Towards a simplified justification of the chosen sample size

Nick J. Broers ^a and Henry Otgaar ^{b, c}

^a Department of Methodology and Statistics, Faculty of Psychology and Neuroscience, Maastricht University, The Netherlands

^b Department of Clinical Psychological Science, Faculty of Psychology and Neuroscience, Maastricht University, The Netherlands

^c Leuven Institute of Criminology, Faculty of Law, KU Leuven, Belgium

Submitted to *Meta-Psychology*. Participate in open peer review by commenting through [hypothes.is](https://tinyurl.com/mp-submissions) directly on this preprint. The full editorial process of all articles under review at Meta-Psychology can be found following this link:

<https://tinyurl.com/mp-submissions>

You will find this preprint by searching for the first authors name.

Abstract

Since the early work of Cohen (1962) psychological researchers have become aware of the importance of doing a power analysis to ensure that the predicted effect will be detectable with sufficient statistical power. APA guidelines require researchers to provide a justification of the chosen sample size with reference to the expected effect size; an expectation that should be based on previous research. However, we argue that a credible estimate of an expected effect size is only reasonable under two conditions: either the new study forms a direct replication of earlier work or the outcome scale makes use of meaningful and familiar units that allow for the quantification of a minimal effect of psychological interest. In practice neither of these conditions is usually met. We propose a different rationale for a power analysis that will ensure that researchers will be able to justify their sample size as meaningful and adequate.

Keywords: Power analysis, common effect size, direct replication, meaningful outcome scales

In the early '60s, Cohen stressed the need of paying attention to the size of the effect that one hoped to demonstrate, and the sample size required to be able to reject the null hypothesis of no effect with an acceptable probability. In a seminal article, Cohen (1962) introduced the now familiar distinction between a small, a medium and a large effect size. By analyzing 70 volumes of the *Journal of Abnormal and Social Psychology* he showed that, given the reported sample sizes, the mean statistical power for demonstrating a small effect was 0.18 and the mean power for demonstrating a medium effect was only 0.48. Only large effect sizes could be detected with what we would now consider to be a minimal acceptable level of statistical power (on average 0.83). Cohen was also interested in the concept of power analysis (Cohen, 1988) and put forward the definitions of what constitutes a small, a medium and a large effect size (for a mean difference expressed as Cohen's d , respectively 0.2, 0.5 and 0.8). Such was the impact of Cohen's work on methodological discourse in academic psychology, that for many his name now immediately evokes an association with the concepts of effect size and power analysis.

However, the increased awareness of the importance of paying equal attention to minimizing both Type I and Type II errors did not readily translate into an improved practice with regard to power analysis. Twenty-four years later, Sedlmeier and Gigerenzer (1987) conducted a follow-up of Cohen's original study on power and effect size. They sampled 64 experiments that were published in the *Journal of Abnormal Psychology*, a successor of the original *Journal of Abnormal and Social Psychology* that Cohen had used for his study. Remarkably, the mean power of these studies for detecting respectively small, medium and large effects was virtually identical to those found in the original Cohen study, with only an average power of 0.21 for detecting small effects, 0.50 for medium effects and 0.84 for detecting large effects. Sedlmeier and Gigerenzer also listed an overview of 11 other studies on power that appeared in the

seventies, all well over a decade after Cohen's initial article. These other studies also demonstrated no change in the inadequate consideration of power. Sedlmeier and Gigerenzer pointed at the predictable consequence of conducting these low powered studies. The well known publication bias that favors studies that produce significant results risks turning many of the research studies into predictable wastes of effort.

In comparison to the original studies that were reviewed by Cohen, the sample of studies examined by Sedlmeier and Gigerenzer did show a clear methodological development in that many more researchers (50% of all studies) employed an alpha correction to guard against inflation of the Type I error rate in the case of multiple testing. For this subsample of more rigorous studies, the power to detect small, medium, or large effects was even smaller than for the sample as a whole (respectively 0.10, 0.28 and 0.72). The suggestion that Type I errors were a more prominent concern to the researchers than Type II errors seemed borne out by the fact that none of the articles made mention of a power analysis in advance of the study, and only 2 of the 64 studies made any reference to power at all. Sedlmeier and Gigerenzer concluded that studies on power (such as those by Cohen) did not seem to have an effect on the power of studies. They predicted that this unwanted state of affairs would not change until journal editors adopt a no-acceptance policy for studies that do not openly reflect on effect size, power, and sample size.

A major boost to this desired change in acceptance policy came in 1999, when a Task Force on Statistical Inference, commissioned by the American Psychological Association, set out the official guidelines that publications in APA journals would henceforth have to follow (Wilkinson & Task Force on Statistical Inference, 1999). Although the lines devoted to recommendations on power analysis formed just a small fraction of the totality of methodological recommendations,

the guidelines clearly stipulated that all research reports must provide a justification for the chosen sample size. In particular, the guideline requires all researchers to “show how effect-size estimates have been derived from previous research and theory in order to dispel suspicions that they might have been taken from data used in the study or, even worse, constructed to justify a particular sample size.” (Wilkinson & Task Force on Statistical Inference, 1999, p.596).

However, notwithstanding this guideline on the need to justify sample size with reference to an a priori power analysis, a survey by Bakker and Wicherts (2011) of 271 psychological articles revealed that, well over a decade after the new APA guideline on power analysis, only 3% of the authors conformed to this rule. In an additional recent study by Bakker et al. (2016), 291 researchers in psychology were interviewed and it was found that decisions on sample size tend to be taken on the basis of either traditional rules of thumb or of intuition. The latter tends to be misleading, as 95% of the respondents systematically underestimated the sample size needed for the detection of a small effect. Not surprisingly, therefore, studies of power that appeared after the publication of the APA guidelines did not show an improvement in the power of reported studies to detect either small or medium effects. Marszalek et al. (2011) showed that inadequate power has in fact become a more prevalent problem due to the rising popularity of more sophisticated multivariate statistical models such as MANOVA and SEM, which tend to require larger sample sizes than the more traditional univariate models. In a similar vein, Maxwell (2001) showed that power to detect small or medium effects remained low, a problem which he showed to be exacerbated by the growing use of multiple testing corrections. As though no official guideline exists, Bakker et al. (2016) concluded their study with the recommendation that all reported studies should include a formal power analysis.

At least nowadays the routine reporting of observed effect sizes seems like a major

improvement of the practice that Cohen exposed in 1962, when awareness of power and effect size seemed to be virtually absent. However, the present-day familiar appearance of the concept of effect size in the results sections is clearly not matched with an equal attention in the method or discussion sections. For example, it is still uncommon to explicitly justify an expected effect size (and hence, the chosen sample size) in the method section (Bakker et al., 2016), or to reflect on the magnitude of an observed effect size in the discussion section (Hoekstra, 2014). This leads to the question of why explicit attention to power analysis and effect size remains absent or superficial in most psychological research, in spite of all the methodological guidelines and the now ubiquitous awareness of its central importance. We believe that the answer is principally related to two key methodological issues. The first of these pertains to the questionable quantitative meaning of effect size in psychological science, and the second to the fact that standardized effect size conflates the size of the effect with the precision with which it was measured. Although the problems surrounding these two issues are not wholly unfamiliar, there is little awareness amongst academic psychologists that taken together, these two issues make understandable why carrying out a power analysis and discussing effect size remain largely mechanical exercises.

The problem of choosing an expected effect size for power analysis

In his original paper on power and effect size, Cohen (1962) defined a medium effect as an effect large enough to be noticeable by the naked eye. As an example, he mentioned the difference in intelligence between people in professional and people in managerial occupations. Without further justification, Cohen next went on to equate such a noticeable difference with a mean

difference equal to half the population standard deviation. Using a standardized IQ test as the operationalization of intelligence, Cohen translated the medium effect into an IQ difference of 8 points (given a population standard deviation of 16). But in fact, of course, the noticeable difference to the naked eye is not an IQ difference, but a difference in intelligence. Unlike IQ, intelligence is a purely theoretical concept without any definite empirical reference. When we reflect on the question of why we perceive one group of people as more intelligent than another, we realize that it is not easy to reduce the answer to a difference on one or a limited set of dimensions. We might conclude that one group tends to use a larger vocabulary, tends to have more sophisticated views on politics, tends to have views on a greater variety of topics, tends to show more creative problem solving, and so on. We abstract from this diffuse collection of experiences by concluding that by and large, one group seems to be more intelligent than another. By postulating this difference in intelligence we are explaining our diverse collection of observations with reference to some hypothetical underlying unidimensional psychological trait. Up to this point, however, ‘intelligence’ is just a qualitative concept allowing for making ordinal comparisons between individuals and groups of individuals.

If we are willing to accept with Cohen’s statement that a “noticeable difference” would translate into a mean difference equal to half the standard deviation of the measurement instrument, there is of course the problem of capturing the theoretical construct with a reliable and valid measurement instrument. For intelligence, any given IQ test forms an empirical operationalization of this theoretical concept. Because we consider a standard IQ test to be a reliable and valid measure of that theoretical concept, we expect that our experience will translate into a demonstrable IQ difference. However that mean IQ difference is not equal to the difference that we perceived by the naked eye. It is one – to some extent arbitrary – way of

demonstrating that perceived difference.

Any given IQ test forms an empirical operationalization of the theoretical concept of intelligence. For an operationalization to be accepted as a scientifically useful way of capturing a theoretical construct, it needs to be validated. The most conventional validation procedure in psychology is probably the practice known as construct validation, first coined by Cronbach and Meehl (1955), in which a network of empirical relations (usually a pattern of correlations) is compared to a corresponding nomological network of theoretical relationships. If the empirical relationships mirror the theoretical ones, the embedded concepts are considered to have demonstrable construct validity. This validation procedure, which has since been further elaborated by Messick (1989), was argued to be flawed in an article by Borsboom et al. (2004). They contended that, although the two terms are often used interchangeably, validation and validity are two very different concepts.

As a simple example of how construct validation can easily yield misleading results, Borsboom et al. suggested a nomological network in which the theoretical construct of intelligence is positively related to the construct of general knowledge and negatively related to the construct of criminal behavior. If an IQ test is then shown to correlate 0.5 with a general knowledge test and -0.4 with number of months spent in prison, the IQ test is said to have good construct validity. However, there are many concepts that will show the pattern of positive and negative relationships shown in this example. Although in reality most nomological networks will specify a greater number of theoretical relationships than in this example, the problem remains that the theoretical relationships are not articulate enough to give us a clue as to what exactly is being measured with the empirical instrument. Indeed, Borsboom et al. conclude that “The typical network in psychology is in terms of higher and lower correlations between

attributes. Such a loose network is unrestrictive and can be satisfied by an indefinite number of attributes besides the intended one.” Borsboom et al. (2004) argued that an alternative way for establishing whether an IQ test really measures some underlying theoretical attribute known as intelligence is, first of all, to take up a realist position and to assume that there exists some dimension along which we can order our subjects from less to more intelligent, and second, to establish that this underlying reality affects the response behavior of subjects on this IQ test. The only way to realize the latter, is to develop a falsifiable theory that specifies how the theoretical construct affects variation on the empirical instrument. If the response behavior of subjects reflects the predicted behavior on the basis of the theory, then per definition the measurement instrument is valid for measuring the theoretical construct. Conversely, “if one does not have an idea how the attribute variations produce variations in measurement outcomes, one cannot have a clue as to whether the test measures what it should measure”. (Borsboom et al., 2004, p. 1068). Another way of phrasing this latter statement would be: If your theory does not allow you to make quantitative predictions about the expression of behavior of individuals on the measurement scale, then any numbers that are produced on the scale cannot provide information on the verisimilitude of the underlying response theory. However, the verbal theories of psychology only produce ordinal predictions, not quantitative predictions. The only theory that allows you to make quantitative predictions of the sort that would be required, is a formal theory.

Although less familiar to the general public, formal theories are no anomaly in the science of psychology. Especially in cognitive psychology, there has long been a subcurrent of mathematical models aiming at quantitative prediction of learning indicators such as response times and retrieval rates. Ever since the early work of Estes on learning (Estes, 1950, 1955), which permitted quantitative predictions of response times in operant conditioning experiments,

progress has been made in formalized theories of such diverse research domains as memory (e.g. Eich, 1982; Murdock, 1989; Brainerd, 2020), linguistics (e.g. Dennis & Humphreys, 2001), reasoning (e.g. Johnson-Laird & Byrne, 1991) and decision theory (most notably, the prospect theory of Kahneman and Tversky (1979), which earned the former the Nobel Prize in Economics). In developmental psychology, the mental stage theory of Piaget has been formalized using catastrophe theory (Van der Maas & Molenaar, 1992).

Notwithstanding the progress that has been made in the development of formal mathematical and computational models of behavior, it is probably fair to state that such progress is to be found at the fringes of psychological science. As Simonsohn (2015) remarks: “Psychological theories are almost exclusively qualitative rather than quantitative – predicting the sign rather than the magnitude of effects” (, p.560). That the science of psychology remains predominantly attached to verbal theories is probably related to Meehl’s observation that psychological concepts tend to be “open” concepts, an “openness arising from the indefinite extensibility of our provisional list of operational indicators of the construct” (Meehl, 1978, p.815). Formal theories require a strict demarcation of the domain of observations to which the theory pertains, as well as a precise definition of theoretical concepts and the functional rules that govern their interaction (Hintzman, 1991). This methodological rigour will inevitably rob the intended theoretical concepts of their open nature, and the gain in clarity and precision will come at the expense of simplification. The resulting formal theory may allow for precise quantitative predictions, but the external validity is likely to be limited.

In contrast, verbal theories are more fuzzy in nature but have the advantage of retaining the richness of its open concepts. Simplification only enters at the stage of operationalization, not in the theory itself. The richness of the verbal theory poses problems as soon as we wish to elicit

predicted effects and measure the predicted consequences. If our theory predicts that certain stimuli presented in certain situations will cause some effect then there exists an indefinite variety of such stimuli and situations that we could actually choose for our specific research design, a freedom of choice that Meehl (1978) called the situation-taxonomy problem. Likewise, the registration of the effect can also be realized in countless different ways, which Meehl (1978) referred to as the response class problem. The richness of the verbal theory transcends a specific choice of stimuli or outcome measure, which implies that there is no clear functional relationship between theory and any specific operationalization, or between different operationalizations of the same concept.

It should therefore come as no surprise that Santor et al. (2006, referenced in Flake & Fried, 2019, January 17) identified 280 different scales for measuring depression, or that the relationship between loneliness and extraversion varied from $r = -0.44$ to $r = -0.19$, depending on which scales were used to measure these constructs (Buecker et al., 2019, referenced in Flake & Fried, 2019, January 17). According to Broers (2021) such variability in reported effect size of the same relationship should not be considered as disturbing, because at the theoretical level there is no quantitative relationship, but only an ordinal one. This follows from the fact that the theory used to predict the relationship is a verbal theory, not a formal theory, and consequently only ordinal relationships can be predicted with the theory. Depending on the measurement scales that were used, the observed effect sizes can have real meaning, but that meaning is not theoretically explained and cannot be generalized beyond the specific research design that was used. There is therefore nothing inherently problematic about different research designs for studying the same relationship producing very different effect size outcomes. The only consistency that the theory predicts is an ordinal one. If the theory predicted a negative

relationship between loneliness and extraversion, then the conceptual replications that produced effect sizes ranging from $r = -0.44$ to $r = -0.19$ have consistently supported this prediction.

10 different studies on the role that confirmation bias plays in sustaining racial prejudice will in practice make use of different stimulus material and outcome scales. This means that the quantitative results can only be compared by converting the unstandardized effect sizes into standardized effect sizes. Finding that the r values of these 10 studies ranged from 0.16 to 0.24 provides an illusion of quantitative information in that the only reason for combining these transformed values is a purely semantic association: all studies focused on the effect of confirmation bias on the sustainment of racial prejudice. But whether the individual magnitudes, let alone the averaged magnitude, conveys any quantitative information on the real existence of an identical psychological effect that we may summarize as confirmation bias is a matter of belief for which no empirical justification has been given. A researcher who aims to conduct another study on the effect of confirmation bias on the sustainment of racial prejudice, using yet other stimulus material and outcome scales, cannot therefore use the average effect size found in the meta-analysis as a theoretically based expectation of the magnitude that will be found in the upcoming study. There is no such theoretical basis for expecting any particular quantity or magnitude, although it would be legitimate to claim that this order of magnitude has been *commonly* found in related studies and to use this as a justification for the use of a particular sample size. In other words: rather than claiming to expect an effect size on theoretical grounds, we would be claiming to expect an effect size on purely empirical grounds. This is the effect size that is commonly found, but there is no way of explaining why this common effect size should have this particular magnitude.

The conflation problem of standardized effect size

It is not possible to use the effect size for determining sample size without specifying the precision with which the effect size will be established (Lenth, 2001). In other words, for any power analysis we will be making use of a standardized effect size. In order to compute a standardized effect we need information both on the absolute size of the effect (i.e., the unstandardized effect size) as well as on the precision of its measurement, expressed by the residual variance of the outcome, which in the case of Cohen's d is established by taking a pooled estimate of the standard deviation. By combining these two sources of information we obtain an effect size estimate that no longer has any dimension. It can therefore seemingly be used for communication about some theoretical effect that was determined with a multitude of different operationalizations across various conceptual replications of the effect. It is not much use to know that the effect of dissonance reduction consisted of a dissonance induced treatment group being willing to pay 10 dollars more than a none dissonant control group, if you are next planning to examine the effect of dissonance reduction on some weight watch scheme, where the outcome is weight measured in pounds. However, by dividing the dollar difference by the pooled estimate of the standard deviation of the dollar scale we might find that the standardized effect size was 0.50. This could be used to transfer the information on effect size to any other study on cognitive dissonance using a different outcome scale. Without additional information, we might compute the desired sample size for the pound difference study on the basis of an expected effect size of 0.50.

However, as Lenth (2001, p.191) claims: "If only a standard effect is sought without regard to how this relates to an absolute effect, the sample size calculation is just pretense". The problem

is that the 0.50 in the dollar study could be reflective of a medium sized absolute effect established with great precision, but it could just as well be a large absolute effect established with a not so very reliable instrument. Funder and Ozer (2019) provide a telling example for the perils of being wrong-footed by a standardized effect size. Using r as an estimate of standardized effect size, we might find that there is a near perfect correlation between age and salary. Many people would quickly conclude that this means that growing old will be accompanied by a sizeable increase in income. However, if all employees will always get about one extra dollar in annual income with each successive year in age, this near perfect consistency would translate in an r close to 1, even though the size of the age effect on annual income is almost zero. Alternatively, salaries that on average show a 20 percent increase with each successive year in age would turn up as a large standardized effect size even in the presence of considerable variability in employee earnings. A large standardized effect size on its own therefore cannot inform on effect size per se, it tells us something about effect size, relative to the precision with which the effect was measured. Apart from improving the reliability of the measurement instrument, that precision can be increased by including suitable covariates in the analysis model and by the choice of a within-subject design instead of a between-subjects design (Mayo, 1978). In assessing the effect of a within-subjects factor participants form their own controls, thereby considerably reducing the unexplained variance and thus increasing the precision with which the experimental effect can be assessed. In addition, increasing the number of levels of an experimental factor tends to increase the proportion of variance that can be explained in the outcome (O'Grady, 1982) These methodological choices do nothing to alter the absolute effect size, but they can dramatically increase the standardized effect size by affecting the precision.

Although this conflation of absolute effect size with precision is well known, many authors

continue to defend the use of the standardized effect size as a means of informing on the strength or even meaningfulness of the effect. Rosenthal and Rubin (1982), for instance, introduced the Binomial Effect Size Display (BESD) to convey a more intuitive grasp on the meaningfulness of an effect size that is expressed as r . Suppose that a study on the effectiveness of a new treatment for preventing heart attacks throws up a correlation of 0.20 between Treatment (conventional vs new) and outcome (heart attack vs no heart attack). Although it would be correct to conclude that the Treatment explains just 4% of the variance in the outcome, this statement would totally muddy the success of the new treatment. By creating a BESD we get a proper look on the consequences of the new vs the conventional treatment, and it turns out that of the 100 patients who received the new treatment 40 developed a heart attack, against 60 patients from the group that received conventional therapy. But where BESD helps us to counteract the simplistic conclusion that an r of 0.20 cannot amount to a meaningful effect, it does nothing to solve the conflation problem. The binomial display rather tells us that what looks like a small effect, could be conceived as psychologically or medically important. What this imaginative display of the effect does is to take away our simplistic preconceptions of what would constitute a small or a large effect. However, the ambiguity surrounding the interpretation of what at first glance looks like a minor effect merely serves to underline the fuzziness that is inherent in the reporting of a standardized effect size.

Another strategy for increasing the interpretability of effect sizes is suggested by Funder and Ozer (2019), who propose the idea of benchmarking. By meta-analyzing effects with a high intuitive meaning to the general public, we can obtain benchmarks for the relative importance of standardized effects. Some examples they mention are an r of 0.10 for the relationship between credibility of communicators and their perceived persuasiveness, and an r of -0.17 between

conscientiousness and sexual infidelity. The importance of such relationships can be gauged by relating these r values to familiar relationships like the fact that males tend to weigh more than females ($r = 0.26$), or the fact that taller persons tend to weigh heavier ($r = 0.44$). Whereas such benchmarks undoubtedly help to convey some intuition regarding the relative magnitude of an effect (an effect of X on Y may be perceived as not very impressive when portrayed as explaining 0.067% of the variance in Y , but may be appraised as meaningful when it becomes clear that it compares to the relationship between Sex and Weight), the fact remains that the same r value may or may not permit an interpretation in terms of meaningfulness. As Blanton and Jaccard (2006) explained, a meaningless difference does not miraculously become meaningful after division by a constant. Whereas an unstandardized effect will be meaningful when expressed in scale units that have a wide familiarity or obvious relevance (such as units on popular IQ scale, or physical metrics like kilo's or seconds), in many psychological studies the outcome scale has arbitrary units that cannot be meaningfully interpreted. The mean difference on a rating scale has no intrinsic meaning, and dividing it by the standard deviation produces a standardized version of that meaningless difference. It can tell us a lot about precision, but nothing about importance.

Using the standardized size of some psychological effect found in a previous study for predicting the standardized effect of a conceptual replication of that study would imply that one does not so much expect the absolute effect size of the previous study to be replicated in the new study, nor that one expects the precision of the original outcome scale to be identical to the precision of the new outcome scale, but rather that one expects the ratio of the absolute effect to the precision with which it is measured to be the same. In the practice of psychological research, there is no theoretical basis for such an expectation.

Considering effect size as a supplement to a significant test result can only be truly important and revealing in its unstandardized form. Kirk (2007) uses an artificial example to demonstrate this. He mentions the case of an imaginary researcher who tests the efficacy of a new drug on the mental functioning of Alzheimer patients. It turns out that the mean difference in IQ is not significant, but the standardized effect, expressed as Hedges' g (a variant of Cohen's d) is 0.89, which by conventional benchmarks would be considered a large effect. This contradictory information is then supplemented by the 95% confidence interval of the mean difference in IQ units, in other words, a confidence interval for the unstandardized effect in meaningful units. This confidence interval runs from a near zero -0.05 to a 24.5 mean IQ difference. Whereas the supposedly large effect reflected by the value of Hedges g cannot tell us much about the psychological relevance of the effect, the interval in unstandardized units clearly does provide us with information on the practical importance of the established effect. It is likely that the new drug has a positive effect on IQ and this effect could be substantial. Further research into the promise of this new drug is warranted.

Power analysis for psychological science

Doing a power analysis to compute the required sample size for a study requires specification of the alpha level (conventionally 5%), the desired power of the test (usually 0.80 or higher) and the expected effect size. Since 1999 the APA requires all researchers to justify their sample size by specifying the expected effect size on the basis of theory and prior research. However, for studies that are based on verbal theories an expected or required effect size for power analysis can only be objectively specified under two conditions. First, when we are dealing with a direct

replication in which the new study aims to make use of the exact same stimulus configuration of an original study, as well as the same outcome scale. Especially in the wake of the replication crisis there have been a notable number of multicenter direct replication studies (e.g. Eerland et al., 2016; Wagenmakers et al., 2016; Bouwmeester et al., 2017; McCarthy et al., 2018; Verschuere et al., 2018). In the case of a direct replication it is perfectly legitimate to use the observed standardized effect size of the original study as the expected value for the standardized effect size of the replication. A second condition under which we will be able to prespecify an effect size for the purpose of determining a required sample size, especially to be found in applied research, is when we are dealing with an outcome scale with meaningful and interpretable measurement units and we are able to specify in advance what difference would amount to a meaningful difference, i.e., a psychologically important effect. As in all research that is based on a verbal psychological theory, there will not be a way of predicting a quantitative effect size, but it is objectively possible to specify the minimal effect size that would be of interest. A study by Burriss et al. (2015, discussed in Lakens, Scheel and Isager, 2018) provides an example of this. Burriss et al theorized that females would display a slight increase of reddishness in the face during the fertile phase of their ovulatory cycle. The idea was that this increase of reddishness was a product of evolution and had the effect of making them more attractive to males. Although it could be experimentally established that the increase of reddishness did indeed objectively occur, the validity of the hypothesis obviously depended on the ability of men to be able to detect the difference in reddishness. Since the detection threshold could be experimentally determined, this provided a smallest effect size of interest in the form of a minimal increase in redness, below which the evolution hypothesis could not be realistic.

Current practice in psychological science conducts a power analysis for an upcoming study on

the basis of standardized effect sizes that were found in previous studies on the same research topic. These previous studies usually employed different stimulus material and different outcome scales and had no theoretical basis for expecting the quantitative result that emerged. Suppose a simple theory stipulates that testosterone gives rise to aggression and that therefore men and women will differ in their reaction to frustration, with men being expected to show more acts indicative of aggression. There are countless ways in which we can manipulate the level of frustration and likewise many different ways in which we can establish acts of aggression. Any given operationalization of frustration and aggression will enable us to establish a quantitative outcome. We may find that the odds for a male to display an act of physical aggression were twice as large as the odds for a female doing so. Or we may find that on average males produced 7 acts of aggression during a specified time interval, against only 3 for women. Alternatively, we may register that males on average took five minutes longer than females to calm down after having been provoked into anger. It is important to realize that while each of these different ways of testing the hypothesized relationship between sex and aggressive responding led to a quantitative outcome that could be used for corroborating the theory, none of these quantitative outcomes was actually predicted by the theory. Because the theory underlying the prediction is a verbal theory, the only outcome that could actually be predicted was the direction of the relationship.

Nonetheless, current practice in doing a power analysis acts as if the quantitative outcomes have actual meaning at the theoretical level, in other words as if at theoretical level there exists some quantitative dimension 'aggressive responding to frustration' that all these different studies tap into and inform about. However, the theory does not assume such a quantitative dimension and the research studies do not analyze the data with the objective of establishing empirical proof

for the existence of such a quantitative dimension. In fact, we might say that the numbers produced in any individual research are little more than an expedient for expressing the hypothesized ordinal relationship. The numbers provide a means of supporting the contention of a specific order relationship. By using a quantitative outcome in any individual study for predicting a quantitative outcome in a new study that employs different stimulus material and different outcomes, we are assuming that the quantitative information found in the original study transcends the confines of the research design and is reflective of some common underlying theoretical reality. In basing our prediction of the effect size in the new study on this untested assumption, we are acting purely on faith. When using the outcome of a meta-analysis where a whole collection of quantitative outcomes are averaged to reveal a “true” effect size, we may be deluding ourselves even harder. At least, as long as we do not have a theory at our disposal that actually predicted a collection of quantitative outcomes across a diversity of different studies, or have not gathered empirical support for such a formal theory.

So what should we do when we wish to conduct a power analysis? If we cannot hope to formulate the effect size that we expect, we must at least formulate the effect size that we want to demonstrate. But on what basis? Admitting that we usually cannot objectively specify an expected effect size for our study does not mean that we should revert to the practice of pre-Cohenian times, when empirical studies were carried out without any explicit consideration of the power to detect an hypothesized effect. Examining a predicted mean difference by contrasting two groups with 10 participants each will be a predictable waste of effort because almost certainly no significant difference will be found. Using two groups of 200 participants each will obviously be more promising, but may bring an unnecessary waste of time and resources, because the expected effect might well have been demonstrated with a considerably

smaller sample. How to find a reasonable optimum, without deluding oneself into thinking that an expected effect size can be derived from previous research in the field?

Cohen was adamant that any research should always be preceded by a power analysis, but he did concede that in practice many researchers do not really know what effect size their study is supposed to yield. In those cases, claimed Cohen, the most reasonable strategy would be to do a power analysis with the presumption of a medium effect. Per definition, this was taken to be an effect noticeable to the naked idea, so even without knowledge derived from prior research, one could defensibly claim to suspect the existence of such an effect. “Those cases”, we contend, are the typical cases in psychological research. So instead of all the pretense of insight into an effect size outcome that we have reason to expect, we can spare ourselves a lot of unnecessary difficulty by simply carrying out power analyses with the default values specified by the software at our disposal. Typically, programs such as G*Power3 (Faul et al, 2007) fill in a medium effect size, a significance level of .05 and a power of 0.95 as default values. Of these, only the prespecified power level can be justifiably altered to suit the researchers’ needs.

Over the years, some discussion has arisen over Cohen’s definition of a medium effect as corresponding to a Cohen’s d of 0.5 or an r of 0.3. (Actually, for a comparison of two groups, the medium Cohen’s d value corresponds with an r of 0.24). For examples, based on an overview of a large collection of meta-analysis studies on social psychological effects, Richard et al (2001) revealed that the mean effect size (expressed in r) for social psychology was only 0.21. Likewise, Gignac and Szodorai (2016) looked into the typical effect size values found in individual differences research and reported a median correlation value of 0.19. Hemphill (2003) examined 380 meta-analytic reviews of either psychological assessment or psychological treatment, and found that of all effect sizes (r) reported, the middle third typically showed correlations in the

range of 0.20 to 0.30. Morris and Fritz (2013) tabulated a cumulative frequency distribution of effect sizes found in memory research, and found the median Cohen's d value to correspond to 0.57 (roughly equal to what Cohen considered a medium effect) and the median partial eta squared value to be equal to 0.18 (way above the .06 value that Cohen considered a medium effect expressed in eta squared). Such overviews of effect size values that are common in a field may give researchers a reason for altering the expected effect size to conform to the magnitude that is typically found in the field. Not because this quantity has some bearing on a "true" effect size that one is interested in, but rather because this quantity is accepted by the research community as being worthy of interest. This is attested to by the fact that many journals have published research results containing this common effect size value. To make sure that it can be revealed in the new study with sufficient level of power, the necessary sample size can be computed. In fact, the prime reason for doing a power analysis should be for the researcher to demonstrate that he or she has made use of a reasonable sample size. If the average effect size in terms of Cohen's d is 0.43 rather than 0.5 in social psychology (corresponding with an r of 0.21), then it follows that for a study on the difference between two independent groups we will be needing two samples of at least 115 each to obtain a power of 0.90 for a two-sided test, rather than two samples of at least 86 that would be needed if the common effect size in the field were 0.5. There is no other basis for expecting this effect size than the simple fact that, when looking around in the field, this is the magnitude we are most likely to encounter – for whatever reason. This means that a researcher who decides to make use of a sample size that seems inadequate to meet a desirable power level for the detection of the common effect size is thereby failing to meet a methodological standard for good research. This could be a reason for outright rejection

of a manuscript showing promising results on the basis of an inadequate sample size, with the request of redoing the study with a more appropriate sample size.

Reporting

Just like there can be only two scenario's where one has a substantive basis for specifying an expected effect size – either when doing a direct replication or when conducting an applied study with use of a meaningful outcome scale – so likewise one can have only two reasons for elaborating on the observed effect size in the discussion section of a manuscript. In a direct replication any change in the magnitude of the original study is worthy of serious reflection, and in the applied study we would obviously like to know whether the effect is psychologically meaningful enough to render it of practical value. Axsom and Cooper (1985), for instance, applied cognitive dissonance theory to the treatment of overweight. They showed that having to exert effort in a pseudo-treatment that was purportedly meant to reduce overweight, and having consciously consented to exert that effort to that end, resulted in a durable mean weight loss of 8 pounds. A difference that could be objectively considered as psychologically meaningful.

In the usual theoretical studies where we had no basis for predicting a particular quantitative outcome, the observed magnitude can likewise only be interpreted in terms of psychological meaningfulness or importance in those cases where the outcome scale itself was psychologically meaningful. This precludes discussion of effects that were established with rating scales, which are inherently meaningless. However, although the employment of meaningful outcome scales gives scope for a richer interpretation of the study result as potentially of practical meaning or not, it does not inform on the proposed psychological mechanism underlying that study result. The theory that proposed the mechanism is inherently qualitative, in the sense of permitting no more than ordinal predictions, and therefore any outcome that could be meaningfully interpreted

as quantitative cannot tell us anything about the theory or its proposed constructs. At best, the quantitative result could be considered as an explanandum worthy of theoretical analysis, leading to the future development of a formal theory that does permit quantitative predictions.

The present convention of reporting standardized effect sizes in the results section which are then not further considered in the discussion section reflects the illusory undertaking of focusing on numbers that do not have any substantive meaning. Researchers know that APA guidelines require them to report effect sizes for the benefit of possible future meta-analyses. But because those effect sizes have no clear quantitative meaning they cannot be related to other such effect sizes found in similar studies, other than on a basis of purely semantic association. Numbers that are treated as numbers can reveal all sorts of delusory patterns, not unlike the meaningful patterns that one often sees in a collection of random numbers. We would do well to treat the numbers for what they are: pseudo-quantitative expedients that help us to pass a verdict over the credibility of a verbal theory that gave rise to the prediction of ordinal patterns. It are those ordinal patterns, not the meaningless magnitudes, that should be reflected on in the discussion section. This is precisely what happens in practice, for good reason.

Conflict of Interest and Funding

We have no conflicts of interest to declare. This study did not receive any specific funding.

Author Contributions

Nick J. Broers and Henry Otgaar both contributed to the concept and writing of this article. Nick Broers drafted the original manuscript and Henry Otgaar provided critical revisions. Both

authors approved the final version of the manuscript for submission.

Open Science Practices

This article earned no Open Science Badges because it forms a commentary on general research practice and does not contain any data or data analyses.

References

- Axsom, D. & Cooper, J. (1985). Cognitive dissonance and psychotherapy: The role of effort justification in inducing weight loss. *Journal of Experimental Social Psychology*, 21, 149-160. [http://dx.doi.org/10.1016/0022-1031\(85\)90012-5](http://dx.doi.org/10.1016/0022-1031(85)90012-5)
- Bakker, M. & Wicherts, J. M. (2011). The (mis)reporting of statistical results in psychology journals. *Behavior Research Methods*, 43, 666–678. <http://dx.doi:10.3758/s13428-011-0089-5>
- Bakker, M., Hartgerink, C. H. J., Wicherts, J. M. & Van der Maas, H. L. J. (2016). Researchers' intuitions about power in psychological research. *Psychological Science*, 27, 1069-1077. <http://dx.doi/10.1177/0956797616647519>
- Blanton, H. & Jaccard, J. (2006). Arbitrary metrics in psychology. *American Psychologist*, 61, 27-41. <http://dx.doi.org/10.1037/0003-066X.61.1.27>
- Borsboom, D., Mellenbergh, G. J., & van Heerden, J. (2004) The concept of validity. *Psychological Review*, 111, 1061-1071. <http://dx.10.1037/0033-295X.111.4.1061>

Bouwmeester, S., Verkoeijen, P. P. J. L., Aczel, B., Barbosa, F., Bègue, L., Branas-Garza, P., ..., Wollbrant, C. E. (2017). Registered replication report: Rand, Greene & Nowak (2012).

Perspectives on Psychological Science, 12(3), 527-542.

<https://doi.org/10.1177/1745691617693624>

Broers, N. J. (2021). When the Numbers Do Not Add Up: The Practical Limits of Stochastologicals for Soft Psychology. *Perspectives on Psychological Science*. 16(4), 698 706. doi: 10.1177/1745691620970557. PMID: 33482071.

Buecker, S., Maes, M., Denissen, J., & Luhmann, M. (2019). Loneliness and the Big Five personality traits: A meta-analysis. *European Journal of Personality*.

<https://doi.org/10.22201/fq.18708404e.2004.3.66178>

Brainerd, C. J. (2020). Deep distortion. *Memory*. <https://doi.org/10.1080/09658211.2020.1844756>

Burriss, R. P., Troscianko, J., Lovell, P. G., Fulford, A. J. C., Stevens, M., Quigley, R., Payne, J., Saxton, T. & Rowland, H. M. (2015). Changes in women's facial skin color over the ovulatory cycle are not detectable by the human visual system. *PLOS ONE*, 10, Article e0130093. doi:10.1371/journal.pone.0130093

Cohen, J. (1962). The statistical power of abnormal-social psychological research: A review. *Journal of Abnormal and Social Psychology*, 65, 145-153. <http://dx.doi.org/10.1037/h0045186>

Cohen, J. (1988). *Statistical power analysis for the behavioral sciences*. (2nd ed.). New York: Academic Press.

- Cronbach, L. J. & Meehl, P. E. (1955). Construct validity in psychological tests. *Psychological Bulletin*, 52, 281-302. PMID: 13245896
- Dennis, S. & Humphreys, M. S. (2001). A context noise model of episodic word recognition. *Psychological Review*, 108, 452-477. PMID: 11381837
- Eerland, A., Sherrill, A. M., Magliano, J. P., Zwaan, R. A., Arnal, J. D., Aucoin, P., . . . Prenoveau, J. M. (2016). Registered replication report: Hart & Albarracín (2011). *Perspectives on Psychological Science*, 11, 158–171.
<https://doi.org/10.1177/1745691615605826>
- Eich, J. M. (1982). A composite holographic associative recall model. *Psychological Review*, 89, 627-661. <http://psycnet.apa.org/doi/10.1037/0033-295X.89.6.627>
- Estes, W. K. (1950). Toward a statistical theory of learning. *Psychological Review*, 57, 97-107. <http://psycnet.apa.org/doi/10.1037/h0058559>
- Estes, W. K. (1955). Statistical theory of spontaneous recovery and regression. *Psychological Review*, 62, 145-154. <http://psycnet.apa.org/doi/10.1037/h0048509>
- Faul, F., Erdfelder, E., Lang, A.-G., & Buchner, A. (2007). G*Power 3: A flexible statistical power analysis program for the social, behavioral, and biomedical sciences. *Behavior Research Methods*, 39, 175–191.
- Flake, J. K., & Fried, E. I. (2019, January 17). Measurement schmeasurement: Questionable measurement practices and how to avoid them. <https://doi.org/10.31234/osf.io/hs7wm>

Funder, D. C. & Ozer, D. J. (2019). Evaluating effect size in psychological research: Sense and nonsense. *Advances in Methods and Practices in Psychological Science*, 2, 156-168.

<https://doi.org/10.1177/2515245919847202>

Gignac, G. E., & Szodorai, E. T. (2016). Effect size guidelines for individual differences researchers. *Personality and Individual Differences*, 102, 74–78.

<https://doi.org/10.1016/j.paid.2016.06.069>

Hemphill, J. F. (2003). Interpreting the magnitudes of correlation coefficients. *American Psychologist*, 58, 78–79. <https://doi.org/10.1037/0003-066X.58.1.78>

Hoekstra, R. (2014). The interpretation of effect size in published articles. *Proceedings of the Ninth International Conference on the Teaching of Statistics*.

https://iaseweb.org/icots/9/proceedings/pdfs/ICOTS9_6B3_HOEKSTRA.Pdf

Hintzman, D. L. (1991). Why are formal models useful in psychology? In W. E. Hockley and S. Lewandowsky (Eds.) *Relating theory and data: Essays on human memory in honor of Bennet B. Murdock*. (pp. 3956). Hillsdale, NJ: Lawrence Erlbaum Associates.

Johnson-Laird, P. N. & Byrne, R. M. J. (1991). *Deduction*, Hillsdale, NJ: Lawrence Erlbaum Associates.

Kahneman, D. & Tversky, A. (1979). Prospect theory: An analysis of decision under risk. *Econometrica*, 47, 263-292. DOI: 10.2307/1914185

Kirk, R. E. (2007). Effect magnitude: A different focus. *Journal of Statistical Planning and Inference*, 137, 1634-1646.

- 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, 259–269. <https://doi.org/10.1177/2515245918770963>
- Lenth, R. V. (2001). Some practical guidelines for effective sample size determination. *The American Statistician*, 55, 187-193.
- Marszalek, J. M., Barber, C., Kohlhart, J. & Holmes, C. B. (2011). Sample size in psychological research over the past 30 years. *Perceptual Motor Skills*, 112, 331-348. <http://dx.doi.org/10.2466/03.11.PMS.112.2.331-348>
- Maxwell, S. E. (2001). The persistence of underpowered studies in psychological research: Causes, consequences, and remedies. *Psychological Methods*, 9, 147-163. <http://dx.doi.org/10.1037/1082-989X.9.2.147>
- Mayo, R. J. (1978). Statistical considerations in analyzing the results of a collection of experiments. *Behavioral and Brain Sciences*, 1, 400-401. <https://doi.org/10.1017/S0140525X00075683>
- McCarthy, R. J., Skowronski, J. J., Verschuere, B., Meijer, E. H., Jim, A., Hoogesteyn, K., ..., Yildiz, E. Registered replication report: Srull & Weijer (1979). *Advances in Methods and Practices of Psychological Science*, 1, 321-336. <https://doi.org/10.1177/2515245918777487>
- Meehl, P. E. (1978). Theoretical risks and tabular asterisks: Sir Karl, Sir Ronald, and the slow progress of soft psychology. *Journal of Consulting and Clinical Psychology*, 46, 806-834. <https://doi.org/10.1016/j.appsy.2004.02.001>

Messick, S. (1989). Validity. In R.L. Linn (Ed.), *Educational measurement* (3rd ed., pp. 13-103). New York: Macmillan.

Morris, P. E., & Fritz, C. O. (2013). Effect sizes in memory research. *Memory*, 21(7), 832-842.

Murdock, B. B. (1989). Learning in a distributed memory model. In C. Izawa (Ed.), *Current issues in cognitive processes: The Tulane Symposium of Cognition*. Hillsdale, N.J.: Erlbaum.

O'Grady, K. E. (1982). Measures of explained variance: Cautions and limitations. *Psychological Bulletin*, 92, 766–777. <https://doi.org/10.1037/0033-2909.92.3.766>

Richard, F. D., Bond Jr., C. F., & Stokes-Zoota, J. J. (2003). One hundred years of social psychology quantitatively described. *Review of General Psychology*, 7, 331-363.
<http://dx.doi.org/10.1037/1089-2680.7.4.331>

Rosenthal, R., & Rubin, D. B. (1982). A simple, general purpose display of magnitude of experimental effect. *Journal of Educational Psychology*, 74, 166-169.
<https://doi.org/10.1037/0022-0663.74.2.166>

Santor, D. A., Gregus, M., & Welch, A. (2006). Eight decades of measurement in depression. *Measurement*, 4, 135–155. https://doi.org/http://www.tandfonline.com/doi/abs/10.1207/s15366359mea0403_1

Sedlmeier, P. & Gigerenzer, G. (1987). Do studies of statistical power have an effect on the power of studies? *Psychological Bulletin*, 105, 309-316. doi: 10.1037/0033-2909.105.2.309

- Simonsohn, U. (2015). Small telescopes: Detectability and the evaluation of replication results. *Psychological Science*, 26, 559-569. PMID: 25800521
- Van der Maas, H. L. J., & Molenaar, P. C. M. (1992). Stagewise cognitive development: an application of catastrophe theory. *Psychological Review*, 99, 395-417. PMID: 1502272
- Verschuere, B., Meijer, E. H., Jim, A., Hoogesteyn, K., Orthey, R., McCarthy, R. J., ..., Yildiz, E. (2018). Registered replication report: Mazar, Amir & Arieli (2008). *Advances in Methods and Practices in Psychological Science*, 1, 299-317.
<https://doi.org/10.1177/2515245918781032>
- Wagenmakers, E.-J., Beek, T., Dijkhoff, L., Gronau, Q. F., Acosta, A., Adams, R. B., Jr., . . . Zwaan, R. A. (2016). Registered replication report: Strack, Martin, & Stepper (1988). *Perspectives on Psychological Science*, 11, 917-928.
<https://doi.org/10.1177/1745691616674458>
- Wilkinson, L. & Taskforce on Statistical Inference (1999). Statistical methods in psychology journals: Guidelines and explanations. *American Psychologist*, 54, 594-604. <http://dx.doi.org/10.1037/0003-066X.54.8.594>