Large enough to matter, too small to be consequential, or insufficient evidence? A unified
framework for more meaningful statistical inference

Adam H. Smiley^{1,2}, Jessica J. Glazier^{1,3}, and Yuichi Shoda^{1*}

¹Department of Psychology, University of Washington, Seattle, WA, 98195

²Department of Psychological Science & Neuroscience, Belmont University,
Nashville, TN 37221

³Department of Psychology, Northeastern University, Boston, MA 02115

AHS, JJG, and YS contributed equally to articulating the unified framework for statistical inference described in this article as well as drafting and revising the manuscript.

Correspondence concerning this article should be addressed to Yuichi Shoda, University of Washington, 119A Guthrie Hall Box 351525, Seattle, WA 98195-1525. Email: yshoda@uw.edu

We thank Stats Atwood, Dan Benjamin, Bryn Bandt-Law, Doris Dai, Léïla Eisner, Brian Flaherty, Jon Gallegos, Jin Goh, Tabea Hässler, Cheryl Kaiser, Kevin King, Randy Lee, Ella Lombard, and Ishika Ray for their helpful comments on earlier versions of this manuscript. This project is supported in part by the National Science Foundation Graduate Research Fellowship Program (NSF GRFP).

This paper has not been published in a peer-reviewed journal. A pre-print of this article can be found at <https://psyarxiv.com/465ha/>.

Abstract

Does the observed effect warrant inferring that the true effect in the population is large enough to matter, or too small to matter? Is the evidence too weak to support either inference? The framework presented here answers these questions by considering a “null region,” rather than a point null hypothesis. It also eliminates concerns about “over-powered” studies and facilitates registered reports. This article describes how to: apply this framework without specialized computer programs whenever confidence intervals can be calculated, operationalize “large enough to matter,” and integrate minimum effects testing, equivalence testing, and traditional null hypothesis testing into a single framework. [100 words]

Keywords: Clinical and Practical Significance, Best Practices for Statistical Inference, Minimum-Effect Testing, Equivalence Testing, Null Hypothesis Significance Testing

Large enough to matter, too small to be of consequence, or insufficient evidence? A unified framework for a more meaningful statistical inference

Traditional null hypothesis significance testing (NHST) deems an observed effect “significant” if it provides strong enough evidence against the counterfactual, “null” hypothesis that there is no effect at all. A major issue with this is that even if an effect is *too small to matter*, it can still be labeled “significant,” an increasingly common scenario due to the availability of larger samples in research making use of “Big Data.” In fact, the use of traditional NHST to define a “finding” is one of the reasons “why most published findings are false” (Ioannidis, 2005). Furthermore, NHST does not provide a way to infer, with sufficient evidence, that the true effects in the population are too small to matter. In this paper, we describe a general framework based on the principles of minimum-effect testing (Fowler, 1985; Murphy & Myors, 1999; Serlin & Lapsley, 1985) and equivalence testing (Dunnett & Gent, 1977; Tryon, 2001; Lakens et al., 2018) to address these shortcomings and allow for more meaningful statistical inference. In addition, we describe a way of using confidence intervals to conduct the tests in our framework without any specialized software. We believe this, combined with the conceptual simplicity and clarity of the framework, has a chance of finally providing a viable alternative, or a supplement, to the traditional NHST, and would like to offer them for consideration by the scientific communities, as well as suggesting research practices facilitated by this framework.

In the introduction of this paper, we first provide context to our proposal by describing how others have addressed the issues with traditional null hypothesis testing. Then, we briefly review how minimum effect-testing has been conceptualized in the past, and detail why conducting MET using confidence intervals will make this approach more versatile and accessible. Next, we describe the practice of equivalence testing, how it can work in tandem with MET, and the

inferences that can be drawn by using them in a unified framework.

Null Hypothesis Significance Testing: In Search of an Alternative

For some research questions, small effects may be of theoretical importance (e.g., Prentice & Miller, 1992; Greenwald et al., 2015). However, in many circumstances it is misleading to refer to very small effects as *significant*, which in common usage means “sufficiently great or important to be worthy of attention; noteworthy; consequential, influential” (Oxford English Dictionary; Oxford University Press, 2000). Adopting a criterion more stringent than the commonly used $p < .05$ (e.g., $p < .005$; Benjamin, et al., 2018) would increase NHST’s evidentiary strength (reflected in the relative odds of alternative vs. null hypotheses, i.e., the Bayes factor), as well as the replicability of statistically significant results. However, even with a different standard, NHST remains fundamentally a test against the null hypothesis of no effect at all. Additionally, the standard practice of conducting null hypothesis testing does not involve *a priori* specification of how effect sizes will be interpreted, leaving room for bias to enter into conclusions drawn based on the size of the observed effect.

With the widely recognized problems associated with NHST such as those described above (see also Greenwald et al., 1996), many have advocated abandoning statistical significance testing altogether in favor of using confidence intervals (e.g., Cohen, 1994; Szucs & Ioannidis, 2017). If interpreted properly, confidence intervals are highly informative. In fact, our suggested approach uses confidence intervals as a central part of its decision criterion. However, in practice, researchers often use confidence intervals only to ask if they contain zero, which is functionally no different from traditional null hypothesis testing (Wasserstein & Lazar, 2016). Furthermore, many researchers have “a gross misunderstanding” of the meaning of confidence intervals (Hoekstra et al., 2014).

One reason for the persistent use of NHST may be the heuristic value of a discrete decision criterion. In fact, research on decision making under uncertainty has found that focusing on a discrete consequential outcome while also describing its degree of uncertainty promotes the most effective use of uncertainty information (e.g., Nadav-Greenberg & Joslyn, 2009). Discrete decision criterion is also useful for stating the results concisely (e.g., “X had a significant effect on Y” is more concise than “the effect of X on Y is such-and-such, with a confidence interval of [lower bound, upper bound]”). Authors and editors may also find it useful to have a discrete criterion for deciding what findings are reasonable to be mentioned in the abstract. Is there a way to preserve these heuristic values without suffering from the myriad problems of NHST? Below is our attempt to answer this question.

The Evolution of Minimum-Effect Testing (MET)

Meehl (1967) argued that when using traditional NHST, increasing precision (e.g., sample size) virtually guarantees that a “statistically significant” result can be found to support any theory even “if the theory is entirely without merit”¹ (p. 111). That is, with NHST, increasing sample size will increase the false positive rate when the true effect is too small to matter, a reason “why most published findings are false” (Ioannidis, 2005). However, this would not be the case if minimum-effect testing (MET) were used instead. If the observed effect is weaker than the smallest possible effect considered to be meaningful, MET will *never* be significant, regardless of sample size. On the other hand, if the observed effect is stronger than the smallest effect of consequence, the larger the sample size, the greater the likelihood that it is significant by MET.

In addition, minimum-effect testing intrinsically accommodates the fact that often there is

¹ Serlin and Lapsley (1985) took this argument a step further, writing in the abstract, “because the psychological null hypothesis is *always false*, increases in precision in psychology always lead to a weaker test of a theory, whereas the converse is true in physics” (p. 73; emphasis added).

a *range* of true effects that are not consequential. This is because MET tests against the null hypothesis that the true effect in the population is within a range of values that are pre-declared as not noteworthy, consequential, or meaningful. It extends the “good-enough principle” advocated by Serlin and Lapsley (1985), in which one first specifies a “belt” around zero and determines if the observed effect falls reliably outside of that range. Murphy and Myers (1999) built upon this idea with minimum-effect tests. In their method of MET, an F test is used to determine if the observed effect is larger than can be reasonably expected from the pre-specified minimally consequential variance accounted for. It will consider an effect significant only when there is sufficient evidence that the observed effect would be unlikely to occur if the true effect in the population was too small to matter. However, in the over twenty years since Murphy and Myers’ article was published, MET has been seldom used for empirical research. We suspect this is largely because of the difficulty associated with conducting MET with F tests, as well as the lack of versatility offered by the particular method proposed by Murphy and Myers. As we demonstrate below, however, we will describe a way to conduct MET (and Equivalence tests) simply using confidence intervals. This will make MET more versatile, as it can be applied to any statistic for which confidence intervals can be calculated (i.e., not just F), and it can be used for both raw and standardized units.

Equivalence Testing: More than “Retaining the Null”

If an effect is deemed to be non-significant by MET (or NHST), one cannot simply conclude that the null hypothesis is true. However, by using Equivalence Testing (EqT), researchers can test against the null hypothesis that the true effect in the population is too large to be considered “equivalent to 0”; rejecting this null hypothesis provides a reasonable basis for inferring that the effect is likely too small to be meaningful (e.g., Lakens et al., 2018). Similar to

our proposed method of MET, EqT can also be conducted using confidence intervals (Dunnett & Gent, 1977; Tryon, 2001; Westlake, 1976; Bauer & Kieser, 1996). Much of the literature on Equivalence Testing emphasizes its implementation for determining if two sample means are equivalent (Tryon, 2001) or if an effect is very close to zero (Lakens et al., 2018). However, the logic of equivalence testing can be generalized to a situation where the goal is to provide a basis for inferring that the true strength of the effect is in *any* prespecified range, as we discuss later in this article.

A Unified Framework that Combines MET and EqT

Both MET and EqT are useful methods on their own, but because they use the same underlying logic, they can naturally be incorporated into an integrated framework. In MET, researchers test if they can infer that the true effect falls *outside* of the range of values that are not of interest. In Equivalence Testing, researchers test if they can infer that the true effect is somewhere *within* the range of values that are not of interest (Lakens et al., 2018). Various terms have been used to refer to this range. In order to use a term applicable to both tests, while emphasizing its similarity to the point null hypothesis in traditional NHST, we refer to it as the “null region.” Using MET and EqT in tandem, researchers can draw one of the following three possible inferences about the true effect: (1) the true effect is *outside* of the null region, (2) the true effect is *within* the null region, or (3) the data *do not provide sufficient evidence* to support either inference. This last inference cannot be made by MET alone, or EqT alone. Maxwell and colleagues (2015) have described that data should be looked at from this perspective, and others (e.g., Lakens et al., 2018) have discussed the relationship between MET and EqT. Building upon their recommendations, in this article, we describe how to integrate MET with EqT to optimize

the informativeness of statistical inference.

Below, we first describe the conceptual framework underlying minimum-effect testing for a directional (one-tailed) hypothesis and then extend it to a non-directional (two-tailed) hypothesis. Second, we describe an overarching framework in which minimum-effect testing is used in conjunction with equivalence testing and discuss how this framework can facilitate more meaningful scientific inference. Third, we demonstrate how to conduct MET and EqT using confidence intervals. Fourth, while we believe each scientific community should develop their own standards for specifying the magnitude of effects that would be considered meaningful for a given research question, we offer a few suggestions that could be helpful. Fifth, we discuss how the results of MET and EqT can be reported succinctly and clearly. And finally, we describe the implications for registered reports and “big data.”

Conceptualizing One-Tailed Minimum-Effect Testing

The rationale for MET is the same as traditional null hypothesis significance testing (NHST), except for the null hypothesis being tested. Thus, we will first review traditional NHST and then describe how MET differs from it and how to use confidence intervals to perform MET.

The logic of traditional NHST is as follows: (1) We make a counterfactual assumption that, in the population, the true effect is 0 (e.g., the difference between the mean of the experimental condition and that of the control condition is 0). This is the null hypothesis (H_0). (2) We then consider the fact that data from any given study are only a subset of the population, and thus the observed value varies from one study to the next and is never *exactly* 0 even if the true effect is 0 in the population. The theoretically expected distribution of the observed values (e.g., sample means) across many repetitions of the study is the sampling distribution shown as the Gaussian “bell” curve shown in panel A of Figure 1. (3) Based on this distribution, we compute the

probability of obtaining the observed value or stronger under the null hypothesis. This is the orange area in Panel A of Figure 1. (4) The size of this area, relative to the entire distribution, is the p-value. Conventionally if the p-value is smaller than the value α , chosen as the criterion for considering that the observed effect (or stronger) is “unlikely enough” (e.g. $p < .05$, or $p < .005$, as suggested by Benjamin et al., 2018), the null hypothesis is rejected and the observed effect is considered statistically significant.

In one-tailed NHST, the effect is considered “significant” when there is strong enough evidence to reject the null hypothesis that the true effect is 0 or lower (when predicting a positive effect) or that the true effect is 0 or higher (when predicting a negative effect). Thus, conceptually, the null hypothesis refers to a null *region* of “0 or lower,” or “0 or higher”, depending on whether a positive or negative effect is predicted. Fortunately, when the observed effect is positive and far enough from 0 to reject the null hypothesis of “absolutely no true effect,” it also allows one to reject all other point null hypotheses that the true value is in the null region of “0 or lower,” because the p values would have been even smaller if the population value was lower than 0. Similarly, when the observed effect is negative and far enough from 0 to reject the null hypothesis of “absolutely no true effect,” it also rejects all other point null hypotheses that the true value is in the null region of “0 or higher.”

Figure 1. Panel A: A sampling distribution of observed values expected under the null hypothesis H_0 that the true effect is zero (i.e., traditional NHST). The orange portion represents the area under the curve above the $1 - \alpha$ quantile (e.g., 0.95) of the sampling distribution. Any observed value in this region is considered unlikely enough given H_0 , and deemed statistically significant by NHST. Panel B: A sampling distribution expected under the MET null hypothesis (H'_0) that the true effect is at a value considered the greatest effect of no interest, NR_U . The green portion represents the area above the $1 - \alpha$ quantile for the sampling distribution under H'_0 . Any observed value in this region is considered significant by MET. Note: In this example the NR_U is shown as larger than the critical value for NHST, but does not necessarily need to be the case.

The rationale for MET is the same as NHST described above except in a one-tailed MET predicting a positive result, the null hypothesis is that, that the true effect in the population is at most a certain value pre-declared by the researcher as the largest that is still too small to matter. We will refer to this value as NR_U (“Upper bound of the Null Region”).² Thus in contrast to the traditional one-tailed NHST for which the null region is “0 or lower” when predicting a positive

² This example could also be conceptualized as a two-tailed MET (described in the next section) with NR_U = the largest effect considered too small to matter and $NR_L = -\infty$. In bounded distributions, such as correlations (–1 to 1) or percentages (0 to 100%), NR_L would be equivalent to the lowest possible value (e.g., –1 or 0%) rather than $-\infty$.

effect, for a one-tailed MET when predicting a positive effect, the null region is “ NR_U or lower.” For example, suppose that a public vaccination program is known be able to prevent a virus epidemic as long as the reproduction number of the virus (the average number of people one infected person will infect) is 3.5 or lower.³ If a study provides good evidence that the reproduction number of the virus in the community of interest is less than 3.5, then it suggests that the current vaccination program is good enough. On the other hand, if the study provides evidence that the reproduction number is greater than 3.5 then urgent action is called for. Then for a study assessing the reproduction number of a new variant, it would be reasonable to set the null range for a one-tailed MET as “reproduction number of 3.5 or lower.” To distinguish this form of null hypothesis from the traditional null hypothesis (H_0) of no effect whatsoever, we refer to it as a *null region hypothesis* (H'_0).

Panel B of Figure 1 is the sampling distribution relevant to MET. It is the expected distribution of the observed reproduction number across many samples randomly drawn from the population where the true reproduction number is NR_U (e.g., the maximum reproduction number of a virus that the current vaccination program can control). When the green region is “small enough” (i.e., smaller than α), then the observed effect is considered significant using MET. Rejecting H'_0 in favor of the prediction of larger effect thus provides evidence that the true effect is large enough to matter. What is “large enough to matter” of course depends on the purpose of research, which should be clearly stated when conducting MET. It is possible that even in the same study, one may report the results of multiple METs with a different null region for each of different purposes, using different values of NR_U as the threshold for “large enough to matter.” Multiple

³ See, for example, Bartsch et al. (2020). It predicts that if the vaccine has 80% efficacy (i.e., a vaccinated group of people will have 80% less disease cases compared to an unvaccinated group), vaccinating 75% of the population is enough to prevent an epidemic due to a virus with the reproduction number of 3.5 or lower, but that a virus with a reproduction number greater than 3.5 will result in an epidemic even in this relatively well vaccinated population.

METs, each with different null regions, can easily be accommodated in publications (see Table 6 as an example).

Conceptualizing Two-Tailed Minimum-Effect Testing

Similar to a one-tailed traditional NHST, one-tailed MET only considers an effect “significant” when it is in *one* direction. In contrast, a two-tailed MET considers two directional hypotheses. For example, in a study assessing the reproduction number of a new virus variant, one hypothesis may be that the reproduction number is 3.5 or higher, which prompts urgent action to add additional measure to prevent an epidemic, as the current vaccination program will not be able to contain the virus. The other hypothesis is that the reproduction number of the new variant is low enough so that the current vaccination program could be safely halted now even though its target population vaccination rate has not been achieved. Suppose this action (halting the current vaccination program) is warranted if the reproduction number of the new variant is 2.5 or lower.

A reproduction number of between 2.5 and 3.5 does not prompt either action, and suggest that no change needs to be made and the current program should continue as planned. Thus we may designate the 2.5 to 3.5 range as the null region, in contrast to the “point null hypothesis” of the traditional NHST. A two-tailed MET tests if a result is significantly higher than the upper bound of the null region (NR_U) or if it is lower than the lower bound of the null region (NR_L). In this scenario, the virus reproduction number of 3.5 corresponds to the value referred to as NR_U in Figure 2, and the reproduction number of 2.5 corresponds to the value referred to as the NR_L in Figure 2.

In another scenario, for example a study examining the relationship between the level of a newly discovered hormone and mental health, any correlation, positive or negative, with a magnitude of 0.1 may be considered notable. In this case, the null region would span from -0.1

(NR_L) to $+0.1(NR_U)$.

In short, the logic of a two-tailed MET is identical to the one-tailed version except it considers two null hypotheses (corresponding to the two Gaussian curves shown in Figure 2).

Figure 2. Two sampling distributions located at NR_L and NR_U , respectively. If an observed effect falls within the shaded region of either distribution (each region = α) it would be considered significant in a two-tailed MET.

Two-tailed MET deems an observed effect significant if it falls in either of the two shaded regions of these curves. That is, two-tailed MET applies the “two one-sided tests” (TOST) procedure, each with a different directional null hypothesis (Schuirmann, 1987; Lakens et al., 2018). For an alpha level of .05, an effect is deemed significant using a two-tailed MET if it is higher than the 97.5% percentile of the sampling distribution expected from a population at NR_U , or if it is lower than the 2.5% percentile of the sampling distribution expected from a population located at NR_L .⁴

⁴ Strictly speaking, if the true population value is at NR_U , there is a non-zero probability of obtaining results that fall in the area on the *left* that is labeled “significant MET.” In most situations, the lower critical value is so much farther away from the population value, this probability is virtually 0, making the overall type I error under the null hypothesis nearly entirely in the area under the curve on the right. However, when NR_L is very close to NR_U the sampling distribution from a population located at NR_U may extend meaningfully below the lower critical value. When NR_L and NR_U are identical (i.e., when MET becomes identical to a traditional two-tailed test using NHST), the area under the

Conceptualizing A Unified Framework

When an observed effect is not significant using a two-tailed MET, it may be tempting to conclude that the true effect is located within the null region (i.e., between NR_L and NR_U). However, as is the case for any failure to reject null hypotheses, a non-significant effect alone should not be considered sufficient support for *accepting* the null hypothesis.

In this context, it would be helpful to explicitly test if the observed effect is significantly higher than NR_L , and *at the same time* significantly lower than NR_U . This process is called equivalence testing (Lakens et al., 2018). If they are both significant, this provides good evidence that the true effect in the population of interest is somewhere within the null region.

As shown in Figure 3, this occurs when the following two conditions are *both* met: (1) the observed value is lower than the $\alpha/2$ quantile (e.g., 2.5 percentile) of the sampling distribution for NR_U , and (2) the observed value is higher than the $1 - (\alpha/2)$ quantile (e.g., 97.5 percentile) of the sampling distribution for NR_L .

curve below the lower critical value will equal the area under the curve above the upper critical value. Thus by setting the lower critical value to the $\alpha/2$ quantile of the sampling distribution from NR_L and at the $1 - \alpha/2$ quantile of the sampling distribution from NR_U , the type I error rate will be no larger than α (when $NR_L = NR_U$) and no smaller than $\alpha/2$. The same reasoning applies under the null hypothesis that the true population value is at NR_L . When the population value is between NR_L and NR_U , the type-I error rate is likely much smaller than .05.

Figure 3. MET and Equivalence testing can be used together. If the observed effect is in the blue (central) region, it is significantly lower than NR_U and significantly higher than NR_L . The green regions (extreme left and extreme right) represent the areas where the effect would be significant by a two-tailed MET, because they are reliably higher than NR_U or reliably lower than NR_L .

When MET is used together with EqT, there are five potential outcomes corresponding to the five regions in Figure 3, starting from the left most region. This is summarized in Table 1.

Region	Region Description	Corresponding hypothetical result in Table 2
Region 1	The observed value is significantly lower than the lower bound of the null region (i.e., NR_L).	A
Region 2	The observed value is not significant by MET nor by EqT.	B, C
Region 3	The observed value is significant by EqT, offering evidence that it is within the null region.	D
Region 4	The observed value is not significant by MET nor by EqT.	E, F
Region 5	The observed value is significantly higher than the upper bound of the null region (i.e., NR_U).	G

Table 1. Descriptions of and decisions for the five regions shown in Figure 3, starting from the left most region.

When the obtained result is not significant by either test (e.g., it is in region 2 or region 4),

researchers should consider the results to be inconclusive and refrain from drawing any conclusions. Note that as sample size increases while the effect size remains the same, Regions 2 and 4, where neither of the tests are significant, will become increasingly narrow.

MET and EqT complement traditional NHST and confidence intervals. For example, results of a hypothetical study assessing the reproduction number of different new virus variants may be reported as follows:

Virus variant	Observed reproduction number	NHST, H_0 : reproduction number = 0	Confidence interval	MET and EqT with $NR_L = 2.5$ and $NR_U = 3.5$		
				MET: reproduction number ≤ 2.5	EqT: $2.5 < \text{reproduction number} < 3.5$	MET: reproduction number ≥ 3.5
A	2.0	$p < .001$	1.7-2.3	$p < .001$	ns	ns
B	2.3	$p < .0005$	2.0-2.6	ns	ns	ns
C	2.7	$p < .0005$	2.4-3.0	ns	ns	ns
D	3.0	$p < .0001$	2.7-3.3	ns	$p < .001$	ns
E	3.3	$p < .00005$	3.0-3.6	ns	ns	ns
F	3.7	$p < .00001$	3.4-4.0	ns	ns	ns
G	4.0	$p < .000001$	3.7-4.3	ns	ns	$p < .001$

Table 2. Hypothetical results of a vaccine efficacy study. ns = not significant.

Conducting Minimum-Effect and Equivalence Testing with Confidence Intervals

MET and EqT can be conducted if the statistic of interest has a method for determining confidence intervals, without explicitly needing to determine the sampling distribution (this is because the computation of confidence intervals already takes sampling distributions into account, see Supplementary Text S1: *The Rationale for using Confidence Intervals to conduct MET* for an explanation). The procedure is described in Table 3.

How to Conduct Minimum-Effect Testing (MET) and Equivalence Testing (EqT) using confidence intervals
<p><u>One-tailed MET, when <i>predicting that the effect is</i> higher than the null region, is significant when:</u></p> <ul style="list-style-type: none"> the confidence interval is entirely above NR_U, the highest population value that is still too low to be of consequence.
<p><u>One-tailed MET when <i>predicting that the effect is</i> lower than the null region, is significant when:</u></p> <ul style="list-style-type: none"> the confidence interval is entirely below NR_L, the lowest population value that is still too high to be of consequence.
<p><u>Two-tailed MET</u> is significant when the result is significantly outside of the null region. i.e.,</p> <ul style="list-style-type: none"> the confidence interval is entirely above NR_U, the highest population value that is still too low to be of consequence. <p><u>or</u></p> <ul style="list-style-type: none"> the confidence interval is entirely below NR_L, the lowest population value that is still too high to be of consequence.
<p><u>Equivalence testing is significant</u> when the result is significantly within the null region. i.e.,</p> <ul style="list-style-type: none"> the confidence interval is entirely below NR_U, the highest population value that is still too low to be of consequence. <p><u>and</u></p> <ul style="list-style-type: none"> the confidence interval is entirely above NR_L, the lowest population value that is still too high to be of consequence.

Table 3. A summary of how to conduct a one-tailed or two-tailed MET and an EqT using confidence intervals. NR_L and NR_U refer to the value(s) of lower and upper bounds of the null region.

Specifying the Null Region

Specifying NR_L and NR_U plays a critical role for MET and EqT. As it should be guided by specific research questions and involves considering many factors, an important challenge for any field of study (or scientific journal) is developing practices and standards for specifying them. For example, clinical psychologists may be most interested in considering clinical significance, but public health researchers may find that utility analysis often makes the most sense. Below, as a guide, we summarize a number of approaches that could be used to specify meaningful NR_L and NR_U values.

Figure 4 presents possible bases for NR_L and NR_U values and describes the rationale for each. They are not meant to be exhaustive; there may be other means of specifying NR_L and NR_U values that have not yet been suggested. We encourage each scientific discipline to develop their own guidelines and practices most applicable to their research question.

Figure 4. This flowchart provides a series of questions researchers could ask when determining the NR_L and NR_U .

Any non-zero true effect is meaningful. In some circumstances, *any* non-zero true effect might be meaningful, in which case NR_L or NR_U would be 0 and MET becomes equivalent to NHST. Researchers, however, should explicitly state and justify why any true non-zero effect would be meaningful.

Utility Analysis. When Murphy and Myors (1999) described their procedure for minimum-effect testing, they suggested the use of utility analysis in determining a minimum effect. Utility analysis involves comparing the costs of an intervention against their benefits measured in “utility based units” (Robinson, 1993). This utility is often described in terms of financial consequences (Schmidt et al., 1979), or in terms of patient outcomes such as quality adjusted life years (Robinson, 1993). Regardless of the units used, if NR_L or NR_U is set at the magnitude of the effect where the cost and benefit of the effect “break even,” a significant MET suggests that the benefit outweighs the costs.

Clinical Significance. In the context of clinical intervention research, traditional statistical significance testing has long been criticized for its failure to accurately reflect clinically meaningful change, because a statistically significant finding is not necessarily clinically meaningful (e.g., Kazdin, 1999). Jacobson and colleagues (1984) first proposed clinical significance in the context of clinical practice as a function of whether a client has moved from the dysfunctional to functional range on the assessment used to track change during treatment. Similarly, in medicine, the concept of *Minimal Clinically Important Difference* (MCID; Guyatt et al., 1987) has been proposed and clinical researchers are urged to make active use of it (e.g., Embry & Piccirillo, 2020; McGlothlin & Lewis, 2014; see King, 2011, for an extensive discussion of related terms, such as Minimal Important Difference (MID), Minimally Important Change (MIC),

Clinically Important Differences (CID), and Smallest Real Difference (SRD)).

Practical Significance. Clinical significance can be thought of as an evolution of a more general concept of practical significance (Pintea, 2010). For the purpose of conducting MET, NR_L and NR_U can be the smallest true effect size that would be of potential utility in *practice* or in the real world (Kirk, 1996).

Ambient Noise Level. For many research questions, for example in research where there is no obvious cost to weigh against outcomes, the approaches above may not be feasible or appropriate in determining NR_L and NR_U . Under such circumstances, one reasonable basis for setting NR_L and NR_U is the “ambient noise level” (Lykken, 1968; also called also called the “crud estimate,” Orben & Lakens, 2020). This refers to the fact that even in extremely large samples, non-zero correlations exist between variables with no obvious relation to each other (e.g., interest in woodworking and birth order). This led Paul Meehl to observe that “In the social sciences and arguably in the biological sciences, ‘everything correlates to some extent with everything else’” (Meehl, 1990, p. 204). That is, because of associations due to historical accidents (e.g., Krugman, 1991) rather than to any intrinsic or theoretically meaningful relationship, the correlation between two variables that are conceptually not related at all will still approach a significant value with an ever-increasing sample size. Similarly, in an experiment, even if the theoretical variable intended to be manipulated has no true effect at all, with an ever-increasing sample size the effects of any actual experimental manipulation will asymptote towards a non-zero value, because it is virtually impossible to manipulate only the theoretical variable of interest and no other variables (i.e., it is impossible to eliminate all confounds).

Several scholars have attempted to estimate the size of the ambient noise level by looking at large correlation matrices. Pearson’s r of .07 (corresponding to Cohen’s d of 0.14) seems to

provide a reasonable approximation for correlations between variables for which there is no theoretical reason to expect any correlation. This could be used as a baseline across disciplines (see Standing et al., 1991, who based this estimate on correlations computed with a sample size over 2000; see also Webster & Starbuck, 1988). Thus, when it is not possible to specify NR_L and NR_U via other methods, NR_L and NR_U could be set at least as large as the ambient noise level. For example, the null region could be set between $r = -.07$ and $r = +.07$ (or $d \leq 0.14$). Doing so can reduce the likelihood of labelling relationships reflecting historical accidents as theoretically meaningful.

***A priori* specification of the Null Region**

We suggest that NR_L and NR_U values be specified and pre-registered prior to data collection/analysis whenever possible. Without specifying NR_L and NR_U prior to data analysis, researchers may be more likely to fall into committing *post hoc* rationalizations of small effect sizes, or *HARKing* (*Hypothesizing After the Results are Known*, Kerr, 1998). Researchers can utilize pre-registration as a way to show that their NR_L and NR_U values were specified *a priori*.

Considerations for Editors and Reviewers

Although we hope that researchers will be unbiased in specifying NR_L and NR_U , it is important to consider the incentives that affect researchers. Many researchers may be motivated to set NR_L and NR_U values to specify an excessively narrow null region, out of concern that setting a wide null region would make it difficult to reject the MET null hypothesis. For this reason, editors and reviewers should consider the appropriateness of the NR_L and NR_U values chosen by researchers. We believe there are a number of potential mechanisms that can be used to address this issue (for an example, see Supplemental Text S2: *Peer-Review*), such that the balance of the authors' and reviewers' incentives, in addition to their commitment to the fundamental goals of

science, will result in constraining the proposed NR_L and NR_U values within a reasonable range. This process will also facilitate a benefit of using MET, namely increasing discourse about effect sizes as well as transparency in the research process.

Reporting Findings

We have provided three examples of how MET may be used in reporting findings. These examples demonstrate how this can be done concisely, and how it can be done in addition to, rather than instead of, the results of traditional NHST.

Example 1

Table 4 shows an example of the (hypothetical) results from a two-group experiment using one-tailed MET and EqT. Specifically, this study tested the hypothesis that when Democrats and Republicans undergo a “fast friends” paradigm to get to know each other (Aron et al., 1997) their level of warmth towards members of the other political party (the “DV” in Table 4) will increase meaningfully (effect size $\geq .2$) in comparison to a control intervention.

Experimental Condition	DV	Effect Size (EF)	Confidence interval	NHST	MET and EqT	
					EqT $-0.2 < \text{true EF} < +0.2$	One tailed MET $0.2 \leq \text{true EF}$
Fast Friends (FF)	3.5		[3.1, 3.9]			
Control (C)	2.4		[2.1, 2.7]			
Difference		.7	[0.3, 1.1]	$p < .001$	ns	$p < .01$

Table 4. The hypothetical results of a two-group experiment analyzed with MET and EqT.

Example 2

Table 5 provides an example of a linear regression from a simulated “big data” study and how results could be interpreted using a two-tailed MET and EqT. $N = 10,000$ and H'_0 s (NR_L and NR_U) vary by predictor variable (see H'_0 column).

Predictor	Standardized beta	SE	t-value	p-value (NHST)	95% CI	H' ₀	MET/EqT
Intercept	0.000	.009	0	1	[−.017, .017]	N/A	N/A
A	0.300	.009	34.44	<.001	[.283, .317]	−.1, .1	Sig. MET
B	0.009	.009	.21	.21	[−.006, .027]	−.1, .1	Sig. EqT
C	0.390	.009	44.79	<.001	[.373, .407]	−.2, .2	Sig. MET
D	0.108	.009	9.18	<.001	[.091, .125]	−.1, .1	n.s.

Table 5. The results of linear regression with $N = 10,000$ and varying H'_0 s by predictor variable.

Example 3

Below, the description of the results reported in Table 4 of Shoda et al. (1990) has been modified to include results of NHST, MET, and EqT simultaneously.

For MET and EqT, two sets of analyses were conducted, corresponding to two possible purposes for which the Shoda et al. (1990) findings are likely to be used. The first purpose is testing the theoretical proposition that individual differences in waiting time in the diagnostic condition reflect something relatively lasting (e.g., cognitive skills, family, or community environment conducive to school achievement). For this purpose, any evidence that the correlations between waiting time and future outcomes are more than what are expected for any randomly chosen pair of variables would be meaningful. Thus, we used the commonly accepted “ambient noise level” of $r = -.07$ and $r = .07$ as the NR_L and NR_U respectively. Against this null hypothesis, the two correlations that were significant by NHST (the ones with asterisks in Table 4) were also significant with a two-tailed MET.

The second purpose for which the present results can be used involves practical consequences of the differences in SAT scores predicted from preschool waiting time. Thus, in this second set of MET and EqT, we based the H'_0 values on practical significance, in particular the expert opinion published in *New York Times* that a difference of as few as 20 points on a student’s SAT could make a difference in college admissions or financial aid (Aviv, 2009). Based on this, we chose a 20-point difference in SAT as the smallest effect of interest (positive or negative). How strong do the correlations need to be in order to predict a 20-point difference in SAT scores if one child waits 5 minutes longer than another? We used the additional information from the same data set reported on page 936 of Mischel, Shoda, & Rodriguez (1989) and determined that if correlations between preschool waiting time and SAT are at least 0.139 for SATV, and 0.145 for SATQ, then a difference of 20 points on SATV and SATQ, respectively, is predicted for every 5 additional minutes of waiting time in preschool. Thus, in this set of MET and EqT, for the observed correlations of SATV and SATQ with preschool waiting time, correlations of ± 0.139 (for SATV) and ± 0.145 (for SATQ) were used as NR_L and NR_U values.

Below (Table 6 in this manuscript) is how Table 4 from the original Shoda et al. (1990) article would be revised, in light of the above. No deletions from the originally published table were made; additions related to MET and EqT are underlined.

Correlations Between Preschool Delay Time and Scholastic Aptitude Test (SAT) Scores				
Measure	Spontaneous ideation		Suggested ideation	
	Rewards exposed	Rewards obscured	Rewards exposed	Rewards obscured
SAT Verbal	.42*	-.12	-.40	-.21
SAT Quant.	.57**	-.31	-.26	-.23
Sample size	35	33	14	12

Table 6. *: $p < .05$, and **: $p < .01$ by NHST with $H_0 : r = 0$. **Bold:** $p < .05$ by two-tailed MET with $NR_L: r = -.139$, $NR_U: r = +.139$ for SATV and $NR_L: r = -.145$, $NR_U: r = +.145$ for SATQ, based on considerations of practical significance. Another set of MET, conducted against the null hypothesis that the effect is no greater than a common estimate of "ambient noise" indicated that the two correlations that were significant by NHST were also significant with a two-tailed MET. None of the correlations were close enough to 0 to be significant by Equivalence Testing using the same H'_0 s.

The following is how the findings may be referred to in the discussion section, in light of MET and EqT. The passage is taken verbatim from page 985 of Shoda et al. (1990). No deletions were made; additions related to MET and EqT are underlined:

In addition, given the smallness of the sample, the obtained coefficients could very well exaggerate the magnitude of the true association. For example, in the diagnostic condition, the 95% confidence interval for the correlation of preschool delay time with SAT verbal score ranges from .10 to .66, and with SAT quantitative score, the confidence interval ranges from .29 to .76. Thus, using a two-tailed MET, only the correlation with SAT quantitative is significant against the null hypothesis that the true magnitude of the correlation is large enough to predict a meaningful practical difference for a college applicant (at least a 20-point difference in SAT score for every additional 5 minutes in preschool waiting time). Although the magnitude of the observed correlation with SAT verbal was large enough to reject the traditional null hypothesis that the true correlation is 0 as well as the null hypothesis that the true correlation is at or weaker than "ambient noise," it was not large enough to reject the null hypothesis that the true correlation is large enough to make a practical difference

for college applicants. In the other conditions, none of the correlations were significant with MET or EqT against either null hypothesis. Thus, overall, except for the correlation with SAT quantitative in the diagnostic condition, the correlations were not large enough to make a practical difference. Nor was it small enough to reject the null hypothesis that the true correlation is no larger than the smallest effect that still matters. In addition, [t]he value and importance given to SAT scores in our culture make caution essential before generalizing from the present study; at the very least, further replications with other populations, cohorts, and testing conditions seem necessary next steps...

Below is the entirety of the Shoda et al. (1990) abstract, edited in light of MET (deletion shown in ~~strikeout~~, and additions are underlined):

Variations of the self-imposed delay-of-gratification situation in preschool were compared to determine when individual differences in this situation may predict aspects of cognitive and self-regulatory competence and coping in adolescence. Preschool children from a university community participated in experiments that varied features of the self-imposed delay situation. Experimental analyses of the cognitive-attentional processes that affect waiting in this situation helped identify conditions in which delay behavior would be most likely to reflect relevant cognitive and attentional competencies. As hypothesized, in those conditions, coherent patterns of statistically significant correlations were found between seconds of delay time in such conditions in preschool and cognitive and academic competence and ability to cope with frustration and stress in adolescence. In the “diagnostic” condition, the correlation between participants’ preschool delay time and SAT quantitative score was significant against the pre-registered null hypothesis that the true correlation is $r = .145$ (or smaller), which is the minimum that may predict a difference of 20 SAT points for every 5 additional minutes of preschool waiting time; the correlation with the verbal score was not strong enough to suggest such practical difference, although it was still significant against the traditional null hypothesis that the true correlation is 0. In the other conditions, the results did not provide sufficient evidence for any conclusion of potential interest; we cannot conclude they are large enough to make a practical difference, nor can we conclude that they are practically equivalent to 0.

As shown by these examples, reporting MET can concisely add important information about the magnitude of observed effects without sacrificing what is already conveyed using traditional NHST.

Registered Reports

MET solves an important challenge for the promising new publication mechanism called the

registered report. Registered reports are a publishing format which conduct peer review prior to data collection. Publication of the results is provisionally accepted as long as the authors closely adhere to the registered methodology (Center for Open Science, n.d.). It minimizes a variety of questionable research practices, such as publication bias and selective reporting of “significant” or hypothesis-confirming results, while making it possible for reviewers to “help authors to improve the protocol or rationale while it is still possible to make changes” (Chambers, 2019). One important question, however, has remained: How does one define that a prediction has been confirmed? Simply stating that predictions are confirmed when there is a statistically significant effect using NHST is problematic because as long as the true effect is not exactly zero, traditional NHST statistical significance is all but guaranteed as the sample size increases (in a two-tailed test). Considering that the prediction is confirmed when the observed effect was larger than the pre-registered minimum predicted effect is also problematic. This is because the likelihood of such an outcome is higher, rather than lower, with a smaller sample size (hence greater standard error), in effect incentivizing researchers to opt for smaller sample sizes.

The proposed framework integrating MET and EqT avoids both of these problems. For example, suppose in a registered report authors predicted that a new housing policy will reduce the number of people experiencing homelessness in the population by at least 30%. Unlike NHST, unless the reduction observed in a study is greater than 30%, the results will never be statistically significant by MET no matter how large the sample is. Smaller sample size and the resultant larger standard error will increase the probability of observing reductions greater than 30% by chance, but because MET takes the larger standard error into account, the false positive rate (i.e., Type I error) will be capped at the chosen level of α (e.g., 0.05). For an example of MET in a registered report, see King et al. (in press).

Implications for Sample Size

Open science practices increasingly recognize the importance of adequately large sample size. At the same time, it has been noted that when using traditional NHST, increasing the sample size will make it more likely to reject the NHST null hypothesis as long as the true effect is not exactly 0. Results of studies using exceptionally large sample sizes can become unintentionally misleading if one only considers whether the NHST p-value is statistically significant. Unlike traditional NHST, however, when using the framework described above, larger sample sizes do not result in misleading conclusions. For example, if the observed effect falls within the null region, it will *never* be significant by MET regardless of sample size. To see why this is the case, note that as the sample size increases and the “width” of sampling distributions decreases in Figure 1, any observed effect between 0 and NR_U will eventually be in the area shown in orange, hence significant by a traditional NHST. However, the same observed values will never be in the green area (values that are significant by MET), no matter how narrow the sampling distributions become. On the other hand, if the observed effect falls outside of the null region, the larger the sample size is, the more likely that it is significant by MET. Similarly, increasing sample size will increase the likelihood of significant EqT only if the effect is within the null region. Another way to describe the consequence of increasing sample size is that it would decrease the region in which the results are inconclusive, allowing researchers to accurately conclude either the effect to be strong enough for a significant MET, or weak enough to be equivalent to 0 (or in the null region, more generally).

Conclusion

With the present framework, researchers can infer, with greater confidence, if an effect is large enough to matter, too small to be of consequence, or if neither inference is warranted. It

avoids many well-documented problems stemming from testing against the null hypothesis of “no effect,” even in a research context in which rejecting such a null hypothesis is akin to a proverbial attack on a “straw man.” And it provides a language for stating the inference very concisely (e.g., “significant by MET”), for example in abstracts where brevity is paramount, while at the same time requiring researchers to explicitly spell out the basis for the inference in the main body of the paper (i.e., confidence intervals and the range of population values pre-declared as meaningful). It does not result in referring to small but non-zero effects as “significant” if they are too small to matter, even in research involving “big data.” Additionally, when an effect is very small, the proposed framework allows researchers to test if the true effect is in fact too small to matter. This can facilitate the publication of “null” results, mitigating the “file drawer” problem in the publication process (a bias against publishing findings that are not statistically significant). It supports the emerging practice of registered reports by providing a much-needed method for stating, and testing, predictions about the true magnitude of an effect. This avoids incentivizing smaller sample size, as would be the case if the criterion is stated in observed (rather than population) values. In sum, we believe that the framework proposed here can lead to improvements in the quality of data-based inferences in scientific research.

References

- Aron, A., Melinat, E., Aron, E. N., Vallone, R. D., & Bator, R. J. (1997). The experimental generation of interpersonal closeness: A procedure and some preliminary findings. *Personality and Social Psychology Bulletin*, 23(4), 363-377.
- Aviv, R. (2009, October 26). Proctor, Please Get Off the Phone. *New York Times*.
- Bartsch, O'Shea, K. J., Ferguson, M. C., Bottazzi, M. E., Wedlock, P. T., Strych, U., McKinnell, J. A., Siegmund, S. S., Cox, S. N., Hotez, P. J., & Lee, B. Y. (2020). Vaccine Efficacy Needed for a COVID-19 Coronavirus Vaccine to Prevent or Stop an Epidemic as the Sole Intervention. *American Journal of Preventive Medicine*, 59(4), 493–503. <https://doi.org/10.1016/j.amepre.2020.06.011>.
- Benjamin, Daniel J, Berger, James O, Johannesson, Magnus, Nosek, Brian A, Wagenmakers, E-J, Berk, Richard, . . . Johnson, Valen E. (2018). Redefine statistical significance. *Nature Human Behaviour*, 2, 6-10.
- Bauer, P., & Kieser, M. (1996). A unifying approach for confidence intervals and testing of equivalence and difference. *Biometrika*, 83(4), 934-937.
- Chambers, C. (2019, September 10). What's next for Registered Reports? Retrieved September 13, 2020, from <https://www.nature.com/articles/d41586-019-02674-6>
- Center for Open Science (n.d.). Registered Reports: Peer review before results are known to align scientific values and practices.
- Cohen, J. (1994). The earth is round ($p < .05$). *American Psychologist*, 49(12), 997.
- Dunnett, C. W., & Gent, M. (1977). Significance testing to establish equivalence between treatments, with special reference to data in the form of 2 x 2 tables. *Biometrics*, 593-602.
- Embry T.W., Piccirillo J.F. (2020). Minimal Clinically Important Difference Reporting in Randomized Clinical Trials. *JAMA Otolaryngology - Head Neck Surgery*, 146(9), 862–863.

- Fowler, R. L. (1985). Testing for Substantive Significance in Applied Research by Specifying Nonzero Effect Null Hypotheses. *Journal of Applied Psychology*, 70, 215–218.
<https://doi.org/10.1037/0021-9010.70.1.215>
- Greenwald, A. G., Banaji, M. R., & Nosek, B. A. (2015). Statistically small effects of the Implicit Association Test can have societally large effects. *Journal of Personality and Social Psychology*, 108, 553–561.
- 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(2), 175–183.
- Hoekstra, Rink, Morey, Richard, Rouder, Jeffrey N, & Wagenmakers, Eric-Jan. (2014). Robust misinterpretation of confidence intervals. *Psychonomic Bulletin & Review*, 21(5), 1157–1164.
- Ioannidis, J. P. (2005). Why most published research findings are false. *PLoS Medicine*, 2(8), e124.
- Jacobson, N. S., Follette, W. C., & Revenstorf, D. (1984). Psychotherapy outcome research: Methods for reporting variability and evaluating clinical significance. *Behavior Therapy*, 15(4), 336–352.
- Kazdin. (1999). The Meanings and Measurement of Clinical Significance. *Journal of Consulting and Clinical Psychology*, 67(3), 332–339.
- Kerr, N. L. (1998). HARKing: Hypothesizing After the Results are Known. *Personality and Social Psychology Review*, 2(3), 196–217.
- King, M. T. (2011). A point of minimal important difference (MID): a critique of terminology and methods. *Expert Review of Pharmacoeconomics & Outcomes Research*, 11(2), 171–184.

- King, K. M., Feil, M. C., Halvorsen, M. A., Seldin, K., Shoda, Y., Lee, C., Smith, G (in press). Negative urgency as a state-level process. Accepted as a Stage 1 Registered Report by *Journal of Personality*. [10.17605/OSF.IO/9DBEJ](https://doi.org/10.17605/OSF.IO/9DBEJ)
- Kirk, R. E. (1996). Practical significance: A concept whose time has come. *Educational and Psychological Measurement*, 56(5), 746-759.
- Krugman, P. R. (1991). *Geography and Trade*, Cambridge: MIT Press.
- Lakens, D., Scheel, A. M., & Isager, P. M. (2018). Equivalence testing for psychological research: A tutorial. *Advances in Methods and Practices in Psychological Science*, 1(2), 259-269.
- Lykken, D. T. (1968). Statistical significance in psychological research. *Psychological Bulletin*, 70(3p1), 151.
- Maxwell, S. E., Lau, M. Y., & Howard, G. S. (2015). Is psychology suffering from a replication crisis? What does “failure to replicate” really mean? *American Psychologist*, 70(6), 487.
- McGlothlin A.E., Lewis, R.J. (2014). Minimal Clinically Important Difference: Defining What Really Matters to Patients. *JAMA*, 312(13):1342–1343.
- Meehl, P. E. (1967). Theory-testing in psychology and physics: A methodological paradox. *Philosophy of science*, 34(2), 103-115.
- Meehl, P. (1990). Why summaries of research on psychological theories are often uninterpretable. *Psychological Reports*, 66, 195-244.
- Mischel, W., Shoda, Y., & Rodriguez, M. I. (1989). Delay of gratification in children. *Science*, 244(4907), 933-938.
- Murphy, K. R., & Myors, B. (1999). Testing the hypothesis that treatments have negligible effects: Minimum-effect tests in the general linear model. *Journal of Applied Psychology*, 84(2), 234.

- Nadav-Greenberg, Limor, & Joslyn, Susan L. (2009). Uncertainty Forecasts Improve Decision Making Among Nonexperts. *Journal of Cognitive Engineering and Decision Making*, 3(3), 209-227.
- Orben, A., & Lakens, D. (2020). Crud (re)defined. *Advances in Methods and Practices in Psychological Science*, 3, 238–247.
- Oxford University Press, publisher. (2000). Oxford English dictionary. Oxford, England: Oxford University Press.
- Pintea, S. (2010). The relevance of results in clinical research: Statistical, practical, and clinical significance. *Journal of Evidence-Based Psychotherapies*, 10(1), 101.
- Prentice, D. A., & Miller, D. T. (1992). When small effects are impressive. *Psychological Bulletin*, 112(1), 160–164.
- Robinson, R. (1993). Cost-utility analysis. *Bmj*, 307(6908), 859-862.
- Schmidt, F. L., Hunter, J. E., McKenzie, R. C., & Muldrow, T. W. (1979). Impact of valid selection procedures on work-force productivity. *Journal of Applied Psychology*, 64(6), 609.
- Schuirmann, D. J. (1987). A comparison of the two one-sided tests procedure and the power approach for assessing the equivalence of average bioavailability. *Journal of Pharmacokinetics and Biopharmaceutics*, 15(6), 657-680.
- Serlin, R. C., & Lapsley, D. K. (1985). Rationality in psychological research: The good-enough principle. *American Psychologist*. 40(1), 73-83.
- Shoda, Y., Mischel, W., & Peake, P. K. (1990). Predicting adolescent cognitive and self-regulatory competencies from preschool delay of gratification: Identifying diagnostic conditions. *Developmental Psychology*, 26(6), 978–986. <https://doi.org/10.1037/0012-1649.26.6.978>.
- Standing, L., Sproule, R., & Khouzam, N. (1991). Empirical statistics: IV. Illustrating Meehl's

- sixth law of soft psychology: everything correlates with everything. *Psychological Reports*, 69(1), 123-126.
- Szucs, D., & Ioannidis, J. P. (2017). When null hypothesis significance testing is unsuitable for research: a reassessment. *Frontiers in Human Neuroscience*, 11, 390.
- Tryon, W. W. (2001). Evaluating statistical difference, equivalence, and indeterminacy using inferential confidence intervals: An integrated alternative method of conducting null hypothesis statistical tests. *Psychological Methods*, 6(4), 371.
- Wasserstein, R. L., & Lazar, N. A. (2016). The ASA's statement on p-values: context, process, and purpose. *The American Statistician*, 70(2), 129-133.
- Webster, J., & Starbuck, W. H. (1988). Theory building in industrial and organizational psychology. *International Review of Industrial and Organizational Psychology*, 3(1), 93-138.
- Westlake, W. J. (1976). Symmetrical confidence intervals for bioequivalence trials. *Biometrics*, 741-744.

Supplementary Text

Section S1: The Rationale for using Confidence Intervals to conduct MET and EqT

MET and EqT require estimating the probability of obtaining the observed value (or more extreme (MET) or less extreme (EqT)) of a statistic of interest (e.g., differences in means, correlations, proportions), randomly drawn from a population in which the null hypothesis H'_0 holds. Unlike the null hypothesis (H_0) of a traditional NHST, MET and EqT use the TOST (two one-sided significance testing), which typically assumes a certain, pre-specified, non-zero value in the population as null hypotheses (H'_{0s}).

For example, one-tailed METs with the alpha level of X% require determining the 100-X percentile (e.g., 95th percentile) of the sampling distribution of the statistic under the null hypothesis H'_0 that the population value is a certain hypothesized non-zero value. Murphy and Myers (1999) derived such sampling distributions for the F statistic. In order to make MET applicable to a wide variety of statistics, however, the present approach makes use of the fact that the sampling distribution under H'_0 is already taken into account when determining confidence intervals. Thus, as described below, we propose that whenever confidence intervals can be determined, it is possible to “reverse engineer” the critical values for use in MET.

We start by first reviewing the logic for estimating the confidence interval of a quantitative variable:

- (1) Suppose we draw many samples of the same size from a population for which the statistic of interest is H'_0 . (Note we are using H'_0 to refer also to the population value of the statistic under the null hypothesis H'_0). Let δ_{95} be the distance between H'_0 and the 95th percentile of the sampling distribution from this population, and δ_5 be the distance between H'_0 and the 5th percentile of the sampling distribution. Thus, the 95th percentile of this distribution is $H'_0 + \delta_{95}$, and the 5th percentile is $H'_0 - \delta_5$.
- (2) Let $\bar{x}_i$ be the observed statistic of interest in sample i drawn from this population (e.g., the sample mean). Then by definition, $\text{Prob}(\bar{x}_i < H'_0 - \delta_5) = 5\%$. It also follows that $\text{Prob}(\bar{x}_i + \delta_5 < H'_0) = 5\%$.
- (3) Let δ_i^*5 be an unbiased estimate of δ_5 based on the data in sample i . If we can assume, at least approximately, that $\text{Prob}(\delta_5 < \delta_i^*5) = \text{Prob}(\delta_5 > \delta_i^*5) = 0.5$, then because $\text{Prob}(\bar{x}_i + \delta_5 < H'_0) = 5\%$, $\text{Prob}(\bar{x}_i + \delta_i^*5 < H'_0) = 5\%$ in the long run. If we define $\bar{x}_i + \delta_i^*5$ as the upper bound of the confidence interval based on sample i , CI_{upper_i} , then $\text{Prob}(CI_{\text{upper}_i} < H'_0) = 5\%$ in the long run.

That is, the probability that the upper bound of the CI is below H'_0 is 5%. In other words, the probability that the entire CI is below H'_0 is 5%.
- (4) Similarly, let δ_i^*95 be an unbiased estimate of δ_{95} based on the data in sample i . Then because from the definition of δ_{95} , $\text{Prob}(H'_0 + \delta_{95} < \bar{x}_i) = 5\%$. Thus, in the long run, $\text{Prob}(H'_0 + \delta_i^*95 < \bar{x}_i) = \text{Prob}(H'_0 < \bar{x}_i - \delta_i^*95) = 5\%$. If we define $\bar{x}_i - \delta_i^*95$ as the

lower bound of the confidence interval based on sample i , CI_{lower_i} , then $\text{Prob}(H'_0 < CI_{lower_i}) = 5\%$ in the long run. That is, the probability that the entire CI is above H'_0 is 5%.

- (5) Thus, it makes sense to define the CI as $[\bar{x}_i - \delta_i^*_{95}, \bar{x}_i + \delta_i^*_5]$. With this definition of the CI, the CI computed in each sample will be entirely above H'_0 in 5% of the samples randomly drawn from the population where H'_0 is true. Similarly, the CI will be entirely below H'_0 in 5% of the samples.

As a result, the probability that the CI will contain the population value is 90%.

Now let us reverse engineer $\delta_i^*_5$ and $\delta_i^*_{95}$. Suppose for a statistic of interest, there is a well-accepted formula for determining the CI based on a given sample. Then, if we define, as described above, $CI_{upper_i} = \bar{x}_i + \delta_i^*_5$ and $CI_{lower_i} = \bar{x}_i - \delta_i^*_{95}$, where CI_{upper_i} and CI_{lower_i} refer to the upper and lower bounds of the CI by applying the well-accepted formula to the data in sample i , we can “reverse engineer” $\delta_i^*_5$ and $\delta_i^*_{95}$ as follows:

$$\begin{aligned} \delta_i^*_5 &= CI_{upper_i} - \bar{x}_i \\ \text{and} \\ \delta_i^*_{95} &= \bar{x}_i - CI_{lower_i} \end{aligned}$$

By the definition of MET as described in the main text, a one-tailed MET is considered significant when $\bar{x}_i > NR_U + \delta_{95}$. But of course, δ_{95} is unknowable empirically. However, as shown above, in the long run, $\bar{x}_i > NR_U + \delta_{95}$ when $\bar{x}_i > NR_U + \delta_i^*_{95}$, or when $\bar{x}_i - \delta_i^*_{95} > NR_U$. That is, a one-tailed MET is significant when $CI_{lower_i} > NR_U$ (i.e., the entire CI is above NR_U).

Now let's us consider a two-tailed MET. A two tailed test is significant when $\bar{x}_i > NR_U + \delta_{95}$ or $\bar{x}_i < NR_L - \delta_5$. As shown above, the first condition is met when $CI_{lower_i} > NR_U$ (i.e., the entire CI is above NR_U). It can be shown similarly that the latter condition is met when $CI_{upper_i} < NR_L$ (i.e., the entire CI is below NR_L). In other words, a two-tailed MET is significant when the entire CI is either above NR_U or when the entire CI is below NR_L .

Note that when the sampling distribution is symmetrical, the above becomes simplified because $\delta_5 = \delta_{95}$. But we did not make that assumption so that our reasoning applies even to statistics for which the sampling distribution is not symmetric.

Section S2: Peer-Review

An important goal of MET is to encourage a discourse on the magnitude of effects. We believe that a peer-review process will help achieve this goal, while also promoting greater transparency of how authors arrive at the null region for use in MET.

To achieve these goals, it is important to consider the incentives that may influence researchers while setting H'_0 . For example, researchers may be motivated to set a “low bar” with null regions very close to zero out of fear that setting the null region to be too large would make it difficult to reject the MET null hypothesis.

There are multiple points at which NR_L and NR_U values can be reviewed and refined in the research/publication process. First, a researcher preregistering their studies and/or analyses can request peer reviews of the proposed H'_0 from colleagues knowledgeable about the domain of research. The names of those who review the H'_0 , along with their reviews and the H'_0 they suggested, should be included as part of a pre-registration and in the eventual publication (with their consent). Both the authors and the reviewers should be mindful that an unreasonably small or large H'_0 may call their judgment into question.

Second, journal editors could request the reviewers of a manuscript reporting MET to specifically comment on the appropriateness of the H'_0 chosen by the authors. If the editor and/or reviewers determines that the H'_0 is inappropriate, then the editor can ask the authors to repeat their MET with a more appropriate value for H'_0 . In this context, manuscript reviewers who suggest an unreasonably large H'_0 need to consider that their own future (or past) work in the same domain would be held to the same standards.

Sample Review Request

To help facilitate the process of requesting an informal peer-review, below is an example of a request a researcher may send to another researcher, requesting peer-review of their H'_0 values prior to data analysis. A request by editors to reviewers of a submitted manuscript may be generated by editing this sample request.

Hello Dr. Garcia,

We are writing you in order to request your peer-approval of the null hypothesis we have chosen for our study. Because of your expertise in technology and learning in public middle schools in large metropolitan areas, we believe that you would be able to effectively evaluate whether our null hypothesis is reasonable.

Our study examines the effect of giving all middle school students their own tablet. We plan to evaluate the utility of doing so by comparing schools that have recently implemented a program that gives students tablets to schools that have not, and assessing students on a standardized math exam at the end of the year.

Because of the great expense of implementing a tablet program, we used utility analysis to determine the size of the effect below which it would be not worthwhile to implement the proposed program. Specifically, we have decided that a 5% increase in the proportion of students who meet the grade level standard for math is appropriate as the smallest effect of

consequence. This will be used as a basis for statistical tests (please see the attached article describing them) determining if the true comparative effectiveness of the free tablet program is large enough to warrant further consideration, because other interventions such as extra tutoring have been shown to result in at least a 5% increase for the same cost.

If you respond to this message with feedback, your name and a summary of your response will be included in a supplement of our paper. Please let us know if you agree that this is a reasonable basis for specifying the smallest effect of consequence. If you do not agree, we would appreciate any feedback you have on how to make this value more appropriate and/or a suggestion of a better value for consideration as the smallest effect of consequence. We value your expertise and believe that your feedback will improve the quality of our statistical inference.

Sincerely,

Adam Smiley, Jessica Glazier, and Yuichi Shoda

Searchable H'_0 Database

Initially practices such as those above will rely on the scientific judgment of authors, reviewers and editors. However, we believe it should be possible for the scientific community to start a searchable database of H'_0 s and accumulate H'_0 s that have been judged reasonable (along with the names of those who provided the judgment) for a particular research question, research goal, and participant populations. The database will also serve as a repository of H'_0 currently considered the minimum for a meaningful result for each relevant context (e.g., research question and goals, participant populations, etc.), along with a record of how often empirical research has rejected them.

As each scientific discipline strives toward establishing a useful database of H'_0 s, we believe that the public discourse about the H'_0 s itself will encourage transparency and communicate the meaning of the observed effects, beyond the question of whether or not it can arise if there was absolutely no true effect. Though the details for the creation of such a database is beyond the scope of this paper, we believe that the scientific community can create and maintain an effective database (or multiple databases).

Avoiding Conflict of Interest

As with peer-review of journal articles, those chosen to review the H'_0 should avoid conflicts of interest. Differences between superiors and subordinates in the power they hold over the other should be considered when determining who should be asked to review the NR_L and NR_U values, as a subordinate might feel pressure accept the value. For example, it would be a conflict of interest if a graduate student were asked to approve of a H'_0 value proposed by a faculty member in their own graduate program, or if an applicant on the academic job market were asked to review a H'_0 value for a faculty member from a university they are seeking to gain employment from. Avoiding

conflicts of interest is important for the quality of research being conducted, as researchers should be able to review proposed H_0 s without experiencing undue pressure. To this end, researchers should use their best judgment when considering potential reviewers in order to avoid such conflicts of interest, and practice transparency and accountability, for example by including a statement declaring no conflict of interest, and disclosing all potential sources of conflict.