Test-Retest Reliability is not a Measure of Reliability or Stability: A Friendly Reminder

Lukas Röseler

Harz University of Applied Sciences and University of Bamberg

Daniel Wolf, Johannes Leder, and Astrid Schütz

University of Bamberg

Author Note

This research was partly funded by a graduate scholarship granted by the state of Saxony-Anhalt, Germany, to Lukas Röseler. The source of funding had no involvement in the study design, data analysis, interpretation of data, writing, or decision to submit.

We thank Jane Zagorski for language editing.

Abstract

We argue that the test-retest reliability coefficient, which is the correlation between a measurement and a repeated measurement using the same diagnostic instrument in the same sample (sometimes referred to as repeatability or falsely referred to as stability), is by itself not an appropriate measure of the reliability of the diagnostic instrument or of the stability of the construct in question. In combination with an actual coefficient of reliability such as Cronbach's alpha, the test-retest reliability coefficient can be used to estimate and compare the stabilities of constructs using a procedure based on the correction for attenuation. However, results from a simulation study showed that classically constructed confidence intervals for the estimator exhibit under-coverage and thus cannot be interpreted correctly.

Keywords: test-retest reliability, repeatability, reliability, validity, stability, correction for attenuation

Test-Retest Reliability is not a Measure of Reliability or Stability: A Friendly Reminder

Reliable measurement is a necessary condition for informative psychological science. An unreliable score (i.e., $r < .5$) is mere measurement error and cannot be correlated with any other variable (e.g., Röseler, Schütz, & Starker, 2019). In spite of its importance, reliability is completely ignored in some lines of research (Parsons, Kruijt, & Fox, 2018) and misunderstood in many others. An even more problematic case involves so-called *test-retest reliability*, for which we assembled this friendly reminder on how to understand and apply it.

We distinguish between stability and reliability where high stability refers to the situation in which a construct's "true" score remains unchanged in relation to other people's scores (i.e., the valid measurement of a stable personality trait), and high reliability refers to a diagnostic instrument's accuracy and its sensitivity to differences between true scores at a given point in time. However, the two terms are often used synonymously. For example, test-retest reliability is also sometimes labelled *stability* (e.g., Döring & Bortz, 2016, p. 465), and sometimes, especially in discussions about replicability, the terms stability and reliability of paradigms and phenomena are used interchangeably (e.g., Nieuwenstein et al., 2015, p. 1).

For example, measures of heart rate that use today's technology are highly reliable, and ongoing changes in interbeat intervals are accurately reflected by changes in the measure. Because the interbeat interval is unstable, test-retest reliability would also be reduced by changes in the true value and would therefore not be 1. By contrast, psychological measurements, such as the measurement of extraversion on the basis of the NEO-FFI, are not a perfect mirror of "true" scores either because they are not perfectly reliable. Thus, using the test-retest reliability as an indicator of stability would lead to an underestimation of stability. Still, according to the definition of personality traits such as extraversion, such traits must be stable.

When people speak of test-retest reliability as a diagnostic instrument's accuracy, they implicitly assume that the construct that is being measured is stable. And when test-retest reliability is used as an indicator of a construct's stability, there is an implicit assumption that the diagnostic instrument is reliable. Thus, test-retest reliability depends on both the reliability of the diagnostic instrument and the stability of the construct.¹ On the basis of this reasoning, we argue that when considered in isolation, the classification and both labels (i.e., test-retest reliability and stability) are misleading. Instead, the term test-retest reliability should be replaced by the term *test-retest correlation*, and it should be used only after considering the reliability of the measurement (e.g., internal consistency) at the different points in time.

Test-Retest Reliability is a not a Good Measure of Reliability

Test-retest reliability can be computed as the correlation of the two measurements when the measurement is conducted twice with a single diagnostic instrument in the same group of people. If the scores obtained by the first and second measurements yield a strong and positive correlation, the diagnostic instrument is assumed to have high test-retest reliability. A widely recognized problem (e.g., Döring & Bortz, 2016, p. 466) is that test-retest reliability does not make sense as a measure of reliability for unstable constructs. This is because the diagnostic instrument can be very accurate, but due to the changeable nature of the construct, consecutive measurements will yield different results in terms of the overall level of test scores but also in the rank ordering of the cases. For example, it is possible to measure the number of heart beats in 1 min with a sophisticated diagnostic instrument such as an ECG and again after some time. Let us assume that the diagnostic instrument is perfectly reliable ($r = 1.00$) because we can count the true value, that is, the correct number of heart beats, without any error. However, the correlation

¹ Note that it actually depends on even more factors, such as the sample and other factors that were comprehensively discussed by Arnold (1996), which we left out for the sake of simplicity.

between two measurements at different points in time would not be 1.00 because the number of heart beats per minute changed (i.e., heart rate is not perfectly stable). The problem that the test-retest reliability is negatively associated with the instability of constructs is probably one of the reasons why there is typically only a small correlation between the test-retest correlation and the internal consistency ($r = .25$; Chmielewski & Watson, 2009).

Especially when measurements are psychological in nature, another problem that can occur when the same diagnostic instruments are administered two times in a row is that participants might remember their answer from the first run. Consider math problems for which participants might remember their answers because they just solved the problems a few minutes ago. A similar situation can occur for self-reports. In these cases, the test-retest correlation can be enormously inflated without revealing anything about the reliability of the diagnostic instrument.

Test-Retest Reliability is not a Good Measure of Stability

The test-retest correlation is not suitable for estimating a diagnostic instrument's reliability, but it is important to ask whether it can provide information about its validity, or more specifically its stability. McCrae, Kurtz, Yamagata, and Terracciano (2011) demonstrated that in comparison with internal consistency, the test-retest correlation provides a better predictor of stability. We argue that this happens because the correlation captures a mixture of reliability and validity. Hence, problems arise when the test-retest reliability is used as a measure of stability. For example, when measuring a stable trait, high construct validity may hold when there is a strong correlation between two measures taken years apart for the same participants. Now, consider a researcher who measured both extraversion and intelligence in a certain sample on two measurement occasions spaced 2 years apart, and say the correlation between the extraversion scores is significantly smaller than the correlation between the intelligence scores.

Can the researcher conclude that intelligence is more stable than extraversion? The not-so-widely recognized problem that we are referring to here is that even when measured constructs are very stable, test-retest reliability is not appropriate for estimating their stability over time because it is affected by the reliability at both measurement occasions. Given that having a reliable measurement (for which test-retest reliability is not a good estimate, as argued above) is the necessary condition for determining whether the measurement itself is valid, bad measurements end up with low validity, that is, low “stability.” The lower correlation that resulted for the extraversion scores between the two points in time in comparison with the intelligence scores does not necessarily mean that intelligence is more stable because the test-retest correlation is affected by the reliability of the measurement, the measurements of intelligence might simply be measured more reliably. If a measurement itself is unreliable, its stability will necessarily be low, independent of the actual stability of the construct. Still, although the degree of stability is an important issue in personality research, the stability of traits is sometimes assumed (e.g., Schmidt, Le, & Ilies, 2003, p. 208) instead of tested. In Table 1, we list examples of reliable and unreliable diagnostic instruments for stable or unstable constructs.

Table 1

Examples that Demonstrate that Reliability is a Necessary Condition for Stability

Reliability	Stability	
	Stable construct	Unstable construct
Higher reliability	Age in years	Heart beats per minute using an ECG
Lower reliability	Extraversion using the NEO-FFI (McCrae & Costa, 2004)	Mood using the PANAS (Watson, Clark, & Tellegen, 1988)

When Does Test-Retest Reliability Serve Which Goals?

Although validity estimators have been discussed and altered to counter the abovementioned problem of a lack of reliability (Chaves & Barber, 1974), this discussion has rarely been transferred to the case of stability (cf. McCrae et al., 2011). The measured test-retest reliability is a function of the reliability of the diagnostic instrument (i.e., the correlation of the measured values with the true values at one point in time, which can be estimated by using alternative approaches such as Cronbach's alpha; Cronbach, 1951), and the stability of the construct (i.e., the correlation of the true values at two points in time). Note, however, that the internal consistency can be larger due to large redundancies (e.g., McCrae et al., 2011, p. 1). The relationship is described in Equation 1.

$$r_{tt} = stability * \sqrt{reliability_{t1} \times reliability_{t2}} \quad (1)$$

Thus, the test-retest reliability can still be a useful index, first, because reliability is a necessary condition for validity. High validity provides evidence of reliability. That is, if the test-retest correlation is high, it is possible to conclude that the diagnostic instrument is reliable *and* that the construct is stable.

Second, test-retest reliability can be considered along with actual reliability. That is, if it is possible to estimate the reliability using a different method than the test-retest reliability, for example, Cronbach's alpha, one can estimate the stability by reorganizing Equation 1. For this purpose, the test-retest reliability can be divided by the square root of the reliabilities' product similar to the procedure proposed by the correction for attenuation (Heise, 1969; Spearman, 1987). By reorganizing Equation 1 in this way, it is possible to estimate the true stability of a construct. Therefore, we call this the *stability estimator* (see Equation 2). If the test-retest reliability is low (e.g., $r_{tt} = .5$) and the measurement is inaccurate (e.g., reliability = .7), the

stability can still be high ($.5/.7 = .71$). We provide an illustration of the relations between stability, reliability, and the resulting test-retest reliability online (<https://l-air.shinyapps.io/relstab/>).

$$\text{stability estimator} = \frac{r_{tt}}{\sqrt{\text{reliability}_{t1} \times \text{reliability}_{t2}}} \quad (2)$$

Comparisons between different constructs' stabilities are usually conducted by using latent modeling. However, for latent modeling, the full data set is necessary. Using the proposed stability estimator, only the reliabilities and the test-retest correlation are necessary. For example, Robins, Fraley, Roberts, and Trzesniewski (2001) applied the correction for attenuation to rank-order stabilities of the Big Five in college students at the ages of 18 and 22 and found that the stability of Openness was higher than the stability of Neuroticism ($\text{stability}_{\text{Openness}} = .92$, $\text{stability}_{\text{Neuroticism}} = .63$, p. 628) even though the test-retest reliabilities of the two traits were very similar ($r_{tt \text{ Openness}} = .88$, $r_{tt \text{ Neuroticism}} = .89$, p. 625). They reported all the necessary values, that is, the internal consistencies at both time points and the test-retest correlation. These values were sufficient for significance testing for differences in stability estimators between all kinds of variables. Such tests would treat the stability estimator as a correlation coefficient, even though it is itself based on three empirical correlations. Assessing the accuracy (in our case: bias and error) of this estimate is thus necessary but has not yet been done to our knowledge.

Accuracy of the Stability Estimator

Computing confidence intervals and comparing the stabilities of two different constructs using the stability coefficient seems to be a simple task at first sight. As the stability estimator is a correlation, it would be reasonable to think that a simple Fisher's z-transformation for confidence intervals and a subsequent z test to compare two stability estimators should be sufficient. Looking more closely, however, it becomes apparent that the estimator is based on

three other estimators, namely, the test-retest correlation and the reliabilities of the two measurements. This might result in common confidence intervals that are too narrow. We ran a simulation study to investigate this concern. We further tested the stability coefficient's bias and error and what the values depend on.

Method

To assess the stability estimator's bias and error as well as the appropriateness of computing 95% confidence intervals using Fisher's z-transformation, we conducted a simulation using R (R Core Team, 2018) version 3.6 and the psych package (Revelle, 2018) version 1.9.12.31. Using a set of predetermined ranges for reliabilities, true stabilities, and sample sizes, we created random data representing true values for two points in time (the correlation between which represented the predetermined true stability) and their corresponding measurement values. In doing so, the correlation between the true value and the corresponding measurement value was represented by the square root of the true reliability. This computation works because, in classical test theory, as reliability is defined as the percentage of the total variance accounted for by true measurement variance, the respective correlation can be derived by extracting the square root of the reliability. Note that when we refer to true stability in the following, we use the correlation of the true values, which might deviate slightly from the determined value. For each combination of parameters, we computed multiple runs that differed due to the use of pseudo-random numbers. The stability estimator was computed from the simulated measurement values that were based on Equation 2 for all combinations for which the determined true stability was as least as high as the determined reliabilities. Stability estimators that were larger than 1 or smaller than -1 were rounded to 1 and -1 for some analyses, respectively. Note again that the reliability was estimated by using the measured and true values and thus represented an approach that

would not be possible using nonsimulated data. The code and the data for the simulation study can be found online (<https://osf.io/jvrn5>).

To evaluate whether the stability estimator could be treated as a correlation, we computed confidence intervals for the estimated stabilities and tested (i.e., counted) how many of these estimated intervals captured the determined true values.

Furthermore, to analyze the influence of each parameter involved in the computation of the stability estimator, we aggregated the respective 500 runs that were based on identical parameters and computed linear regressions with the simulation parameters (i.e., true reliability, empirical stability, and sample size) for (a) the difference between the stability estimator and the true stability (i.e., systematic deviance from the true values or bias) and (b) the standard deviations of the estimate for 500 runs (i.e., unsystematic deviance from the true values or error), respectively. The regression analyses were computed for the full simulated data (no exclusions due to estimates greater than 1 or smaller than -1). In this way, we aimed to assess the extent to which the stability estimator's bias and error depended on the simulation parameters' true reliability, true stability, and sample size. A perfect estimator would have neither bias nor error. Our use of bias and error is illustrated in *Figure 1*.

Figure 1. Examples of error and bias used in the simulation study. Error bars represent variances.

Results and Discussion

For 13.40% of the cases, no confidence interval could be computed (e.g., a stability estimator > 1 , computed according to Equation 2). In 45.14% of the remaining cases (39.12% of all cases), the confidence interval of the stability estimator captured the true stability, indicating that the confidence intervals were too narrow. This confirmed our concerns that the stability estimator consisted of three other estimates itself and that the error variances could have accumulated. Further analyses revealed that the coverage did not depend on the determined sample size but on the determined reliability and stability: High reliability and low stability were associated with greater coverage. Nevertheless, the coverage was below 95% across all cases. Results are presented in *Figure 2*.

Figure 2. Coverage of the 95% confidence intervals for the stability estimators for sample size, reliability, and stability. Dashed lines represent the 95% confidence intervals. The data were not aggregated, and no cases were excluded.

The regression analyses in which we used the simulation parameters to predict the stability estimator's bias and error revealed that the empirical stability estimator's error but not its bias was related to all the varied parameters. There were significant main effects of all parameters on the error (i.e., the standard deviation of the stability estimator), but there were no effects of any of the parameters on the bias (i.e., the difference between the true value and the stability estimator). Not surprisingly, the stability estimator's error was smaller for large sample sizes, large reliabilities, and large stabilities. The use of the stability estimator—especially for large sample sizes with a reliability and a stability close to 1—can thus be recommended. However, once the true stability reaches 1, the estimator will be larger than 1 in half of the cases, and it will no longer be possible to compute the confidence intervals. Note that the results also depend on the determined simulation parameters (i.e., values for the reliability, true stability, and sample size) and are based on the assumption that the relationships are linear, although the influence of the sample size clearly follows the law of large numbers. An overview of the results

is presented in *Figure 3* and Table 2. The code used for the simulation can be found online (<https://osf.io/jvrn5>).

Table 2

Simulation Parameters and their Impact on the Estimator's Bias and Error

Predictor	Min	Max	Step size	Dependent variable: Bias	Dependent variable: Error
Intercept	—	—	—	$b_0 = -1.62\text{E-}4, t(116) = -0.09, p = .931$	$b_0 = 0.23, t(116) = 20.09, p < .001$
True reliability	.5	.7	.1	$b_1 = -3.45\text{E-}3, t(116) = -1.62, p = .107$	$b_1 = -0.16, t(116) = -12.02, p < .001$
True stability	.7	1	.1	$b_2 = 2.02\text{E-}3, t(116) = 1.30, p = .197$	$b_2 = -.01, t(116) = -7.00, p < .001$
Sample size	100	1000	100	$b_3 = 7.83\text{E-}7, t(116) = 1.30, p = .198$	$b_3 = -5.98\text{E-}5, t(116) = -16.28, p < .001$

Note. For each possible parameter combination, 500 simulations were conducted and aggregated for the bias and error assessments.

Figure 3. Error and bias of the stability estimator by sample size, reliability, and stability. The absence of bias is supported by the median of 0. Error is represented by the range of the boxplots. The data were not aggregated, and no cases were excluded.

Limitations

Using the proposed stability estimator to estimate the correlation between the true values at two points in time yielded a value that should not be treated as a correlation when computing confidence intervals. If a scale's internal consistency is extremely high (i.e., close to 1), the stability estimator will likely be larger than 1 and will prohibit further processing, such as the computation of confidence intervals. Alternatively, correlations can be rounded to 1 or -1. Moreover, confidence intervals that are computed by using Fisher's z-transformation will provide underestimated boundaries (i.e., intervals that are too narrow and that yield under-coverage). In almost every case, researchers tend to treat the stability estimator as though it is more precise than it actually is.

Recommendations

In light of the arguments presented, we advise researchers to avoid using the misleading terms test-retest reliability and stability. Instead, we suggest that a more precise terminology be used. We recommend the term test-retest correlation for the correlation of two measurements using the same diagnostic instrument at two points in time (alternatively, one can stick with the term repeatability), whereas the term stability concerns an aspect of the validity of a trait. A construct's true stability itself is based merely on theoretical assumptions and can hardly be assessed. As a solution, we provide the stability estimator. We could not find evidence that the stability estimator is systematically biased, but we warn that the classic computation of confidence intervals leads to under-coverage (i.e., intervals that are too narrow). In cases where repeatability is mentioned as a kind of reliability, we suggest that researchers note that repeatability is not actually a measure of reliability.

Conflicts of Interest

The authors declare that there were no conflicts of interest with respect to the authorship or the publication of this article.

Author Contributions

- LR created the first versions of the shinyapp, the simulation study, and the manuscript.
- LR and DW revised the shinyapp and the simulation study.
- All authors revised and approved the final manuscript.

Data, Material, and Online Resources

The simulation code and data and the shinyapp code are available at <https://osf.io/jvrn5>.

The shinyapp is available at <https://l-air.shinyapps.io/relstab/>.

References

- Arnold, M. E. (1996). Influences on and Limitations of Classical Test Theory Reliability Estimates. *Paper Presented at the Annual Meeting of the Southwest Educational Research Association, New Orleans.*
- Chaves, J. F., & Barber, T. X. (1974). Cognitive strategies, experimenter modeling, and expectation in the attenuation of pain. *Journal of Abnormal Psychology*, 83(4), 356–363.
<https://doi.org/10.1037/h0036789>
- Chmielewski, M., & Watson, D. (2009). What is being assessed and why it matters: The impact of transient error on trait research. *Journal of Personality and Social Psychology*, 97(1), 186–202. <https://doi.org/10.1037/a0015618>
- Cronbach, L. J. (1951). Coefficient alpha and the internal structure of tests. *Psychometrika*, 16(3), 297–334. <https://doi.org/10.1007/BF02310555>
- Döring, N., & Bortz, J. (2016). *Forschungsmethoden und Evaluation in den Sozial- und Humanwissenschaften* (5. vollständig überarbeitete, aktualisierte und erweiterte Auflage). *Springer-Lehrbuch*. Berlin, Heidelberg: Springer. Retrieved from <http://dx.doi.org/10.1007/978-3-642-41089-5> <https://doi.org/10.1007/978-3-642-41089-5>
- Heise, D. R. (1969). Separating Reliability and Stability in Test-Retest Correlation. *American Sociological Review*, 34(1), 93. <https://doi.org/10.2307/2092790>
- McCrae, R. R., & Costa, P. T. (2004). A contemplated revision of the NEO Five-Factor Inventory. *Personality and Individual Differences*, 36(3), 587–596.
[https://doi.org/10.1016/S0191-8869\(03\)00118-1](https://doi.org/10.1016/S0191-8869(03)00118-1)
- 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 : An Official Journal of the Society for Personality and Social Psychology, Inc, 15(1), 28–50. <https://doi.org/10.1177/1088868310366253>

Nieuwenstein, M. R., Wierenga, T., Morey, R. D., Wicherts, J. M., Blom, T. N.,

Wagenmakers, E. M., & van Rijn, H. (2015). On making the right choice: A meta-analysis and large-scale replication attempt of the unconscious thought advantage. *Judgment and Decision Making*, 10(1), 1–17. Retrieved from <http://journal.sjdm.org/14/14321/jdm14321.html>

Parsons, S., Kruijt, A.-W., & Fox, E. (2018). Psychological Science needs a standard practice of reporting the reliability of cognitive behavioural measurements. Advance online publication. <https://doi.org/10.17605/OSF.IO/6KA9Z>

R Core Team (2018). R [Computer software]. Vienna, Austria: R Foundation for Statistical Computing. Retrieved from <https://www.R-project.org/>

Revelle, W. (2018). psych: Procedures for Personality and Psychological Research, Version = 1.8.3 [Computer software]: Northwestern University, Evanston, Illinois, USA. Retrieved from <https://CRAN.R-project.org/package=psych>

Robins, R. W., Fraley, R. C., Roberts, B. W., & Trzesniewski, K. H. (2001). A longitudinal study of personality change in young adulthood. *Journal of Personality*, 69(4), 617–640. <https://doi.org/10.1111/1467-6494.694157>

Röseler, L., Schütz, A., & Starker, U. (2019). Cognitive Ability does not and cannot Correlate with Susceptibility to Anchoring Effects. Advance online publication. <https://doi.org/10.31234/osf.io/bnsx2>

Schmidt, F. L., Le, H., & Ilies, R. (2003). Beyond alpha: An empirical examination of the effects of different sources of measurement error on reliability estimates for measures of individual-

differences constructs. *Psychological Methods*, 8(2), 206–224. <https://doi.org/10.1037/1082-989X.8.2.206>

Spearman, C. (1987). The Proof and Measurement of Association between Two Things. *The American Journal of Psychology*, 100(3), 441–471. Retrieved from <https://www.jstor.org/stable/1422689>

Watson, D., Clark, L. A., & Tellegen, A. (1988). Development and validation of brief measures of positive and negative affect: The PANAS scales. *Journal of Personality and Social Psychology*, 54(6), 1063–1070. <https://doi.org/10.1037/0022-3514.54.6.1063>