

Psychology will be a Much Better Science When We Change the Way We Analyze Data

Geoffrey R. Loftus
University of Washington

In 1964, I entered the field of Psychology because I believed that within it dwelt some of the most fundamental and challenging problems of the extant sciences. Who could not be intrigued, for example, by the relation between consciousness and behavior, or the rules guiding interactions in social situations, or the processes that underlie development from infancy to maturity?

Today, in 1996, my fascination with these problems is undiminished. But I've developed a certain angst over the intervening thirty-something years—a constant, nagging feeling that our field spends a lot of time spinning its wheels without really making all that much progress. This problem shows up in obvious ways—for instance, in the regularity with which findings seem not to replicate. It also shows up in subtler ways—for instance, one doesn't often hear Psychologists saying, "Well this problem is solved now; let's move on to the next one" (as, for example, Johannes Kepler must have said over three centuries ago, after he had cracked the problem of describing planetary motion).

I've come to believe that at least part of this problem revolves around our *tools*—particularly the tools that we use in the critical domains of data analysis and data interpretation. What we do, I sometimes feel, is akin to trying to build a violin using a stone mallet and a chain-saw. The tool-to-task fit is not all that good, and as a result, we wind up building a lot of poor-quality violins.

My purpose here is to elaborate on these issues. In what follows, I will summarize our major data-analysis and data-interpretation tools, and describe what I believe to be amiss with them. I will then offer some suggestions for change.

The writing of this manuscript was supported by an NIMH grant to the author. I thank Nelson Cowan, David Irwin, Gerd Gigerenzer, Elizabeth Loftus, Michele Nathan, and especially Emanuel Donchin for their very insightful and helpful comments on early versions of this manuscript. Correspondence should be addressed to: gloftus@u.washington.edu.

The Universality of Null Hypothesis Significance Testing

The vast bulk of data analysis and data interpretation in the social and behavioral sciences is carried out using a set of techniques collectively known as *null hypothesis significance testing* (NHST). The logic of NHST goes as follows.

1. One begins with the hypothesis that some independent variable will have an effect on some dependent variable. At its most general level, the hypothesis can be expressed as:

$$\text{not } (\mu_1 = \mu_2 = \dots = \mu_J) \quad (\text{Eq. 1})$$

where $\mu_1 \dots \mu_J$ are the means¹ of J population distributions of the dependent variable that correspond to J levels of the independent variable. Equation 1 is generically referred to as an "alternative hypothesis," or H_1 .

2. An experiment is conducted in which J random samples of the dependent variable are obtained—one sample for each level of the independent variable. This yields the observed sample means $M_1 \dots M_J$, which are estimates of the population means $\mu_1 \dots \mu_J$. Generally it is not true that $M_1 = M_2 = \dots = M_J$; i.e., there will always be *some* differences among the sample means. What needs to be determined is whether the observed differences among the sample means are due only to errors in measurement, and/or to noise, or whether they are due, at least in part, to corresponding differences among the population means. Informally, the investigator needs to provide a convincing argument that the observed effect is "real."

3. To this end, the investigator endeavors to compute the probability (known as p)² of ob-

¹ Actually NHST can be applied to any population parameter. I use means here because means are by far the parameter of greatest concern in social science experiments.

² A brief note of clarification is in order here. The entity I have referred to as " p ", which is computed from the data, is Fisher's *exact level of significance* (Fisher,

serving differences among the M 's as great as those that were actually observed given that, in fact, the J population means are all equal, i.e., that

$$\mu_1 = \mu_2 = \dots = \mu_J \quad (\text{Eq. 2})$$

This hypothesis that the population means are equal is referred to as the "null hypothesis," or H_0 .

4. Based on a comparison of the computed value of p with a criterion value of p known as α (is usually set at 0.05), the investigator makes a binary decision. If p is less than α then the investigator makes a strong decision to "reject H_0 " in favor of H_1 (a decision which is usually phrased as, "the observed effect is statistically significant"). If p is greater than α , then the investigator makes a weak decision to "fail to reject H_0 ." Two types of errors can be committed in this setting. One commits a *Type-I error* when one incorrectly rejects a true null hypothesis. If the null hypothesis is true, then the probability of a Type-I error is, by definition, equal to α (typically .05). One commits a *Type-II error* when one incorrectly fails to reject a false null hypothesis. If the null hypothesis is false, a Type-II error is committed with probability β . In general, we do not know the value of β because we have no information, and no assumptions, regarding the values of the actual μ 's, given that the null hypothesis is false. *Statistical power* is defined to be $(1 - \beta)$, which is interpreted as the probability of correctly rejecting a false null hypothesis. Because β is not generally known, neither is power.

5. Finally, based on a series of such decisions—that is, rejecting, or failing to reject a series of null hypotheses—the investigator tries to

1925; 1935). The related entity known as " α ", described below, is the Neyman-Pearson *probability of at Type-I error*, which is decided upon before the data are collected. The Pearson and Neyman-Pearson approaches to statistics are quite different but, as splendidly described by Gigerenzer, Swijtink, Porter, Daston, Beatty, & Kruger (1989, pp. 78-109), they have been "hybridized" over the years and it is the resulting mish-mash that has been almost universally taught as "the statistical method" over the past half century. A detailed analysis of this issue is beyond the scope of this article, but Gigerenzer et al. make a convincing case that the confusion of the two approaches is responsible for much of what has gone astray with modern statistical practice. The description of NHST that I provide here—and which I will inveigh against—is, essentially, a description of this commonly used hybridized approach.

make sense of the data set, no matter how complex the data set might be.

Although variants of this procedure constitute the primary means of making conclusions from the vast majority of psychology experiments, I do not believe that it is a fruitful way of interpreting data or understanding psychological phenomena. On the contrary, I believe that reliance on NHST has channeled our field into a series of methodological cul-de-sacs, and it's been my observation over the years (particularly over my four years as editor of *Memory & Cognition*) that conclusions made entirely or even primarily based on NHST are at best severely limited, and at worst highly misleading. Below I will articulate the reasons for these beliefs.

I'm by no means the first to issue such charges. Periodically a book or an article will appear, decrying the enormous reliance we place on NHST³. Sadly, however, although such airings of the issues occasionally attract attention, they have not (up until now, anyway) impelled widespread action⁴. They have been carefully crafted and put forth for consideration, only to just kind of dissolve away in the vast acid bath of existing methodological orthodoxy.

I have two goals in this article: first to summarize some of the major problems with NHST, and second to suggest some alternative techniques for extracting more insight and understanding from a data set.

³ A sample of these writings is, in chronological order: Tyler (1931); Jones (1955); Nunnally (1960); Rozeblum (1960); Grant (1962); Bakan (1966); Meehl (1967); Lykken (1968); Carver (1978); Meehl (1978); Berger & Berry (1988); Hunter & Schmidt (1989); Gigerenzer, Swijtink, Porter, Daston, Beatty, & Kruger (1989); Rosnow & Rosenthal (1989); Cohen (1990); Meehl (1990); Loftus (1991); Carver (1994); Cohen (1994); Loftus & Masson (1994); Maltz (1994); Schmidt (1996); and Schmidt & Jones (1997).

⁴ This situation may finally be changing. Symposia at both the 1996 APS convention and the 1996 APA convention have aired the shortcomings of NHST as the primary data-analysis technique in the social sciences; see also Shea (1996). An APA task force has been set up to study the value of NHST. At least one journal editor has tried to discourage NHST (Loftus, 1993a, which provoked the caustic observation by Greenwald, Gonzalez, Harris, & Guthrie, 1994, that *every* recent empirical article in Loftus's journal, *Memory & Cognition*, has used NHST nonetheless).

Six Things to not Like about NHST

In this section, I will articulate six major problems with NHST. As I have indicated, they have been described before. But they bear repeating, and it is useful to consider them in concert.

(1) *The Usual Impossibility of a Typical Null Hypothesis*

NHST usually revolves around the testing of a null hypothesis that couldn't really be true to begin with. To illustrate, suppose one presented subjects with digits on a computer screen and were interested in effects of stimulus duration on the subsequent recall of the digits. One might run an experiment in which digit strings were shown for one of five exposure durations ranging from say 10 to 100 ms, measuring proportion of correctly recalled digits in each condition. In the usual hypothesis-testing framework, one would establish the following null hypothesis:

$$\mu_1 = \mu_2 = \mu_3 = \mu_4 = \mu_5 \quad (\text{Eq. 3})$$

where the μ 's refer to the population means of percent digits recalled for each of the five exposure durations. Note that the logic of NHST demands that the equal signs in Equation 3 mean *equal to an infinite number of decimal places*. If one weakens this requirement such that "=" means "pretty much equal," then one must add additional assumptions specifying what is meant by "pretty much." Although the mathematical machinery for doing this has been worked out (e.g., Greenwald, 1975; Hays, 1973; Serlin, 1993), it is rarely (if ever) implemented in practice.

This null hypothesis of identical population means cannot be literally correct. As Paul Meehl (1967) has pointed out, "Considering...that everything in the brain is connected with everything else, and that there exist several 'general state-variables' (such as arousal, attention, anxiety and the like) which are known to be at least *slightly* influenceable by practically any kind of stimulus input, it is highly unlikely that *any* psychologically discriminable situation which we apply to an experimental subject would exert literally *zero* effect on any aspect of performance." Alternatively, the μ 's can be viewed as measurable values on the real-number line. Any two of them being identical implies that their difference (also

a measurable value on the real-number line) is exactly zero—which has a probability of zero⁵.

Accordingly, differences in exposure duration *must* lead to performance differences even if such differences are small, and the relevant question is not really, "Are there any differences among the population means?" Rather, the relevant questions are: (1) How *big* are the differences, (2) are they big enough for the investigator to care about; and if so, (3) what pattern do they form? In short, testing the null hypothesis of Equation 2 can't provide new information. All it can do is to indicate whether there is enough statistical power to detect whatever differences among the population means must be there to begin with. As I have noted elsewhere (Loftus, 1995), rejecting a typical null hypothesis is like rejecting the proposition that the moon is made of green cheese. The appropriate response would be: "Well yes, ok...but so what?"

(2) *"Significance" vs the Underlying Pattern of Population Means*

A finding of "statistical significance" only constitutes evidence (and vague evidence at that, as we shall see) that a null hypothesis of the sort embodied in Equation 2 is false. Such a finding provides no information about form of the *underlying pattern of population means* which is presumably what's important for making scientific conclusions.

There are several ways of dealing with this problem. First, one can address it using post-hoc tests, whereby reject/fail-to-reject decisions are made about particular pairs of means. But post-hoc tests have problems. First, within the hypothesis-testing framework, the more such tests are carried out, the greater is the probability of committing at least one Type-I error. It is possible to adjust for this problem, but only at the expense of raising the probability of committing a Type-II error. Second, and more generally, post hoc tests focus on specific *pairs* of means,

⁵ One caveat is in order here. There are *some* experiments in which a null hypothesis could genuinely be true (see, e.g., Frick, 1995, for a discussion of this topic). A good example (attributable to Greenwald, et al., 1996) is a *qualitative* null hypothesis such as that a defendant in a murder case is actually the murderer. In such a case, the null hypothesis could certainly be true, and rejecting it (say based on DNA matches) would be a meaningful conclusion. However, these kinds of experiments are the exception rather than the rule in the social sciences.

which, when there are more than two conditions, is only an indirect way of assessing the entire *pattern* of means.

A second means of addressing the problem is via planned comparisons. Use of planned comparisons entails first generating a pattern of *weights*—one weight per experimental condition—that constitute the prediction of some experimental hypothesis about the overall pattern of population means. The correlation between the weights and the observed sample means then constitutes (essentially) a measure of how good the hypothesis is. Use of planned comparisons can be an informative and efficient process. The problem with planned comparisons, however, is that in practice, they are rarely used. I will have more to say about planned comparisons in a later section.

(3) Power

The third problem has to do with lack of attention to statistical power. NHST has come to revolve critically around the avoidance of Type-I errors. This occurs mainly because the probability of a Type-I error (α) can be computed. In contrast, the probability of a Type-II error (β) and concomitantly power ($1 - \beta$) usually cannot be computed. This is because computation of β and power requires a specific, quantitative hypothesis (e.g., $\mu_2 = \mu_1 + 10$ in a two-condition experiment), and such quantitative hypotheses are exceedingly rare in the social sciences.

Lack of power analysis is particularly troublesome when an investigator bases a conclusion on the proposition that some null hypothesis is true (rather than making the logically correct decision of "failing to reject the null hypothesis"). In such a case one cannot be sure whether there are in fact relatively small differences among the population means (a conclusion that can only be made when there is high power) or whether there may be large differences among population means that are undetected (a conclusion that is implied by low power). With high power, an investigator may be justified in accepting the null hypothesis "for all intents and purposes." The lower the power, the less acceptable is such a conclusion.

As noted, lack of power analysis often stems from the lack of quantifiable alternative hypotheses that characterizes the social sciences in general, and Psychology in particular. Nonetheless, there are ways of conveying the overall state of statistical power in some experiment (particularly through use of confidence intervals, as will be illustrated in an example below).

(4) The Artificial "Effects/Non-Effects" Dichotomy

A related problem is not, strictly speaking, a problem in the logic of NHST. Rather, it is a problem that arises because investigators, like all humans, are averse to making decisions that are both complicated and weak, such as, "we fail to conclude that the null hypothesis is false." Rather, people prefer simple strong decisions such as, "the null hypothesis is true." This fact of human nature fosters an artificial dichotomy that revolves around the arbitrary nature of the .05 criterion level.

Most people, if pressed, will agree that there is no essential difference between finding, say, that $p = .050$ and finding that $p = .051$. However, investigators, journal editors, reviewers, and scientific consumers often forget this and behave as if the .05 cutoff were somehow *real* rather than arbitrary. Accordingly, the world of perceived psychological reality tends to get divided into "real effects" ($p \leq .05$) and "non-effects" ($p > .05$). Statistical conclusions about such "real effects" and "non-effects" made in Results sections then somehow get sanctified and transmuted into conclusions that endure into Discussion sections and beyond, where they insidiously settle in and become part of our discipline's general knowledge structure⁶. The mischief thereby stirred up is incalculable. For instance, when one experiment shows a significant effect ($p \leq .05$), and an attempted replication shows no significant effect ($p > .05$), a "failure to replicate" is proclaimed. Feverish activity ensues, as method sections are scoured and new experiments run, in an effort to understand the circumstances under which the effect does or does not show up—and all because of an arbitrary cutoff at the .05 level. No wonder there is an epidemic of "conflicting" results in psychological research! This state of affairs is analogous to a chaotic phenomenon in which small initial differences lead to enormous differences in the eventual outcomes. In the case of data analysis, chaos is inimical to understanding, and it is more appropriate that similar results (e.g., $p = 0.050$ and $p = 0.051$) yield similar conclusions rather than entirely different conclusions.

(5) The Hypothesis-Testing Tail Wags the Theory-Construction Dog

Central to NHST is computation of p , the probability of a Type-I error. However, this

⁶ See Rosenthal & Gaito, 1963, for an interesting empirical demonstration of this assertion.

probability is difficult or impossible to compute unless routine, simplifying assumptions are made about the nature of the psychological processes under consideration. These assumptions then insinuate themselves into, and become integral to much of psychological theory. Some of the most common such assumptions are these:

1. The dependent variable is a linear combination of a collection of "effects"—effects due to the independent variables, to the interactions among the independent variables, and to various sources of "error."
2. The errors are distributed normally.
3. The variances of the error distributions are equal across various conditions.

Thus, the nature of the data-analysis technique generally dictates the nature of psychological theory which, in turn, engenders strong biases against formulating theories incorporating other perhaps more realistic and/or interesting assumptions. Accordingly, psychological theory becomes generic linear-model theory and a lot of potential for insight is lost. The situation might be summarized as, "Off the shelf, assumption produce off the shelf conclusions."

To illustrate this problem, suppose an experiment is designed to investigate whether the rate of forgetting depends on degree of original learning. In this experiment, word lists are taught to subjects whose recall is tested following intervals that vary from zero to five days. There are two groups of subjects. The *high-learning* subjects are allowed to learn the lists to a criterion level of 100% correct, while *low-learning* subjects learn the lists to a lower criterion level. The major data from this experiment are *forgetting curves* of the sort shown in Figure 1: here, memory performance is plotted as a function of retention interval.

The default data-analysis procedure in such an experiment would be to carry out a two-way analysis of variance (ANOVA). Let us suppose that the experimental power is sufficiently great that the ANOVA reveals statistically significant main effects of both learning level and retention interval, along with a significant interaction. It is evident in Figure 1 that the form of the interaction indicates a shallower "slope" for the low-learning than for high-learning condition (i.e., the vertical difference between the curves becomes smaller with increasing retention interval). The typical conclusion issuing from these observations would be that forgetting is slower following low learning than following high learn-

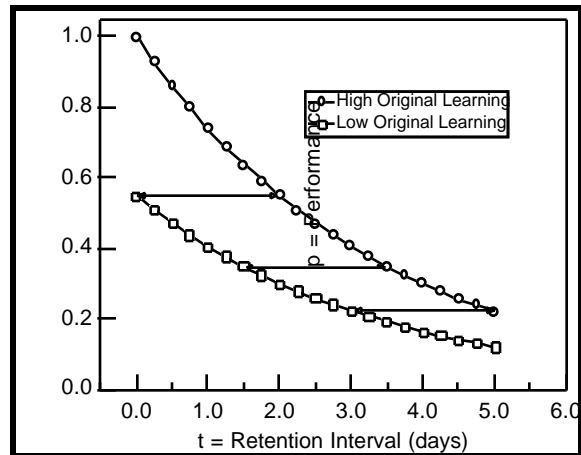


Fig. 1. Forgetting curves following high and low learning. Horizontal lines are meant to indicate that the two curves are horizontally parallel.

ing (this logic was used by Slamecka & McElree, 1983⁷).

This standard data-analysis procedure, along with the concomitant conclusion, would mask a very interesting regularity in the data: As is indicated by the horizontal lines on the figure, the *horizontal difference* between the high-learning and low-learning forgetting curves is *constant*. As shown by Loftus (1985a; 1985b; see also Loftus & Bamber, 1990) such horizontal equality is, under very general assumptions, a necessary and sufficient condition to infer *equal* high- and low-learning forgetting rates. Indeed, the Figure-1 curves were generated from exponential decay equations of the form:

$$\text{Low learning: } p = e^{-0.3t}$$

$$\text{High learning: } p = e^{-0.3(t+2)}$$

where p is performance and t is forgetting time (in days). Here, the equal forgetting rates are expressed by the same exponential decay parameter (0.3) in both equations.

⁷ Actually Slamecka & McElree found no significant interaction between degree of original learning and retention interval, and hence concluded that forgetting rate did not depend on degree of original learning. However, different data sets using the same general paradigm show different forms of interactions that would lead variously to the conclusion that higher-learning forgetting is faster, that higher-learning forgetting is slower and that high- and low-learning forgetting do not differ (a meta-finding that should, in and of itself, provoke suspicion that something is fundamentally amiss). In any event, it is the logic of the data-analysis technique, not the conclusion, that is primarily at issue here.

Condition	Exposure Duration (ms)							
	20	50	80	110	140	170	200	230
NVE/HU	0.287	0.503	0.843	1.005	1.468	1.664	2.102	2.257
VE/HU	0.461	1.192	1.399	2.360	3.008	3.236	3.908	4.649
NVE/LU	0.099	0.536	1.192	1.461	1.626	2.048	2.657	2.874
VE/LU	0.683	1.475	2.822	3.747	4.863	5.397	6.861	7.849

^aNotes: NVE: No Verbal Encoding.
VE: Verbal Encoding. HU: High Uncertainty
LU: Low Uncertainty

In this example, the principal finding is (or should be) that the horizontal difference between the two curves is constant. It is this finding that would imply no difference between high- and low-learning forgetting rates. However, investigators working within the hypothesis-testing framework would tend to miss this critical regularity, or would dismiss it, for at least two reasons. First, it is not immediately obvious how a standard significance test could be carried out that is relevant to the finding⁸, and without a significance test a finding is not generally deemed "valid." Second the logic of the linear model within which standard ANOVA is couched focuses on differences between the dependent variable at a fixed level of the independent variable (vertical differences), rather than differences between the independent variable at a fixed level of the dependent variable (horizontal differences).

(6) NHST Provides Only Imprecise Information about the Validity of the Null Hypothesis

The final problem which has been hammered at by Bayesian statisticians⁹ for decades, is this. By convention, one rejects some null hypothesis when

$$(1) p(\text{observed data} \mid \text{null hypothesis}) < 0.05,$$

while rejecting the null hypothesis implies the conclusion that

$$(2) p(\text{null hypothesis} \mid \text{observed data}) \text{ is small.}$$

(What else could be meant by the phrase "reject the null hypothesis"?) But without additional

⁸ Which isn't to say such a test couldn't be invented; it simply isn't part of general statistical knowledge or (what is probably more important) part of present statistics computer packages.

⁹ e.g., Berger & Berry (1988); Winkler (1993); see also Cohen (1994).

information, there is no logical basis for concluding the validity of (2) given the finding embodied in (1). Indeed, the probability of the null hypothesis given the data could be shown to be *anything* given suitable assumptions about the prior (pre-data) probability that the null hypothesis is true. Without specific assumptions about this prior probability, the exact probability of the null hypothesis given the observed data is unknown. In short, the common belief that the precise quantity, ".05", refers to anything meaningful or interesting is illusory.

Alternatives

I now suggest four (by no means mutually exclusive) alternatives to traditional NHST. My major goal in making each of these suggestions is simple and modest: it is to increase our ability to understand what a data set is trying to tell us. These techniques are not fancy or esoteric. They're just sensible.

(1) Plot Data Rather than Presenting Them as Tables-Plus-F-and-p-Values

In a previous article (Loftus, 1993b) I described some fictional data collected by a fictional psychologist named Jennifer Loeb. The story went as follows. Loeb was interested in memory for visual material and carried out a task in which visual stimuli were displayed and then recalled. There were three independent variables in Loeb's experiment, all varied at the time of stimulus presentation, which were: Stimulus exposure duration (eight values, ranging from 20 to 230 ms), verbal encoding (prohibited or required), and stimulus spatial uncertainty (high or low).

Based on a specific theory, Loeb had three predictions. First she predicted her performance measure to be a linear function of stimulus exposure duration. Second, she predicted the slope of this function to be higher with verbal encod-




Fig. 2. Loeb's data

ing. Third, she predicted the slope to be higher with lower spatial uncertainty.

Table 1, which shows a common way of presenting data like Loeb's, lists the values of her performance measure for all $8 \times 2 \times 2 = 32$ conditions. Accompanying NHST results—long compendia of F-ratios and p-values corresponding to the main ANOVA plus subsidiary tests—are generally provided as part of the text (often spanning many tedious pages).

Most people find such tabular-cum-text data presentation difficult to assimilate. That isn't surprising. Decades of cognitive research, plus millennia of common sense, teach us that the human mind isn't designed to integrate information that is presented in this form. There's too much of it, and it can't be processed in parallel¹⁰.

An alternative way of presenting Loeb's data is shown in Figure 2. Here, performance is plotted as a function of exposure duration for

¹⁰ One might argue that tables are useful when a reader needs exact values of the data points. However, such situations are quite rare, and when they do occur, exact values are obtainable from the investigators—a process that is particularly easy in these days of electronic communication.




Fig. 3. Loeb's data (plotted with 95% confidence intervals)

the verbal conditions (top panel) and the non-verbal conditions (bottom panel). The two curves within each panel are for the low-uncertainty and high-uncertainty conditions. Best fitting linear functions are drawn through the data points. With the data presented like this, one can acquire in a glance—or at most, a couple of glances—the same information that it would have taken practically forever to get out of Table 1. A picture really *is* worth 1000 words.

(2) Provide Confidence intervals

The Figure-2 plot indicates that Loeb's obtained pattern of sample means confirms her predictions pretty well. The curves are generally linear, verbal encoding yields a higher slope than no verbal encoding, and low uncertainty yields a higher slope than high uncertainty.

However, this plot provides no indication of the sort of *error variance* that is typically included as part of an ANOVA. Figure 3 remedies this deficiency: It shows the same data along with 95% confidence intervals placed around the sample means (computed, in this within-subjects design, as described by Loftus & Masson, 1994).

Figure 3 illustrates my second suggestion, which is to put confidence intervals around all sample statistics that are important for making

conclusions. Figure 3 provides most of the crucial information about Loeb's data. The observed pattern of sample means provides the best estimate of the underlying pattern of population means upon which conclusions should be based. The provision of confidence intervals allows the reader to assess the degree of statistical power—the smaller the confidence intervals, the greater the power. Power, in contrast to its profoundly convoluted interpretation within the hypothesis-testing framework, can here be simply interpreted as an indication of how seriously the observed pattern of sample means should be taken as a reflection of the underlying pattern of population means.

One could further analyze these data by, for instance, computing the *slope* for each of the four verbal encoding x uncertainty conditions for each subject. These four slopes could then be plotted along with *their* confidence intervals. One could go still further by, for instance, computing mean slope *differences* along with *their* confidence intervals. Such procedures correspond to graphic illustrations of various kinds of interactions. Creative use of such procedures allows one to jettison NHST entirely.

Confidence intervals as a guide to "accepting the null hypothesis"

The provision of confidence intervals is particularly useful when one wants to accept some null hypothesis “for all intents and purposes.” Another fictional data set¹¹ involves a clinician whom I will call “Christopher Sanders.” In this account, Sanders developed a clinical technique to decrease agoraphobia that is cheaper than the generally used “Standard technique.” Sanders ran a simple experiment to compare his technique with the Standard technique. In this experiment, forty agoraphobic individuals were randomly assigned to be administered either the Standard or the Sanders treatment. A year later, each individual's agoraphobia was assessed on a 10-point scale. Sanders' hope was that there would be no difference between the two treatments, in which case his treatment, being cheaper, would presumably be preferred to the Standard treatment.

Sanders got his hoped-for result, and reported it thusly: “The mean agoraphobia scores of the Standard and the Sanders groups were 5.05 and 5.03. The difference between the two groups was not statistically significant, $p > .05$.” Sanders went on to conclude that the cheaper

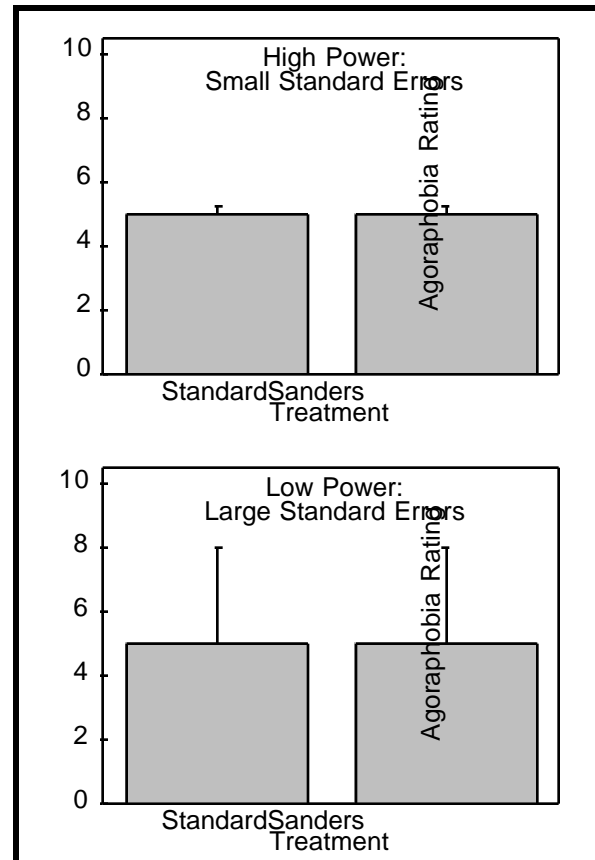


Fig. 4. Two possible outcomes of Sanders's experiment, illustrating high and low power.

Sanders technique was therefore the preferred one.

What is implied by the phrase “not significantly different” in Sanders's report? We can't tell, because Sanders provided no indication of statistical power, evaluation of which would be critical for justifying his accepting the null hypothesis of no treatment difference.

Sanders' report is consistent with many possibilities, two of which are presented in Figure 4. Given the top panel outcome, where small confidence intervals reflect high experimental power, Sanders' acceptance of the null hypothesis would be reasonable: Here it is readily apparent that the actual difference between the Sanders and Standard population means must be quite small. In contrast, given the bottom panel outcome, where large confidence intervals reflect low statistical power, Sanders' acceptance of the null hypothesis would be unconvincing: the actual population mean difference between the two treatments could vary widely.

There is a noteworthy epilogue to this story: When pressed, over drinks, at a professional conference, Sanders further defended his acceptance

¹¹ Also introduced by Loftus (1993b).

of the null hypothesis by pointing out that the actual t-value he obtained in his t-test was very small—"Practically zero!" he declared proudly. And so it was that Sanders committed the common error of equating smallness of the test statistic with permissibility of accepting the null hypothesis. Ironically, it can be easily shown that, given a particular mean difference, the smaller the t-value, the lower is power—and hence, the *less* appropriate it is to accept the null hypothesis.

Plus ça change, plus c'est la même chose.

The notion of using confidence intervals in this way—as a guide to accepting a null hypothesis "for all intents and purposes"—is not new. Thirty-four years ago the following suggestion appeared in the pages of the *Psychological Review*.

"In view of our long-term strategy of improving our theories, our statistical tactics can be greatly improved by shifting emphasis away from overall NHST in the direction of statistical estimation. For example [when testing a presumed null pre-treatment difference between two groups, an investigator] would do better to obtain a...confidence interval for the pre-treatment difference. If the interval is small and includes zero, [the investigator] is on fairly safe ground; but if the interval is large, even though it includes zero, it is immediately apparent that the situation is more serious. In both cases, H_0 would have been accepted" (Grant, 1962).

Confidence intervals and statistical significance

A section on confidence intervals would be incomplete without a discussion of the question: "How do you make a decision about whether some variable has an effect by just looking at confidence intervals?"¹²

The answer is: You can't¹³. This, I would assert, is an advantage rather than a disadvantage

of using plots plus confidence intervals rather than depending on NHST. I have argued above that a major difficulty with NHST is that it reduces data sets into a series of effect/no effect decisions and that this process, artificial as it is, leads the field astray in many ways. It imposes the illusion of certainty on a domain that is inherently ambiguous. Simply showing data, with confidence intervals, provides a superset of the quantitative information that is provided by a hypothesis-testing procedure, but it does not foster the false security embodied in concrete decision that is based on a foundation of sand.

(3) Meta-Analysis and Effect Size

An increasingly popular technique is that of meta-analysis (e.g., Schmidt, 1992, 1996; Rosenthal, 1995). Meta-analysis entails considering a large number of independent studies of some phenomenon (for example, gender differences in spatial ability) and (essentially) averaging the observed effects across studies to arrive at a overall effect. This technique is particularly useful when two conditions are being compared, but trickier when the question under investigation involves more than two conditions (see Richardson, 1996, for a discussion of this issue).

To illustrate, consider the question of gender differences in spatial ability. Suppose that, in fact, there is some difference in spatial ability between the population of males and the population of females. The direction and magnitude of this difference might be investigated in many (say 25) separate studies. Each individual study produces some observed gender difference (presumably the mean male spatial score minus the mean female spatial score) that constitutes that experiment's estimate of the population mean difference. Meta-analysis would entail averaging all 25 reported differences, thereby arriving at a single estimated difference that is (roughly speaking) five times as accurate as any of the individual estimates.

One problem is: What actual number should be extracted from each of the individual studies? One could simply use the raw difference, whatever it might be. The problem with this is that different studies presumably use somewhat different measures of spatial ability, different ways of carrying out the experiment, different populations of subjects, and so on; thus the raw measures wouldn't be comparable across the studies.

experiment the relation between confidence interval pattern and statistical significance is not immediately apparent.

¹² In numerous statistics classes and in other forums in which I have discussed this issue, a question that I can absolutely count on is: In a two-condition situation, what is the relation between error-bar overlap, and the reject/fail to reject decision?

¹³ At least not usually. In a two group design, nonoverlapping 95% confidence intervals implies that a two-tailed, $\alpha = .05$ t-test would lead to a conclusion of "statistically significant." However if the error bars do overlap, and/or if there are more than two conditions in the

The typical solution is to compute *effect size* which, in its simplest incarnation, is the mean difference observed in a given experiment divided by the obtained estimate from that experiment of the population standard deviation. It is these effect sizes that are then averaged to arrive at an overall estimate of the population effect size.

(4) *Planned Comparisons (Contrasts)*

My final suggestion, which I alluded to earlier, is to carry out planned comparisons on a data set. The use of planned comparisons is a technique that has been strongly advocated, and clearly described by many¹⁴. In my opinion, however, it is a technique that is surprisingly underutilized.

Briefly, carrying out a planned comparison involves the following steps.

1. First, one generates a *quantitative hypothesis* (even a relatively simple one will do) about the underlying pattern of population means corresponding to the conditions in some experiment. Suppose, for example, that one were investigating the relation between problem-solving time and alcohol consumption. An experiment is designed in which subjects are assigned to one of five groups that differ in terms of amount of alcohol consumed—0, 1, 2, 3, or 4 oz—and problem-solving time is measured for each group. A hypothesis to be tested is that problem-solving time increases linearly with number of ounces of consumed alcohol.

2. Next, hypothesis in hand, one generates *weights*—one weight per condition—that correspond to the hypothesized pattern of means. One constraint is that the weights must sum to zero; thus, in this example, the weights representing the linearity hypothesis would be: -2, -1, 0, 1, 2.

3. The experiment is carried out and the means—in this example, the five mean problem-solving times—are computed.

4. Finally, one (essentially) computes an over-conditions *correlation* between the weights and the sample means. The magnitude of the Pearson r^2 that emerges reflects the goodness of the hypothesis. Various other more sophisticated procedures can also be carried out, the nature of which are beyond the scope of this article.

¹⁴See, for example, Abelson (1995), Hays (1973), Loftus & Loftus (1988), and Rosenthal & Rosnow (1985).

Conclusions

I've tried to provide a variety of reasons why NHST, as typically utilized, is barren as a means of transiting from data to conclusions. I've tried to provide some examples of techniques—these are standard techniques, not bizarre or fancy ones—to replace standard NHST. These techniques are: First to *plot* the data rather putting them in tabular form, second to put *confidence intervals* around important sample statistics, third to use *meta-analysis*, and fourth to use *planned comparisons*. These techniques are all designed to assist in the ultimate goal of understanding what it is that some data set is trying to tell us.

This article is entitled “Psychology will be a much better science when we change the way we analyze data.” I hope that my arguments make it clear why I believe this to be true. I believe that in order for any science to progress satisfactorily, its primary data-analysis techniques must provide genuine insight into whatever phenomenon its practitioners set out to investigate. The primary data analysis technique of Psychology—NHST—does not, as I've tried to demonstrate, meet this criterion.

Acquisition of insight is often difficult in the social sciences, which are cursed with large numbers of uncontrollable variables, and hence error variance that has to be dealt with *somehow*. I believe that, historically, social scientists have embraced NHST procedures because they provide the appearance of objectivity. These procedures may indeed be objective in the sense that they provide rules for making scientific decisions. But objectivity is not, alas, sufficient for insight. I believe that these rules provide only the *illusion* of insight—which is worse than providing no insight at all.

References

- Abelson, R. (1995). *Statistics as principled argument*. Hillsdale, NJ: Erlbaum.
- Bakan, D. (1966). The test of significance in psychological research, *Psychological Bulletin*, 66, 423-437.
- Berger, J.O. & Berry, D.A. (1988). Statistical analysis and the illusion of objectivity, *American Scientist*, 76, 159-165.
- Carver, R.P. (1978). The case against statistical significance testing. *Harvard Educational Review*, 48, 378-399.
- Carver, R.P. (1993). The case against statistical significance testing, revisited. *Journal of Experimental Education*, 61, 287-292.

- Cohen, J. (1990). Things I have learned (so far). *American Psychologist*, 45, 1304-1312
- Cohen, J. (1994). The earth is round ($p < .05$). *American Psychologist*, 49, 997-1003.
- Fisher, R.A. (1925). *Statistical methods for research workers*. Edinburgh: Oliver and Boyd.
- Fisher, R.A. (1935). *The design of experiments*. Edinburgh: Oliver and Boyd.
- Freedman, D.A., Rothenberg, T. & Sutch, R. (1983). On Energy Policy Models. *Journal of Business and Economic Statistics*, 1, 24-36.
- Frick, R.W. (1995). Accepting the null hypothesis. *Memory & Cognition*, 23, 132-138.
- Gigerenzer, G., Swijtink, Z., Porter, T., Daston, L., Beatty, J., & Kruger, L.. (1989) *The Empire of chance: How probability changed science and everyday life*. Cambridge England: Cambridge University Press
- Grant, D.A. (1962). Testing the null hypothesis and the strategy and tactics of investigating theoretical models. *Psychological Review*, 69, 54-61.
- Greenwald, A.G. (1975). Consequences of prejudice against the null hypothesis. *Psychological Bulletin*, 83, 1-20.
- Greenwald, A.G., Gonzalez, R., Harris, R.J., & Guthrie, D. (1996) Effect Sizes and p-values: What should be reported and what should be replicated? *Psychophysiology*, 33, 175-183.
- Hays, W. (1973). *Statistics for the Social Sciences (second edition)*. New York: Holt.
- Hunter, J.E. & Schmidt, F. (1990). *Methods of meta-analysis: Correcting error and bias in research findings*. Newbury Park, CA: Sage.
- Jones, L.V. (1955). Statistics and research design. *Annual Review of Psychology*, 6, 405-430. Stanford: Annual Reviews, Inc.
- Loftus, G.R. (1985). Evaluating forgetting curves. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 11, 396-405 (a).
- Loftus, G.R. (1985). Consistency and confoundings: Reply to Slamecka. *Journal of Experimental Psychology: Learning, Memory and Cognition*, 11, 817-820 (b).
- Loftus, G.R. (1991). On the tyranny of hypothesis testing in the social sciences. *Contemporary Psychology*, 36, 102-105 (1991)
- Loftus, G.R. (1993). Editorial Comment. *Memory & Cognition*, 21, 1-3. (a)
- Loftus, G.R. (1993). Visual data representation and hypothesis testing in the microcomputer age. *Behavior Research Methods, Instrumentation, and Computers*, 25, 250-256. (b)
- Loftus, G.R. (1995). Data analysis as insight. *Behavior Research Methods, Instrumentation and Computers*, 27, 57-59.
- Loftus, G.R. & Bamber, D. (1990) Weak models, strong models, unidimensional models, and psychological time. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 16, 916-926.
- Loftus, G.R. and Masson, M.E.J. (1994) Using confidence intervals in within-subjects designs. *Psychonomic Bulletin & Review*, 1, 476-490.
- Lykken, D.T. (1968). Statistical significance in psychological research. *Psychological Bulletin*, 70, 131-139.
- Maltz, M.D. (1994). Deviating from the mean: The declining significance of significance. *Journal of research in crime and delinquency*, 31, 434-463.
- Meehl, P.E. (1967). Theory testing in psychology and physics: A methodological paradox. *Philosophy of Science*, 34, 103-115.
- Meehl, P.E. (1978). Theoretical risks and tabular asterisks: Sir Karl, Sir Ronald and the slow process of soft psychology. *Journal of Consulting and Clinical Psychology*, 46, 806-834.
- Meehl, P.E. (1990). Why summaries of research on psychological theories are often uninterpretable. *Psychological Reports, Monograph Supplement 1-V66*.
- Nunnally, J. (1960). The place of statistics in psychology, Educational and *Psychological Measurement*, 20, 641-650.
- Oakes, M.L. (1986). *Statistical inference: A commentary for the social and behavioral sciences*. New York: Wiley.
- Richardson, J.T.E. (1996). Measures of effect size. *Behavior Research Methods, Instrumentation, & Computers*, 28, 12-22.
- Rozeboom, W.W. (1960). The fallacy of the null-hypothesis significance test. *Psychological Bulletin*, 57, 416-428.
- Rosenthal, R. (1995). Writing meta-analytic reviews. *Psychological Bulletin*, 118, 183-192.
- Rosenthal, R. & Gaito, J. (1963). The interpretation of levels of significance by psychological researchers. *Journal of Psychology*, 55, 33-38.
- Rosenthal, R. & Rosnow, R.L. (1985). *Contrast analysis: Focused comparisons in the analysis of variance*. Cambridge, England: Schmidt, Frank (1996). Statistical significance testing and cumulative knowledge in Psychology: Implications for training of researchers. *Psychological Methods*, 1, 115-129.
- Rosnow, R.L. & Rosenthal, R. (1989). Statistical procedures and the justification of knowledge in psychological science. *American Psychologist*, 44, 1276-1284.
- Schmidt, Frank (1996). Statistical significance testing and cumulative knowledge in Psychology: Implications for training of researchers. *Psychological Methods*, 1, 115-129.
- Schmidt, Frank & Hunter, J. (1997). Eight false objections to the discontinuation of significance testing in the analysis of research data. In L. Harlow & S. Mulaik (Eds.), *What if there were no significance testing*. Hillsdale, NJ: Erlbaum.
- Serlin, R.C. (1993). Confidence intervals and the scientific method: A case for Holm on the range. *Journal of Experimental Education*, 61, 350-360.

- Shea, C. (1996). Psychologists debate accuracy of "significance test". *The Chronicle of Higher Education*, August, 16.
- Slamecka, N.J. & McElree, B. (1983). Normal forgetting of verbal lists as a function of their degree of learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 9, 384-397.
- Tyler, R.W. (1935). What is statistical significance? *Educational Research Bulletin*, 10, 115-118, 142.
- Winkler, R.L. (1993). Bayesian statistics: An overview. In G. Kerens, and C. Lewis (Eds.) *A Handbook for Data Analysis in the Behavioral Sciences: Statistical Issues*, Hillsdale NJ: Erlbaum.