Use, Knowledge and Misconceptions of Effect Sizes in Psychology

Elizabeth Collins^{1*} & Roger Watt¹

1. Department of Psychology, University of Stirling, UK

* Corresponding author: Elizabeth Collins, Department of Psychology, University of Stirling,
Stirling, FK9 4LA, elizabeth.collins@stir.ac.uk

This pre-print has been submitted for publication but has not been peer-reviewed.

Abstract

An effect size is a value which quantifies the size of a phenomenon of interest. They have not historically been widely used in psychology, although this has improved over time with ongoing encouragement from organisations such as the American Psychological Association. In this study, 247 psychologists were surveyed to examine effect size use, barriers to use, and effect size knowledge measured using quantitative and qualitative questions. We found high levels of effect size use, with only 9.3% of participants never using effect sizes. One third of participants scored full marks on a true-false knowledge test, but many misconceptions were apparent when participants were asked to define the term effect size. Engaging in open science appears to have some positive influence, but does not consistently relate to increased knowledge. Barriers to use of effect sizes include lack of incentives and lack of knowledge. We argue that successful statistical reform requires a combination of institutional incentives (the ‘stick’ approach), such as journal regulations for statistics reporting, but also personal incentives (the ‘carrot’ approach), such as demonstrating how effect sizes can easily provide deeper insight into a researcher’s data.

Keywords

effect size, new statistics, open science, statistical reform

Introduction

Any quantitative study in Psychology benefits from being given two types of evaluation. First, the effect-size question: how big is the effect of an intervention, experimental manipulation, or observational variable? Second, the uncertainty of the effect-size question: how reliable is the finding? Null hypothesis significance testing (NHST), used throughout quantitative studies in Psychology, conflates the two into a single evaluation which compares the sample effect size with its standard error to produce a dichotomous outcome. However, the resulting p-value is a reductionist statistic that gives no detail about any specific hypotheses aside from the null, and provides no information about the size or direction of any effects that have been found (Cohen, 1990; Cumming, 2014).

Calculating and reporting effect sizes provides multiple benefits. Primarily, it quickly translates evidence into a useful measure of size, beyond NHST's limited response to "can we reject the null hypothesis?", providing much more information to an interested reader. Findings can therefore be more easily interpreted with regards to contextual importance and impact, sometimes referred to as 'practical significance'. Effect sizes also encourage researchers to critically evaluate results generated using NHST, as a statistically significant finding may not necessarily translate to an effect size large enough to be of theoretical or practical interest. Furthermore, reporting effect sizes facilitates the computation of meta-analyses (combining estimates from multiple studies), which are crucial for increasing the accuracy of knowledge about any particular phenomenon (Cumming, 2014).

Recommendations to adopt effect sizes, either in addition to NHST, or to replace NHST altogether, are not new (Cohen, 1965; Cohen, 1990). Towards the end of the 20th century, the American Psychological Association (APA) Task Force on Statistical Inference issued a report stating that academic psychologists should "*always present effect sizes for primary outcomes... We must stress again that reporting and interpreting effect sizes in the context of previously reported effects is essential to good research*" (Wilkinson, 1999). The APA subsequently published the 5th edition of their Publication Manual with similar text, strongly encouraging the use of effect sizes (APA, 2001). In the wake of psychology's replication crisis, many researchers have re-emphasised the need for statistical reform, which includes using effect sizes to support or replace NHST (Cumming, 2014; Wasserstein, Schirm and

Lazar, 2019). Reviews indicate that effect size reporting appears to be increasing over time (Peng, Chen, Chiang and Chiang, 2013).

It must also be acknowledged that encouraging, or enforcing, the reporting of effect sizes alone is not sufficient to improve how data is evaluated. If effect-sizes are of importance for evaluating studies, it is necessary that they are known, understood properly and evaluated with insight. There is a considerable difference between the statement that “A has a significant effect on B” and the statement that “variation in A accounts for X% of the variation in B”. Reviews of the literature indicate that effect sizes, even when reported, are often not examined within the discussion section of the same article, meaning that they are frequently not used to draw conclusions about the topic under investigation. When mentioned, often evaluations are often brief and frequently rely on standardised benchmarks for what constitutes a small, medium, or large effect to summarise results (Sun, Pan and Wang, 2010; Funder and Ozer, 2019). This approach to evaluation has been widely criticised as often standardised benchmarks are unsuited to the typical range of effect sizes found in any specific psychological context (Bakker et al., 2019; Correll et al., 2020). Psychological researchers need to have a sturdy foundation of knowledge about effect sizes both to enable them to successfully interpret and compare effects, but also to compute and report them correctly in the first instance (Flora, 2020). Without sufficient knowledge, the risk remains that effect sizes will eventually be misused or misunderstood, which has been a longstanding issue with NHST (Haller and Krauss, 2002; Lyu, Peng and Hu, 2018).

Psychological Researchers and Effect Sizes

Despite many reviews of effect size reporting in the published literature, only a small amount of research has directly surveyed psychologists about their use and knowledge of effect sizes. One survey of 472 Spanish academic psychologists found that 87.1% claimed to know about effect sizes, although only 68.4% could name an effect size statistic (Badenes-Ribera et al., 2016). While 40.7% of participants reported using effect sizes ‘quite often’, nearly a quarter of the sample reported using effect sizes never, or very infrequently. A subsequent replication, which surveyed 159 Italian psychologists, found even less familiarity with effect sizes: 81.8% of participants claimed to know about them, but only 44.7% could name an effect size statistic (Badenes-Ribera et al., 2018). Within this study, only 35.9% of participants reported using effect sizes ‘quite often’, with nearly one third of the sample using

effect sizes never or ‘very little’. As yet, there is little evidence that psychologists have a firm understanding of effect sizes.

The Present Study

Our intention was to further explore the use and knowledge in a global sample of psychologists, considering potential barriers to effect size use and measuring knowledge through a combination of quantitative and qualitative questions. If knowledge is high, then encouraging or enforcing effect size reporting and evaluation is a sensible idea. The risk, however, is that the reporting of effect sizes becomes another unreflective step in the process of data analysis. Barriers to effect size use must also be identified in order to successfully change behaviour in psychology. In light of psychology’s replication crisis, we have also examined whether engagement in any kind of psychological reform (e.g. the open science movement) has had a positive influence on effect size use and knowledge.

Methods

This study received ethical approval from the General University Ethics Panel at the University of Stirling.

Sampling

Due to the exploratory nature of this study, a power analysis was not deemed suitable for sample size planning. Instead, the intention was to collect as many responses as possible. Participants were recruited using convenience sampling during a four week period in 2020, with the study advertised on Twitter and promoted through academic mailing lists. All self-identifying psychologists involved in psychological research in any location were eligible to take part in this study, including PhD and Masters-level students.

Materials

Data was collected via a Qualtrics (Qualtrics, Provo, UT) questionnaire, found on the OSF page for this manuscript (Collins, 2021). Participants were asked whether they currently calculate effect sizes in their research (*yes*, *not always*, or *no*); and were asked to explain their own reasons for not using effect sizes if they responded with *no* or *not always*. Participants who reported calculating effect sizes were also asked about their software preferences.

To examine knowledge of effect sizes, all participants were first asked to define the term ‘effect size’ in their own words. Participants were then presented with a series of five statements about effect sizes and asked to label them as true or false, based on similar research into p-value and confidence interval knowledge (e.g. Haller and Krauss, 2002; Hoekstra et al., 2014).

In line with criticisms of the True-False testing method (see Miller and Ulrich, 2016), this survey added a third “I don’t know” response option for each statement, to reduce false correct responses as a result of guessing. Participants were also presented with a combination of correct and incorrect statements. The statements presented to participants were as follows:

1. Effect sizes express the magnitude of the influence one variable has on another
2. A larger sample size will result in a larger (stronger) effect size

3. If you are doing high quality research, your aim is to find the largest effect sizes possible
4. A statistically significant p-value corresponds to a medium or large effect size
5. A small effect size indicates that the null hypothesis should fail to be rejected

Statement one provides a definition of effect size which is true, while statements two to five present false information about effect sizes regarding the relationship between effect size and sample size, effect size and research quality, and effect size and statistical significance.

Statement development was influenced by past research including the survey of Badenes-Ribera et al. (2016), Kelley and Preacher's tutorial paper (2012), and criticisms raised by Morris (2020).

Participants were subsequently asked to report the values they perceived to be small, medium and large effect sizes in either Pearson's r , Cohen's d , or another measure of their choosing. They were given free-text space to provide any explanations or comments about their decision. This data was primarily collected to assess how many participants still use Cohen's benchmarks (Cohen, 1992) to measure effect sizes, and identify any evidence of critical thinking about context.

The survey ended with minimal demographic questions concerning academic job role, location, field of psychology, and psychological research reform engagement, measured with the question "*are you actively engaged with any elements of the current movement towards improving psychological science, such as replication, pre-registration, new statistics, producing open data, the Society for the Improvement of Psychological Science (SIPS), or anything similar?*". For brevity, this variable is titled "open science?" throughout this manuscript.

Procedure

All participants gave informed consent via an electronic check box. Participants were first asked to define effect size, so that later questions did not influence their knowledge of the concept. They were subsequently asked about their use and knowledge of effect sizes, as detailed in Materials.

Data Handling

The data is available on the OSF here: osf.io/5hu8m (Collins, 2021). 326 participants started the survey, but 75 responses were removed because the participant did not move beyond the consent page. Four further participants were removed due to only answering a random selection of questions and providing no complete responses to any section of the survey. The resulting data set consisted of data from 247 participants.

Data Analysis

Quantitative data was analysed descriptively, calculating frequencies, averages and associated statistics, and examining demographic differences. The estimated effect size values were cleaned to remove 9 outliers, all of which were instances of using overly large integers such as 2, 5 & 8 as measures of Cohen's d .

Qualitative data examining barriers to effect size use was analysed using inductive content analysis (Drisko and Maschi, 2016), treating each question response as an individual meaning unit. Frequencies for each category were then counted to identify common patterns.

Definitions of effect size were first deductively analysed to categorise each definition as acceptable, incorrect, or shows some understanding. An 'acceptable' definition was defined as one which correctly and clearly defined effect size as a measure of the magnitude of a phenomenon of interest, without being too restrictive (e.g. explaining effect size as the size of a difference between two groups), or being too vague (e.g. "size of an effect"). Kelley and Preacher's (2012) definition of effect size, "*a quantitative reflection of the magnitude of some phenomenon*" (p. 137) and Watt and Collins' (2019) definition, "*a numerical description of the strength of the relationship [between variables]*" (p. 70) were used to guide the categorisation process.

Incorrect definitions were characterised by describing other concepts (e.g. p-values) or making clear mistakes such as explaining that an effect size is a measure of reliability. The category '*shows some understanding*' captures all definitions that demonstrated some awareness of an effect size while either being too specific or lacking sufficient clarity to be categorised as *acceptable*. This category was further divided into 'overly specific' and

‘vague’ definitions, and a further inductive content analysis was used to identify patterns of misconceptions.

The categorisation process was conducted in full by EC, and a random 20% of the data set was independently rated by RW. 88% inter-rater agreement was identified after independent categorisation, which was revised through discussion to 100% agreement.

Participants

The final data set consisted of 247 participants from a variety of academic roles, locations, and fields of psychology. Participants were predominantly based in the UK ($n = 153$) or the USA ($n = 34$), with the remaining participants coming from Germany (6), Ireland (6), Australia (5), Canada (4), Italy (3), Sweden (3), Serbia (2), Spain (2), and one each from Belgium, Finland, France, The Netherlands, New Zealand and Singapore.

A wide range of fields were represented within the sample, the largest of which were health psychology ($n = 51$), cognitive psychology (25), neuropsychology (23), social psychology (22), clinical psychology (20), forensic psychology (16) and developmental psychology (10). The other fields reported by one or more participants were: affective, applied, autism, behavioural, community, comparative, counselling, cyber, decision, educational, environmental, evolutionary, experimental, legal, methodological, multidisciplinary, occupational, personality, psychobiology, psycholinguistics, quantitative, and sports psychology.

Table 1 presents the breakdown of job roles and open science engagement within the sample. More than half of participants were early career researchers at the PhD or postdoctoral level. Nearly half of participants reported being engaged in some way with psychological reform (46.96%).

Table 1.

Job Role and Open Science Engagement, $n=247$

Demographic Categories	Frequency	
	<i>N</i>	%
Job Role^a		
MSc Student	16	6.48

USE, KNOWLEDGE AND MISCONCEPTIONS OF EFFECT SIZES

Research or Teaching Assistant (no PhD)	10	4.05
PhD Student or equivalent trainee	98	39.68
Postdoctoral Researcher	36	14.58
Lecturer or Senior Lecturer	45	18.20
Professor	16	6.48
Other ^b	5	2.02
Prefer not to say	3	1.22
Missing	18	7.29
Open Science?		
Yes	116	46.96
No	104	42.11
Prefer not to say	9	3.64
Missing	18	7.29

^a Participants were presented with a list of UK academic job roles, with accompanying suggestions on equivalent foreign positions to allow them to self-identify their position, or the option to select 'other' and report their title if unsure. ^b The category "other" includes one clinical psychologist, one occupational psychologist and part-time lecturer, one psychology practitioner, one research and training lead, and one participant who did not supply further details.

Results

Using Effect Sizes

A high proportion of the sample reported some use of effect sizes in their research, as shown in Table 2. Only 9.3% of the sample indicated that they do not ever calculate effect sizes.

Table 2.

Frequencies of Reported Effect Size Use, with Open Science Comparisons

Using Effect Sizes	Yes (%)	Not always (%)	No (%)	Missing (%)
Sample	175 (70.9)	48 (19.4)	23 (9.3)	1 (0.4)
Open Science?				
Yes	95 (54.3)	17 (35.4)	4 (17.4)	0
No	65 (37.1)	23 (47.9)	16 (69.6)	0
Prefer not to say	5 (2.9)	3 (6.3)	1 (4.3)	0
Missing	10 (5.7)	5 (10.4)	2 (8.7)	1 (100)

The most frequently used software for effect size calculation was SPSS, reported 93 times. SPSS was closely followed by R (mentioned 79 times). Many other options were listed by participants, in order of frequency: online calculators (21 times), Excel (16), by-hand calculations (14), JASP (12), G*Power (11), Jamovi (9), Matlab (7), MPlus (7), Stata (7), ESCI (2), SAS (2), BrawStats (1), Minitab (1) and RevMan (1).

Not Using Effect Sizes

Content analysis identified four explanations for not personally using effect sizes: lack of knowledge, lack of incentive, conditional use, and not doing quantitative research. All four categories were relevant to both *never* and *not always* using effect sizes. The categories and sub-categories are presented in Table 3 along with frequencies and illustrative quotes. The two most common explanations were a lack of personal knowledge, and not using effect sizes out of habit.

Table 3.

Why Psychological Researchers Don't, or Don't Always, Use Effect Sizes

USE, KNOWLEDGE AND MISCONCEPTIONS OF EFFECT SIZES

Categories	Frequency			Illustrative Quotes
	No	Not Always	Total	
<i>Lack of Knowledge</i>				
Personal knowledge	3	8	11	<i>“I'm not totally sure of my statistical knowledge”</i>
Collective knowledge of complex statistics	2	5	7	<i>“Debate over how to do this in linear regression (e.g. some people have told me beta is an effect size but others have said this isn't true)”</i>
<i>Lack of Incentive</i>				
Bad habit	4	7	11	<i>“Mostly habit, but I am starting to calculate effect sizes more frequently”</i>
Not required	3	1	4	<i>“...feel as though it is more work and often not required. Although, some journals now require effect sizes to be calculated”</i>
<i>Conditional Use</i>				
Depends on audience	1	3	4	<i>“Depends who I am writing for - academic journal, lay audience and what the purpose of the research is”</i>
Depends on project	1	3	4	<i>“Depends on what I am observing. If it is an intervention I would use.”</i>
Informed statistical preferences	-	5	5	<i>“Sometimes they're not the most effective way to communicate the information and I instead use graphics and Bayes factors”</i>
				<i>“Calculate it if the p value is significant”</i>
				-
Other condition	-	3	3	<i>“I tend to include them when they are automatically generated in statistical analyses”</i>
				-
				<i>“as directed by supervisor”</i>
<i>Not doing quantitative research</i>	4	1	5	<i>“My research is predominantly qualitative. But I am familiar with the term”</i>

True-False Knowledge Measure

The mean score for the five true-false effect size statements was 3.99, or four statements out of five correctly answered. Knowledge score increased with job ranking, and participants engaging in some aspect of research reform scored higher than participants who did not (4.27 versus 3.76), as shown in Table 4.

Table 4.*Mean and Median True-False Knowledge Scores for the Sample*

Participants	Mean (SD) Score	Median Score
Full Sample	3.99 (1.15)	4
Effect Size Use		
Yes	4.2 (0.95)	4
Not Always	3.49 (1.47)	4
No	3.41 (1.33)	4
Job Role		
Professor	4.69 (0.48)	5
Lecturer	4.56 (0.66)	5
Postdoctoral Researcher	4.17 (0.91)	4
PhD Student	3.88 (1.20)	4
Research/Teaching Assistant*	3.8 (1.229)	4
MSc Student	2.75 (1.29)	3
Other	3.4 (1.52)	4
Prefer not to say	2.33 (1.16)	3
Open Science?		
Yes	4.27 (1.07)	5
No	3.76 (1.19)	4
Prefer not to say	3.33 (0.87)	4

*without a PhD

It can be seen in Figure 1 that more than one third of participants (40.89%) scored full marks. A further 28.75% of participants scored four out of five, with most other participants scoring two or three. Only two participants scored zero.

Figure 1.*Distribution of True-False Statement Scores out of Five*

The breakdown of responses for each statement is shown in Table 5. For each of the five statements, at least three quarters of the sample chose the correct answer. Statements 1 and 2 had the highest proportions of incorrect responses, at 18.22% and 15.39% respectively. Uncertainty was highest for Statement 5, with 12.96% of participants opting for “I don’t know”.

Table 5.

Response Frequencies for True-False Knowledge Statements.

Statement	True		False		I don’t know		Missing	
	n	%	n	%	n	%	n	%
1 Effect sizes express the magnitude of the influence one variable has on another ^a	176	<i>71.26</i>	45	<i>18.22</i>	13	<i>5.26</i>	13	<i>5.26</i>
2 A larger sample size will result in a larger (stronger) effect size	38	<i>15.39</i>	176	<i>71.26</i>	20	<i>8.10</i>	13	<i>5.26</i>
3 If you are doing high quality research, your aim is to find the largest effect sizes possible	25	<i>10.12</i>	184	<i>83.81</i>	25	<i>7.69</i>	13	<i>5.26</i>
4 A statistically significant p-value corresponds to a medium or large effect size	8	<i>3.24</i>	207	<i>83.81</i>	19	<i>7.69</i>	13	<i>5.26</i>
5 A small effect size indicates that the null hypothesis should fail to be rejected	11	<i>4.45</i>	191	<i>77.33</i>	32	<i>12.96</i>	13	<i>5.26</i>

^a Statement 1 was the only statement where the correct response was “true”.

Effect Size Estimation

50 participants made their estimations of effect size in Pearson’s r , compared to 145 participants opting for Cohen’s d . Seven participants chose ‘other’, with two opting for partial eta squared, one each for Cohen’s f^2 , Cramer’s V and ‘neuroimaging bits’, and two providing no further information.

Participant estimations of effect size clearly aligned with Cohen’s historic benchmarks (Cohen, 1988) for both popular measures of effect size, as shown in Tables 6 and 7. There

was no noticeable difference between estimations made by participants engaging in some aspect of research reform, compared to those not engaging with it. Note that all reported means have standard deviations between 0.01 and 0.18.

Table 6

Participant Estimates of Effect Size (ES) Using Pearson's r

	Small ES			Medium ES			Large ES		
	Mean	Median	Mode	Mean	Median	Mode	Mean	Median	Mode
Sample	0.13	0.1	0.1	0.34	0.3	0.3	0.57	0.5	0.5
Open Science?									
Yes	0.13	0.1	0.1	0.33	0.3	0.3	0.53	0.5	0.5
No	0.14	0.1	0.2	0.37	0.4	0.5	0.64	0.7	0.7

Table 7

Participant Estimates of Effect Size (ES) Using Cohen's d

	Small ES			Medium ES			Large ES		
	Mean	Median	Mode	Mean	Median	Mode	Mean	Median	Mode
Sample	0.19	0.2	0.2	0.45	0.5	0.5	0.72	0.8	0.8
Open Science?									
Yes	0.2	0.2	0.2	0.45	0.5	0.5	0.73	0.8	0.8
No	0.18	0.2	0.2	0.45	0.5	0.5	0.71	0.8	0.8

Despite estimations themselves not differing between open science (OS) and non-OS participants, follow-up questions established that 11 non-OS participants expressed uncertainty about correctly remembering effect size values compared to three OS participants. Additionally, 22 OS participants highlighted the importance of context or critiqued relying on benchmarks compared to seven non-OS. However, six other OS participants commented that they follow Cohen's guidelines, demonstrating that reform engagement does not always translate to critical thinking about effect size interpretation.

Definitions of Effect Size

Participants were asked to define effect size in their own words. 28.9% of definitions were categorised as *acceptable*, 56.2% as *shows some understanding*, and 14.9% as *incorrect*. Rater 1 (EC) then further divided all definitions categorised as *shows some understanding* into two sub-categories: *vague definitions* and *overly specific definitions*. The final results are shown in Table 8.

Acceptable definitions

70 definitions were categorised as *acceptable* based on the criteria defined in the Methods section. Examples of acceptable definitions include: “*effect size measures the strength of the relationship between variables. different metrics exist, e.g. depending on the kind of variables involved*” and “*magnitude of effect for a phenomena, e.g. extent of difference between two groups*”. Note that the second example mentions a difference between two groups, but due to the obvious context of this being provided as an example, the definition was not categorised as overly specific.

Table 8.

Categories and Frequencies for Definitions of Effect Size.

		Acceptable	Shows Some Understanding		Incorrect
		<i>n (%)</i>	<i>Vague</i> <i>n (%)</i>	<i>Overly Specific</i> <i>n (%)</i>	<i>n (%)</i>
Sample		70 (28.9)	56 (23.1)	80 (33.1)	36 (14.9)
Open Science?					
	Yes	44 (62.9)	26 (46.4)	32 (40.0)	13 (36.1)
	No	23 (32.9)	23 (41.1)	38 (47.5)	18 (50.0)
	Prefer not to say	1 (1.4)	2 (3.6)	3 (3.8)	3 (8.3)
	Missing	2 (2.9)	5 (8.9)	7 (8.8)	2 (5.6)

Incorrect definitions

36 definitions were categorised as *incorrect*. Multiple participants confused effect sizes and other concepts, such as statistical power – e.g. “[*an effect size is*] *the power of your analysis*” – and degrees of freedom: “*an effect size is degrees of freedom*”. Other participants incorrectly defined effect sizes as related to NHST, such as “*the value that adds meaningfulness to a p-value*”, or the more muddled “*how statistically significant something is in relation to participant numbers*”. Some participants also demonstrated a misunderstanding of how effect size should be interpreted. For example, one participant added the following incorrect detail to their definition: “*if effect size is high, then you are more likely to be able to make reliable predictions than if it is low, in which case it is useless for making predictions and has little practical use as a result*”. Participants who gave incorrect definitions also mentioned a variety of other ideas, such as effect size being a binary indicator of an intervention working or not, effect sizes being numbers which compare populations, and effect sizes being numbers which show how important an effect is. While the latter idea

sounds plausible, it is important to remember that an effect size alone is simply a measure of magnitude: it may allow the importance of the effect to be evaluated, but does not do this on its own.

Definitions categorised as shows some understanding

Any definition that demonstrated an awareness of effect sizes but did not meet the criteria to be classed as *acceptable* was put into this category ($n=136$), which was then further divided into *vague* or *overly specific* definitions.

Vague definitions were characterised by a lack of detail, typically offering short descriptions such as “*effect size is the size of a statistical effect*”, “*magnitude of a phenomenon*”, or even more simply, “*measure of the effect*”. Most of these definitions correctly noted that an effect size relates to magnitude, but did not provide enough clarity or detail to be categorised as acceptable.

Contrastingly, overly specific definitions were typically correct, but defined effect size as only existing within a specific context. The recurrent ideas in these definitions were analysed inductively with content analysis, the results of which are detailed in Table 9. Misconceptions were grouped together based on common ideas, forming three categories: *effect sizes measure a difference*, *effect sizes measure experiments*, and *effect sizes are connected to statistical testing*. Two further repeated misconceptions were identified, along with four unique ideas. Frequencies and examples are all provided in Table 9. Note that some definitions specified more than one misconception, and so frequencies add up to more than the total number of definitions provided.

Table 9.

Misconceptions About Effect Size With Frequencies of Occurrence and Example Definitions

Misconception	Frequency	Example Quote
Effect sizes measure a difference	53	
Difference (general)	8	<i>“the difference in parameters one is interested in”</i>
Group difference (general)	13	<i>“Magnitude of differences between groups”</i>

USE, KNOWLEDGE AND MISCONCEPTIONS OF EFFECT SIZES

Two-group difference	20	<i>“Effect sizes reflect the magnitude of any given difference between two groups”</i>
Difference between conditions	8	<i>“Effect size means how much of a difference there is between conditions”</i>
Difference between means	4	<i>“a standardized estimate of the difference between your means”</i>
Effect sizes measure experiments	11	
Relates to experiments	4	<i>“A numerical measure of the scale/level of an effect observed in an experiment”</i>
Relates to interventions	2	<i>“The size of the change in the outcome following an intervention”</i>
Relates to manipulations	3	<i>“A measure of the strength of the effect of your manipulation on outcome variables”</i>
Relates to treatments	2	<i>“the proportion/percentage of variance due to the treatment”</i>
Effect sizes are connected to statistical testing	10	
Linked to inferential tests	8	<i>“The magnitude of the effect of a statistical test”</i>
Linked to significant p-values	2	<i>“When a significant association has been found between a dependant & independent variable, the effect size gives an indication of how much (how strongly) the independent variable changes the dependent variable”</i>
Further misconceptions	21	
Effect sizes are standardised metrics	15	<i>“tells us the magnitude of a reported effect using a standardized way”</i>
Effect sizes measure the strength of an interaction	2	<i>“Explains the strength of an interaction”</i>
Other ^a	4	-

^a Four unique ideas were categorised as other: effect sizes are unstandardised values, Cohen’s *d* is the only effect size; effect sizes demonstrate the real world effect, effect sizes provide meaning, and effect sizes measure clinical significance.

A wide range of misconceptions about what an effect size is were demonstrated by participants. Several misconceptions indicate that knowledge may be constrained by personal experience or discipline, such as participants who described an effect size as being related to treatment, interventions, manipulations and experiments. A few participants demonstrated the misconception that effect sizes are simply an addition to null hypothesis significance testing,

whether that be as a connection to particular inferential tests, or more specifically believing that effect sizes should be used to explain statistically significant findings.

Discussion

This study collected data on effect size use and knowledge from 247 psychological researchers from various locations and fields of psychology. More than 90% of participants indicating using effect sizes in this sample, suggesting that effect size use may be increasing in comparison to earlier research (Badenes-Ribera et al., 2018). Only 9.3% of participants reported never calculating effect sizes. We identified two key barriers to effect size use: lack of knowledge, and lack of incentive. In part, lack of knowledge was linked to the perceived complexity of some effect size calculations, and the lack of consensus on what to report for complex analyses. This will continue to be a constraint until more accessible software and materials become available. However, lack of knowledge is also partly the responsibility of individuals. In an academic culture dominated by NHST (Ioannidis, 2019), it is unsurprising that many individuals have not educated themselves more on other statistics, because they are not necessary for academic success in traditional circles. It is likely that a lack of incentive is both a barrier to effect size use at the collective level (i.e. journals), and a personal barrier to learning about effect sizes in the first place.

Our data indicates that basic knowledge of effect sizes is quite high in this sample, with more than one third of participants scoring full marks on the true-false statements presented in this study. Nonetheless, many participants still incorrectly perceive that effect sizes increase with sample size, and nearly 20% of the sample incorrectly marked a true definition of effect size as false, demonstrating that there is still plenty of progress to be made regarding effect size knowledge. The misconceptions that emerged in definitions of effect size (shown in Table 9) imply that many researchers have adopted the small bits of information they need in order to continue with their research, but do not necessarily understand the variety of effect size measures that exist or how they might be used.

It must be acknowledged that our sample does not generalise well to the wider psychology community. Firstly, the sample is skewed towards Western countries, with no African or South American representation, and only one participant in Asia (Singapore). It is known that many governments in these regions favour research productivity over quality, meaning that

there are often fewer incentives or opportunities to engage with new statistics (e.g. Onie, 2020), and we expect that effect size use and knowledge is lower in these communities. Secondly, a large proportion of the sample identified as being engaged with psychological reform in some way, which is unlikely to be an accurate representation of the psychology community as a whole (both within and outside Western countries). There appears to be a slight advantage of belonging to the open science community (note that we use this term broadly, to capture any participants who self-reported engaging in some kind of psychological reform behaviour). Average knowledge scores were higher, and two thirds of ‘acceptable’ rated definitions came from participants in this demographic group. We suggest that many participants in this group have more exposure to effect sizes through processes such as power analysis and meta-analysis, both championed as part of psychological reform post-replication crisis. Despite this, the overall differences between open science and non-open science participants are small, and indeed are hardly seen at all in the estimation data, where several open science participants commented that they rely on Cohen’s guidelines.

Outside the open science world, effect sizes are often simply a chore, especially with the ongoing prioritisation of NHST. Journal regulations will continue to increase their use by providing the ‘stick’ of a carrot-and-stick approach, which arguably is better than nothing at all, but this will not translate to effect sizes being understood. In our survey, two thirds of participants made mistakes on the true-false knowledge statements, and misconceptions of effect size are still high. We expect that there is less knowledge and higher levels of misconceptions in the wider psychology community, due to the demographic limitations previously discussed. Effect sizes risk becoming a check-box exercise, open to miscalculation, misuse and misunderstanding.

We believe that true statistical reform occurs when psychologists are shown the value of the statistical tools that they are using. NHST use, $p < 0.05$ in particular, comes with the increased likelihood of publication in many journals. As an example of how there can be intrinsic value to effect-sizes that extends beyond power analysis and meta-analysis, consider a variant on the binomial effect size display (BESD) (Randolph and Edmondson, 2005). For a scatter plot, the BESD splits the data into 4 quadrants by the median value in each dimension and counts how many participants are in each quadrant. BESD then reports this as the proportions, so that 50:50 would correspond to no effect and 100:0 would correspond to a full effect. For a

group design (such as a t-test), the groups define the split in one dimension and the overall median of the continuous measure defines the other.

Examples of the BESD concept are shown in Figure 2: the left panel shows two examples of no effect (both show a BESD of 50:50), compared to the right-side, which shows two graphs both with a BESD of 33:67 (equivalent to approximately $r = 0.4$).

Figure 2.

Examples of the BESD Concept, Which Compares the Distribution of Data Across Quadrants

Note. Graphs on the left demonstrate no effect (50:50), while both graphs on the right demonstrate an effect of approximately $r = 0.4$ (33:67).

We suggest the BESD is even more useful when modified to report the number of participants instead of proportions, using the total sample size as the denominator. For instance, no effect would be reported as $0/n$. Looking at either graph in the right-hand panel in Figure 2, we can see there are 10 more participants in the bottom left than the top left quadrant, and there are 10 more participants in the top right compared to the bottom right quadrant. This translates to 20 ($10 + 10$) displaced participants, or $20/60$ (where $n = 60$). This

is much more intuitive than the original BESD of 33:67, quickly offering insight into both ‘how much’ of an effect exists within the data, and how much data that there is in total.

Conclusions

In conclusion, our sample demonstrates high effect size use, and generally good knowledge of effect sizes. However, this must all be viewed within the constraints of the sampling bias towards Western psychologists and those engaged in some way with psychological reform. We believe that effect size use and understanding is therefore lower in the wider psychology population than we have found here. Journal requirements are a powerful way to influence behaviour, but more work needs to be done to encourage psychologists to expand their knowledge of effect sizes. Psychologists need both institutional and personal incentives to widen their statistical knowledge. We have commented on the value of the BESD as one example of how effect sizes can be valuable for researchers, and we also suggest that journal author guidelines make space for effect size tutorials and educational materials, as this makes them easily accessible for researchers.

References

- American Psychological Association. (2001). Publication manual of the American Psychological Association (5th ed.). Washington, DC: American Psychological Association.
- Badenes-Ribera, L., Frias-Navarro, D., Pascual-Soler, M., & Monterde-i-Bort, H. (2016). Knowledge level of effect size statistics, confidence intervals and meta-analysis in Spanish academic psychologists. *Psicothema*, 28(4), 448-456.
- Badenes-Ribera, L., Frias-Navarro, D., Iotti, N. O., Bonilla-Campos, A., & Longobardi, C. (2018). Perceived statistical knowledge level and self-reported statistical practice among academic psychologists. *Frontiers in psychology*, 9, 996.
- Bakker, A., Cai, J., English, L., Kaiser, G., Mesa, V., & Van Dooren, W. (2019). Beyond small, medium, or large: Points of consideration when interpreting effect sizes. *Educational Studies in Mathematics*, 102(1), 1-8. <https://doi.org/10.1007/s10649-019-09908-4>
- Cohen, J. (1990). Things I have learned (so far). *American Psychologist*, 1990, 45[12], 1304-1312
- Cohen, J. (1965). Some statistical issues in psychological research. *Handbook of clinical psychology*, 95-121.
- Collins, E. (2021, July 1). 2021 Understanding and Misunderstanding Effect Size. Retrieved from osf.io/5hu8m
- Correll, J., Mellinger, C., McClelland, G. H., & Judd, C. M. (2020). Avoid Cohen's 'small', 'medium', and 'large' for power analysis. *Trends in cognitive sciences*, 24(3), 200-207. <https://doi.org/10.1016/j.tics.2019.12.009>
- Cumming, G. (2014). The new statistics: Why and how. *Psychological science*, 25(1), 7-29. <https://doi.org/10.1177/0956797613504966>
- Drisko, J. W., & Maschi, T. (2016). *Content analysis*. Pocket Guides to Social Work R.
- Flora, D. B. (2020). Thinking about effect sizes: From the replication crisis to a cumulative psychological science. *Canadian Psychology/Psychologie canadienne*. <https://doi.org/10.1037/cap0000218>
- 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(2), 156-168. <https://doi.org/10.1177/2515245919847202>
- Haller, H., & Krauss, S. (2002). Misinterpretations of significance: A problem students share with their teachers. *Methods of Psychological Research*, 7(1), 1-20.
- Hoekstra, R., Morey, R. D., Rouder, J. N., & Wagenmakers, E. J. (2014). Robust misinterpretation of confidence intervals. *Psychonomic bulletin & review*, 21(5), 1157-1164. <https://doi.org/10.3758/s13423-013-0572-3>

Ioannidis, J. P. (2019). What have we (not) learnt from millions of scientific papers with P values?. *The American Statistician*, 73(sup1), 20-25.
<https://doi.org/10.1080/00031305.2018.1447512>

Kelley, K., & Preacher, K. J. (2012). On effect size. *Psychological methods*, 17(2), 137.
<https://doi.org/10.1037/a0028086>

Lyu, Z., Peng, K., & Hu, C. P. (2018). P-value, confidence intervals, and statistical inference: A new dataset of misinterpretation. *Frontiers in Psychology*, 9, 868.
<https://doi.org/10.3389/fpsyg.2018.00868>

Miller, J., & Ulrich, R. (2016). Interpreting confidence intervals: A comment on Hoekstra, Morey, Rouder, and Wagenmakers (2014). *Psychonomic bulletin & review*, 23(1), 124-130.
<https://doi.org/10.3758/s13423-015-0859-7>

Morris, P. H. (2020). Misunderstandings and omissions in textbook accounts of effect sizes. *British Journal of Psychology*, 111(2), 395-410. <https://doi.org/10.1111/bjop.12401>

Onie, S. (2020, November 3). Redesign open science for Asia, Africa and Latin America. *Nature*. <https://doi.org/10.1038/d41586-020-03052-3>

Peng, C. Y. J., Chen, L. T., Chiang, H. M., & Chiang, Y. C. (2013). The impact of APA and AERA guidelines on effect size reporting. *Educational Psychology Review*, 25(2), 157-209.
<https://doi.org/10.1007/s10648-013-9218-2>

Randolph, J. J., & Edmondson, R. S. (2005). Using the binomial effect size display (BESD) to present the magnitude of effect sizes to the evaluation audience. *Practical Assessment, Research, and Evaluation*, 10(1), 14. <https://doi.org/10.7275/zqwr-mx46>

Sun, S., Pan, W., & Wang, L. L. (2010). A comprehensive review of effect size reporting and interpreting practices in academic journals in education and psychology. *Journal of Educational Psychology*, 102(4), 989. <https://doi.org/10.1037/a0019507>

Wasserstein, R. L., Schirm, A. L., & Lazar, N. A. (2019). Moving to a world beyond “p<0.05”. *The American Statistician*, 73:sup-1, 1-19.
<https://doi.org/10.1080/00031305.2019.1583913>

Watt, R., & Collins, E. (2019). *Statistics for Psychology: A Guide for Beginners (and Everyone Else)*. SAGE.

Wilkinson, L. (1999). Statistical methods in psychology journals: Guidelines and explanations. *American psychologist*, 54(8), 594. <https://doi.org/10.1037/0003-066X.54.8.594>