Reliability, Replicability, and Validity: A Meta-Scientific Analysis

Kristian E. Markon

University of Iowa

Author Note

Correspondence concerning this article should be addressed to Kristian E. Markon, Department of Psychological and Brain Sciences, University of Iowa, Iowa City, IA 52242; kristian-markon@uiowa.edu

Abstract

The role of measurement error in replicability of psychological findings has become of increasing interest, with some researchers suggesting it is critical to replicability, and others arguing that it is likely secondary to other effects on generalizability of findings. This work examined the relationship between reliability and effect size in published registered replication (i.e., many-labs) projects (313 samples from 44 studies) and a meta-analysis (16 samples from 8 studies). Among multiple-item registered replication designs, at lower reliabilities effect size was near zero regardless of reliability; at greater reliabilities (above approximately 0.80), effect size appeared to increase with reliability for some effects but not others. However, among the broader set of registered replication studies, including single-item designs, number of items was not associated with greater effect size, and in fact decreased with measure length. In the meta-analysis, greater reliability was associated with greater validity coefficients throughout the range of reliability. Results point to the importance of measurement precision in replicability of psychological findings, but also to the importance of precision per se and not proxies such as measure length.

As replicability has become an increasingly critical focus in the behavioral sciences, attention has turned to the role of measurement quality in generalizability of results (e.g., Shaw, et al., 2020; Stanley & Spence, 2014; Zuo, Xu, & Milham, 2019). In theory, all other things being equal, results obtained using more reliable measurements should generalize more: reliability can be conceptualized in terms of measurement variability across sampling occasions (Brennan, 2001), and classical psychometric theory indicates empirical associations should be attenuated when the measurements involved have smaller reliabilities. As such, one argument is that effects should generalize more across study replications when measurements are more reliable, either in initial report samples or in replication samples (e.g., Shaw, et al., 2020; Stanley & Spence, 2014). From this perspective, reliability and measure quality has important consequences for replicability, either through stochastic generalizability effects, or unaccounted degrees of freedom afforded to researchers through "measure hacking" (Flake & Fried, 2020; Marshall, et al., 2000).

Traditional reliability estimates can provide misleading impressions of measurement error, however, due to misspecified assumptions, misestimation, or other factors. These types of considerations raise questions about how strongly traditional measurement error indices signal generalizability and replicability in practice. For example, traditional internal consistency reliabilities do not necessarily predict indices of score validity and generalizability such as temporal stability, heritability, or cross-observer agreement, even when test-retest reliabilities do (Clarke, 1956; Kasten & Nevo, 2008; McCrae, et al., 2011). It has similarly been noted that single-item measures often exhibit psychometric and empirical properties comparable to longer measures (e.g., Bowling, 2005; Konstabel, et al., 2017; Robins, Hendin, & Trzesniewski, 2001; Wanous, Reichers, & Hudy, 1997), even though classical and modern psychometric theory both suggest score reliability, and therefore validity, should increase with test length, all other things being equal.

One explanation for these observations is that many approaches to reliability and measurement information, especially those based on internal structural considerations, conflate error with valid unique variance (McCrae & Möttus, 2019). That is, individual items likely comprise valid unique variance that is predictively informative, in the sense of future responses or external criteria, even if it is not informative

about contemporaneous responses to other items. In this way, a measure comprising significant valid unique variance may appear to have low reliability by certain indices, yet be reliable with regard to unique factors and produce generalizable effects. Conversely, theory also suggests that multidimensional measures with lower reliability in a certain sense can have greater validity, if they are constructed so as to represent the measurand of interest more completely (Little, Lindenberger, & Nesselroade, 1999). These types of observations have led to conclusions that measurements are both more reliable than has been traditionally assumed, and that reliability as indexed by many traditional statistics is less central to replicability of effects in practice than might be concluded based on traditional theory and assumptions (McCrae, et al., 2011).

The role of measurement error in the replication crisis, at least in the social sciences, has also been questioned due to the complex relationship between measurement error, induced biases, and the relative effects of measurement error and other factors in generalizability of estimated effects (Loken & Gelman, 2017). From this perspective, stochastic errors of measurement per se are less realistic moderators of effect size estimates than other factors more closely related to variability in true processes underlying response, design, or analysis. This may be especially true if, as has been noted in the cognitive science literature (Hedge, Powell, & Sumner, 2018), many indices of measurement error are focused on individual differences, which might not be directly pertinent to experimental designs focused on maximizing group differences. Quantifying the error of a group difference estimate may be more relevant in such cases than the error in an individual difference estimate (although they can be considered in the same framework; Markon, 2017).

The purpose of this study was to empirically examine the role of measurement error in effect size estimation and replicability within psychological studies, to quantify the extent to which traditional measurement error metrics are related to generalizability of effects in a prototypic registered replication scenario. Registered replication projects (Klein, et al., 2018) provide useful opportunities to do so: large numbers of effect size estimates are preregistered across a variety of study sites, for a variety of effects of interest, which should provide sufficient opportunity for measurement precision to vary across measures, samples, and effects. This allows for estimation, within a preregistered design, of how

differing measurement properties affect effect size estimation.

A meta-analysis also conducted to examine the relationship between reliability indices and validity. Multiple-effect, multiple-lab registered replications have advantages of preregistration, of admitting potential heterogeneity in effects, but existing projects tend to be relatively homogenous in certain respects. Measures are often specific to a given target effect, for instance, are often narrow in scope, and are shorter compared to some measures (e.g., Shaw, et al. 2020). A meta-analysis of empirical reliability-validity relationships permits examination of a wider range of measures and effects, to determine how much findings from many-labs data generalizes to other types of measures and effects. A meta-analysis of studies reporting reliability-validity associations was therefore conducted.

Study 1: Registered Replication Data

Method

Samples

Samples were obtained from two sets of registered replication projects: the Many Labs 2 Project (Klein, et al., 2018), and all registered replication projects listed as published (i.e., Stage 2) through the Association for Psychological Science Registered Replication Reports Project as of September 2021 (APS, 2021). These were chosen as they represent a range of typical social psychological, cognitive, and personality effects, comprise a diversity of samples and settings, and have been examined with regard to moderators and other variables.

Analyses of the Many Labs 2 effects were preregistered (<https://osf.io/jf3ar/>); samples from the other APS registered replication projects were added later to increase the number of effects included in analyses. The preregistration protocol was applied to the entire set of studies. Results using the preregistered set of Many Labs 2 samples alone are similar to those for the entire set and are presented separately in the supplement (<https://osf.io/7wkpv/>).

Effects and Analyses

Per the preregistered protocol, reliabilities were calculated for primary effects where the dependent variable was not modeled directly as a repeated measure in effect estimation (for example, dependent variable items not modeled directly in a mixed effect model), and could be represented as a

simple additive composite of items (for example, a sum, mean, or percent); ratios and other similar nonadditive dependent variables were not included (this criteria corresponded to all Many Labs 2 studies except 10, 19, 8, 24, and 28, and all of the AMPS many labs projects). Three forms of reliability statistic were calculated for each included sample and measure when possible: α , ω_h and ω_{tot} .

Single-item measures were handled in two ways. First, analyses were conducted randomly imputing average interitem correlations, squared general factor item loadings, and communalities of items from multi-item scales (corresponding to analyses of α , ω_h and ω_{tot} , respectively), using simulation-based methods to model uncertainty in the single-item reliabilities (for example, in estimating confidence intervals and average estimates). Second, analyses were conducted omitting the single-item scales, and analyzing the number of items alone as a measurement error index.

Primary analyses of the relationship between reliability and effect size estimate were conducted using simple correlations (for instance, Pearson's r or rank correlation). Linear modeling was also used, without sample and effect of interest as random predictors, with both as predictors, and with each included alone. A primary hypothesis of interest was that effect size estimates would be greater in absolute magnitude when reliability is greater. A secondary hypothesis was that effect size estimates would specifically increase with the square root of the reliability, following classical formulations of attenuation due to unreliability (e.g., Bobko, Roth, & Bobko, 2001).

Results

Multiple-Item Measures

A scatterplot of effect size versus reliability is presented in Figure 1 (results presented here focus on ω_t ; results for other reliability indices were similar and are presented in Figure S1). As is evident from the figure, there was a tendency for effect size to increase with increased reliability, although this relationship was largely restricted to the upper end of the reliability range. For ω_t below approximately 0.80, there was little or no relationship between reliability and effect size, with effect size uniformly low regardless of reliability. For values above that, there was heterogeneity in reliability-effect size relationships, with greater reliability being associated with greater effect size for some effects, and effect size being near zero regardless of reliability for other effects (Horn's $D = 36372408$ for test of

heteroscedasticity, $p < .0001$). Overall the correlation between effect size and reliability was 0.368 (Pearson; 0.239 Kendall; $p < 0.0001$ for both).

As an exploratory analysis of heterogeneity in reliability effect-size relationships, a 2-class mixture model was fit to the multiple-item measures, under the hypothesis that one class would represent null effects, with no relationship between reliability and effect size, and another class would represent valid effects, with a positive relationship between reliability and effect size. This account was largely supported by the analyses, with one class characterized by lower average reliability (0.758), near-zero average effect size (0.009), and no relationship between reliability and effect size ($r=0.017$); and the other, higher average reliability (0.888), positive average effect size (1.10), and positive relationship between reliability and effect size ($r=0.61$). This mixture model is plotted in Figure 2; the null class was the more prevalent estimated class, representing 0.725 of the observed effect sizes; the valid class was the less prevalent estimated class, representing 0.275 of the data.

Single-Item Measures

Effect sizes for single-item measures are also presented in Figure 1, at the average estimated reliability for single-item measures given the other observed multiple-item measures (that is, the average estimated reliability of single items from the multiple-item measures). The mean and median estimated single-item reliabilities were 0.372 and 0.249, respectively, with a variance of 0.077 (in comparison, the mean and median multiple-item composite reliabilities were 0.794 and 0.825, with a variance of 0.020). Effect sizes for the single-item measures varied substantially, with mean and medians of 0.571 and 0.379, with a variance of 0.516. The effect sizes for the single-item measures were typically greater than those for the multiple-item measures (the mean and median d for multiple-item measures were 0.308 and 0.081, with a variance of 0.433; $W = 301199$, $p < .0001$, for difference in location between the single-item and multiple-item reliabilities).

Number of Items

Figure 3 plots the distribution of effect sizes by number of items in the measure. As can be seen, the shorter measures had larger effect sizes (Kendall's $r = -0.168$, $p < .0001$). This pattern was partially accounted for by greater heterogeneity among the single-item measures, although the two-item

measures had less heterogeneity in effect sizes than the single-item measures and also relatively large effect sizes compared to longer measures.

Linear Models of Effect Size and Reliability

Variance Associated with Effect, Sample, and Effect x Reliability Interaction. Table 1 presents the results of regression models of the relationship between effect size, reliability, and sources of variance in reliability, for multiple-item measures. Per preregistered analyses, effect and sample were modeled as random effects; many models did not converge, however, so they were treated as fixed effects instead. As can be seen from Table 1, effect of interest was a significant predictor of effect size, as was the reliability-effect interaction, but sample was not. This suggests the primary predictor of effect size, aside from reliability, was the effect of interest being studied; the sample from which a reliability and estimate was derived was not a significant predictor of effect size per se. The significant reliability-effect interaction is consistent with heterogeneity in effect size-reliability associations (Figure 1), such that reliability was associated with effect sizes for some effects, but was less related or unrelated to effect size for other effects.

Models in which multiple and single-items were jointly modeled, with single-item reliabilities being randomly imputed based on the multiple-item measures, resulted in similar conclusions. Reliability, effect of interest, and the effect-reliability interaction significantly predicted effect size (Supplementary Table S2; permutation $p = 0.000$, 0.000 , and 0.011 respectively), but sample did not ($p=0.0.361$).

Models Predicted by Unreliability Attenuation Theory. In contrast to predictions based on attenuation of validity corrections, the square root of the reliability did not add significantly to prediction of effect size beyond reliability itself (permutation $p = 0.7265$). Moreover, in models with a square root term, the regression weight for the square root of the reliability was negative, in the opposite direction as theory (β for the square root term = -14.91 ; β for the linear term = 10.67 ; intercept = 5.06), suggesting that the square root term was largely capturing nonlinearity in the data in the form of increasing relationships between reliability and effect size at larger values of the former.

Study 2: Meta-Analysis

Method

Samples

Studies and samples were included if they met the following criteria: report of association statistics between reliability and validity coefficients (e.g., the correlation between coefficient alphas and correlations with external criteria), or a primary focus on reliability-validity relationships with report of reliability and validity coefficients from which reliability-validity associations could be calculated (e.g., a paper that has as a primary focus reliability-validity relationships, with presentation of empirical reliability and validity coefficients, but no explicit calculation of correlation between those coefficients). Common examples of excluded papers were those discussing theoretical reliability-relationships without reporting empirical reliability and validity coefficients, or those reporting on differential reliability and validity of measures without discussion of reliability-validity associations across measures (e.g., a paper that compares the subscales of two measures in their psychometric and predictive properties, but that does not examine reliability-validity associations across subscales). Papers were identified by treeing forward and backward through the literature, and by searching for phrases such as “relationship[correlation] between reliabilities[alphas, dependabilities] and validities”, and “psychometric properties and validity”.

Based on the result of this search, 30 reliability-validity correlations from 16 samples in 8 studies were identified. 21 of these involved internal consistency statistics, and 9 involved short-term dependability (i.e., test-retest) statistics. Most of the validity criteria, representing 22 of the reliability-validity correlations, were observer ratings of some form; 5 were academic performance variables; and the remainder reflected miscellaneous other criteria. One was from a dissertation, and the remainder were in published articles (two of these were themselves meta-analyses). Most of the reported reliability-validity correlations were taken over samples of measures, but one was taken over samples of persons, and two were taken over some combination. The reliability-validity correlations were based on 56.7 reliability-validity criteria pairs on average ($sd=65.8$, median 30), ranging from 7 to 261. Studies included in the meta-analysis are listed in the supplement.

Analyses

Given the heterogeneity in measures and variables of focus in the included studies, random effects models were assumed *a priori* (using the metafor package for R). As each reliability-validity

correlation reflected a correlation of correlations, usually taken over measures, the sampling variance of each study used as a weight was modeled as the sum of the variances due to number of measures and due to number of persons. In cases where the sample size differed between validity and reliability coefficients, average sample size was used. Unweighted descriptive statistical approaches were also adopted for comparison.

Results

Overall, the mean reliability-validity correlation was 0.259, the median was 0.298, and the standard deviation was 0.225. The overall estimate under a weighted random effects model was 0.312. In general, results supported psychometric theory in that larger reliabilities were associated with larger validity coefficients ($p < 0.0001$ against a null hypothesis of zero, according to both Wilcoxon rank sum tests and random effects model tests).

A plot of the reliability-validity correlations is presented in Figure 4, where reliability-validity correlations are presented as a function of average reliability in the study, and separately for internal consistency and test-retest statistics. Although the reliability-validity correlations in this set of studies was larger for dependabilities compared to internal consistencies (mean $r = 0.353$ and 0.218 for dependabilities and internal consistencies, respectively), and larger in samples with greater mean reliability (Kendall's tau = 0.241), these differences were not significant (for the test of difference between dependabilities and internal consistencies, $W = 123$, $p = 0.205$ for the Wilcoxon rank sum test and $z = 1.106$, $p = 0.269$ for the random effects model test; for the test of association with mean reliability, $z = 1.830$, $p = 0.067$ for the test of Kendall's tau and $z = 0.907$, $p = 0.365$ for the random effects model test). Consistent with the registered replication studies, there appeared to be a slight increase in reliability-validity correlations in samples with mean reliabilities greater than approximately 0.80.

Discussion

In both the registered replication analyses and meta-analysis, reliability was related to validity in general: the magnitude of the overall relationship was similar in the two analyses — 0.368 and 0.312 Pearson correlations in each, depending on assumptions. Moreover, in both analyses, there was some

evidence for increase in the magnitude of the reliability-validity relationship at high levels of reliability. Overall, results across the two studies support classical emphasis on measurement in understanding observed relationships between psychological variables, and underscore the importance of reliability and measurement error in the generalizability of those observations.

Different authors have variously argued that stochastic measurement error might or might not play a significant role in replicability in the psychological sciences (e.g., Loken & Gelman, 2017; Shaw, et al., 2020). Results presented here arguably support both positions in different ways. In both the registered replication analyses and meta-analysis, reliability did significantly predict effect size in general. However, in the registered replication studies, this was true only above a certain level of reliability and only for certain effects. Results of the registered replication analyses underscore that a nonnull true relationship is necessary for observed associations to be moderated by reliability: a “true” null effect will not be elevated to a nonnull effect by virtue of improved measurement precision. Moreover, there was little evidence that measure length per se — often assumed to be related to greater measurement precision, all other things being equal — predicted greater effect size estimates, as observed length was actually inversely related to effect size in the studies examined.

It has been noted in many contexts that shorter measures, sometimes even single items, can produce reliable measurements – this is a tenet of modern item response theory (Embretson, 1996), and has been noted empirically by many authors (e.g., Cooke, Michie, Hart, & Hare, 1999; Edelen & Reeve, 2007). However, even in item response theory, reliability generally is assumed to increase with measure length, and in general measurement science, length is commonly assumed a priori to be a proxy for measure quality (e.g., Shaw, et al., 2020; Smith, McCarthy, & Anderson, 2000). The results presented here reinforce that, at least with regard to criterion-related effect size, measurement precision is key, and not length per se.

Relatedly, some authors have argued for an emphasis on dependabilities as estimates of reliability, noting that valid unique variance, such as that in items, can be conflated with error variance in internal-structural estimates of reliability (e.g., McCrae, R. R., & Möttus, 2019). Results of the registered replication studies are consistent with this perspective, in that measure length was not predictive of

greater effect size. The meta-analytic results were also partially consistent with this, in that the dependabilities were associated with larger reliability-validity associations. However, this difference between internal consistency and dependability statistics was not significant, so it is unclear how much these results would generalize to a larger set of persons, measures, and effects.

There were important differences between the registered replication analyses and meta-analysis. Most prominent, perhaps, is the difference in reliability-validity associations at low levels of reliability, with associations being essentially zero at lower levels of reliability in the registered replication analyses, and positive in the meta-analysis. This is likely due to the nature of the effects being examined in the two sets of studies: the registered replication studies emphasized replication or lack thereof of possibly null effects, but the meta-analytic studies leveraged likely nonnull effects to characterize reliability-validity relationships. Large true effects likely afford more room for observed variability in associations due to unreliability attenuation than small true effects.

Limitations

One significant limitation of the registered replication analyses is that reliabilities were not available for the single-item measurements. As a result, it is impossible to know if they would have exhibited similar patterns as the multiple-item measurements, had their reliabilities been available, or if the single-item measures or their corresponding effects of interest were different in important ways. It is unlikely that the single-item measures were associated with greater effect sizes on average than the multiple-item measures because they were shorter per se. The single-item reliabilities implied by the multiple-item measures did vary substantially, ranging from 0.031 to 0.971, so it is plausible that some of the single-item measurements associated with larger effect sizes were also relatively reliable. Results point to the need for further characterization of the psychometric properties of measures, including those that are relatively brief and used in experimental designs.

Another limitation is that although the registered replication studies and meta-analytic studies were focused on different domains, within each studies tended to be relatively homogenous in focus. The registered replication studies examined in these analyses, for example, were relatively focused on social and cognitive phenomena, and the meta-analytic studies tended to focus on personality and academic

achievement. It is plausible that other patterns of relationships might be observed in other areas, such as biological psychology. Relatedly, most of the registered replication studies examined here utilized experimental designs, which might exhibit other patterns than those focused on observational designs.

Implications for Measurement Decisions and Development

Fundamentally, these results demonstrate that measurement precision is important to making generalizable and replicable inferences about psychological effects. This precision per se, rather than proxies for it, is key, as is the validity of the effect being examined. In many ways, the current results underscore traditional perspectives that reliable measurements, and informative measures, are necessary but not sufficient in making generalizable inferences. Establishing measurement reliability helps avoid Type I inferential errors by demonstrating estimated unreliability does not account for null or quasi-null effect estimates, as well as Type II inferential errors, by demonstrating that an effect is present when reliable measures are used.

Although the current results underscore that single-item measures can produce relatively reliable measurements, some care is required in interpreting what those measurements are reliable with reference to. Shorter measures, or measures that are otherwise limited in scope (e.g., time or content), are arguably more likely to be constrained in the breadth of aspects of a measurand they are able to measure. Being reliable with reference to some subattribute of a measurand (e.g., hopelessness) does not mean that other subattributes (e.g., concentration problems, dysphoria, demotivation) of a more general measurand (e.g., depression) are also being measured reliably; adequate representation of the variable domain is also important for accurate estimation of its relationships with other variables (Little, Lindenberger, & Nesselroade, 1999). Research on intraindividual variation and measurement noninvariance also suggests that there may be individual or group differences in how measurands are best measured (e.g., some individuals might be more dysphoric when depressed, others more demotivated), suggesting that breadth may be necessary to adequately capture some constructs in a generalizable sense across populations. Reliability with reference to variables involved in the specific effect of interest is necessary, where the effect of interest is understood with regard to potential variation across individuals and groups.

These types of considerations might be broadened further to consider the measurand of interest in a particular design. Recently proposed frameworks suggest that reliability can be seen in terms of the information a measurement provides about a parameter of interest (Markon, 2017). Traditionally these parameters of interest have been values of an attribute of interest that vary across individuals (e.g., negative mood, political attitude, memory ability), but it may be advantageous to think in terms of other parameters a measurement might be providing information about, such as values of an attribute that differ across groups. To the extent that an optimal measure of group difference might be different from an optimal measure of an individual difference (e.g., due to interest in an experimental design), different choices of measures might be made. This would require a broadened conception of reliability, one that encompasses other forms of difference than individual differences, or other types of parameters.

References

- Association for Psychological Science (2021). *Registered Replication Reports*.
<https://www.psychologicalscience.org/publications/replication/ongoing-projects>
- Bobko, P., Roth, P. L., & Bobko, C. (2001). Correcting the effect size of d for range restriction and unreliability. *Organizational Research Methods*, 4(1), 46-61.
- Bowling, A. (2005). Just one question: If one question works, why ask several? *Journal of Epidemiology & Community Health*, 59(5), 342-345. <https://doi.org/10.1136/jech.2004.021204>
- Brennan, R. L. (2001). An essay on the history and future of reliability from the perspective of replications. *Journal of Educational Measurement*, 38, (4), 295-317. doi:10.1111/j.1745-3984.2001.tb01129.x.
- Clarke, H. W. (1956). *An experimental investigation of theorems relating to the structure and content of rating instruments* [Dissertation]. Ohio State University.
- Cooke, D. J., Michie, C., Hart, S. D., & Hare, R. D. (1999). Evaluating the screening version of the Hare Psychopathy Checklist—Revised (PCL:SV): An item response theory analysis. *Psychological Assessment*, 11(1), 3-13. <https://doi-org.proxy.lib.uiowa.edu/10.1037/1040-3590.11.1.3>
- Edelen, M. O., & Reeve, B. B. (2007). Applying Item Response Theory (IRT) Modeling to questionnaire development, evaluation, and refinement. *Quality of Life Research*, 16, 5-18.
- Embretson, S. E. (1996). The new rules of measurement. *Psychological Assessment*, 8(4), 341-349. <https://doi.org/10.1037/1040-3590.8.4.341>
- 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>
- Hedge, C., Powell, G., & Sumner, P. (2018). The reliability paradox: Why robust cognitive tasks do not produce reliable individual differences. *Behavior Research Methods*, 50(3), 1166-1186. <https://doi.org/10.3758/s13428-017-0935-1>
- Kasten, R., & Nevo, B. (2008). Exploring the relationship between interrater correlations and validity of peer ratings. *Human Performance*, 21(2), 180-197. <https://doi.org/10.1080/08959280801917792>

- Klein, R. A., Vianello, M., Hasselman, F., Adams, B. G., Adams, R. B., Alper, S., Aveyard, M., Axt, J. R., Babalola, M. T., Bahník, Š., Batra, R., Berkics, M., Bernstein, M. J., Berry, D. R., Bialobrzeska, O., Binan, E. D., Bocian, K., Brandt, M. J., Busching, R., ... Nosek, B. A. (2018). Many Labs 2: Investigating variation in replicability across samples and settings. *Advances in Methods and Practices in Psychological Science*, 1(4), 443–490. <https://doi.org/10.1177/2515245918810225>
- Konstabel, K., Lönnqvist, J.-E., Leikas, S., García Velázquez, R., Qin, H., Verkasalo, M., & Walkowitz, G. (2017). Measuring single constructs by single items: Constructing an even shorter version of the “Short Five” personality inventory. *PLOS ONE*, 12(8), e0182714. <https://doi.org/10.1371/journal.pone.0182714>
- Little, T. D., Lindenberger, U., & Nesselroade, J. R. (1999). On selecting indicators for multivariate measurement and modeling with latent variables: When “good” indicators are bad and “bad” indicators are good. *Psychological Methods*, 4(2), 192–211.
- Loken, E., & Gelman, A. (2017). Measurement error and the replication crisis. *Science*, 355(6325), 584–585. <https://doi.org/10.1126/science.aal3618>
- Markon, K. E. (2017). *A generalized definition of reliability based on Lindley information*. <https://psyarxiv.com/vgpfb/>
- Marshall, M., Lockwood, A., Bradley, C., Adams, C., Joy, C., & Fenton, M. (2000). Unpublished rating scales: A major source of bias in randomised controlled trials of treatments for schizophrenia. *British Journal of Psychiatry*, 176(3), 249–252. doi:10.1192/bjp.176.3.249
- McCrae, R. R., Kurtz, J. E., Yamagata, S., & Terracciano, A. (2011). Internal consistency, retest reliability, and their implications for personality scale validity. *Personality and Social Psychology Review*, 15(1), 28–50. <https://doi.org/10.1177/1088868310366253>
- McCrae, R. R., & Möttus, R. (2019). What personality scales measure: A new psychometrics and its implications for theory and assessment. *Current Directions in Psychological Science*, 28(4), 415–420. doi:10.1177/0963721419849559
- Robins, R. W., Hendin, H. M., & Trzesniewski, K. H. (2001). Measuring global self-esteem: Construct validation of a single-item measure and the Rosenberg Self-Esteem Scale. *Personality and Social*

Psychology Bulletin, 27(2), 151-161.

Shaw, M., Cloos, L. J. R., Luong, R., Elbaz, S., & Flake, J. K. (2020). Measurement practices in large-scale replications: Insights from Many Labs 2. *Canadian Psychology*, 61(4), 289–298.

<https://doi.org/10.1037/cap0000220>

Smith, G. T., McCarthy, D. M., & Anderson, K. G. (2000). On the sins of short form development.

Psychological Assessment, 12(1), 102-111.

Stanley, D. J., & Spence, J. R. (2014). Expectations for replications: Are yours realistic? *Perspectives on Psychological Science*, 9(3), 305–318. <https://doi.org/10.1177/1745691614528518>

Wanous, J. P., & Reichers, A. E. (1997). Overall job satisfaction: How good are single-item measures?

Journal of Applied Psychology, 82(2), 247-252.

Zuo, X-N., Xu, T., & Milham, M. P. (2019). Harnessing reliability for neuroscience research. *Nature Human Behavior*, 3, 768-771.

Figure 1

Note. Values plotted are effect size (d) versus reliability estimates (ω_t), for the total set of studies. Single-item effect size estimates are shown as blue squares at the mean values of the estimated single-item reliabilities (0.372).

Figure 2

Note. Values are effect size plotted as a function of reliability. Individual estimates are colored according to class membership in the 2-class mixture model discussed in the text. The sizes of points reflects the uncertainty of classification, with larger points reflecting greater uncertainty. The ellipses reflect the magnitude and direction of within-class covariances of the two classes.

Figure 3

Note. Values are effect size (d) versus number of items, across all measures, for the total set of studies.

Figure 4

Note. Values are reliability-validity correlations plotted as a function of mean sample reliability. Internal consistency reliabilities are plotted as filled circles; dependabilities are plotted as open squares.

Table 1: ANOVA of Linear Model of Effect Size

	<i>df</i>	<i>SS</i>	<i>MS</i>	<i>F</i>	<i>p</i>
<i>Reliability</i>	1	32.529	32.529	410.343	0.000
<i>Effect</i>	13	168.869	12.990	163.861	0.000
<i>Sample</i>	211	11.520	0.055	0.689	0.995
<i>Reliability x Effect</i>	9	2.139	0.238	2.998	0.004
<i>Residual</i>	320	25.368	0.079		

Note. Values in table are from ANOVA of the fixed-effect linear model predicting effect size from reliability (ω_i) estimate, effect of interest, sample, and the reliability-effect interaction, in the total set of studies, for multiple-item measure studies only. p-values are permutation p-values, obtained by comparing the observed F statistics to the distribution of F statistics obtained through 10000 permutations of the outcome variable (effect size d) with regard to the predictors. Tables for the preregistered studies only, and with single-item reliabilities imputed, are in the supplement.