Data analysis as insight: Reply to Morrison and Weaver

GEOFFREY R. LOFTUS

University of Washington, Seattle, Washington

I begin by underscoring a fundamental point of agreement between Morrison and Weaver and me: that data analysis should entail designing whatever set of techniques optimizes insight into whatever question the experiment was designed to address, rather than implementing some rote set of preordained rules and regulations. Against this shared philosophical backdrop, I then (1) reiterate problems with hypothesis testing, and (2) address some of the quite pertinent issues that Morrison and Weaver raise with respect to computation of confidence intervals.

Morrison and Weaver (1995) end their commentary by noting that “programs to plot results . . . are not in themselves an adequate substitute for thought” (p. 56). Although this assertion sounds like a truism, it is not. Rather, it is a legitimate response to one aspect of a trend about which numerous psychologists (including me) have become increasingly concerned over the years—that of data analysis as no-brainer. I believe, as I have asserted elsewhere, that this trend is largely rooted in an almost exclusive reliance on hypothesis-testing procedures which, as they are modally used, amount to reducing a rich and complex data set to an impoverished and uninformative string of significant/nonsignificant decisions.

Throughout Morrison and Weaver’s remarks is woven the theme that such gratuitous information reduction is unacceptable—that data analysis should be carried out and reported in such a way as to optimally exploit the information carried by a data set. For example, they note that if confidence intervals are reported, there is no need to make the standard homogeneity of variance assumption; one can individualize confidence intervals for individual conditions. Likewise, they note that if a critical conclusion depends on an over-conditions comparison of *variances*, then confidence intervals of the *variances* should be reported; that if infrequently reported statistics (like skewness) provide insight about the question being addressed, then such statistics should be computed and reported. And so on.

In short, there is a fundamental basis of agreement between Morrison and Weaver and me: data analysis should entail designing whatever set of techniques optimizes insight into whatever question the experiment was designed to address, rather than implementing some rote set of preordained rules and regulations (e.g., “do the ap-

propriate ANOVA and leave it at that”). To the degree that the general attitude embraced by Morrison and Weaver enters the methodological zeitgeist, they, and I, and many others will rejoice. Given our shared attitudes, our points of disagreement seem relatively unimportant.

With this foundation in mind, let me turn to some specific points of agreement and disagreement.

ON THE VALUE OF HYPOTHESIS TESTING

As I have emphasized, Morrison and Weaver and I espouse the same fundamental goal, which might be expressed as “creative and flexible data analysis as a route to optimal insight.” Morrison and Weaver assert that the plot-plus-error-bar (PPE) method favored by Loftus (1993c; see also Loftus, 1991, 1993a, 1993b) is burdened with various difficulties that ought to render it subsidiary to standard hypothesis-testing procedures. Below, I will address some of Morrison and Weaver’s complaints about the PPE method. First, however, I would like to reiterate two remarks I have made in the past about fundamental problems with hypothesis testing.

Implausible Null Hypotheses

A null hypothesis is a specific statement about relations among a set of population means (e.g., that all population means are identical to one another). In almost all cases, a null hypothesis (at least in the social sciences) can be assumed a priori to be false (e.g., no set of population means is identically equal to an infinite number of decimal places). Thus, a test of the null hypothesis, contrary to conventional wisdom, does not usually test the null hypothesis. What it usually *does* test is whether there is sufficient statistical power to detect whatever violation of the null hypothesis must exist.

As an example of this issue, consider Morrison and Weaver’s analysis of linear functions. Morrison and Weaver state that “we would happily [believe in the validity of fundamentally linear functions] if only a line or two of text were added stating, for example, that 95% of

The writing of this article was supported by an NIMH grant to the author. Correspondence concerning it may be sent to Geoffrey R. Loftus, Department of Psychology, University of Washington, Seattle, WA 98195 (e-mail: gloftus@u.washington.edu).

the exposure $\times$ uncertainty interaction could be accounted for by the linear component and that the t value for this component had a probability of $<.01$. In this case, one p value would be worth a thousand pictures” (p. 54).

There are two quite separate suggestions here. The first—that planned comparisons should be used as a basis for describing the percent variance accounted for by some effect—is, in my opinion, an excellent one; as I have stated elsewhere (Loftus, 1993b), I believe that planned comparisons are puzzlingly underutilized in the social sciences. However, what about reporting that $p < .01$? This small p value asserts that, given the null hypothesis of zero correlation between a set of linear weights and a corresponding set of population means, the observed degree of linearity has a probability of less than 1%. But the assumption of an exactly zero correlation is entirely implausible to begin with. So what’s the point of considering anything that depends on its validity? It’s like saying that the moon’s observed color is highly unlikely given that the moon is made of green cheese. True, but so what?

Miscast Objectivity

Morrison and Weaver (p. 54) consider assessment of two not-quite-parallel observed linear functions described by Loftus (1993c). They observe that the functions themselves in conjunction with the error bars around the sample means strongly suggest the underlying (population) functions to be likewise nonparallel. However, Morrison and Weaver describe a plausible scenario (of variances covarying with means) by which the underlying functions might indeed be parallel. They ask, “Might not two readers seeing the same data be led to conflicting conclusions?”

Well . . . yes. And, indeed, it is the lot of many data sets reported in the social sciences to be similarly ambiguous. Morrison and Weaver seem to imply, however, that a simple ANOVA would clear up this ambiguity, since the ANOVA would cleanly result in a “reject” or a “don’t reject” decision.

There are several interrelated problems here. The first, noted above, is that *rejecting* a null hypothesis (in this case, a null hypothesis of parallel curves) generally doesn’t tell you anything you didn’t know before, as the null hypothesis can almost always be assumed a priori to be false. However, what if the ANOVA results in a failure to reject the null hypothesis? As we have all been taught, this is essentially a “no-decision” outcome. By convention, we couldn’t reject the null hypothesis; nor, by logic, could we infer it to be true. The problem, however, is that many investigators would (either explicitly or implicitly) accept the null hypothesis anyway. Whether we like it or not, a statistical analysis technique has sociological consequences. A consequence of using hypothesis testing is that most humans—even humans who are sophisticated scientists schooled in statistical reasoning—can’t seem to accept a “no-decision” state,

succumbing instead to the irresistible temptation to divide effects into those that exist (when statistical significance has been found) and those that don’t (when it has not).

COMPUTING APPROPRIATE CONFIDENCE INTERVALS

Morrison and Weaver have performed a very useful service in underscoring a number of oversimplifications in the computation of confidence intervals as described by Loftus (1993c). In general, these problems have been addressed (although certainly not entirely solved) in an article that Morrison and Weaver acknowledge (Loftus & Masson, 1994). Here I briefly describe some of Loftus and Masson’s points that are relevant to Morrison and Weaver’s arguments.

Repeated Measures Designs

Loftus (1993c) restricted his remarks to entirely between-subject designs, wherein computation of confidence intervals was straightforward and embodied in the equation,

$$CI = M_j \pm \sqrt{\frac{MS_w}{n_j}} [\text{criterion } t(df_w)],$$

where M_j is the mean of Condition j , n_j is the number of observations in Condition j , MS_w is the (pooled) mean square within conditions, and df_w is degrees of freedom within.

Loftus and Masson argue that any kind of statistical analysis (e.g., ANOVA or PPE) is fundamentally used to elucidate the underlying *pattern* of population means. This reasoning allowed them to propose a confidence interval in an entirely within-subject design; here $MS_{S \times C}$, the mean square due to subject $\times$ condition interaction is substituted for MS_w , and, accordingly, the formula becomes

$$CI = M_j \pm \sqrt{\frac{MS_{S \times C}}{n}} [\text{criterion } t(df_{S \times C})],$$

where n is now the total number of subjects. Loftus and Masson show that this confidence interval, while not appropriate for inferring the value of any *single* population mean is, like its between-subject counterpart, appropriate for inferring how much faith one can put in the *pattern* of sample means as a reflection of the underlying pattern of population means.

Homogeneity of Variance

Morrison and Weaver assert that the homogeneity of variance assumption is rarely correct. Although they do not say so, this is an argument for abandoning the ANOVA; however, heterogeneity of variance poses no problem for computation of confidence intervals. In a between-subject design, the confidence interval around M_j becomes

$$CI_j = M_j \pm \sqrt{\frac{MS_{Wj}}{n_j}} [\text{criterion } t(n_j - 1)],$$

where MS_{Wj} is the mean-square error computed within Condition j only. Loftus and Masson demonstrate an analogous confidence interval for within-subject designs, which is

$$CI_j = \sqrt{\frac{\text{estimator}_j}{n}} [\text{criterion } t(n - 1)],$$

where n is the total number of subjects, and “estimator _{j} ” is

$$\left(\frac{J}{J - 1} \right) \left(MS'_{Wj} - \frac{MS_{S \times C}}{J} \right).$$

Here, J is the number of conditions and MS'_{Wj} is the variance within Condition j after overall subject variation has been removed (see Loftus & Masson, 1994, for details).

Multifactor Within-Subject Designs

Morrison and Weaver point out that in multifactor completely within-subject designs, there are numerous subject $\times$ effect error terms, and hence numerous possible confidence intervals. Morrison and Weaver acknowledge Loftus and Masson's point that if the error terms are sufficiently similar, they can be pooled—that is, the design can be treated as a one-way design (with $J \times K \times L \times \dots$ conditions) and a single confidence interval, based on $MS_{S \times ABC \dots}$, can be computed with little loss of validity. If the various error terms are quite different from one another, then computation of a single confidence interval would be, as Morrison and Weaver stress, a misleading procedure. Note, of course, that in such an event the error-term differences would likely themselves become a phenomenon of some interest. (Why, the investigator might ask, for example, would subjects be less consistent over levels of Factor 1 than over levels of Factor 2?)

Mixed Designs

A mixed design is more complicated as there would be no reason to suppose that the within- and between-subject error terms would be the same. Again, Loftus and Masson suggest several means of addressing this problem. First, if appropriate, one could compute whatever confidence interval is most germane to the critical conclusion. Second, one could compute multiple confidence intervals corresponding to the various effects.

Third, using some a priori model (e.g., a linear scanning model applied to a Sternberg short-term memory scanning task) one can sometimes reduce a mixed design (e.g., set size within subjects $\times$ stimulus type between subjects, with RT as the dependent variable) to a pure design (e.g., a between-subject design entailing different stimulus types and *slope* as the dependent variable).

ET ALIA

Finally, a couple of minor miscellaneous comments:

PPE Is Not Close to Hypothesis Testing

In considering the arbitrariness of choosing to use, say, a standard error versus a 95% confidence interval, Morrison and Weaver pose the question “is the PPE not coming dangerously close to hypothesis testing?” (p. 56). The answer is: “No.” Yes, there is a correspondence between the alpha level used in hypothesis testing and a confidence interval's probability level. But, for reasons I have alluded to above and elsewhere, hypothesis testing and confidence intervals entail different kinds of basic logic (the former testing the plausibility of a data set given some null hypothesis, and the latter directly describing a pattern of population parameters).

Mea Culpa

Morrison and Weaver are correct (p. 55). I made a computational error in describing the fictional Lowry data. The t value should have been reported as simply “ $t < 1.0$.”

REFERENCES

- LOFTUS, G. R. (1991). On the tyranny of hypothesis testing in the social sciences. *Contemporary Psychology*, **36**, 102-105.
- LOFTUS, G. R. (1993a). Editorial comment. *Memory & Cognition*, **21**, 1-3.
- LOFTUS, G. R. (1993b, November). *On the overreliance on hypothesis testing in the social sciences*. Presented at the 34th Annual Meeting of the Psychonomic Society, Washington, DC.
- LOFTUS, G. R. (1993c). A picture is worth a thousand p -values: On the irrelevance of hypothesis testing in the computer age. *Behavior Research Methods, Instrumentation, & Computers*, **25**, 250-256.
- LOFTUS, G. R., & MASSON, M. E. J. (1994). Using confidence intervals in within-subject designs. *Psychonomic Bulletin & Review*, **1**, 476-490.
- MORRISON, G. R., & WEAVER, B. (1995). Exactly how many p values is a picture worth? A commentary on Loftus's plot-plus-error-bar approach. *Behavior Research Methods, Instruments, & Computers*, **27**, 52-56.

(Manuscript received May 24, 1994;
revision accepted for publication July 19, 1994.)