How observed effects in behavioral science should (not) be reported

¹Erich H. Witte & ²Frank Zenker

¹Institute for Psychology, University of Hamburg, Von-Melle-Park 5, 20146

Hamburg, Germany; ORCID: 0000-0002-9446-3211; ²Dpt. of Philosophy, Boğaziçi University, 34342 Istanbul, Turkey; ORCID: 0000-0001-7173-7964;

*corresponding author: frank.zenker@boun.edu.tr

Abstract: Empirical studies in behavioral science typically report observed effects by relying on standardized effect size measures (e.g., Cohen's d). These measures relate the observed mean difference to the observed standard deviation. The latter thus functions as an empirical measurement unit. Since the observed standard deviation is conceptually related to the quality of measurement, the interpretation and the meta-analytically aggregation of standardized observed effects require a transparent account of the measurement scale reliability, the measurement error, the quality of the empirical setting, and how it is standardized. Even top-tier journals, however, do not demand that published studies provide such an account. This arguably places a cumulative science of human behavior farther out of reach. Standardized observed effects, therefore, should be reported alongside a transparent error account.

Keywords: effect size measure, error account; measurement; standardization; standard deviation; statistical power; meta-analysis; theory construction

How observed effects in behavioral science should (not) be reported

1. Introduction

Already the APA task force on statistical inference (Wilkinson et al., 1999, 599) had stressed that “reporting and interpreting effect sizes in the context of previously reported effects is essential to good research,” because “[i]t enables readers to evaluate the stability of results across samples, designs, and analyses. Reporting effect sizes also informs power analyses and meta-analyses needed in future research.” Moreover, a statistically significant effect that was observed under sufficient statistical power can often stand proxy for a *theoretically predicted* effect. For both meta-analytical and theory construction research, therefore, the size of observed effect is arguably “the most important outcome of empirical studies” (Lakens, 2013, 1).

The effect size measures that behavioral science publications most frequently rely on are Cohen’s (1977) d and Pearson’s (1895) r (see Schäfer & Schwarz, 2019). (Additional measures are reviewed in Grissom & Kim, 2012.) The d and the r measure are *standardized* effect size measures. For both measures report the observed mean difference ($m_1 - m_0$) relative to—or in the standard of—the observed standard deviation (s) (Baguley, 2009). What a typical behavioral science publication today reports thus is a standardized observed effect.

The interpretation of object-level standardized observed effects, as well as their meta-analytical aggregation, requires a transparent account of the *errors* a study is subject to. Providing this account requires explaining, among others, how the size of observed effect depends on the measurement scale reliability, the measurement error, the empirical setting, and how it is standardized. (For the contrasting error-theoretical approach, see Mayo, 1996; Mayo & Spanos, 2011.) Even top-tier behavioral science journals, however, do not require that published studies provide this account, nor do authors normally provide it themselves.

This, we argue, places a cumulative science of human behavior farther out of reach. Our main recommendation is that standardized observed effects should be reported *alongside* a transparent error account. We begin by reviewing the measurement-theoretical background (Sect. 2), then make our case using hypothetical yet *realistic* data from educational research (Sect. 3), followed by a discussion (Sect. 4) and conclusions (Sect. 5). Several formal details are relegated to the Appendix.

2. Standardized effect size measures and measurement

2.1 Cohen's *d*-measure

In experimental studies, Cohen's (1977) *d*-measure (see 1) states the difference ($m_1 - m_0$) between the observed means in the treatment (m_1) and the control group (m_0) relative to the pooled standard deviation (s). The latter quantity states the average difference between each data point and m_x ($x = 0, 1$) for both groups (Grissom & Kim, 2012; Lakens, 2013; Schäfer & Schwarz, 2019; Schulze, 2004).

$$d = (m_1 - m_0) / s \quad (1)$$

Although the sample of a study is ideally distributed evenly across both groups ($n_1 = n_0$), the *d*-measure can—up to a 90 : 10 % partitioning of the sample—be transformed without distortion into Pearson's (1895) point-biserial correlation measure r_{bis} (see 2) typically used in correlational studies (McGrath & Meyer, 2006).

$$r_{\text{bis}} = d / \sqrt{d^2 + 4} \quad (2)$$

Or alternatively:

$$d = r_{\text{bis}} / \sqrt{[N_1 \times N_0 / (N_1 + N_0)^2]}. \quad (3)$$

Because the relevant properties of the *d*-measure do under the above condition transfer to the *r*-measure, we can focus without loss of generality on the *d*-measure. What follows, therefore, holds for all uses of a standardized effect size measure (see Grissom & Kim, 2012) to report an observed effect.

2.2 Conjoint measurement

A measurement score (X_i) on a measurement scale is the sum of a *true score* (T_i) and a *measurement error* (e_i): $X_i = T_i + e_i$. According to the law of large numbers, if T_i and e_i are *independent* and e_i is *normally distributed* (such that its mean is $m_e=0$), then the sum of errors ($\sum e_i$) in a sample will decrease as the sample size (n) increases, because as the number of measurement scores increases, the instances of e_i to both sides of T_i will cancel out (Lord & Novick, 1968). Since the law of large numbers applies to the observed sample mean ($m_x \approx m_T$), the *observed variance* ($s^2_{x_i}$) thus depends on the

variance of the *true scores* (s^2_{Ti}) and the variance of the *errors* (s^2_{ei}). Consequently, if s^2_{Ti} and s^2_{ei} are independent, then both terms can be added: $s^2_{Xi} = s^2_{Ti} + s^2_{ei}$. In a sufficiently large random sample, therefore, s^2_{Xi} estimates the variance of the true scores in a *population* (s^2_{Xi-POP}). Notice that the square root of the variance ($\sqrt{s^2}$) is the standard deviation (s).

For contingent reasons, it is unavoidable that several independent studies of the same phenomenon display differences concerning how members of a given population are drawn into a sample. These differences entail that the true scores in any given pair of samples are prone to vary in different ways. In each sample, therefore, a different degree of *sampling error* (associated with the true scores) biases m_x away from the population mean (m_{POP}) (Borsboom, 2005; Lord & Novick, 1968). Although a given instance of m_x thus need not match m_{POP} , the law of large numbers guarantees that, as further *independent* instances of m_x (randomly sampled from the same population) are added, the estimate of m_{POP} continuously improves. In brief, if the underlying measurement scores are normally distributed, then adding sufficiently many independent measurement instances of m_x provides a robust basis to estimate m_{POP} .

2.3 Measurement quality

The measurement instances of m_x and s_x that independent studies of the same phenomenon observe thus underly what publications report as standardized *object-level* observed effects. Their interpretation and meta-analytical aggregation, however, is challenged by the fact measurements in behavioral science frequently occur by fiat (Flake & Fried, 2020), because measurements here typically occur in the absence of a *theoretically motivated* measurement unit, i.e., a unit derived from a theoretical construct postulated by a substantial theory of human behavior (Andreas, 2021; Rakover, 2022, cf. Sherry, 2011). The main reason is that the available theoretical structures are less well-developed. Indeed, several authors speak of a regular theory crisis (Eronen & Bringmann, 2021; Eronen & Romeijn, 2020; Feynman, 1974; Gigerenzer, 1998; Irvine, 2021; Meehl, 1978; Miłkowski, Hohol & Nowakowski, 2019; Oberauer & Lewandowsky, 2019; van Rooij & Baggio, 2020).

Absent a theoretically motivated measurement unit, the one plausible and available unit to standardize several independently observed object-level effects is the *observed standard deviation* (s), which thus functions as an *empirical* measurement unit. Conceptually, the use of s as a measurement unit presupposes a *transformation*

of raw measurement scores into elements of a distribution of a sample of measurement scores, i.e., elements of an uncertain statistic. Raw measurement scores as elements relative to a measurement theory thus become elements of a formally describable distribution relative to a statistical theory (Baguley, 2009). For instance, if a set of raw measurement scores is embedded into a normal distribution, then a d -value does *not* represent a score on a measurement scale but a position on the x -axis of this distribution, called a z -value. The relation between an individual raw measurement score and a theoretical construct, therefore, is *indirect* because it is mediated by s .

Conceptually, the s -value reflects the *quality of measurement* because it quantifies the measurement uncertainty that is associated with $(m_1 - m_0)$. As stated, a pair of independent studies of the same phenomenon are prone to arrive at different observed s -values already because of how observations vary naturally. Additionally, the observed s -value depends on measurement-related assumptions that crucially determine whether $(m_1 - m_0)$ is *clearly* observable. A clear observation of $(m_1 - m_0)$ generally requires *homogeneous* measurements of m_x , as indicated by a *small* s -value. Whereas a *large* s -value indicates measurements of m_x that are *heterogeneous*.

Particularly a *small* d -value that goes along with a *large* s -value entails that measurements are *insufficient* to confidently exclude the possibility that factors *besides* those used theoretically model m_1 account for $(m_1 - m_0)$, factors that include various errors. As object-level standardized observed effects are to be interpreted or meta-analytically aggregated, therefore, the main challenge is emphatically not that error-prone measurements *underdetermine* a theoretical construct (which is true), nor that error-prone measurements cannot *probabilistically support* a theoretical construct (which is false). The main challenge is that, if a small d -value goes along with a large s -value, then the measurements of m_x *keep* from discerning $(m_1 - m_0)$ clearly behind the metaphorical veil that a large s -value is.

Because recent *meta*-meta-analyses strongly suggest that the ‘small d , large s ’-pattern is *representative* of standardized observed effects in behavioral science (Schauer & Hedges, 2020; Olsson-Collentine, Wicherts & van Assen, 2020; Linden & Hönekopp, 2021), the crucial insight is that the observed s -value can be reduced if a study’s measurement-related assumptions are being made transparent.

2.4 Measurement-related assumptions

A measurement outcome is generally subject to changes in measurement-related assumptions. In Einsteinian relativity theory, for instance, time measurement scores depend on whether the measurement instrument is stationary or in motion (Galison, 2004). This dependency breaks with Newtonian physics, where time measurement scores have absolute meaning. In relativity theory, time measurement scores instead depend on the frame of reference and the inertial system where a measurement occurs. Compared to time measurements with a stationary clock, for instance, measurements with a clock mounted on a train moving at light speed would display a *time dilation*.

As we proceed to show, a change in measurement-related assumptions can similarly result in an *effect size dilation*. Particularly a small and heterogeneous observed effect (small d , large s) can thus change into a large and homogeneous observed effect (large(r) d , small s). Crucially, this change of pattern is a direct consequence of how the *errors* that a study is subject to are accounted for (Morris & DeShon, 2002), and thus a consequence of differences in measurement quality.

Stating that a difference between a pair of error accounts entails a difference in measurement quality may appear trivial. The implications for the development of a cumulative science of human behavior are nevertheless significant, because already a quick glance at the top journals in behavioral science suffices to recognize that publications typically report standardized observed effects *without* a transparent error account. From what a typical publication in behavioral science reports, therefore, the effect size dilation is effectively *indiscernible*.

The absence of a transparent error account challenges the meta-analytical integration of independent object-level standardized effects into a population estimate as much as it challenges the development of empirically adequate theoretical constructs that predict a population effect. In both cases, the main problem is the same. Although the random factors that influence an empirical setting cannot be known exhaustively, they must in any two samples be assumed to affect the measurement of $(m_1 - m_0)$ in *different* ways. Hence, even if statistical power is maximal, s never takes the value zero. Particularly if a *small* d -value goes along with a large(r) s -value, therefore, the value of $(m_1 - m_0)$ must remain *vastly* uncertain. This uncertainty fully transfers to a meta-analytically estimated population effect size.

As for theory construction, a good theoretical construct must postulate an empirically adequate version of $(m_1 - m_0)$ *independently* of s , because s does not reflect how m_x is theoretically modeled but how well $(m_1 - m_0)$ is measured. A

theoretical construct “is empirically adequate exactly if what it says about the observable things and events in the world is true—exactly if it ‘saves [or captures] the phenomena’” (van Fraassen, 1980, 12). The smaller ($m_1 - m_0$), therefore, the more homogenous must the underlying measurements be to clearly discern particularly the theoretically predicted m_1 . Otherwise, it is difficult to see how an empirically adequate construct could be identified at all. The mean of a random distribution ($m_0=0$), of course, is in no way easier to discern; it is merely easier to “predict.”

We proceed to demonstrate the relevance of a transparent error account using an example of hypothetical, yet *realistic* data from educational research on the impact of a new teaching method on the average student’s IQ score. Because what follows is orthogonal to the controversy over the *construct validity* of IQ (Cook, Campbell & Day, 1979; Cronbach & Meehl, 1955; Shadish, Cook & Campbell, 2002; Schimmack, 2022), replacing the IQ construct with any other construct that is measured without a theoretically motivated measurement unit cannot affect the *conceptual* point we are about to make.

3. Transparent error accounting

3.1 The impact of a new teaching method on IQ scores

A central question for educational research concerns the impact of the duration of school education (*independent variable*) on the average student’s standard IQ-score (*dependent variable*). Relative to the standard IQ measurement scale—where the population mean is conventionally set to $m_{POP} = 100$ IQ points and the standard deviation is $s_{POP} = 15$ IQ points—Ritchie & Tucker-Drob’s (2018) meta-analysis estimates that one year of school education increases the average student’s standard IQ-score by 1.197 IQ points (ibid.). After dividing it by s_{POP} , the meta-analytically estimated effect of one year of school education thus is $d = 1.197 / 15 = 0.08$ IQ points. For six years, we find $d = 0.48$ IQ points, for ten years $d = 0.80$ IQ points, and for thirteen years $d = 1.04$ IQ points. Notice that these d -values roughly correspond to how Cohen (1977) defined the terms *small*, *medium*, *large*, and *very large* effect.

Although Ritchie and Tucker-Drob (2018) discuss causal considerations and the role of moderators (see our Sect. 4), it suffices for our purposes to assume that some (undefined) new teaching method exerts a *larger* effect on the average student’s IQ score than a traditional teaching method. To quantify this effect, the average IQ score of students taught with a traditional method (*control group*) must be compared

to that of students taught with the new method (*treatment group*). The one relevant observation, therefore, is the mean difference ($m_1 - m_0$) in IQ scores as standardized by s . Assume, then, that three researchers independently conduct control group studies in three countries, grant that all observations are statistically significant ($p < 0.05$), and that the typically large samples available in educational research yield well-powered ($(1 - \beta) > 0.95$) observations (Hattie, 2009). Finally, using $d = (m_1 - m_0) / s_{Cn}$, where ‘ C_n ’ stands for ‘country n ’, assume the standardized observed effects are $d_{C1} = 0.20$, $d_{C2} = 0.50$, $d_{C3} = 1.0$ IQ points. (Ritchie and Tucker-Drob’s (2018, 5) meta-analysis reveals a rather similar variation of the observed effect.)

Since the d -values differ across the three countries, the impact of the teaching method on the average student’s IQ scores appears to differ *locally*. This appearance, however, can be explained away by studying how different error accounts standardize the observed ($m_1 - m_0$) by the observed s . This explanation is possible in the first place because—as a measurement unit—the observed s lacks *theoretical* meaning (see Sect. 2). All three researchers, therefore, are free to construct this unit according to their *preferred* error account. The only relevant constraint is that each error account must explain the variation of the observed effect by factors *other* than those which model the observed effect theoretically.

The explanation rests on the conceptual truth (Cohen, 1977; Baguley, 2009) that a description of a standardized observed effect, a comparison of this effect across different statistically standardized measurement scales, as well as a significance test for this effect, entail a transition from the original measurement unit ($s_{POP} = 15$ IQ points), via the empirically induced statistical measurement unit s , to a t -statistic of the form:

$$t = [(m_1 - m_0) / s \times \sqrt{(n_1 \times n_2) / (n_1 + n_2)}] \quad (4)$$

The main consequence is that a t -test for a standardized observed effect against random *requires* a standardization by s , and that the statistical power of this t -test depends on the specific standardization a researcher has chosen. (This is more fully explained in the Appendix).

We proceed to summarize how four progressively more sophisticated error accounts make different measurement-related assumptions transparent. In each account, the measurement error, as well as various theoretically non-modeled factors

that influence the teaching method's impact on the average student, are represented by *different* standardization factors, factors that are traced by indexing the d -measure (see Table 1). The d -measure draws the s -value that an IQ-test manual states for the population (researcher one). The d' -measure draws on local s -values (researcher two). The d^* -measure eliminates the variation of students' true individual IQ scores; the d^{**} -measure is standardized to the s -value in country two, and the d'^{**} -measure to that in country three (researcher three).

3.2 Four progressively more sophisticated error accounts

Assume that researcher one treats data in country one by standardizing the observed means (m_0, m_1) to $s_{POP} = 15$ IQ points (the measurement unit of the standard-IQ scale), thereby treating $(m_1 - m_0)$ as an element of the population distribution as reported in an IQ-test manual (Ritchie & Tucker-Drob, 2018). This is equivalent to the study in country relying on a *random* sample of students. A publication, therefore, would report an observed effect of $dc_1 = (m_1 - m_0) / sc_1 = 3 / 15 = 0.20$ IQ points *without* providing an error account, thus describing the normal case in behavioral science today.

Further assume that the studies in countries two and three rely on samples of students who are implicitly *selected for* IQ by the progressively more demanding admission requirements that apply in a secondary and a high school system.

Country one: *random* sample; *no* selection for IQ

Country two: *secondary* school students; *medium* selection for IQ

Country three: *high* school students; *strong* selection for IQ

Because the three different IQ-selection conditions influence students' IQ scores, $sc_1 = 15$ IQ points becomes an *invalid* estimate for sc_2 and sc_3 . Therefore, as researcher two acknowledges, an explanation why the observed effect differs in each country requires a standardization of the local d -values to a *common* measurement unit. This standardization is conceptually equivalent to measuring m_x with a single measurement instrument to which the same amount of measurement uncertainty is associated.

Because the one plausible and available measurement unit is the observed standard deviation sc_n , assume the locally observed s -values are $sc_1 = 15$, $sc_2 = 6$, $sc_3 =$

3 IQ points. (Notice that these s -values are realistic.¹) After the local measurement units are standardized to $s_{C2} = 6$, researcher two obtains the standardization factors $s_{C1}^* = 15 / 6 = 2.50$ IQ points and $s_{C3}^* = 3 / 6 = 0.50$ IQ points. This standardization entails that the new teaching method now has the *same* effect throughout, namely $dc_1 = (0.20 \times 2.50) = dc_2 = dc_3 = (1.0 \times 0.50) = 0.50$ IQ points.

Although researcher two's error account is more sophisticated than that of researcher one, the locally observed s -values remain induced from observations. The status of s , therefore, is *purely* empirical. But this status changes as researcher three reasons on theoretical grounds that the dispositional variance of the students' true *individual* IQ-scores is irrelevant for purposes of developing a theoretical construct that predicts $(m_1 - m_0)$ for the *average* student. The sampling error associated with student's true individual IQ scores, therefore, can be eliminated from the s -value. Conceptually, this elimination is equivalent to an observed s -value being *corrected* on the basis of theoretical considerations. On the assumption that the measurement scale reliability is 0.80 (see the Appendix, formula 5), the observed effect size again increases—from $dc_1 = 0.20$, $dc_2 = 0.50$, $dc_3 = 1.0$ IQ points—to $d^*_{C1} = 0.27$, $d^*_{C2} = 1.12$, $d^*_{C3} = 2.73$ IQ points, yet continues to vary locally.

Since the measure $(m_1 - m_0)$ is constant under relevant formal transformations, and is in each country the same, researcher three speculates that the observed effect varies because the empirical settings are in each country differentially well-controlled against the influence of random factors. Conceptually, this speculation is equivalent to assuming a *local implementation condition* for the teaching method. Researcher three's error account consequently standardizes both the measurement unit and the measurement error to the implementation condition of country two. Conceptually, this is equivalent to using a single relative measure on a common measurement scale. The observed effect—which now comes to $dc_1^{**} = 1.13$, $dc_2^{**} = 1.12$, and $dc_3^{**} = 1.11$ —thus *ceases* to vary locally.

¹ In a recent crowdsourcing replication study, Landy et al. (2020) coordinated fifteen labs on using samples of around $n_1 = n_2 = 300$ to conduct independent well-powered original and replication t -tests. The ratio of the largest and smallest s -values for an original and a replication t -test came to 4.97 vs. 5.00 (study 1), 3.48 vs. 2.85 (study 2), 2.28 vs. 2.53 (study 3), and 2.71 vs. 2.58 (study 4). These s -values are very similar to those used in our example.

Researcher three can moreover calculate the *expected* effect size increase for an implementation condition with the *smallest* s -value (country three). The observed effect thus comes to $d'^{**} = 2.70$ IQ points, corresponding to an effect of $2.70 \times 15 = 40.50$ IQ point (!) on the standard-IQ scale. For country one, one would find $2.75 \times 15 = 41.25$ IQ points, and for country two $2.73 \times 15 = 40.95$ IQ points.

Formal details omitted from this section are found in the Appendix. We proceed to formulate the upshot and discuss implications for meta-analytical research, theory construction, and the replication crisis in behavioral science.

Table 1: Five versions of the d -measure relying on different standardization factors

Country	Standardized effect size measure versions				
	d	d'	d^*	d^{**}	d'^{**}
	$(m_1 - m_0)/s_{\text{POP}}$	$(m_1 - m_0)/s_{C2}$	$(m_1 - m_0)/s_{e-Cn^*}$	$(m_1 - m_0)/s_{C2}^{**}$	$(m_1 - m_0)/s_{C3}^{**}$
1	$d_{C1} = 0.20$	$d'_{C1} = 0.50$	$d_{C1}^* = 0.27$	$d_{C1}^{**} = 1.13$	$d'_{C1}^{**} = 2.75$
2	$d_{C2} = 0.50$	$d'_{C2} = 0.50$	$d_{C2}^* = 1.12$	$d_{C2}^{**} = 1.12$	$d'_{C2}^{**} = 2.73$
3	$d_{C3} = 1.00$	$d'_{C3} = 0.50$	$d_{C3}^* = 2.72$	$d_{C3}^{**} = 1.11$	$d'_{C3}^{**} = 2.70$

4. Discussion

4.1 Transparent error accounting

Compared to $d_{C1} = 0.20$, $d_{C2} = 0.50$, and $d_{C3} = 1.0$ IQ points (as observed by researcher one who provided no error account), the observed effect did significantly increase as various measurement-related assumptions were made transparent. In country two, for instance, the observed effect eventually came to $d^{**} = 1.12$ IQ points, or what Cohen calls a *very large* effect. On the standard IQ-scale (where $s_{\text{POP}} = 15$ IQ points), $d^{**} = 1.12$ IQ points corresponds to an effect of 16.80 IQ points. On the same scale, the initial *non-standardized* effect of $(m_1 - m_0) = 3.00$ IQ points—or what Baguley (2009) calls *simple* effect size—thus increases by a factor of 5.6. Under the best-controlled

implementation condition (country three), the effect increases even to 41 IQ points, or by a factor of 13.

Without a transparent error account, what compared to the *medium* observed effects that are typically reported in educational research (Ritchie & Tucker-Drob, 2018; Hattie, 2009) is a *significantly* increased effect size would simply remain “hidden.” Crucially, the different ways of accounting for the various errors discussed in Sect. 3 did exclusively draw on *data-based* standardization factors (e.g., the local s -values themselves, the influence of students’ true individual IQ scores on s , and the local implementation condition), factors that are *independent* of how $(m_1 - m_0)$ is modeled theoretically. An error account, therefore, inevitably reflects a researcher’s best understanding of how an observed effect is influence by factors that are *not* used to model it theoretically.

As indicated, recent *meta-meta*-analyses (Schauer & Hedges, 2020; Olsson-Collentine, Wicherts & van Assen, 2020; Linden & Hönekopp, 2021) strongly suggest that our starting point—a *small* and *heterogeneous* standardized observed effect (i.e., small d , large s)—is a *representative* pattern for reported object-level effects in behavioral science. With each progressively more sophisticated error account, this pattern eventually changed to that of a *very large* and *homogeneous* observed effect. At the same time, $(m_1 - m_0)$ remained constant. The change in pattern, therefore, can only be explained by how $(m_1 - m_0)$ is standardized by s , i.e., by how well $(m_1 - m_0)$ is measured.

Because the number of *plausible* error accounts entails the same number of distinct standardization factors, one may even inquire into the meaningfulness of assuming that the new teaching method has a *single* true effect. Absent a theoretically motivated measurement unit, moreover, a single “right” effect size measure cannot be had either, because a reasoned choice among available standardization factors always favors s . And s simply varies with the details of a specific error account. All plausible standardization factors, therefore, depend on the error account as a background assumption.

4.2 Implications

The forgoing provides a sufficient reason to report object-level standardized observed effects exclusively *alongside* a transparent error account. Conversely, not providing a transparent error account severely challenges the development of a theoretical

construct that predicts an *empirically adequate* population effect, because object-level standardized observed effects that deserve to be modeled theoretically must in the first place *exceed* the size of the standard measurement error. Observed effects that satisfy this criterion, however, tend to be highly heterogeneous (Linden & Hönekopp, 2021), suggesting that they are underlain by *vastly* different measurement units (Flake & Field, 2020). Without providing a transparent error account, therefore, sufficiently large and homogenous object-level standardized observed effects (that are needed to develop theoretical constructs that predict empirically adequate population effects) simply cannot be identified.

This arguably places a cumulative science of human behavior farther out of reach. By contrast, where it is provided a transparent error account harbors empirical and statistical information to *separate* ($m_1 - m_0$) from the random influences on an empirical setting that are represented by s . This particularly renders an *empirically adequate* theoretically predicted instance of ($m_1 - m_0$) more clearly discernable. A sophisticated meta-analytical approach aiming to estimate a population effect of interest, therefore, would *only* include topically related object-level standardized observed effects that are sufficiently homogeneous. Instead, current meta-analytical approaches typically employ a moderator analysis to explain the heterogeneity of object-level effects *away* (Schauer, & Hedges, 2020). But rarely are any promising moderators identified and more rarely yet verified (Linden & Hönekopp, 2021).

Concerning theory construction, using moderators to explain the heterogeneity of standardized observed object-level effects away ignores that the *correspondence* between a theoretically predicted effect postulated by a theoretical construct, on one hand, and a population effect estimated by a meta-analysis, on the other, simply cannot be evaluated unless an empirically adequate theoretical construct can be distinguished with *sufficient confidence* from its empirically inadequate counterpart. Standardized effect sizes, however, can serve to *draw* this distinction only if a small s -value indicates that uncontrolled factors exert merely a negligibly small influence on measurements. A sophisticated meta-analytical approach, therefore, should *correct* published object-level standardized observed effects according to how “disturbed” their empirical setting are. But without a transparent error account this is impossible.

Recognizing the importance of a transparent error account also offers insights that may improve our understanding of the causes of the replication crisis in behavioral science (Olsson-Collentine, Wicherts & van Assen, 2020; Schauer &

Hedges, 2020; Szucs & Ioannidis, 2017). The assumption that a reported object-level standardized observed effect matches a population effect in size suffices to post-hoc determine a study's observed statistical power. On the assumption that the sample size is constant, the larger the study's observed effect, the larger its statistical power. Observed statistical power, moreover, is strongly correlated with the probability of replicating a true observed effect,² and the population effect *is* the true effect that an empirically adequate theoretical construct must predict. Identifying an empirically adequate theoretical construct, therefore, requires that studies achieve sufficient statistical power.

In behavioral science, however, *median* observed statistical power for studies that observe *small* ($d = 0.20$), *medium* ($d = 0.50$) and *large* standardized effects ($d = 0.80$) is estimated as, respectively, $(1 - \beta) = 0.12, 0.44, 0.73$ (Szucs & Ioannidis, 2017). Actual samples thus are regularly too small to reasonably expect that a replication study duplicates a published originally observed effect. Although already a simple error account can suffice to decrease the observed s -value—thereby increasing both the d -value and the $(1 - \beta)$ -value—the minimum sample size that is required to observe a specific effect under sufficient statistical power can be calculated a priori (*before* data collection), only if a theoretical construct had been used to predict ($m_1 - m_0$) (Neyman & Pearson, 1967). In behavioral science, however, predictive constructs are rare. Low median observed statistical power, therefore, is also a consequence of the broad absence of predictive constructs.

4.3 Measurement precision

A *homogeneous* standardized observed effect, as we saw, presupposes a small s -value. Obtaining a small s -value requires a reduction, and ultimately the elimination, of *most* random influences on measurements. This eliminative strategy happens to describe how cosmologists currently study the early universe's evolution. As the microwave and submillimeter skies are scanned, the systematic residual polarization effects of

² A replication study is subject to a regression effect, the size of which is inversely correlated with the sample size. Observed effects in replication studies, therefore, are expectably smaller than in original studies (Fiedler & Prager, 2018). The most basic factor in determining the replicability of originally observed effects, of course, is the base-rate of replicable observed effects in a field (Bird, 2020).

interest are generally estimated separately from random noise (Bouchet et al., 2020, 9). By “cleaning” observational data of this noise, the underlying measurements become so precise (i.e., the s -value is so small) that even a *very small* effect is clearly discernable.

Thinking of random noise as an element to be eliminated starkly contrasts with thinking of moderators as elements that a meta-analysis seeks to identify rather than eliminate. Applying this eliminative strategy to the phenomena that are studied in behavioral science thus requires the near-elimination of *theoretically non-modeled* influences on a dependent variable of interest. A possible way forward is to examine the residuals of hierarchical linear models that differentiate measurements according to a study’s *units of research* (e.g., in our example, the country, the school type, the class size). The measurement unit, therefore, can be standardized according to these units of research (Raudenbush & Bryk, 2002).

In our example, additional measurements would be required to determine how these units of research affect the measurement error associated with the individual IQ scores in a sample. Once the variance that each such unit of research accounts for is eliminated, the residuals and a consequently reduced measurement error would yield new standardization factors, thus allowing to eliminate additional kinds of errors from the implementation error. This would differentiate s and at the same time decrease it. An analysis of the residuals, therefore, would yield a further improved estimate of the new teaching method’s effect on the average student’s IQ score.

5. Conclusion

Whether the object of evaluation is an intervention upon a dependent variable, an observed correlation, or a theoretical prediction, the crucial properties of an observed effect are its *size* and its *heterogeneity*. With standardized effect size measures, these two properties depend on the observed mean difference ($m_1 - m_0$) and the observed standard deviation (s). The latter quantity depends on factors such as the measurement scale reliability, the measurement error, the quality of an empirical setting, and how it is standardized. These factors are relevant in the first place because the measurement scales in behavioral science typically lack a theoretically motivated measurement unit. The one plausible and available measurement unit, therefore, is s . Yet the observed s -value varies with the specific of a researcher’s error account. Hence, as object-level standardized observed effects that are published *without* a transparent error account

shall be interpreted or meta-analytically aggregated, the population effect that an empirically adequate theoretical construct must predict typically remains *vague*. For the development of a cumulative science of human behavior, however, this vagueness, presents a self-inflicted, avoidable drawback. For this reason, standardized observed effects should exclusively be reported alongside a transparent error account.

Acknowledgements/contributions/author-declaration

DELETED. The author(s) declare(s) that they have no conflict of interest concerning the publication of this manuscript. N.N. wrote a first draft, which N.N. edited. Both author(s) approved the final submitted version. N.N. acknowledges funding from DELETED.

References

- Andreas, H. (2021). Theoretical terms in science. In: Zalta, E. N. (ed.), *The Stanford Encyclopedia of Philosophy* (Fall 2021 Edition).
<https://plato.stanford.edu/archives/fall2021/entries/theoretical-terms-science/>
- Baguley, T. (2009). Standardized or simple effect size: What should be reported? *British Journal of Psychology*, 100, 603-617.
<https://doi.org/10.1348/000712608X377117>
- Borsboom, D. & Mellenbergh, G. J. (2002). True scores, latent variables, and constructs: A comment on Schmidt and Hunter. *Intelligence*, 30, 505–514.
[https://psycnet.apa.org/doi/10.1016/S0160-2896\(02\)00082-X](https://psycnet.apa.org/doi/10.1016/S0160-2896(02)00082-X)
- Bird, A. (2020). Understanding the replication crisis as a base rate fallacy. *The British Journal for the Philosophy of Science*, 72 (4), 965-993.
<https://doi.org/10.1093/bjps/axy051>
- Borsboom, D. (2005). *Measuring the mind. Conceptual Issues in Contemporary Psychometrics*. New York: Cambridge University Press.
- Bouchet, F. R. et al. (2020). Planck2018 results. *Astronomy & Astrophysics* 641, A1
<https://doi.org/10.1051/0004-6361/201833880>
- Cook, T. D., Campbell, D. T., & Day, A. (1979). *Quasi-experimentation: Design & analysis issues for field settings* (Vol. 351). Boston, MA: Houghton Mifflin.
- Cohen, J. (1977, rev. ed). *Statistical power analysis for the behavioral sciences*. London: Academic Press.

- Cronbach, L. J., & Meehl, P. E. (1955). Construct validity in psychological tests. *Psychological Bulletin*, 52(4), 281-302. <https://doi.org/10.1037/h0040957>
- Dietrich, C. F. (1991²). *Uncertainty, calibration and probability*. Bristol: Adam-Hilger.
- Eronen, M. I., & Bringmann, L. F. (2021). The theory crisis in psychology: How to move forward. *Perspectives on Psychological Science*, 16(4), 779-788. <https://doi.org/10.1177/1745691620970586>
- Eronen, M. I., & Romeijn, J. W. (2020). Philosophy of science and the formalization of psychological theory. *Theory & Psychology*, 30(6), 786-799. <https://doi.org/10.1177/0959354320969876>
- 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. <https://doi.org/10.3758/BF03193146>
- Fraassen, B. van (1980). *The Scientific Image*. Oxford: Oxford University Press.
- Feynman, R. P. (1974). Cargo cult science. *Engineering and Science*, 37(7), 10-13. <https://resolver.caltech.edu/CaltechES:37.7.CargoCult>
- Fiedler, K., & Prager, J. (2018). The regression trap and other pitfalls of replication science—illustrated by the report of the Open Science Collaboration. *Basic and Applied Social Psychology*, 40(3), 115-124. <https://doi.org/10.1080/01973533.2017.1421953>
- Flake, J. K., & Fried, E. I. (2020). Measurement schmeasurement: Questionable measurement practices and how to avoid them. *Advances in Methods and Practices in Psychological Science*, 3(4), 456-465. <https://doi.org/10.1177/2515245920952393>
- Grissom, R. J., & Kim, J. J. (2012). *Effect sizes for research*. New York: Routledge.
- Galison, P. (2004). *Einstein's clocks, Poincaré's maps: empires of time*. New York: WW Norton & Company.
- Gigerenzer, G. (1998). Surrogates for theories. *Theory & Psychology*, 8 (2), 195–204. <https://doi.org/10.1177/0959354398082006>
- Hattie, J. A. C. (2009). *Visible learning. A synthesis of over 800 meta-analyses relating to achievement*. New York: Routledge.

- Irvine, E. (2021). The role of replication studies in theory building. *Perspectives on Psychological Science*, 16(4), 844-853.
<https://doi.org/10.1177/1745691620970558>
- Landy, J. et al. (2020). Crowdsourcing hypothesis tests: making transparent how design choices shape research results. *Psychological Bulletin*, 146(5), 451–479.
<http://dx.doi.org/10.1037/bul0000220>
- Lakens, D. (2013). Calculating and reporting effect sizes to facilitate cumulative science: a practical primer for t-tests and ANOVAs. *Frontiers in Psychology*, 4, article 863. <http://dx.doi.org/10.3389/fpsyg.2013.00863>
- Linden, A. H., & Hönokopp, J. (2021). Heterogeneity of research results: A new Perspective from which to assess and promote progress in psychological science. *Perspectives on Psychological Science*, 16(2), 358–376.
<https://doi.org/10.1177/1745691620964193>
- Lord, F. M. & Novick, M. R. (1968). *Statistical theories of mental test scores*. Reading, MA: Addison-Wesley.
- Mayo, D. G. (1996). *Error and the growth of experimental knowledge*. Chicago: University of Chicago Press.
- Mayo, D. G., & Spanos, A. (2011). Error statistics. In: P.S. Bandyopadhyay & M.R. Forster (Eds.). *Handbook of Philosophy of Science. Vol. 7: Philosophy of Statistics* (pp. 153–198). Amsterdam: Elsevier.
- McGrath, R. E., & Meyer, G. J. (2006). When effect sizes disagree: the case of r and d . *Psychological Methods*, 11, 386-401. <https://doi.org/10.1037/1082-989x.11.4.386>
- Miłkowski, M.; Hohol, M., & Nowakowski, P. (2019). Mechanisms in psychology: The road towards unity? *Theory & Psychology*, 29(5), 567-578.
<https://doi.org/10.1177/0959354319875218>
- Morris, S. B., & DeShon, R. P. (2002). Combining effect size estimates in meta-analysis with repeated measures and independent-groups designs. *Psychological Methods*, 7(1), 105-125. <https://doi.org/10.1037/1082-989x.7.1.105>
- Neyman, J., & Pearson, E. S. (1967). *Joint Statistical Papers*. Cambridge: University Press.
- Oberauer, K., & Lewandowsky, S. (2019). Addressing the theory crisis in psychology. *Psychonomic Bulletin & Review*, 26(5), 1596-1618.
<https://doi.org/10.3758/s13423-019-01645-2>

- Olsson-Collentine, A., Wicherts, J. M., & van Assen, M. A. (2020). Heterogeneity in direct replications in psychology and its association with effect size. *Psychological Bulletin*, 146(10), 922-940. <http://dx.doi.org/10.1037/bul0000294>
- Pearson, K. (1895). Notes on regression and inheritance in the case of two parents. *Proceedings of the Royal Society of London*, 58, 240–242. <https://doi.org/10.1098/rspl.1895.0041>
- Rakover, S. S. (2020). Why has the field of psychology not developed like the natural sciences?. *Journal of Mind & Behavior*, 41(3&4), 247–266. https://www.researchgate.net/profile/Sam-Rakover/publication/348390535_Why_Has_the_Field_of_Psychology_Not_Developed_Like_the_Natural_Sciences/links/5ffc5cb8299bf140888adcfd/Why-Has-the-Field-of-Psychology-Not-Developed-Like-the-Natural-Sciences.pdf
- Raudenbush, S. W. Bryk, A. S. (2002). *Hierarchical linear models: Applications and data analysis methods*, second edition. Newbury Park, CA: Sage.
- Ritchie, S. J. & Tucker-Drob, E. M. (2018). How much does education improve intelligence? A meta-analysis. *Psychological Science*, 29(8), 1358-1369. <https://doi.org/10.1177/0956797618774253>
- Sherry, D. (2011). Thermoscopes, thermometers, and the foundations of measurement. *Studies in History and Philosophy of Science Part A*, 42(4), 509-524. <https://doi.org/10.1016/j.shpsa.2011.07.001>
- Schauer, J. M. & Hedges, L. V. (2020, April, 9). Assessing heterogeneity and power in replications of psychological experiments. *Psychological Bulletin*, 146(8), 701-719. <https://doi.org/10.1037/bul0000232>
- Schäfer, Th. & Schwarz, M.A. (2019). The meaningfulness of effect sizes in psychological research: differences between sub-disciplines and the impact of potential biases. *Frontiers in Psychology*, 10, 813. <https://doi.org/10.3389/fpsyg.2019.00813>
- Schimmack, U. (2022). The validation crisis in psychology. *Meta-Psychology*, 5, 1-9. <https://doi.org/10.15626/MP.2019.1645>
- Schulze, R. (2004). *Meta-Analysis. A comparison of approaches*. Göttingen: Hogrefe & Huber.
- Shadish, W. R., Cook, T. D., & Campbell, D. T. (2002). *Experimental and quasi-experimental designs for generalized causal inference*. Boston, MA: Houghton, Mifflin and Company.

- Szucs, D. & Ioannidis, J.P.A. (2017). Empirical assessment of published effect sizes and power in the recent cognitive neuroscience and psychology literature. *PloS Biol*, 15(3), e2000797. <https://doi:10.1371/journal.pbio.2000797>
- van Rooij, I., & Baggio, G. (2020). Theory before the test: How to build high-verisimilitude explanatory theories in psychological science. *Perspectives on Psychological Science*, 16(4), 682–697. <https://doi.org/10.1177/1745691620970604>
- Wilkinson, L., & Task Force on Statistical Inference, American Psychological Association, Science Directorate (1999). Statistical methods in psychology journals: Guidelines and explanations. *American Psychologist*, 54(8), 594–604. <https://doi.org/10.1037/0003-066X.54.8.594>

APPENDIX

1. Standardization

Assume that researcher one, who uses a convenience sample of students in country one, obtains $d_{C1} = (m_1 - m_0) / s = 0.20$ IQ points on the standard IQ scale, and acts *as if* data were sampled *randomly* by setting $s_{C1} = s_{POP} = 15$ IQ points, i.e., the s -value that an IQ-test manual states for a random sample (Ritchie & Tucker-Drob, 2018).

Researcher one's error account thus merely acknowledges that each measured deviation from m_x ($x = 0, 1$) is subject to a measurement error of $s_{C1} = 15$ IQ points. A publication, therefore, would report an IQ-score increase for the average student of $d_{C1} \times s_{C1} = 0.20 \times 15 = (m_1 - m_0) = 3$ IQ points. This represents a *non-standardized* effect, or what Baguley (2009) calls *simple* effect.

Further, assume that the sample in country two consists exclusively of students who are implicitly *selected* for their IQ by the admission requirements of a *secondary* school system. Call this a *medium* IQ-selection condition. For country three, by contrast, assume a *strong* IQ-selection condition because the sample there consists exclusively of *high*-school students.

Country one: *random* sample; *no* selection for IQ

Country two: *secondary* school students; *medium* IQ-selection

Country three: *high*-school students; *strong* IQ-selection

Because these different IQ-selection conditions must be thought to influence students' IQ scores, $s_{C1} = 15$ IQ points (as used by researcher one) becomes invalid as an estimate of s_{C2} and s_{C3} . To correct for the different IQ-selection conditions, researcher two can obtain a *common* measurement unit by using the local s -values $s_{C1} = 15$, $s_{C2} = 6$, and $s_{C3} = 3$ IQ points. Thus, $s_{C1} = 15$ IQ points is standardized to $s_{C1}^* = s_{C1} / s_{C2} = 15 / 6 = 2.50$ IQ points, $s_{C2} = 6$ IQ points is standardized to $s_{C2}^* = 1$ IQ point, and $s_{C3} = 3$ IQ points is standardized to $s_{C3}^* = s_{C3} / s_{C2} = 3 / 6 = 0.50$ IQ points. The observed effects, therefore, can be expressed in the common unit of s_{C2} . To trace this standardization, d changes to d' .

$$d'_{C1} = ((m_1 - m_0) / s_{C1}) \times s_{C1}^* = 3 / 15 \times 2.50 = 0.50 \text{ IQ points}$$

$$d'_{C2} = ((m_1 - m_0) / s_{C2}) \times s_{C2}^* = 3 / 6 \times 1.00 = 0.50 \text{ IQ points}$$

$$d'_{C3} = ((m_1 - m_0) / s_{C3}) \times s_{C3}^* = 3 / 3 \times 0.50 = 0.50 \text{ IQ points}$$

Standardizing the d -value to the common measurement unit $s_{C2} = 6$ IQ points thus results in the *same* observed effect size in each country, namely 0.50 IQ points. This shows that $s_{C1} = 15$ had inflated $(m_1 - m_0)$ in country one (no IQ-selection), whereas $s_{C3} = 3$ IQ points had deflated it in country three (strong IQ-selection). However, although the d -measure's numerator, $(m_1 - m_0)$, is the same as in the d' -measure, standardizing a d -value to a common measurement unit affects the *statistical power* of a t -test. Given $n_1 = n_0 = 100$ and $\alpha = 0.05$ (one-sided), for instance, the statistical power of a t -test for $d_{C1} = 0.20$ IQ points is $(1 - \beta) = 0.41$, for $d_{C2} = 0.50$ IQ points it is $(1 - \beta) = 0.97$, and for $d_{C3} = 1.0$ IQ points we find $(1 - \beta) \approx 1.00$ (Faul et al., 2007).

Despite observing the *same* effect of 0.50 IQ points throughout, the $(1 - \beta)$ -error differs *locally*, because—regardless of how observations are transformed into a t -statistic—the t -value is defined as $t = d \times (\sqrt{n_1 + n_0} / 2)$ (Cohen, 1977, 67). This definition leaves s implicit (i.e., $s = 1$). (The statistical theory underlying a t -test (Neyman & Pearson, 1967) thus *downplays* the role of raw measurement scores.) As raw IQ measurement scores are transformed into a t -statistic, of course, the t -value continues to depend on the d -value, which itself depends on s . Provided the α -error and n are constant, therefore, the t -value and the statistical power of a t -test likewise depend on s . A t -test against random for $d_{C1} \times s_{C1} = 0.20 \times 15 = (m_1 - m_0) = 3$ IQ points, as reported by researcher one, thus *requires* a standardization (cf. Baguley, 2009).

2. Sampling error and error variance

Because a theoretical construct should predict the effect for the *average* student, assume that researcher three avoids differentiating the teaching method's observed effect according to students' *individual* observed IQ scores. How the true IQ scores in a sample vary thus becomes *theoretically independent* of the new teaching method. Consequently, the sampling error that is associated with the sample's mean true score (m_{Ti}) can be *subtracted* from the mean of students' individual observed IQ scores (m_{Xi}), i.e., $m_{Xi} - m_{Ti}$. Since $X_i - T_i = e_i$ (see Sect. 2), researcher three thus retains only the measurement error that is associated with the observed effect on the *average* student's IQ score (s_{ei}).

Researcher three must consequently recalculate the local error variances (s^2_{e-Cn}), as well as the local standard deviations (s_{e-Cn}) (Borsboom & Mellenbergh, 2002; Lord & Novick, 1968). On the assumption that the measurement reliability of the IQ-scale—itsself a property of the scale rather than of the sample—comes to $r_{XX'} = 0.80$, then given

$$s^2_{e-Cn}^* = (s^2_{Cn} - r_{XX'} \times s^2_{Cn}), \text{ where } r_{XX'} = s^2_T / s^2_{Cn} \quad (5)$$

one finds

$$s^2_{e-C1}^* = (15^2 - 0.80 \times 15^2) = 125, \text{ where } s_{e-C1}^* = \sqrt{125} = 11.18 \text{ IQ points}$$

$$s^2_{e-C2}^* = (6^2 - 0.80 \times 6^2) = 7.20, \text{ where } s_{e-C2}^* = \sqrt{7.20} = 2.68 \text{ IQ points}$$

$$s^2_{e-C3}^* = (3^2 - 0.80 \times 3^2) = 1.20, \text{ where } s_{e-C3}^* = \sqrt{1.20} = 1.10 \text{ IQ points}$$

Once the sampling error is subtracted—which yields the measure $d^* = (m_1 - m_0) / s_{e-Cn}^*$ —the observed effects changes to $d^*_{C1} = 3 / 11.18 = 0.27$, $d^*_{C2} = 3 / 2.68 = 1.12$, and $d^*_{C3} = 3 / 1.10 = 2.73$ IQ points. In each country, therefore, the observed effect increases in both size and heterogeneity.³

So far, researcher three's error account merely acknowledges the measurement error that is associated with an *individual* IQ score (e_i). But observed differences from a sample mean reflect the variation of the *true* IQ scores in a sample. Researcher three's error account, therefore, fails to acknowledge that the locally different IQ-selection conditions entail that the true individual IQ scores in a sample are subject to an *additional* variance (s^2_T). (See formula 5, where s^2_{Cn} is subtracted from $r_{XX'} \times s^2_{Cn}$, but $r_{XX'} = s^2_T / s^2_{Cn}$). However, the failure to acknowledge s^2_T can only explain why the

³ Recall that our example presupposes a *between*-subjects design and that, as the IQ-scores are extracted from the observed variance, the residuals that are obtained before administering a treatment do quantify the teaching method's *average* effect without sampling bias (see main text, Sect. 2). In a *within*-subjects design, by contrast, and provided students' prior intelligence (before treatment) is controlled for, the *d*-value for the average effect would expectably be *twice* that of a between-subjects design (see Schäfer & Schwarz, 2019).

d^* -measure increases the observed effect's size and its heterogeneity. But it cannot explain why the error variance remains large.

3. Implementation condition

To explain the large error variance, researcher three speculates that each school system is *in its own way* subject to random influences (e.g., because of varying levels of achievement, motivation, classroom conditions, or other random factors). Since the variation of student's true IQ scores is already eliminated from sc_n , any remaining error variance that may account for the observed effect's heterogeneity (beyond random influences) can be explained only by two factors: the *measurement error* (s_e) that is associated to student's IQ-scores (e_i) and a local *implementation condition*. This implementation condition may reflect, for instance, that students or teachers accept the new teaching method to different degrees.

What the heterogeneity of the observed effect thus suggests is that the local implementation condition *interacts* with the teaching method's effect on the average student's IQ score. To postulate this interaction is conceptually equivalent to thinking of the teaching method's implementation as a transformation of a theoretical construct into a measurable praxis. Researcher three's error account must hence acknowledge a local *implementation error*. However, there is no straightforward way to separate the implementation error from the measurement error (s_e). All uncontrolled factors that go beyond individual-level influences, therefore, continue to be represented squarely by the measurement error.

Finally, even if the standard IQ measurement scale had the same reliability in each country, its measurement unit would continue to be *theoretically uninformed*. Absent an alternative way of identifying how the local implementation error varies—which would allow for the differences between education systems to be standardized away—the details of how the implementation error interacts with the observed IQ scores must depend on the local education system. As each local education system thus provides a distinct *quasi*-experimental setting (to which a distinct measurement quality is associated), each provides *its own* theoretical standard. For this reason, the local implementation condition can only be measured locally, but not theoretically.

4. Theoretical prediction

Once the observed effect is differentiated according to the local implementation condition, one obtains the d^{**} -measure for a *standardized normal* quasi-experimental setting. To arrive at the d^{**} -measure, one may eliminate students' true IQ scores from the d -measure. Alternatively, when starting from the d^* -measure, one may treat a *medium* implementation condition as a “calibrated” quasi-experimental condition. The differentiation according to the local implementation condition thus allows for three local (country-specific) effects to be theoretically predicted.

Provided country two is treated as the *average* standard country (*medium* implementation condition), one obtains $d_{C2}^* = d_{C2}^{**} = 1.12$ IQ points, where $s_{e2}^* = 2.68$ IQ points. Since $s_{e1}^* > s_{e2}^*$, if s_{e2}^* standardizes $s_{e1}^* = 11.18$ IQ points, then the theoretically predicted effect in country one increases from $d_{C1}^* = 0.27$ to $d_{C1}^{**} = 1.13$ IQ points, because $s_{e1}^{**} = s_{e1}^* / s_{e2}^* = 11.18 / 2.68 = 4.17$, and $d_{C1}^{**} = d_{C1}^* \times s_{e1}^{**} = 0.27 \times 4.17 = 1.13$ IQ points. For country two, since $s_{e1}^* = s_{e2}^*$, one obtains $d_{C1}^{**} = 1.13 \approx d_{C2}^{**} = 1.12$ IQ points. And for country three, since $s_{e3}^* < s_{e2}^*$, one obtains a theoretically predicted effect of $d_3^{**} = d_3^* \times 0.41 = 2.72 \times 0.41 = 1.12$ IQ points, because $s_{e3}^{**} = s_{e3}^* / s_{e2}^* = 1.10 / 2.68 = 0.41$ IQ points.

Relative to a *medium* implementation condition, therefore, one obtains *roughly the same* theoretically predicted effect throughout, namely $d_{C1}^{**} = 1.13 \approx d_{C2,C3}^{**} = 1.12$ IQ points. The slight inequality between d_{C1}^{**} and $d_{C2,C3}^{**}$ suggests the presence of *inhibiting* influences on the effect in country one and *facilitating* influences in countries two and three. In any case, the IQ score of the average student would be theoretically predicted to increase by about the same amount throughout.

Relative to the *best* implementation condition (country three), finally, the standardization factors are $s_{e3}^{**} = 1.10$ IQ points and $s_{e2}^{**} / s_{e3}^{**} = 2.68 / 1.10 = 2.44$ IQ points. Because $(m_1 - m_0)$ is now standardized by s_{C3} —thus neutralizing the standardization by s_{C2} —the IQ score of the average student would on the standard IQ measurement scale be theoretically predicted to increase by some 41 IQ points ($d^{**} = 16.80 \times 2.44 = 40.93$ IQ points).

Word count: 8312