting p -values in a continuous manner, as such an approach stands on debatable mathematical grounds. However, mathematics are not the primary problem with p -values: researchers’

overly simplistic interpretations *are* the problem. As many have observed, the still-too-common practice of dismissing a study if the p -value does not meet the .05 criterion can lead to problematic research practices and erroneous interpretations of research findings (e.g. Meehl, 1990; Cohen, 1994; Ziliak & McCloskey, 2008; Simmons, Nelson, & Simonsohn, 2011; Nosek, Spies, & Motyl, 2012). Consequently, whenever a researcher obtains a p -value that falls just short of the “bright line” .05 criterion, he or she is motivated to use so-called “ p -hacking” techniques (see Gelman & Loken, 2013; Head et al., 2015), also called “researcher degrees of freedom” (Simmons, Nelson, & Simonsohn, 2011, p. 1359). A common example of this is to continue data collection until the p -value falls below .05, and then to stop data collection immediately. Since such practices contribute to the epidemic of false-positives (Ioannidis, 2005) that ultimately undermines confidence in social science research (Open Science Collaboration, 2015), researchers need to cease engaging in them. But researchers often perceive that the incentive structure (e.g. getting hired, achieving tenure) is structured in favor of those with the most publications (Nosek, Spies, & Motyl, 2012; also see Kunda, 1990).

A stepwise classification procedure like the one listed above serves to de-emphasize the importance of the p -value, thereby removing a major motivation for researchers to engage in “ p -hacking” techniques. By shifting the focus away from dichotomous yes/no decisions as the ultimate outcome of social science research, the system proposed here will therefore contribute to better research practices. Although it is admittedly quite subjective, critics should consider whether a strict cutoff of .05 is any better. Though flawed, primarily in that the p -value for a sample size of 25 provides a very different strength of evidence than the same value for a sample size of 250, this proposed system maintains three specific advantages: 1) the sample size limit draws attention to the outsized impact of sample size on a p -value, 2) it reflects the shades of gray that exist in the researcher’s decision space (degree of evidence, rather than effect/no effect), and 3) combined with the following recommendations, a stepwise classification system for p -values can reduce the motivation for social

science researchers to “*p*-hack”—even *without* forcing them to abandon a procedure that has, despite its flaws, been a part of statistical training in the field for decades!

Every psychologist should realize that people are rarely converted to better behavior by leveling accusations that a previous behavior was wrong, and then advocating the immediate adoption of a completely new behavior instead. So, although this proposal may be questionable on the mathematical and logical grounds of interpreting a [flawed] frequentist procedure in a Bayesian manner, it has the distinct advantage of accounting for the aspects of our nature that encourage us to stick with what has worked in the past, even if a better way is apparent! It is to this part of our nature that reform efforts must appeal in order to win the hearts and minds of our colleagues who remain entrenched in the use of NHST.

Moving away from the magic cutoff. If the present proposal to change the interpretation of *p*-values (coupled with the subsequent recommendations to include further statistical procedures) is adopted widely and immediately, I suspect that social scientists a few decades hence may find themselves rarely using NHST procedures. If this suspicion is correct, then a gentle move away from dichotomous *p*-value-based inferences may accomplish what 50 years worth of illustrious reformers’ repeated calls to abolish NHST could not: the end of error-prone ‘automatic’ scientific inference!

Toward meta-analytic thinking

From the magic cutoff to the magic toolbox. But the above proposal is incomplete, without also shifting attention toward Cumming and Finch’s meta-analytic thinking. When considering the results described in the Medication X example, a good scientist would probably wonder about the specifics, such as how big the mean difference actually is, or how reliable the effect is. Well, I would like to know those factors as well! Sadly, there is no single statistical procedure that can give us this information. Gerd Gigerenzer has frequently (Gigerenzer, 2004; Gigerenzer, Krauss, & Vitouch, 2004; Gigerenzer & Marewski, 2015) described the variety of statistical procedures as comprising a toolbox,

from which we can draw an appropriate method to help us answer questions about our work. This analogy helps researchers understand that no one-size-fits-all statistic exists (Cohen, 1994).

More on Medication X. So, how big *is* the difference between the group who took Medication X and the waitlist group? The mean number of words recalled by people in the waitlist group was 7.733; the mean number of words recalled by people in the experimental group (those who took Medication X) was 9.7. Okay, now we know that the mean difference was almost 2, *and* that the experimental group had a higher mean number of words recalled. We didn't know either fact from looking only at the t value and whether or not the p -value was smaller than the .05 cutoff.

But knowing that the mean difference is 2 isn't as helpful as we might like. Did everyone in the waitlist group remember less than 9 words, or did they recall anywhere from 3 to 30 words? This information would certainly affect the interpretation of this result! So, reporting the SD helps to put that mean difference into perspective. The SD for the waitlist group is 3.258; the SD for the medication group is 3.131. Reporting the means and the SDs alike can help the reader make sense of the results; running a t test and obtaining an exact two-tailed p -value of .02 can help us to make a stronger case that random chance does not constitute a good explanation for the results.

So far, we have evidence that the null of “no difference” is most likely false. The specifics show that for a sample of $n = 60$, *if* the null hypothesis that Medication X does not affect memory is true, there is only a 2% chance that the observed difference occurred due to sampling error. This is better than knowing only that $t = 2.384$ and $p < .05$, but it is still an incomplete presentation of the data.

Confidence intervals. Many are familiar with the use of confidence intervals. A confidence interval is closely related to the p -value, in that a 95% CI that includes zero is not significant at the customary alpha level of .05; a 95% CI that excludes zero *is* significant at the .05 level.

However, a 95% CI still makes salient a feature of the data that a p -value does not: the CI indicates that, based on the data obtained in the sample, the mean difference would be expected to fall

between the lower and upper bounds in 95 out of 100 replications. This underscores the uncertainty inherent in sampling, and in research more generally (Cumming & Finch, 2001). Knowing that the 95% CI of the difference in this example falls between 0.315 and 3.618 gives us a further indication of the precision of the effect, as well as a reasonable estimate of the actual population value.

Effect sizes. A major weakness of NHST procedures is that significance is contingent upon sample size. Cohen, a legend in the matter of effect size, wrote frequently and persuasively about the need for effect size tests (e.g. Cohen, 1994) and guidelines for interpreting those effect sizes (e.g. Cohen, 1988). Ellis (2009) wrote a brief but extremely useful guide to the necessity of using effect sizes, as well as a good overview of how to interpret common effect-size measures like Cohen's d . Jussim, Crawford, Anglin, Stevens, and Duarte (2016) argue that including effect size measures can ameliorate zealous researchers' temptation to make overly exuberant claims. Ziliak and McCloskey (2008; 2009; 2013) also make an impassioned and compelling plea to think about effect size and practical impact, protesting that p -values measure precision rather than what they call "oomph."

Different measures of effect size are appropriate for different kinds of tests. Cohen's d or Glass' Δ are appropriate measures for a one-sample or independent-samples t test (note that these are inappropriate for repeated-measures tests!). For an ANOVA, eta-squared (or partial eta-squared) is a common effect size measure, though some instead prefer omega-squared. For a chi-square test for independence or chi-square goodness-of-fit test, the phi-coefficient or Cramér's V are proper effect size measures. For correlations, Pearson's r is itself a measure of effect size, as is the coefficient of determination, R^2 .

I will use Cohen's d for the Medication X example. Cohen's d is 0.615 for this dataset, a medium effect by Cohen's interpretation guidelines. This gives us further evidence that the effect is real, nontrivial, and replicable. Since Cohen's d is also immune to sample size concerns, it is difficult to game this test—making it a good check on any conclusions drawn as a result of an NHST procedure.

However, even effect size isn't a perfect measure of practical impact. Ziliak and McCloskey (2009) use a clear analogy to explain why practical impact is usually the most important consideration: imagine that your mother wants to lose as much weight as possible, so she asks you to recommend a diet pill. After some research, you find that dieters who use a pill called "Oomph" lose an average of 20 pounds, plus or minus 14, and that dieters who use a pill called "Precision" lose an average of 5 pounds, plus or minus half a pound. A file with sample data illustrating that "Precision" yields a larger Cohen's d and a larger value of t (and, consequently, a smaller p -value) than "Oomph" is available at <https://osf.io/369a2/>. In this analogy, Ziliak and McCloskey's "sizeless scientist" would determine that the pill called Precision is superior to the pill called Oomph, *even though Oomph is virtually certain to result in more weight loss!* Ziliak and McCloskey blast such a conclusion: "Mom wants to lose weight, not gain precision" (Ziliak & McCloskey, 2009, p. 2304).

Assumptions. I will also mention at this point that a researcher should check the assumptions underpinning the t test. Statisticians will surely chide me for waiting until now, but as the self-serving proverb goes: better late than never! Levene's test for equality of variances yields a non-significant result: $F(1) = 0.004, p = .951$. This shows that the assumption of equality of variance is upheld. A Shapiro-Wilk test of normality shows a W of 0.910 and p of .015 for the waitlist group, and a W of 0.967 with associated p of .461 for the experimental group. This demonstrates a violation of the assumption of normality for the waitlist group, but not for the group that took Medication X.

If this were a real study, instead of a demonstration using made-up data, that violation of the normality assumption might be a bit of a concern, but research has shown that the t test tends to be robust against such issues (Lumley, Diehr, Emerson, & Chen, 2002; Sullivan & D'Agostino, 1992), even with an n of 5 or smaller (de Winter, 2013)! One test to address concerns about the violation of normality is the Mann-Whitney U test, which does not assume normality of the distributions for the sample populations. This test still shows a significant difference between the two groups ($U = 280.5, p$

= .012, two-tailed), so despite the violation of the assumption of normality, our t -based interpretation holds: we can still rule out chance as a good explanation for the effect.

Figures: It is also desirable to depict the result in the form of a graph or table, in order to quickly and clearly display the essential, basic information about the data. Gigerenzer, Krauss, & Vitouch (2004) argue that it is often sufficient to include such information graphically, as an appropriate display can communicate the essential information far more quickly and efficiently than a cascade of words and numbers ever could. In that spirit, then, I offer the following graph of the results explained here:

Figure 1. Words recalled by experimental condition.

Bayesian tests: A final step is to conduct a Bayesian independent-samples t test (for more detail about this procedure, see Rouder, Speckman, Sun, Morey, & Iverson, 2009). This is preferable to Student's t test because it evaluates the relative likelihoods of two different hypotheses, as explained above. The outcome of a Bayesian test—that one hypothesis is X times more likely than the other hypothesis, given the observed data—is also clear and intuitively simple to understand. The non-directional BF_{10} value is 2.697, which indicates that the observed data were 2.7 times more likely to occur if taking Medication X affects memory than if taking Medication X does not impact memory. [If

I were predicting that Medication X would *increase* the number of words recalled, I'd run a one-tailed test, and the corresponding BF_{10} would be stronger at 5.309].

Interpretation. Now, what can we take from all this information? Medication X *appears* to be better for your memory than doing nothing at all. However, the experimental design still leaves the results subject to objections about the possibility of placebo effects, experimenter-behavior effects, the limitations of between-subjects rather than within-subjects designs for analyzing questions like this, etc. So we have some limited evidence that Medication X can improve your memory, but the quality of that evidence is, frankly, not very high. Based on the procedure described above, I would remain highly skeptical of the efficacy of Medication X to improve memory.

Statistics are no substitute for science

To ascertain whether Medication X is truly effective at improving memory, we need more—and better-controlled—experiments (see e.g. Lykken, 1968). We need to implement a double-blind procedure, we should test the same individuals before and after taking Medication X, and we should ideally test Medication X against a variety of alternatives, including mnemonic strategies that have been shown to improve recall (like acrostics and the method of loci) and other medications and herbs that are claimed to improve recall. And if we truly want to see if Medication X affects cognitive function on the whole, we need to use more tests than just testing participants' ability to memorize a list of words.

If these better-controlled tests show that Medication X is at least as effective as other memory-improving medications, *and* is at least as effective as mnemonic strategies, *then* we can conclude that Medication X is an effective memory-enhancing agent. To fully participate in Cumming & Finch's meta-analytic thinking, we need to consider the results of additional experiments that test questions such as these. Without such higher-quality evidence, our conclusions about Medication X should be *extremely* tentative! Statistics are no substitute for scientific thought; advanced statistical analyses of a

poorly-designed experiment will still give you little useful information about the subject. What is true of computer programs is true of any mathematical procedure: garbage in, garbage out.

Hawk-eyed readers may object to the suggestion that they should engage in the regular reporting of this wide variety of statistical procedures, since these procedures sport a good deal of overlap (Anderson, submitted manuscript). That is, in the absence of extremely large sample sizes, researchers are unlikely to find a p -value of .001 with a trivial Cohen's d of .07. However, such objectors are likely overlooking the fact that the weakness is less a matter of the statistical procedure itself, and more so in how people (mis-)use them to support a conclusion. Each test advocated here maintains a unique advantage over the other procedures; to review:

- The p -value associated with NHST procedures provides information that helps us rule out sampling error/chance as an explanation for any differences in recall between groups. This evidence is not conclusive in favor of a particular alternative hypothesis, but it can still be useful for a preliminary study of moderate sample size.
- Reporting means and SDs provides a straightforward way to communicate the central tendency and variability observed in the data, and should help readers to better understand what the data are showing about the effect under investigation.
- Confidence intervals give an estimate of precision and of the probable population value.
- Effect sizes are immune to artificial inflation due to excessive sample size, and thereby give a more honest assessment of the actual real-world influence of the variable(s) being studied.
- A good graph or table makes the important summary information more visually salient.
- And, finally, the Bayes factor provides an easy-to-understand outcome of a hypothesis test.
- A good scientist will design a strong experiment, then will consider *all* of the above analyses when drawing a conclusion about the study's results.

Objectors should also consider that some experimental articles have already implemented many of the suggestions herein: in a paper investigating decision making via experimental methodology, Hilbig, Michalkiewicz, Castela, Pohl, and Erdfelder (2015) made use of all of the statistical information recommended in the present manuscript, with the exception of 95% confidence intervals and graphs. For each of their major analyses, Hilbig et al. presented the NHST results [e.g. t value], exact p -value, effect size measure, and BF_{10} . In testing recognition as a decision strategy, Basehore and Anderson (2016) presented the NHST results [t or F value], exact p -value, effect size measure, and 95% CI, but did not include the BF_{10} or a graph of the results (but for the graphs, see <https://osf.io/9tsxt/>).

Though some readers may be developing a headache at the thought of running all of these different analyses every time they test something, it's not as difficult it may seem: I conducted all the above analyses (except for the Mann-Whitney U test) in under 2 minutes, using a single statistical program! After uploading a .csv file with the made-up data used for this manuscript (this file is freely available for download in .csv format at <https://osf.io/369a2/>), the statistical program JASP will: run the independent-samples t test, test equality of variances, test normality, show the descriptive statistics, give the 95% confidence interval, and display the mean difference—all in only a few clicks of the mouse! In just a few additional clicks, JASP also displays an appropriate effect size measure, generates a plot of the means with 95% CI bars, and runs a Bayesian independent-samples t test.

While JASP (version 0.8.0.1 as of January 2017) is not a fully-fleshed-out program yet, many users will find it very helpful, especially for fulfilling most of the directives herein. And, since it is a relatively lightweight program that is available free of cost for Windows, MacOS, and Linux (see <http://jasp-stats.org>), there is little reason not to use JASP to communicate information such as that recommended above.

The bottom line

A results section that adheres to the recommendations I've made here can be done primarily through the use of JASP. Using the recommendations listed above, a write-up for the Medication X experiment used throughout this manuscript would look something like this:

“...An independent-samples t test showed that people who took Medication X recalled significantly more words in a free-recall test ($M = 9.7$, $SD = 3.131$) than members of a waitlist group who took no medication ($M = 7.733$, $SD = 3.258$), $t(58) = 2.384$, $p = .02$ (two-tailed). This yields strong evidence that chance is not a good explanation for the observed difference. Though a Shapiro-Wilk test showed that the waitlist group violated the assumption of normality, a Mann-Whitney U test still revealed strong evidence against chance ($U = 280.5$, $p = .012$). Cohen's d showed a medium effect of 0.615, and a Bayesian independent-samples t test showed some limited support for the hypothesis that Medication X affects memory ($BF_{10} = 2.697$). Collectively, these tests provide evidence that taking Medication X daily for six weeks does yield a reliable improvement on a recall test.

“As seen in Figure 1, participants who adhered to the Medication X regimen recalled about 2 more words on average than people who took no medication ($M_{\text{diff}} = 1.967$, 95% CI [0.315, 3.618]). However, this only compares people who took Medication X to people who took no medication at all. Further experiments were therefore necessary to rule out confounds such as placebo effects...”

Remember that more (and better!) experiments are necessary before accepting this as a legitimate effect. In the absence of sound experimental design, all the statistical procedures in the world mean nothing! But following the current recommendations, in combination with multiple well-designed studies, will contribute to more trustworthy and replicable results.

Implications

Following the recommendations advanced here will require a culture change among journal editors, since a number of journals still publish single experiments and exploratory, huge- n studies as being more authoritative than the evidence warrants. It will also require a culture change among

researchers: let's not give journal editors any shoddy, single-study, *p*-value-based articles from which to choose, even if an extra publication would beef up our CVs! Finally, it will require a culture change among hiring committees, since doing science properly will necessarily result in fewer publications.

But we must remember that quantity does not equal quality! If we want to do science properly, all of us will have to learn how to interpret a wider variety of statistical procedures. We will also have to become comfortable with the practice of conducting multiple studies to test our results, meaning that each of us will therefore generate fewer publications. But these changes are ultimately a small price to pay for improving scientific practice in psychology!

Postscript:

Earlier, I mentioned the straw man argument as a logical fallacy. *Ad hominem* attacks are also a logical fallacy, but they frequently remain convincing. Readers may suspect that my advocacy of statistical reform stems in part from personal frustration at having had manuscripts rejected. But this suspicion, though understandable, would be incorrect: I have personally submitted one academic manuscript to the Society for Judgment and Decision Making's interdisciplinary publication, *Judgment and Decision Making*; it was published in May 2016. So the present work has nothing to do with personal frustration, as I have thus far been quite fortunate! Rather, my criticism stems from my experiences instructing an undergraduate statistics course, and also from my conviction that judicious recommendations can improve statistical practices immediately—and *without* having to upset the figurative apple cart by completely overturning current methods!

There are more resources on these matters than I could hope to acknowledge here. So, aside from the various works cited above, I would like to express my gratitude to Andrew Gelman for his supremely helpful statistics blog at <http://www.andrewgelman.com>, which (in combination with the various knowledgeable and inquisitive commenters) provided a wealth of information on the subject of improving statistical practice.

I would also like to thank Richard B. Anderson for being a helpful guide through the morass of statistical practice in psychology, including providing helpful comments on a previous draft of this manuscript. I would also like to thank Dr. Anderson's Complex Cognition Lab at Bowling Green State University for agreeing to host a presentation of these ideas, and Christopher B. Arnold and Brent A. Lang for being sounding boards for some of the ideas presented herein. Unbeknownst to them, Jonathan Baron, Daniel Goldstein, and two anonymous reviewers of Basehore and Anderson (2016) provided helpful feedback that planted the seed that ultimately led to this manuscript—I am also grateful to them for planting that seed!

I would also like to acknowledge the work of the R Foundation for Statistical Computing for providing the powerful free and open-source, cross-platform [program R](#), the Free Software Foundation for creating the free-as-in-freedom, cross-platform, SPSS-compatible [statistical program PSPP](#), and to E. J. Wagenmakers and the team that made [JASP](#) available as a free and open-source, cross-platform, user-friendly statistics program.

My recommendations in this article do not necessarily reflect the views of any of the above individuals or organizations.

References

- Anderson, R. B. (2017). *The p-value as a rival to the Bayes factor*. Manuscript in preparation.
- Bakan, D. (1966). The test of significance in psychological research. *Psychological Bulletin*, 66, 423-437.
- Baschore, Z. & Anderson, R. B. (2016). The simple life: New experimental tests of the recognition heuristic. *Judgment and Decision Making*, 11, 301-309.
- Carver, R. P. (1978). The case against statistical significance testing. *Harvard Educational Review*, 48, 378-399.
- Cohen, J. (1988). *Statistical power analysis for the behavioral sciences* (2nd ed.). Hillsdale, NJ: Erlbaum.
- Cohen, J. (1994). The earth is round ($p < .05$). *American Psychologist*, 49, 997-1003.
- Cowles, M. & Davis, C. (1982). On the origins of the .05 level of statistical significance. *American Psychologist*, 37, 553-558.
- Cumming, G. (2014). The new statistics: Why and how. *Psychological Science*, 25, 7-29. doi: 10.1177/0956797613504966
- Cumming, G. & Finch, S. (2001). A primer on the understanding, use, and calculation of confidence intervals that are based on central and noncentral distributions. *Educational and Psychological Measurement*, 61, 532-574.
- de Winter, J. C. F. (2013). Using the Student's *t*-test with extremely small sample sizes. *Practical Assessment, Research & Evaluation*, 18(10), 1-12. Available online: <http://pareonline.net/getvn.asp?v=18&n=10>.
- Ellis, P. D. (2009). Thresholds for interpreting effect sizes. Hong Kong Polytechnic University. Accessed February 2015, at http://www.polyu.edu.hk/mm/effectsizefaqs/thresholds_for_interpreting_effect_sizes2.html
- Etz, A. & Vandekerckhove, J. (2016). A Bayesian perspective on the reproducibility project: psychology. *PLoS One*, 11, 1-12. doi: 10.1371/journal.pone.0149794.
- Gelman, A. & Loken, E. (2013, unpublished manuscript). The garden of forking paths: Why multiple comparisons can be a problem, even when there is no “fishing expedition” or “p-hacking” and the research hypothesis was posited ahead of time. Downloaded January 2017 from http://www.stat.columbia.edu/~gelman/research/unpublished/p_hacking.pdf
- Gigerenzer, G. (2004). Mindless statistics. *The Journal of Socio-Economics*, 33, 587-606. doi: 10.1016/j.socec.2004.09.033
- Gigerenzer, G., Krauss, S., & Vitouch, O. (2004). The null ritual: What you always wanted to know about significance testing but were afraid to ask. In D. Kaplan (ed.), *The Sage Handbook of Quantitative Methodology for the Social Sciences* (p. 391-408), accessed December 2016 at http://www.polyu.edu.hk/mm/effectsizefaqs/thresholds_for_interpreting_effect_sizes2.html
- Gigerenzer, G. & Marewski, J. (2015). Surrogate science: The idol of a universal method for scientific inference. *Journal of Management*, 41, 421-440. doi: 10.1177/0149206314547522
- Gilbert, D. T., King, G., Pettigrew, S., & Wilson, T. D. (2016). Comment on “Estimating the reproducibility of psychological science.” *Science*, 351, 1037. doi: 10.1126/science.aad7243
- Hayek, F. A. (1942). Scientism and the study of society (Part I). *Economica*, 9, 267-291. Downloaded January 2017 from <http://www.jstor.org/stable/2549540>

- Head, M. L., Holman, L., Lanfear, R., Kahn, A. T., & Jennions, M. D. (2015). The extent and consequences of p-hacking in science. *PLoS Biology*, 13(3), 1-15. doi: 10.1371/journal.pbio.1002106
- Hilbig, B. E., Michalkiewicz, M., Castela, M., Pohl, R. F., & Erdfelder, E. (2015). Whatever the cost? Information integration in memory-based inferences depends on cognitive effort. *Memory & Cognition*, 43, 659-671. doi: 10.3758/s13421-014-0493-z
- Hubbard, R. & Armstrong, J. S. (2005). Why we don't really know what "statistical significance" means: A major educational failure. Draft. Downloaded December 2016 from <http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.399.1349&rep=rep1&type=pdf>
- Ioannidis, J. P. A. (2005). Why most published research findings are false. *PloS Medicine*, 2, 0696-0701. doi: 10.1371/journal.pmed.0020124
- Jeffreys, H. (1961). *Theory of Probability*, 3rd ed. Oxford, UK: Oxford University Press.
- Jussim, L., Crawford, J. T., Anglin, S. M., Stevens, S. T., & Duarte, J. L. (2016). Interpretations and methods: Towards a more effectively self-correcting social psychology. *Journal of Experimental Social Psychology*, 66, 116-133. doi: 10.1016/j.jesp.2015.10.003
- Kruschke, J. K. (2013). Bayesian estimation supersedes the *t* test. *Journal of Experimental Psychology: General*, 142, 573-603. doi: 10.1037/a0029146
- Kruschke, J. K. & Liddell, T. M. (in press). The Bayesian new statistics: Hypothesis testing, estimation, meta-analysis, and power analysis from a Bayesian perspective (Preprint version of November 15, 2016). *Psychonomic Bulletin & Review*. doi: 10.2139/ssrn.2606016. Downloaded January 2017 from <https://osf.io/3eewv/>
- Kunda, Z. (1990). The case for motivated reasoning. *Psychological Bulletin*, 108, 480-498.
- Levine, T. R., Weber, R., Hullett, C., Park, H. S., & Massi Lindsey, L. L. (2008). A critical assessment of null hypothesis significance testing in quantitative communication research. *Human Communication Research*, 34, 171-187. doi: 10.1111/j.1468-2958.2008.00317.x
- Luce, R. D. (1988). The tools-to-theory hypothesis. Review of G. Gigerenzer & D. J. Murray, "Cognition as intuitive statistics." *Contemporary Psychology*, 33, 582-583.
- Lumley, T., Diehr, P., Emerson, S., & Chen, L. (2002). The importance of the normality assumption in large public data sets. *Annual Review of Public Health*, 23, 151-169. doi: 10.1146/annurev.publhealth.23.100901.140546
- Lykken, D. T. (1968). Statistical significance in psychological research. *Psychological Bulletin*, 70, 151-159.
- Meehl, P. E. (1990). Why summaries of research on psychological theories are often uninterpretable. *Psychological Reports*, 66, 195-244.
- Morgan, P. L. (2003). Null hypothesis significance testing: Philosophical and practical considerations of a statistical controversy. *Exceptionality*, 11, 209-221.
- Nosek, B. A., Spies, J. R., & Motyl, M. (2012). Scientific utopia: II. Restructuring incentives and practices to promote truth over publishability. *Perspectives on Psychological Science*, 7, 615-631. doi: 10.1177/1745691612459058
- Open Science Collaboration (2015). Estimating the reproducibility of psychological science. *Science*, 349, aac4716: 1-8. doi: 10.1126/science.aac4716
- Roediger, H. L. (2004). What happened to behaviorism. *APS Observer*. Accessed December 2016 from <http://psychologicalscience.org/observer/what-happened-to-behaviorism#.WGRgnve3PqA>

- Rosnow, R. L. & Rosenthal, R. (1989). Statistical procedures and the justification of knowledge in psychological science. *American Psychologist*, 44, 1276-1284.
- Rouder, J. N., Speckman, P. L., Sun, D., Morey, R. D., & Iverson, G. (2009). Bayesian t tests for accepting and rejecting the null hypothesis. *Psychonomic Bulletin & Review*, 16, 225–237.
- Rucci, A. J. & Tweney, R. D. (1980). Analysis of variance and the “second discipline” of scientific psychology: A historical account. *Psychological Bulletin*, 87, 166-184.
- Scheibehenne, B., Rieskamp, J., & Wagenmakers, E. J. (2013). Testing adaptive toolbox models: A Bayesian hierarchical approach. *Psychological Review*, 120, 39-64. doi: 10.1037/a0030777
- Simmons, J. P., Nelson, L. D., & Simonsohn, U. (2011). False-positive psychology: Undisclosed Flexibility in data collection and analysis allows presenting anything as significant. *Psychological Science*, 22, 1359-1366. doi: 10.1177/0956797611417632
- Sullivan, L. M., & D’Agostino, R. B. (1992). Robustness of the *t* test applied to data distorted from normality by floor effects. *Journal of Dental Research*, 71, 1938-1943.
- Stevens, S. S. (1960). The predicament in design and significance. *Contemporary Psychology*, 9, 273-276.
- Thompson, B. (2004). The “significance” crisis in psychology and education. *The Journal of Socio-Economics*, 33, 607-613. doi:10.1016/j.socec.2004.09.034
- Trafimow, D., & Marks, M. (2015). Editorial. *Basic and Applied Social Psychology*, 37, 1-2. doi: 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.
- Watson, J. B. (1913). Psychology as the behaviorist views it. *Psychological Review*, 20, 158-177.
- Wetzels, R. & Wagenmakers, E.-J. (2012). A default Bayesian hypothesis test for correlations and partial correlations. *Psychonomic Bulletin & Review*, 19, 1057-1064. doi: 10.3758/s13423-012-0295-x
- Wilkinson, L., & Task Force for Statistical Inference, APA Board of Scientific Affairs (1999). Statistical methods in psychology journals: Guidelines and explanations. *American Psychologist*, 54, 594-604.
- Ziliak, S. T. & McCloskey, D. N. (2008). *The Cult of Statistical Significance: How the Standard Error Costs Us Jobs, Justice, and Lives*. Ann Arbor: University of Michigan Press.
- Ziliak, S. T. & McCloskey, D. N. (2009). The cult of statistical significance. Presentation at the Joint Statistical Meetings, Section on Statistical Education, August 2009. Downloaded December 2016 from <http://www.deirdremccloskey.com/docs/jsm.pdf>
- Ziliak, S. T. & McCloskey, D. N. (2013). We agree that statistical significance proves essentially nothing: A rejoinder to Thomas Mayer. *Econ Journal Watch*, 10, 97-107.