
INVITED COMMENTARY

International Journal of Sports Physiology and Performance, 2006;1:50-57
© 2006 Human Kinetics, Inc.

Making Meaningful Inferences About Magnitudes

Alan M. Batterham and William G. Hopkins

A study of a sample provides only an estimate of the true (population) value of an outcome statistic. A report of the study therefore usually includes an inference about the true value. Traditionally, a researcher makes an inference by declaring the value of the statistic statistically significant or nonsignificant on the basis of a *P* value derived from a null-hypothesis test. This approach is confusing and can be misleading, depending on the magnitude of the statistic, error of measurement, and sample size. The authors use a more intuitive and practical approach based directly on uncertainty in the true value of the statistic. First they express the uncertainty as confidence limits, which define the likely range of the true value. They then deal with the real-world relevance of this uncertainty by taking into account values of the statistic that are substantial in some positive and negative sense, such as beneficial or harmful. If the likely range overlaps substantially positive and negative values, they infer that the outcome is unclear; otherwise, they infer that the true value has the magnitude of the observed value: substantially positive, trivial, or substantially negative. They refine this crude inference by stating qualitatively the likelihood that the true value will have the observed magnitude (eg, very likely beneficial). Quantitative or qualitative probabilities that the true value has the other 2 magnitudes or more finely graded magnitudes (such as trivial, small, moderate, and large) can also be estimated to guide a decision about the utility of the outcome.

Key Words: clinical significance, confidence limits, statistical significance

Researchers usually conduct a study by selecting a sample of subjects from some population, collecting the data, then calculating the value of a statistic that summarizes the outcome. In almost every imaginable study, a different sample would produce a different value for the outcome statistic, and of course none would be the value the researchers are most interested in—the value obtained by studying the entire population. Researchers are therefore expected to make an inference about the population value of the statistic when they report their findings in a scientific

Batterham is with the School of Health and Social Care, University of Teesside, Middlesbrough, UK. Hopkins is with the Division of Sport and Recreation, AUT University, Auckland 1020, New Zealand.

journal. In this invited commentary we first critique the traditional approach to inferential statistics, the null-hypothesis test. Next we explain confidence limits, which have begun to appear in publications in response to a growing awareness that the null-hypothesis test fails to deal with the real-world significance of an outcome. We then show that confidence limits alone also fail and then outline our own approach and other approaches to making inferences based on meaningful magnitudes.

The Null-Hypothesis Test

The almost universal approach to inferential statistics has been the null-hypothesis test, in which the researcher uses a statistical package to produce a P value for an outcome statistic. The P value is the probability of obtaining any value larger than the observed effect (regardless of sign) if the null hypothesis were true. When $P < .05$, the null hypothesis is rejected and the outcome is said to be statistically significant.

In an effort to bring meaning to this deeply mysterious approach, many researchers misinterpret the P value as the probability that the null hypothesis is true, and they misinterpret their outcomes accordingly. Cohen,¹ in the classic article “The Earth Is Round ($p < .05$),” summed up this confusion by concluding that significance testing “does not tell us what we want to know, and we so much want to know what we want to know that, out of desperation, we nevertheless believe that it does!”^(p997)

Readers might also wonder what is sacred about $P < .05$. The answer is nothing. Fisher,² the statistician and geneticist who championed the P value, correctly asserted that it was a measure of strength of evidence against the null hypothesis and that “we shall not often go astray if we draw a conventional line at 0.05.”^(p80) It was not Fisher’s intention that $P < .05$ be used as an absolute decision rule. Indeed, he would almost certainly have endorsed the suggestion of Rosnow and Rosenthal that “surely God loves the 0.06 nearly as much as 0.05.”^{3(p1277)}

Regardless of how the P value is interpreted, hypothesis testing is illogical, because the null hypothesis of no relationship or no difference is always false—there are no truly zero effects in nature. Indeed, in arriving at a problem statement and research question, researchers usually have good reasons to believe that effects will be different from zero. The more relevant issue is not whether there is an effect but how big it is. Unfortunately, the P value alone provides us with no information about the direction or size of the effect or, given sampling variability, the range of feasible values. Depending, *inter alia*, on sample size and variability, an outcome statistic with $P < .05$ could represent an effect that is clinically, practically, or mechanistically irrelevant. Conversely, a nonsignificant result ($P > .05$) does not necessarily imply that there is no worthwhile effect, because a combination of small sample size and large measurement variability can mask important effects. An overreliance on P values might therefore lead to unethical errors of interpretation. As Rozeboom⁴ stated, “Null-hypothesis significance testing is surely the most bone-headedly misguided procedure ever institutionalized in the rote training of science students.”^(p35)

Confidence Intervals

In response to these and other critiques of hypothesis testing, confidence intervals are beginning to appear in research reports. The strict definition of the confidence interval is hotly debated, but most if not all statisticians would agree that it is the range within which we would expect the value of the statistic to fall if we were to repeat the study with a very large sample. More simply, it is the likely range of the true, real, or population value of the statistic. For example, consider a 2-group comparison of means in which we observe a mean difference of 10 units, with a 95% confidence interval of 6 to 14 units (or lower and upper confidence limits of 6 and 14 units). In plain language, we can say that the true difference between groups could be somewhere between 6 and 14 units, where *could* refers to the probability used for the confidence interval.

A confidence interval alone or in conjunction with a *P* value still does not overtly address the question of the clinical, practical, or mechanistic importance of an outcome. Given the meaning of the confidence interval, an intuitively obvious way of dealing with magnitude of an outcome statistic is to inspect the magnitudes covered by the confidence interval and thereby make a statement about how big or small the true value could be. The simplest scale of magnitudes we could use has 2 values: positive and negative for statistics such as correlation coefficients, differences in means and differences in frequencies, or greater than unity and less than unity for ratio statistics such as relative risks. If we apply the confidence interval for an outcome statistic to such a 2-level scale of magnitudes, we can make 1 of 3 inferences: The statistic could only be positive, the statistic could only be negative, or the statistic could be positive and negative. The first 2 inferences are equivalent to the value of the statistic being statistically significantly positive and statistically significantly negative, respectively, whereas the third is equivalent to the value of the statistic being statistically nonsignificant. We illustrate these 3 inferences in Figure 1.

This equivalence of confidence intervals and statistical significance is a well-known corollary of statistical first principles, and we will not explain it further here, but we stress that confidence intervals do not represent an advance on null-hypothesis testing if they are interpreted only in relation to positive and negative values or, equivalently, the zero, or null, value.




Figure 1 — Only 3 inferences can be drawn when the possible magnitudes represented by the likely range in the true value of an outcome statistic (the confidence interval, shown by horizontal bars) are determined by referring to a 2-level (positive and negative) scale of magnitudes.

Magnitude-Based Inferences

The problem with a 2-level scale of magnitude is that some positive and negative values are too small to be important in a clinical, practical, or mechanistic sense. The simplest solution to this problem is use a 3-level scale of magnitude: substantially positive, trivial, and substantially negative, defined by the smallest important positive and negative values of the statistic. In Figure 2 we illustrate a crude approach to the various magnitude-based inferences for a statistic with substantial values that are clinically beneficial or harmful.

There is no argument about the inferences shown in Figure 2 for confidence intervals that lie entirely within 1 of the 3 levels of magnitude (harmful, trivial, and beneficial). For example, the outcome is clearly harmful if the confidence interval is entirely within the harmful range, because the true value could only be harmful (where *could only be* refers to a probability somewhat more than the probability level of the confidence interval). A confidence interval that spans all 3 levels is also relatively easy to deal with: The true value could be harmful and beneficial, so the outcome is unclear, and we would need to do more research with a larger sample or with better measures to resolve the uncertainty. But how do we deal with a confidence interval that spans 2 levels—harmful and trivial, or trivial and beneficial? In such cases the true value could be harmful and trivial but not beneficial, or it could be trivial and beneficial but not harmful. Situations like these are bound to arise, because a true value is sometimes close to the smallest important value, and even a narrow confidence interval will sometimes overlap trivial and important values. It would therefore be a mistake to conclude that the outcome was unclear. For example, a confidence interval that spans beneficial and trivial values is clear in the sense that the true value could not be harmful.




Figure 2 — Four different inferences can be drawn when the possible magnitudes represented by the confidence interval are determined by referring crudely to a 3-level (beneficial, trivial, and harmful) scale of magnitudes.

One option for dealing with a confidence interval that spans 2 regions is to make the inference *not harmful*, but the reader will then be in doubt as to which of the other 2 outcomes (trivial or beneficial) is more likely. A preferable alternative is to declare the outcome as having the magnitude of the observed effect, because in almost all studies the true value will be more likely to have this magnitude than either of the other 2 magnitudes. Furthermore, there is an understandable tendency for researchers to interpret the observed value as if it were the true value.

We regard the approach summarized in Figure 2 as crude, because it does not distinguish between outcomes with confidence intervals that span a single magnitude level and those that start to overlap into another level. Furthermore, the researcher can make incorrect inferences comparable to the type 1 and type 2 errors of null-hypothesis testing: An outcome inferred to be beneficial could be trivial or harmful in reality (a type 1 error), and an outcome inferred to be trivial or harmful could be beneficial (a type 2 error). We therefore favor the more sophisticated and informative approach illustrated in Figure 3, in which we qualify clear outcomes with a descriptor⁵ that represents the likelihood that the true value will have the observed magnitude. The resulting inferences are content-rich and would surely qualify for what Cohen referred to as “what we want to know.” They are also free of the burden of type 1 and type 2 errors. An article in the current issue of this journal features this approach.⁶

The inferences shown in Figure 3 are still incomplete, because they refer to only 1 of the 3 magnitudes that an outcome could have, and they simplify the likelihoods into qualitative ranges. For studies with only 1 or 2 outcome statistics, researchers could go 1 step farther by showing the exact probabilities that the true effect is harmful, trivial, and beneficial in some abbreviated manner (eg, 2%, 22%,




Figure 3 — In a more informative approach to a 3-level scale of magnitudes, inferences are qualified with the likelihood that the true value will have the observed magnitude of the outcome statistic. Numbers shown are the quantitative chances (%) that the true value is harmful, trivial, or beneficial.

and 76%, respectively, as illustrated in Figure 3), then discussing the outcome using the appropriate qualitative terms for 1 or more of the 3 magnitudes (eg, very unlikely harmful, unlikely trivial, probably beneficial). Hopkins and colleagues have experimented with this approach in recent publications.⁶⁻⁹ For more than a few outcome statistics, this level of detail will produce a cluttered report that might overwhelm the reader. Nevertheless, the researcher will have to calculate the quantitative probabilities for every statistic in order to provide the reader with only the qualitative descriptors. Statistical packages do not produce these probabilities and descriptors without special programming, but spreadsheets are available to produce them from the observed value of the outcome statistic, its *P* value, and the smallest important value of the statistic (see <http://newstats.org/generalize.html>) or from raw data in straightforward controlled trials.¹⁰

The quantitative approach to likelihoods of benefit, triviality, and harm can be further enriched by dividing the range of substantial values into more finely graded magnitudes. Cohen^{11(pp24,83)} devised default thresholds for dividing values of various outcome statistics into *trivial*, *small*, *moderate*, and *large*. Use of these or modified thresholds (see <http://newstats.org/effectmag.html>) allows the researcher to make informative inferences about even unclear effects; for example, “although the effect is unclear, any benefit or harm is at most small.” Compare this statement with “the effect is not statistically significant” or “there is no effect ($P > .05$).”

Other Approaches to Inferences

Authors who promote the use of confidence intervals usually encourage researchers in an informal fashion to interpret the importance of the magnitudes represented by the interval.^{eg,12,13} We also found several who advocate calculating the chances of clinical benefit using a value for the minimum worthwhile effect,^{14,15} but we know of only 1 published attempt by mainstream statisticians or researchers to formalize the inference process in anything like the way we describe here. Guyatt et al¹⁶ argued that a study result can be considered positive and definitive only if the confidence interval is entirely in the beneficial region. This position is understandable for expensive treatments in health-care settings, but in general we believe it is too conservative. We also believe that the 95% level is too conservative for the confidence interval; the 90% level is a better default, because the chances that the true value lies below the lower limit or above the upper limit are both 5%, which we interpret as very unlikely.⁵ A 90% level also makes it more difficult for readers to reinterpret a study in terms of statistical significance.¹⁷ In any case, a final decision about acting on an outcome should be made on the basis of the quantitative chances of benefit, triviality, and harm, taking into account the cost of implementing a treatment or other strategy, the cost of making the wrong decision, the possibility of individual responses to the treatment (see below), and the possibility of harmful side effects. For example, the ironic “what have we got to lose?” would be the appropriate attitude toward an inexpensive treatment that is almost certainly not harmful, possibly trivial, and possibly beneficial (0.3%, 55%, and 45%, respectively), provided there is little chance of harmful individual responses and harmful side effects.

Some readers will be surprised to learn that there is a thriving statistical counter-culture founded on probabilistic assertions about true values. Bayesian statisticians,

as they are known, make an inference about the true value of a statistic by combining the value from a study with an estimate of a probability distribution representing the researcher's belief about the true value before the study.^{eg.18} Bayesians contend that this approach replicates the way we assimilate evidence, but quantifying prior belief is a major hurdle.¹⁸ Meta-analysis provides a more objective quantitative way to combine a study with other evidence, although the evidence has to be published and of sufficient standard. The approach we have presented here is essentially Bayesian but with a "flat prior"; that is, we make no prior assumption about the true value. The approach is easily applied to the outcome of a meta-analysis.

Where to From Here?

Some researchers may argue that making an inference about magnitude requires an arbitrary and subjective decision about the value of the smallest important effect, whereas hypothesis testing is more scientific and objective. We would counter that the default scales of magnitude promulgated by Cohen are objective in the sense that they are defined by the data and that for situations in which Cohen's scales do not apply, the researcher has to justify the choice of the smallest important effect. We concur with Kirk¹⁹ that researchers themselves are in the best position to justify the choice and that dealing with this issue should be an ethical obligation. In any case, magnitudes are implicit even in a study designed around hypothesis testing, because estimating sample size for such a study requires a value for the smallest important effect, along with an arbitrary choice of type I and type II statistical error rates (usually 5% and 20%, respectively, corresponding to an 80% chance of statistical significance at the 5% level for the smallest effect). Studies designed for magnitude-based inferences will need a new approach to sample-size estimation based on acceptable uncertainty. A draft spreadsheet has been devised for this purpose and will be presented at the 2006 annual meeting of the American College of Sports Medicine.²⁰ Sample sizes are approximately one third of those based on hypothesis testing, for what seems to be reasonably acceptable uncertainty.

Making an inference about magnitudes is no easy task: It requires justification of the smallest worthwhile effect, extra analysis that stats packages do not yet provide by default, a more thoughtful and often difficult discussion of the outcome, and sometimes an unsuccessful fight with reviewers and the editor. On the other hand, it is all too easy for a researcher to inspect the *P* value that every statistical package generates and then declare that either there is or there is not an effect. We might therefore have to wait a decade or two before the tipping point is reached and magnitude-based inferences in one form or another displace hypothesis testing. We will need the help of the gatekeepers of knowledge—peer reviewers of manuscripts and funding proposals, journal editors, ethics-committee representatives, funding-committee members—to remove the shackles of hypothesis testing and to embrace more enlightened approaches based on meaningful magnitudes.

A final few words of caution. . . . Statistical significance, confidence limits, and magnitude-based inferences all relate to the 1 and only true or population value of a statistic, not to the value for an individual in that population. Thus, a large-enough sample can show that a treatment is, for example, almost certainly beneficial, yet the treatment can be trivial or even harmful to a substantial propor-

tion of the population because of individual responses. We are exploring ways of making magnitude-based inferences about individuals.²¹

References

1. Cohen J. The earth is round ($p < .05$). *Am Psychol.* 1994;49:997-1003.
2. Fisher RA. *Statistical Methods for Research Workers*. London, UK: Oliver and Boyd; 1950.
3. Rosnow RL, Rosenthal R. Statistical procedures for the justification of knowledge in psychological science. *Am Psychol.* 1989;44:1276-1284.
4. Rozeboom WW. Good science is abductive, not hypothetico-deductive. In: Harlow LL, Mulaik SA, Steiger JH, eds. *What if There Were No Significance Tests?* Mahwah, NJ: Lawrence Erlbaum; 1997:335-392.
5. Hopkins WG. Probabilities of clinical or practical significance. *Sportscience.* 2002;6. Available at: sportsci.org/jour/0201/wghprob.htm.
6. Hamilton RJ, Paton CD, Hopkins WG. Effect of high-intensity resistance training on performance of competitive distance runners. *Int J Sports Physiol Perform.* In press.
7. Van Montfoort MC, Van Dieren L, Hopkins WG, Shearman JP. Effects of ingestion of bicarbonate, citrate, lactate, and chloride on sprint running. *Med Sci Sports Exerc.* 2004;36:1239-1243.
8. Petersen CJ, Wilson BD, Hopkins WG. Effects of modified-implement training on fast bowling in cricket. *J Sports Sci.* 2004;22:1035-1039.
9. Paton CD, Hopkins WG. Combining explosive and high-resistance training improves performance in competitive cyclists. *J Strength Cond Res.* In press.
10. Hopkins WG. A spreadsheet for analysis of straightforward controlled trials. *Sport-science.* 2003;7. Available at: sportsci.org/jour/03/wghtrials.htm.
11. Cohen J. *Statistical Power Analysis for the Behavioral Sciences*. Hillsdale, NJ: Lawrence Erlbaum; 1988.
12. Braitman LE. Confidence intervals assess both clinical significance and statistical significance. *Ann Intern Med.* 1991;114:515-517.
13. Greenfield MLVH, Kuhn JE, Wojtys EM. A statistics primer: confidence intervals. *Am J Sports Med.* 1998;26:145-149.
14. Shakespeare TP, Gebiski VJ, Veness MJ, Simes J. Improving interpretation of clinical studies by use of confidence levels, clinical significance curves, and risk-benefit contours. *Lancet.* 2001;357:1349-1353.
15. Froehlich G. What is the chance that this study is clinically significant? a proposal for Q values. *Eff Clin Pract.* 1999;2:234-239.
16. Guyatt G, Jaeschke R, Heddle N, Cook D, Shannon H, Walter S. Basic statistics for clinicians: 2. interpreting study results: confidence intervals. *Can Med Assoc J.* 1995;152:169-173.
17. Sterne JAC, Smith GD. Sifting the evidence—what's wrong with significance tests. *Br Med J.* 2001;322:226-231.
18. Bland J, Altman D. Bayesians and frequentists. *Br Med J.* 1998;317:1151-1160.
19. Kirk RE. Promoting good statistical practice: some suggestions. *Educational and Psychological Measurement.* 2001;61:213-218.
20. Hopkins WG. Sample sizes for magnitude-based inferences about clinical, practical or mechanistic significance [abstract]. *Med Sci Sports Exerc.* In press.
21. Pyne DB, Hopkins WG, Batterham A, Gleeson M, Fricker PA. Characterising the individual performance responses to mild illness in international swimmers. *Br J Sports Med.* 2005;39:752-756.