Running Head: POOR RELIABILITY AND EXCELLENT VALIDITY

When High Reliability (Assessed in non-Clinical Samples) May Indicate Poor Diagnostic Utility: A Cautionary Note

Nachshon Meiran

Ben-Gurion University of the Negev, Beer-Sheva, Israel

Author's note:

Nachshon Meiran, Department of Psychology and Zlotowski Center for Neuroscience, Ben-Gurion University of the Negev, Beer-Sheva, Israel. This research was supported by a grant from the Israel Science Foundation.

I wish to thank Joseph Tzelgov for helpful comments and Stephanie Knipprath for English proofreading.

Correspondence concerning this article should be sent to the author, nmeiran@bgu.ac.il.

Abstract

Objective: Neuropsychological tests are sometimes used to discriminate between non-pathological and pathological populations, and their reliability is frequently assessed using non-pathological samples. The objective of this note is to alert to the fact that, for these tests, high reliability as assessed in non-clinical samples *may* imply poor diagnostic-utility.

Method: Rational analysis and a mathematical proof were used.

Results: The results indicate a reciprocal relationship between reliability, assessed in non-clinical samples, and validity/diagnostic-utility such that when reliability is low, validity and diagnostic-utility may be high.

Conclusions: Diagnostic utility should be of prime consideration, and for that reason, reliability estimates, if done, should be assessed in samples including both pathological and non-pathological participants.

Keywords: reliability, validity, neuropsychological tests

When High Reliability (Assessed in non-Clinical Samples) May Indicate Poor Diagnostic Utility: A Cautionary Note

Neuropsychologists use tests for a variety of purposes. One is documenting the relative positioning of an individual along a continuum that varies in non-clinical populations (e.g., intelligence tests). Another is in distinguishing pathology from non-pathology. In this note, I discuss what may appear as a paradoxical relationship between reliability and validity that applies *only* to the latter kind of tests. I will demonstrate it by means of a hypothetical (admittedly, simplistic) case. Let us assume that I developed a test to diagnose a specific neuropsychological impairment such as memory impairment, and this test's score ranges between zero and two. Moreover, aside from occasional random measurement error, all the healthy controls score two on the test, while the amnesic participants all score either zero or one. This test proves highly sensitive and specific. Because a large non-clinical sample is easy to find, I compute this test's reliability in that sample, and realize that it is close to zero¹. In the remainder of this note, I will show why high reliability (as assessed in non-clinical samples) might be a bad sign. Actually, several excellent tests are designed with the logic of being easy for healthy participants so that nearly all of them reach ceiling performance (E. Russell, 2001; Uttl, 2005). One example out of many is the question "What country are we in?" taken from the Mini-Mental Status Examination used to detect dementia (Folstein, M.F., Folstein, S.E., & McHugh, 1975).

Why Does High Reliability Assessed in Non-Clinical Samples Indicates a Potential Problem?

In the following discussion, I employ classical test theory, since most of the reports on tests employ measures taken from this theory, and since this is still the test theory with which most psychology practitioners are familiar (Childs & Eyde, 2002, but see Reise & Waller, 2009; see Appendix, presenting the issue within modern Item Response Theory).

According to classical test theory, reliability sets the upper bound on validity (Lord, Novick, & Birnbaum, 1968). Yet, there are several exceptions to this rule (e.g., Loevinger, 1954; Nimon, Zientek, & Henson, 2012). The current note discusses an additional such case. Specifically, I argue that an optimal test to discriminate pathology from non-pathology is characterized by a bi-modal distribution, with one mode for each sub-population, and with minimal overlap between populations.

Specifically, I distinguish between true variance and variance associated with measurement error. Thus, the total observed test variance, σ^2_{Total} , is decomposed into $\sigma^2_{\text{True-Between}}$, the between group (pathological vs. non-pathological) variance in true score, $\sigma^2_{\text{True-Within}}$, the within-group variance in true score, and σ^2_{Error} , the measurement error,

such that $\sigma^2_{\text{Total}} = \sigma^2_{\text{True-Between}} + \sigma^2_{\text{True-Within}} + \sigma^2_{\text{Error}}$. Using this notation, the (discriminant or criterion-related) validity of this test is expressed as the proportion of True-Between variance out of the Total variance:

$$(1) \eta^2 = \sigma^2_{\text{True-Between}} / (\sigma^2_{\text{True-Between}} + \sigma^2_{\text{True-Within}} + \sigma^2_{\text{Error}})$$

Reliability (treated generically), as computed within a single (e.g., non-clinical) population, is thus:

$$(2) \rho_{\text{tt-Single population}} = \sigma^2_{\text{True-Within}} / (\sigma^2_{\text{True-Within}} + \sigma^2_{\text{Error}})$$

Note that Equation 2 does not contain $\sigma^2_{\text{True-Between}}$, which is what I am truly interested in detecting. More generally, the equation acknowledges the well-known fact that reliability depends on the population in which it is being estimated (e.g., Loevinger, 1954; Mellenbergh, 1996).

Most importantly, as $\sigma^2_{\text{True-Within}}$ decreases, validity (η^2) increases (because the denominator in Equation 1 shrinks down) and reliability ($\rho_{\text{tt-Single population}}$) decreases (because the numerator in Equation 2 shrinks down). Actually, when $\sigma^2_{\text{True-Within}}$ equals zero (e.g., all the non-amnesic participants' true score is two), validity is maximal and the single-population reliability is zero. Figure 1 demonstrates this point for a particular set of values. In this demonstration, validity dropped to about one-fifth of its value when reliability increased from zero to .80!

Figure 1: A demonstration of the reciprocal relation between validity and reliability, when reliability is assessed within a single (e.g., non-clinical) population.

This point can also be demonstrated by the common practice of interpreting a score as reflecting pathology when $|Z| > 2.0$ (e.g., Crawford, 2003). Since the standard score is the distance from the mean in standard deviation units, when the variance (and the standard deviation, in the non-clinical population) is small, even slight numerical discrepancies from the norm would be sufficient to detect pathology. Moreover, because the variance equals $\sigma^2_{\text{True-Within}} + \sigma^2_{\text{Error}}$, it would be minimized by minimizing $\sigma^2_{\text{True-Within}}$.

So far, I treated the issue from a purely theoretical point of view. In the sections below, I deal with the question whether this issue has any relevance to actual practices, namely – in tests that are in actual use, in their interpretation, and in how reliability is being assessed.

Is this Problem Reflected in Current Tests?

When considering some published meta-analyses, the answer to this question appears to be “possibly yes”. Table 1 (added here for demonstrative purposes and not intended to provide a definitive summary of existing evidence) presents meta-analytic reliabilities and validities. The reliabilities are taken from a recent meta-analysis (Calamia, Markon, & Tranel, 2013). The validities are taken from other meta-analyses that were chosen under the constraints that (a) the comparison is between a pathological group and a non-pathological group, (b) that results are presented for separate tests, and (c) that the meta-analytic reliability of these tests has been reported by Calamia et al.

Table 1: Demonstrative Meta-Analytic Reliability and Validity (Standardized Mean Difference Relative to Healthy Controls) of Neuropsychological Tests*

	Rel.	AD	ADHD	Schiz.	BPD	Euthymic	Rank
WAIS Block Design	.83	1.70	.41	.90			0.49 (5)
WAIS Comprehension	.77	1.50		1.29			0.52 (6)
WAIS Digit Span	.79	1.38	.66	.86			0.53 (7)
WAIS Digit Symbol	.85	1.72	.85	1.59	.84	.47	0.46 (4)
WAIS Information	.91	1.56		.73			0.86 (13)
WAIS Similarities	.77	1.60	.62	.88			0.62 (9)
WAIS Vocabulary	.90	1.19	.50	.94			0.73 (11)
CVLT Short, free recall	.65			.66		.73	0.57 (8)

CVLT Long, free recall	.74		.75		.71	0.81 (12)
AVLT/CVLT Total T1-5	.74/.75				.90	0.14 (1)
Trails - A	.66	.45	.80	.71	.52	0.89 (14)
Trails-B	.77	.64	.91	.99	.78	0.37 (3)
Boston Naming	.80	2.13	1.01			0.21 (2)
WCST Perseverative	.62	.36		.88	.76	0.64 (10)

---

* *Rel.=Reliability, taken from Calamia et al. (2013, Table 2). Values represent the model-based estimated test-retest reliability of healthy participants of age=40, with zero test-retest delay. AD=Alzheimer's Disease. Values are taken from a meta-analysis comparing AD to healthy aging (Christensen, Hadzi-Pavlovic, & Jacomb, 1991, Table 2, when AD is pooled across all severity levels). ADHD=Attention deficit and Hyperactivity Disorder (Frazier, Demaree, & Youngstrom, 2004, Table 2). Schiz.=First episode schizophrenia (Mesholam-Gately, Giuliano, Goff, Faraone, & Seidman, 2009, Tables 4,5,6,8,9). BPD=Bipolar disorder (Arts, Jabben, Krabbendam, & van Os, 2008, Table 2), Euthymic= Euthymic with BPD (Robinson et al., 2006, Table 2). WAIS=Wechsler Adult Intelligence Test; AVLT=Rey Auditory Verbal Learning Test; CVLT=California Verbal Learning Test; T1-5=Trials 1-5; WCST=Wisconsin Card Sorting Test. Rank: (1) The validities were rank-ordered separately for each pathology; (2) These ranks were normalized by dividing the ranks by the number of tests for which validities were available for the given pathology; (3) The values represent the average normalized ranks (low=good validity). The rank-ordering of these values appear in brackets (1=best). Note that the reliability values presented in this table are different from those in the theoretical analysis below, yet the two systems are consistent with one another.*

Since intelligence tests (the WAIS subsets in Table 1) were developed in order to make discriminations within the healthy population (see more below), their reliability is expected to be relatively high, and this is true for the values in Table 1 (average reliability for WAIS subsets=.83, vs .71 for the non WAIS tests). Interestingly, the best three tests in terms of their validity ("Rank" in table 1) are *non-WAIS* tests. Moreover, the average validity Rank of the non-WAIS tests (.50) is *better* than that of the WAIS subsets (.60).

Nonetheless, these discrepancies in validity might be due to the fact that the domains tapped by the non-WAIS tests are more important for the specific pathologies appearing in Table 1. A better comparison would be between tests that are nearly identical in terms of content, yet show different reliabilities. The only two tests reported

by Calamia et al. (2013) that nearly meet these requirements are the AVLT and CVLT that have been shown to have similar albeit non-identical content (e.g., Stallings, Boake, & Sherer, 1995), and have some indices that differ in their meta-analytic reliability. The relevant indices are Trial 1, Trial 5, and recognition memory, where CVLT indices have higher reliability than the equivalent AVLT indices (.505, .655, and .621, as compared with .454, .579, and .284, respectively). Luckily, I found one meta-analysis in which the AVLT and CVLT results were separated. It indicates that, as predicted, the AVLT shows better validity, .81, 1.11, and .69, versus .65, .86, and .44 in CVLT for Trial 1, Trial 5, and recognition, respectively (Mesholam-Gately, et al., 2009). The results are thus at minimum compatible with the claim that high reliability (as assessed in a non-clinical sample) may really be problematic, not only theoretically, but also in practice.

Is the Practice of Assessing Reliability within Non-Pathological Samples Prevalent?

A recent meta-analysis (Calamia, et al., 2013, Table 10) shows that most of the neuropsychological test reliability reports in the literature that they reviewed were based on estimates taken from *non-clinical* samples, with the proportion being roughly 60-65% (Calamia, personal email, March 26, 2015), and with a higher proportion for non-IQ tests, for which this practice is most problematic.

Does this Practice Lead to False Conclusions?

Uttil (2005) administered memory tests of various difficulty levels (i.e., number of items to be remembered) using a relatively large sample of non-pathological college students and showed that the tests suffer from severe ceiling effects. Despite acknowledging the use of memory tests for making diagnostic decisions (p. 460), Uttil argued that ceiling effects have severe adverse effects on both reliability and validity, and concluded that "The findings demonstrate that tests afflicted by severe ceiling effects are not suitable for many of their original purposes because they cannot measure differences among individuals" (p. 466).

Uttil's point is correct when using the tests to detect memory giftedness, for example. However, the conclusions are wrong when the tests are used to detect and impairment. This can be demonstrated regarding the 3-word recall task from the MMSE (Schultz-Larsen, Kreiner, & Lomholt, 2007). This study shows that over 80% of the non-demented participants recalled two or three words, with almost 50% recalling all three words, thus indicating that the task is trivial for non-demented people. Yet, of those participants who were declared demented in the last wave of the study, about 92% either did not recall any word or recalled just a single word.

How Should Reliability be Assessed?

Because validity (or diagnostic utility) is of prime importance, reliability estimates should not only be accurate, but also favor more valid tests. For this to happen, reliability estimates should be conducted in samples representing both pathological and non-pathological populations. As implied from existing testing standards (American Education Research Association, American Psychological Association, & National Council on Measurement in Education, 2014; Evers et al., 2013), a proper examination of reliability should be conducted on a sample comprising persons from both the pathological and the non-pathological populations. When reliability is examined in this fashion (especially when the proportion of participants with pathology approaches .5, as needed to maximize $\sigma^2_{\text{True-Between}}$), it would not result in such a paradox. This is seen in how reliability is defined in Equation 3:

$$(3) \rho_{\text{tt-All populations}} = (\sigma^2_{\text{True-Within}} + \sigma^2_{\text{True-Between}}) / (\sigma^2_{\text{True-Within}} + \sigma^2_{\text{True-Between}} + \sigma^2_{\text{Error}})$$

This equation shows that reliability increases (rather than decreases) with increasing validity (as defined in Equation 2), since an increase in validity implies an increase in $\sigma^2_{\text{True-Between}}$, which also contributes to the denominator in Equation 3.

Two Types of Tests

Neuropsychologists used (at least) two types of tests. Some tests show large variance within the non-pathological population. The best examples are the very widely used intelligence tests (Butler, Retzlaff, & Vanderploeg, 1991; Rabin, Barr, & Burton, 2005). Indeed, general intelligence has originally been defined as *the ability showing the largest variance within the non-pathological population* (the first unrotated factor, Spearman, 1904). Intelligence tests are vital tools in the diagnosis of deviant intelligence, in ruling out intellectual disability when indications of a more circumscribed pathology are observed, and in predicting future job performance. Yet, intelligence tests show the largest $\sigma^2_{\text{True-Within}}$. As such, they may be sub-optimal tools in detecting *highly specific* performance impairments such as amnesia, alexia, agnosia, etc. Other tests, showing ceiling effects (Uttl, 2005), produce small variance (and reliability) in the non-pathological range and are thus more optimally suited for the diagnosis of specific impairments.

The present conclusions stem naturally from well-known facts. One is that reliability is influenced by range restriction (e.g., Bartlett & Frost, 2008; E. W. Russell, 1992). The other is the general principle suggested by Carey and Gottesman (1978) that “reliability calculated for one base rate cannot be generalized to other situations where the base rate differs appreciably” (p. 1,456). Here, it means that the reliability computed in a non-clinical sample cannot be generalized to clinical settings in which some people have a pathology and others not.

I wish to end with a warning. As seen in Equations 1 and 2, the true within-population variance is not the only factor influencing reliability. Two other factors are the error variance and the true between-population variance. Thus, one should bear in mind that while high reliability *may* be a bad sign in *some* cases, low reliability is not necessarily a good sign because it may reflect small true between-group variance or a high measurement error.

References

- American Educational Research Association, American Psychological Association, & National Council on Measurement in Education (2014). *Standards for educational and psychological testing*. Washington, DC: American Educational Research Association.
- Arts, B., Jabben, N., Krabbendam, L., & van Os, J. (2008). Meta-analyses of cognitive functioning in euthymic bipolar patients and their first-degree relatives. *Psychological Medicine*, 38, 771–785.
<http://doi.org/10.1017/S0033291707001675>
- Baker, F. B. (2001). *The basics of Item Response Theory. Second edition*. For full text: <http://ericae.net/irt/baker>. Retrieved from <http://eric.ed.gov/?id=ED458219>
- Bartlett, J. W., & Frost, C. (2008). Reliability, repeatability and reproducibility: analysis of measurement errors in continuous variables. *Ultrasound in Obstetrics and Gynecology*, 31, 466–475. <http://doi.org/10.1002/uog.5256>
- Butler, M., Retzlaff, P. D., & Vanderploeg, R. (1991). Neuropsychological test usage. *Professional Psychology: Research and Practice*, 22, 510–512.
<http://doi.org/10.1037/0735-7028.22.6.510>
- Calamia, M., Markon, K., & Tranel, D. (2013). The robust reliability of neuropsychological measures: Meta-Analyses of test–retest correlations. *The Clinical Neuropsychologist*, 27, 1077–1105.
<http://doi.org/10.1080/13854046.2013.809795>
- Carey, G., & Gottesman, I. I. (1978). Reliability and validity in binary ratings: Areas of common misunderstanding in diagnosis and symptom ratings. *Archives of General Psychiatry*, 35, 1454–1459.
<http://doi.org/10.1001/archpsyc.1978.01770360058007>
- Childs, R. A., & Eyde, L. D. (2002). Assessment training in clinical psychology doctoral programs: What should we teach? What do we teach? *Journal of Personality Assessment*, 78, 130–144. http://doi.org/10.1207/S15327752JPA7801_08
- Christensen, H., Hadzi-Pavlovic, D., & Jacomb, P. (1991). The psychometric differentiation of dementia from normal aging: A meta-analysis. *Psychological Assessment: A Journal of Consulting and Clinical Psychology*, 3, 147–155.
<http://doi.org/10.1037/1040-3590.3.2.147>

- Embretson, S. E. (1996). The new rules of measurement. *Psychological Assessment*, 8(4), 341–349. <http://doi.org/10.1037/1040-3590.8.4.341>
- Evers, A., Muñiz, J., Hagemester, C., HstmæLingen, A., Lindley, P., Sjöberg, A., & Bartram, D. (2013). Assessing the quality of tests: revision of the EFPA review model. *Psicothema*, 25(Número 3), 283–291.
- Frazier, T. W., Demaree, H. A., & Youngstrom, E. A. (2004). Meta-Analysis of Intellectual and Neuropsychological Test Performance in Attention-Deficit/Hyperactivity Disorder. *Neuropsychology*, 18(3), 543–555. <http://doi.org/10.1037/0894-4105.18.3.543>
- Loevinger, J. (1954). The attenuation paradox in test theory. *Psychological Bulletin*, 51, 493–504. <http://doi.org/10.1037/h0058543>
- Lord, F. M., Novick, M. R., & Birnbaum, A. (1968). *Statistical theories of mental test scores*. Oxford, England: Addison-Wesley.
- Mellenbergh, G. J. (1996). Measurement precision in test score and item response models. *Psychological Methods*, 1, 293–299. <http://doi.org/10.1037/1082-989X.1.3.293>
- Mesholam-Gately, R. I., Giuliano, A. J., Goff, K. P., Faraone, S. V., & Seidman, L. J. (2009). Neurocognition in first-episode schizophrenia: A meta-analytic review. *Neuropsychology*, 23, 315–336. <http://doi.org/10.1037/a0014708>
- Nimon, K., Zientek, L. R., & Henson, R. K. (2012). The assumption of a reliable instrument and other pitfalls to avoid when considering the reliability of data. *Frontiers in Psychology*, 3, 102. <http://doi.org/10.3389/fpsyg.2012.00102>
- Rabin, L. A., Barr, W. B., & Burton, L. A. (2005). Assessment practices of clinical neuropsychologists in the United States and Canada: A survey of INS, NAN, and APA Division 40 members. *Archives of Clinical Neuropsychology*, 20, 33–65. <http://doi.org/10.1016/j.acn.2004.02.005>
- Reise, S. P., Ainsworth, A. T., & Haviland, M. G. (2005). Item Response Theory fundamentals, applications, and promise in psychological research. *Current Directions in Psychological Science*, 14, 95–101. <http://doi.org/10.1111/j.0963-7214.2005.00342.x>
- Reise, S. P., & Waller, N. G. (2009). Item Response Theory and clinical measurement. *Annual Review of Clinical Psychology*, 5, 27–48. <http://doi.org/10.1146/annurev.clinpsy.032408.153553>
- Robinson, L. J., Thompson, J. M., Gallagher, P., Goswami, U., Young, A. H., Ferrier, I. N., & Moore, P. B. (2006). A meta-analysis of cognitive deficits in euthymic patients with bipolar disorder. *Journal of Affective Disorders*, 93, 105–115. <http://doi.org/10.1016/j.jad.2006.02.016>
- Russell, E. (2001). Toward an explanation of Dodrills observation: High neuropsychological test performance does not accompany high IQs. *The Clinical Neuropsychologist*, 15, 423–428. <http://doi.org/10.1076/clin.15.3.423.10267>

- Russell, E. W. (1992). Reliability of the Halstead Impairment Index: A simulation and reanalysis of Matarazzo et al. (1974). *Neuropsychology*, 6, 251–259.
<http://doi.org/10.1037/0894-4105.6.3.251>
- Schultz-Larsen, K., Kreiner, S., & Lomholt, R. K. (2007). Mini-Mental Status Examination: Mixed Rasch model item analysis derived two different cognitive dimensions of the MMSE. *Journal of Clinical Epidemiology*, 60, 268–279.
<http://doi.org/10.1016/j.jclinepi.2006.06.007>
- Spearman, C. (1904). "General Intelligence," objectively determined and measured. *The American Journal of Psychology*, 15, 201–292.
- Stallings, G., Boake, C., & Sherer, M. (1995). Comparison of the California Verbal Learning Test and the Rey Auditory Verbal Learning Test in head-injured patients. *Journal of Clinical and Experimental Neuropsychology*, 17, 706–712.
<http://doi.org/10.1080/01688639508405160>
- Strauss, E., & Spreen, O. (1991). *A compendium of neuropsychological tests*. New York: Oxford University Press.
- Thomas, M. L. (2011). The value of Item Response Theory in clinical assessment: A review. *Assessment*, 18, 291–307. <http://doi.org/10.1177/1073191110374797>
- Uttl, B. (2005). Measurement of individual differences Lessons from memory assessment in research and clinical practice. *Psychological Science*, 16, 460–467.
<http://doi.org/10.1111/j.0956-7976.2005.01557.x>

Appendix

Modern test theory, namely Item-Response Theory (Baker, 2001; Embretson, 1996; Reise, Ainsworth, & Haviland, 2005) has made a similar point a long time ago. In this theory, the probability of a person to score high on an item is a function of the person's ability and item difficulty. An important item (and also test) characteristic is *Information*, which describes the precision with which a given item assesses ability. Item information is expressed as the reciprocal of the standard error of measurement (Baker, 2001, Equation 6-1).

Unlike in classical test theory, where reliability (and the standard error of measurement) is assumed to be constant across item-difficulty levels, in Item-Response Theory, Information is highest around the item's difficulty (e.g., Reise et al., 2005), and with entire tests, around the center of the ability scale of the items comprising the test. For these reasons, an optimal test to use for discrimination should mainly consist of items with difficulty around the cutoff value used for discrimination (Baker, 2001, p. 110). In the present case, these are items that are very easy for healthy controls and difficult for brain-damaged persons. Finally, given the fact that test Information (estimated using Item-Response Theory) and test reliability (estimated using classical test theory) are closely related (e.g., Thomas, 2011), this implies that an ideal such test would tend to have low within-population reliability.

Notes:

1. I ran a small-scale simulation on 300 "healthy participants", in which one or two items were not recorded with $P = .10$ (measurement error). In this simulation, the test-retest correlation turned out to be $-.03$.
2. Here I used the familiar sum-of-squares (SS) decomposition, according to which $SS_{\text{total}} = SS_{\text{True-between}} + SS_{\text{True-within}} + SS_{\text{Error}}$. The values represented in the main text are obtained when dividing both sides of the equation by N .