The Validation Crisis in Psychology

Ulrich Schimmack

University of Toronto Mississauga

Contact: Ulrich.schimmack@utoronto.ca

The Validation Crisis in Psychology

Abstract

In this commentary on the state of validation research in psychology, I review Cronbach and Fiske's (1955) seminal article and point out that the term is widely used, but researchers rarely follow their recommendations. Most important, construct validation requires specification of a nomological net, which could be done with a structural equation model and construct validity should be quantified, which could be done by means of factor loadings in an SEM measurement model.

The Validation Crisis in Psychology

8 years ago, psychologists started to realize that they have a replication crisis. Many published results do not replicate in honest replication attempts that allow the data to decide whether a hypothesis is true or false.

The replication crisis is sometimes attributed to the lack of replication studies before 2011. However, this is not the case. Most published results were replicated successfully. However, these successes were entirely predictable from the fact that only successful replications would be published (Sterling, 1959). These sham replication studies provided illusory evidence for theories that have been discredited over the past eight years by credible replication studies.

New initiatives that are called open science are likely to improve the replicability of psychological science in the future, although progress towards this goal is painfully slow.

This blog post addresses another problem in psychological science. I call it the validation crisis. Replicability is only one necessary feature of a healthy science. Another necessary feature of a healthy science is the use of valid measures. This feature of a healthy science is as obvious as the need for replicability. To test theories that relate theoretical constructs to each other (e.g., construct A influences construct B for individuals drawn from population P under conditions C), it is necessary to have valid measures of constructs. However, it is unclear which criteria a measure has to fulfill to have construct

validity. Thus, even successful and replicable tests of a theory may be false because the measures that were used lacked construct validity.

Construct Validity

The classic article on “Construct Validity” was written by two giants in psychology; Cronbach and Meehl (1955). Every graduate student of psychology and surely every psychologist who published a psychological measure should be familiar with this article.

The article was the result of an APA task force that tried to establish criteria, now called psychometric properties, for tests to be published. The result of this project was the creation of the construct “Construct validity”

The chief innovation in the Committee’s report was the term construct validity. (p. 281).

Cronbach and Meehl provide their own definition of this construct.

Construct validation is involved whenever a test is to be interpreted as a measure of some attribute or quality which is not “operationally defined” (p. 282).

In modern language, construct validity is the relationship between variation in observed test scores and a latent variable that reflects corresponding variation in a theoretical construct (Schimmack, 2010).

Thinking about construct validity in this way makes it immediately obvious why it is much easier to demonstrate predictive validity, which is the relationship between observed tests scores and observed criterion scores than to establish construct validity, which is the relationship between observed test scores and a latent, unobserved variable. To demonstrate predictive validity, one can simply obtain scores on a measure and a criterion and compute the correlation between the two variables. The correlation coefficient shows the amount of predictive validity of the measure. However, because constructs are not observable, it is impossible to use simple correlations to examine construct validity.

The problem of construct validation can be illustrated with the development of IQ scores. IQ scores can have predictive validity (e.g., performance in graduate school) without making any claims about the construct that is being measured (IQ tests measure whatever they measure and what they measure predicts important outcomes). However, IQ tests are often treated as measures of intelligence. For IQ tests to be valid measures of intelligence, it is necessary to define the construct of intelligence and to demonstrate that observed IQ scores are related to unobserved variation in intelligence. Thus, construct validation requires clear definitions of constructs that are independent of the measure that is being validated. Without clear definition of constructs, the meaning of a measure reverts essentially to “whatever the measure is measuring,” as in the old saying “Intelligence is whatever IQ tests are measuring. This saying shows the problem of research with measures that have no clear construct and no construct validity.

In conclusion, the challenge in construct validation research is to relate a specific measure to a well-defined construct and to establish that variation in test scores are related to variation in the construct.

What are Constructs

Construct validation starts with an assumption. Individuals are assumed to have an attribute, today we may say personality trait. Personality traits are typically not directly observable (e.g., kindness rather

than height), but systematic observation suggests that the attribute exists (some people are kinder than others across time and situations). The first step is to develop a measure of this attribute (e.g., a self-report measure "How kind are you?"). If the test is valid, variation in the observed scores on the measure should be related to the personality trait.

A construct is some postulated attribute of people, assumed to be reflected in test performance (p. 283).

The term "reflected" is consistent with a latent variable model, where unobserved traits are reflected in observable indicators. In fact, Cronbach and Meehl argue that factor analysis (not principle component analysis!) provides very important information for construct validity.

We depart from Anastasi at two points. She writes, "The validity of a psychological test should not be confused with an analysis of the factors which determine the behavior under consideration." We, however, regard such analysis as a most important type of validation. (p. 286).

Factor analysis is useful because factors are unobserved variables and factor loadings show how strongly an observed measure is related to variation in an unobserved variable; the factor. If multiple measures of a construct are available, they should be positively correlated with each other and factor analysis will extract a common factor. For example, if multiple independent raters agree in their ratings of individuals' kindness, the common factor in these ratings may correspond to the personality trait kindness, and the factor loadings provide evidence about the degree of construct validity of each measure (Schimmack, 2010).

In conclusion, factor analysis provides useful information about construct validity of measures because factors represent the construct and factor loadings show how strongly an observed measure is related to the construct.

It is clear that factors here function as constructs (p. 287).

Convergent Validity

The term convergent validity was introduced a few years later in another seminal article on validation research by Campbell and Fiske (1959). However, the basic idea of convergent validity was specified by Cronbach and Meehl (1955) in the section "Correlation matrices and factor analysis"

If two tests are presumed to measure the same construct, a correlation between them is predicted (p. 287).

If a trait such as dominance is hypothesized, and the items inquire about behaviors subsumed under this label, then the hypothesis appears to require that these items be generally intercorrelated (p. 288)

Cronbach and Meehl realize the problem of using just two observed measures to examine convergent validity. For example, self-informant correlations are often used in personality psychology to demonstrate validity of self-ratings. However, a correlation of $r = .4$ between self-ratings and informant ratings is open to very different interpretations. The correlation could reflect very high validity of self-ratings and modest validity of informant ratings or the opposite could be true.

If the obtained correlation departs from the expectation, however, there is no way to know whether the fault lies in test A, test B, or the formulation of the construct. A matrix of intercorrelations often points out profitable ways of dividing the construct into more meaningful parts, factor analysis being a useful computational method in such studies. (p. 300)

A multi-method approach avoids this problem and factor loadings on a common factor can be interpreted as validity coefficients. More valid measures should have higher loadings than less valid measures. Factor analysis requires a minimum of three observed variables, but more is better. Thus, construct validation requires a multi-method assessment.

Discriminant Validity

The term discriminant validity was also introduced later by Campbell and Fiske (1959). However, Cronbach and Meehl already point out that high or low correlations can support construct validity. Crucial for construct validity is that the correlations are consistent with theoretical expectations.

For example, low correlations between intelligence and happiness do not undermine the validity of an intelligence measure because there is no theoretical expectation that intelligence is related to happiness. In contrast, low correlations between intelligence and job performance would be a problem if the jobs require problem solving skills and intelligence is an ability to solve problems faster or better.

Only if the underlying theory of the trait being measured calls for high item intercorrelations do the correlations support construct validity (p. 288).

Quantifying Construct Validity

It is rare to see quantitative claims about construct validity. Most articles that claim construct validity of a measure simply state that the measure has demonstrated construct validity as if a test is either valid or invalid. However, the previous discussion already made it clear that construct validity is a quantitative construct because construct validity is the relation between variation in a measure and variation in the construct and this relation can vary. If we use standardized coefficients like factor loadings to assess the construct validity of a measure, construct validity can range from -1 to 1.

Contrary to the current practices, Cronbach and Meehl assumed that most users of measures would be interested in a “construct validity coefficient.”

There is an understandable tendency to seek a “construct validity coefficient. A numerical statement of the degree of construct validity would be a statement of the proportion of the test score variance that is attributable to the construct variable. This numerical estimate can sometimes be arrived at by a factor analysis” (p. 289).

Cronbach and Meehl are well-aware that it is difficult to quantify validity precisely, even if multiple measures of a construct are available because the factor may not be perfectly corresponding with the construct.

Rarely will it be possible to estimate definite “construct saturations,” because no factor corresponding closely to the construct will be available (p. 289).

And nobody today seems to remember Cronbach and Meehl’s (1955) warning that rejection of the null-hypothesis, the test has zero validity, is not the end goal of validation research.

It should be particularly noted that rejecting the null hypothesis does not finish the job of construct validation. The problem is not to conclude that the test “is valid” for measuring- the construct variable. The task is to state as definitely as possible the degree of validity the test is presumed to have (p. 290).

One reason why psychologists may not follow this sensible advice is that estimates of construct validity for many tests are likely to be low (Schimmack, 2010).

The Nomological Net – A Structural Equation Model

Some readers may be familiar with the term “nomological net” that was popularized by Cronbach and Meehl. In modern language a nomological net is essentially a structural equation model.

The laws in a nomological network may relate (a) observable properties or quantities to each other; or (b) theoretical constructs to observables; or (c) different theoretical constructs to one another. These “laws” may be statistical or deterministic.

It is probably no accident that at the same time as Cronbach and Meehl started to think about constructs as separate from observed measures, structural equation model was developed as a combination of factor analysis that made it possible to relate observed variables to variation in unobserved constructs and path analysis that made it possible to relate variation in constructs to each other. Although laws in a nomological network can take on more complex forms than linear relationships, a structural equation model is a nomological network (but a nomological network is not necessarily a structural equation model).

As proper construct validation requires a multi-method approach and demonstration of convergent and discriminant validity, SEM is ideally suited to examine whether the observed correlations among measures in a multi-trait-multi-method matrix are consistent with theoretical expectations. In this regard, SEM is superior to factor analysis. For example, it is possible to model shared method variance, which is impossible with factor analysis.

Cronbach and Meehl also realize that constructs can change as more information becomes available. It may also occur that the data fail to provide evidence for a construct. In this sense, construct validation is an ongoing process of improved understanding of unobserved constructs and how they are related to observable measures.

The example from the natural sciences was the initial definition of gold as having a golden color. However, later it was discovered that the pure metal gold is actually silver or white and that the typical yellow color comes from copper impurities. In the same way, scientific constructs of intelligence can change depending on the data that are observed. For example, the original theory may assume that intelligence is a unidimensional construct (g), but empirical data could show that intelligence is multi-faceted with specific intelligences for specific domains.

Ideally this iterative process would start with a simple structural equation model that is fitted to some data. If the model does not fit, the model can be modified and tested with new data. Over time, the model would become more complex and more stable because core measures of constructs would establish the construct validity, while peripheral relationships may be modified if new data suggest that theoretical assumptions need to be changed.

When observations will not fit into the network as it stands, the scientist has a certain freedom in selecting where to modify the network (p. 290).

Too often psychologists use SEM only to confirm an assumed nomological network and it is often considered inappropriate to change a nomological network to fit observed data. However, structural equation modeling is as useful for testing of an existing construct as exploration of a new construct.

However, given the lack of construct validation research in psychology, psychology has seen little progress in the understanding of such basic constructs such as extraversion, self-esteem, or wellbeing. Often these constructs are still assessed with measures that were originally proposed as measures of these constructs, as if divine intervention led to the creation of the best measure of these constructs and future research only confirmed their superiority. Instead many claims about construct validity are based on conjectures rather than empirical support by means of nomological networks. This was true in 1955. Unfortunately, it is still true over 50 years later.

For most tests intended to measure constructs, adequate criteria do not exist. This being the case, many such tests have been left unvalidated, or a finespun network of rationalizations has been offered as if it were validation. Rationalization is not construct validation. One who claims that his test reflects a construct cannot maintain his claim in the face of recurrent negative results because these results show that his construct is too loosely defined to yield verifiable inferences (p. 291).

Given the difficulty of defining constructs and finding measures for it, even measures that show promise in the beginning might fail to demonstrate construct validity later and new measures should show higher construct validity than the early measures. However, psychology shows no development in measures of the same construct. The most widely used measure of self-esteem is still Rosenberg's scale from 1965 and the most widely used measure of wellbeing is still Diener et al.'s scale from 1984. It is not clear how psychology can make progress, if it doesn't make progress in the development of nomological networks that provide information about constructs and about the construct validity of measures.

Cronbach and Meehl are clear that nomological networks are needed to claim construct validity.

To validate a claim that a test measures a construct, a nomological net surrounding the concept must exist (p. 291).

However, there are few attempts to examine construct validity with structural equation models (Connelly & Ones, 2010; Zou, Schimmack, & Gere, 2013). [please share more if you know some]

One possible reason is that construct validation research may reveal that authors initial constructs need to be modified or their measures have modest validity. For example, McCrae, Zonderman, Costa, Bond, and Paunonen (1996) dismissed structural equation modeling as a useful method to examine the construct validity of Big Five measures because it failed to support their conception of the Big Five as orthogonal dimensions with simple structure.

Recommendations for Users of Psychological Measures

The consumer can accept a test as a measure of a construct only when there is a strong positive fit between predictions and subsequent data. When the evidence from a proper investigation of a

published test is essentially negative, it should be reported as a stop sign to discourage use of the test pending a reconciliation of test and construct, or final abandonment of the test (p. 296).

It is very unlikely that all hunches by psychologists lead to the discovery of useful constructs and development of valid tests of these constructs. Given the lack of knowledge about the mind, it is rather more likely that many constructs turn out to be non-existent and that measures have low construct validity.

However, the history of psychological measurement has only seen development of more and more constructs and more and more measures to measure this increasing universe of constructs. Since the 1990s, constructs have doubled because every construct has been split into an explicit and an implicit version of the construct. Presumably, there is even implicit political orientation or gender identity.

The proliferation of constructs and measures is not a sign of a healthy science. Rather it shows the inability of empirical studies to demonstrate that a measure is not valid or that a construct may not exist. This is mostly due to self-serving biases and motivated reasoning of test developers. The gains from a measure that is widely used are immense. Thus, weak evidence is used to claim that a measure is valid and consumers are complicit because they can use these measures to make new discoveries. For example, Bosson et al. (2000) showed weak convergent validity of the self-esteem IAT with other implicit measures, but Greenwald and Farnham (2001) ignored this evidence and claimed that the self-esteem IAT has construct validity as a measure of implicit self-esteem. Twenty year later, there is still no evidence for this claim (Falk et al., 2015; Schimmack, 2019).

Conclusion

Just like psychologist have started to appreciate replication failures in the past years, they need to embrace validation failures. Some of the measures that are currently used in psychology are likely to have insufficient construct validity. If this was the decade of replication, the 2020s may become the decade of validation. Maybe this is overly optimistic, given the lack of improvement in validation research since Cronbach and Meehl (1955) outlined a program of construct validation research. Ample citations show that they were successful in introducing the term, but psychologists failed in adopting the rigorous practices they were recommending. It is time to change this and establish clear standards of construct validation that psychological measures should meet. Most important, validity has to be expressed in quantitative terms to encourage competition for developing new measures of existing constructs with higher validity.

References

- Bosson, J. K., Swann, W. B., Jr., & Pennebaker, J. W. (2000). Stalking the perfect measure of implicit self-esteem: The blind men and the elephant revisited? *Journal of Personality and Social Psychology*, 79(4), 631-643. <http://dx.doi.org/10.1037/0022-3514.79.4.631>
- Campbell, D. T., & Fiske, D. W. (1959). Convergent and discriminant validation by the multitrait-multimethod matrix. *Psychological Bulletin*, 56(2), 81-105. <http://dx.doi.org/10.1037/h0046016>
- Connelly, B. S., & Ones, D. S. (2010). An other perspective on personality: Meta-analytic integration of observers' accuracy and predictive validity. *Psychological Bulletin*, 136(6), 1092-1122. <http://dx.doi.org/10.1037/a0021212>
- Cronbach, L. J., & Meehl, P. E. (1955). Construct validity in psychological tests. *Psychological Bulletin*, 52(4), 281-302. <http://dx.doi.org/10.1037/h0040957>
- Greenwald, A. G., & Farnham, S. D. (2000). Using the Implicit Association Test to measure self-esteem and self-concept. *Journal of Personality and Social Psychology*, 79(6), 1022-1038. <http://dx.doi.org/10.1037/0022-3514.79.6.1022>
- McCrae, R. R., Zonderman, A. B., Costa, P. T., Jr., Bond, M. H., & Paunonen, S. V. (1996). Evaluating replicability of factors in the Revised NEO Personality Inventory: Confirmatory factor analysis versus Procrustes rotation. *Journal of Personality and Social Psychology*, 70(3), 552-566. <http://dx.doi.org/10.1037/0022-3514.70.3.552>
- Sterling, T. D. (1959). Publication decisions and their possible effects on inferences drawn from tests of significance— or vice versa. *Journal of the American Statistical Association*, 54(285), 30-34. doi:10.2307/2282137
- Zou, C., Schimmack, U., & Gere, J. (2013). The validity of well-being measures: A multiple-indicator–multiple-rater model. *Psychological Assessment*, 25(4), 1247-1254. <http://dx.doi.org/10.1037/a0033902>