On the Reliability of Implicit Measures: Current Practices and Novel Perspectives

Maarten De Schryver, Sean Hughes, Jan De Houwer, and Yves Rosseel

Ghent University

Author Note

MDS, SH, JDH, and YR, Department of Psychology, Ghent University, Belgium.

Preparation of this paper was made possible by Grant BOF16/MET_V/002 of Ghent

University to JDH. Correspondence concerning this article should be sent to

maarten.deschryver@ugent.be or sean.hughes@ugent.be.

Abstract

During the last two decades a new class of indirect measurement procedures has emerged that has been used widely in psychological science. These procedures were developed in order to circumvent the limitations of self-reports and crack open the hidden world of ‘implicit’ cognition. Yet despite their popularity there seems to be no general framework that constrains (or guides) the way in which we think about the reliability of implicit measures. Instead there are far too many (subjective) researcher degrees of freedom when it comes to deciding how to assess, interpret, and even report reliability. In this paper we introduce such a framework and argue that it can be used by novel and seasoned researchers alike to estimate and interpret the reliability of their implicit measures. Our approach draws on Latent Variable Modeling and the idea of parceling in order to approximate reliability (in the sense of consistency, equivalence, stability) and test the assumptions that underlie those approximations. We close by discussing the implications of our framework for the conceptualization of reliability in particular and for implicit cognition research more generally.

Keywords: Reliability, implicit measures, implicit cognition

On the Reliability of Implicit Measures: Current Practices and Novel Perspectives

At the turn of the century, a measurement revolution swept through psychological science altering much of the methodological, theoretical, and empirical landscape. Until then researchers interested in assessing people's attitudes, beliefs, and personality characteristics typically relied on *direct measurement procedures* such as semantic differential scales, feeling thermometers, and questionnaires. Although useful, these procedures were rooted in the assumption that people have both introspective access to, as well as the opportunity and motivation to accurately report on, their psychological attributes. The problem, however, is that this assumption often fails to hold (e.g., when measuring socially sensitive issues or in situations where the demand to respond in specific ways is high).

With the introduction and refinement of *indirect measurement procedures* researchers sought to circumvent these problems and crack open the hidden world of 'implicit' cognition. Broadly speaking, indirect procedures seek to (a) circumvent a person's ability to strategically control their behavior as well as (b) capture psychological processes, attributes, or content in ways that do not depend on introspective access. A vast and ever-growing library of indirect procedures have now been developed and applied to issues both in- and outside of psychology (see Gawronski & Payne, 2011). The most influential of these procedures include the Implicit Association Test (IAT; Greenwald, McGhee, & Schwartz, 1998), evaluative and semantic priming (e.g., Fazio, 2001), and the Affective Misattribution Procedure (AMP; Payne, Cheng, Govorun, & Stewart, 2005) (for more on indirect procedures see Gawronski & De Houwer, 2014; Nosek, Hawkins, & Frazier, 2011). The outcomes of indirect measurement procedures are typically referred to as implicit measures.¹

¹ In-line with De Houwer and Moors (2010), we define a measurement procedure as "direct" or "indirect" based on the way in which the measurement context is arranged to capture the behavior of interest. We define the outcome derived from measurement procedures as implicit or explicit based on the automaticity features of the process by which a psychological attribute influences measurement outcomes.

The introduction of any new type measure in psychology sparks questions about its psychometric properties. Implicit measures were no different, with researchers raising questions about their validity (e.g., Greenwald, Poehlman, Uhlmann, & Banaji, 2009) and reliability (e.g., Nosek, Greenwald, & Banaji, 2007). Indeed, it would not be too much of a stretch to suggest that reliability has been a persistent and thorny issue in implicit cognition research. This is because the vast majority of implicit measures seem to suffer from unacceptably low levels of reliability especially when compared to their explicit counterparts (see Cunningham, Preacher, & Banaji, 2001; Fazio & Olson, 2003; Gawronski, LeBel, & Peters, 2007; Uhlmann, Pizarro, & Bloom, 2008). Although there are a number of exceptions (e.g., the IAT and AMP) reliability estimates often range from “abysmally low (Bosson, Swann, & Pennebaker, 2000) to moderate (Kawakami & Dovidio, 2001)” (p.572) and are argued to have serious knock-on effects for cumulative scientific progress (LeBel & Paunonen, 2011; but see De Schryver, Hughes, Rosseel, & De Houwer, 2015). Yet despite this there is no general framework that constrains (or at least guides) the way in which we think about the reliability of implicit measures and specifies how the reliability of those measures can be determined. Instead there seems to be far too many (subjective) researcher degrees of freedom when it comes to deciding how to assess reliability, whether to even report it, and how to interpret it.

Looking at Reliability with a Fresh Pair of Eyes

The current paper has two goals. First, it aims to provide a quick primer for those interested in the concept of reliability and its relationship to implicit measures. We fully recognize that some readers will already be familiar with some of the statistical and psychometric concepts outlined in the initial sections of this paper. Yet many are not: a recent review of graduate training in psychology revealed that measurement theory and practice is often ignored in doctoral programs and that only a minority of students know how to apply

the methods of reliability correctly (Aiken, West, & Millsap, 2008; Graham, 2006). Thus we anticipated that many readers would benefit from a brief primer on the concept of reliability.

Our second (and main) goal is to introduce a general framework that can be used by novel and seasoned researchers alike to estimate and interpret the reliability of their implicit measures. If we are to have cumulative scientific progress then we need to be able to place trust in the accuracy of our measures. We also need procedures that can reproduce those measures time and time again. As we will argue, both the accuracy and reproducibility of measures is closely tied to the issue of reliability. Too many researcher degrees of freedom during the data collection and reporting stages can undermine our confidence in the accuracy and replicability of findings (e.g., Simmons, Nelson, & Simonsohn, 2011). Just as researchers may (either knowingly or unknowingly) adopt practices that artificially inflate the statistical significance of their findings (*'p-hacking'*) they might also adopt practices that inflate the reliability estimates they obtain (*'r-hacking'*). Given the prominent role that reliability plays in every stage of scientific evaluation it seems that we need clear guidelines and recommendations that explicate and constrain researcher degrees of freedom. With this in mind, we introduce a perspective on reliability that forces researchers to test the basic assumptions that underpin their reliability estimates and to clarify the individual and environmental factors that influence those estimates. Such an analysis leads to a heuristic framework for reviewing and evaluating reliability in implicit cognition research.

Our paper is organized in the following manner. Before we tackle reliability we first highlight a common conceptual confusion between procedures and tests (Part I). Here we argue that tests should not be equated with procedures, and that when we speak about reliability we are referring to the outcomes obtained from a particular test (taken by a specific sample) which has been administered in a wider context. In order to understand reliability we first need to understand how these three (sample, test, and context) parameters interact with

one another. In Part II we shift our attention to reliability and consider how it has traditionally been understood from a (strict) classical perspective. Consistent with others (e.g., Borsboom, 2005) we argue that when one tries to interface this traditional perspective with reality it quickly starts to break down and leads to the realization that reliability is approximated in three different ways (consistency, equivalence, and stability). Each of these approximations explores a different facet of reliability and involves manipulating the test-specific, general context, and individual parameters in a certain fashion. In Part III we propose that certain conceptual problems can be circumvented by drawing on the concepts of Latent Variable Modeling (LVM) and the idea of parceling. Although parceling has been subject to some controversy (Little, Cunningham, Shahar and Widaman, 2002) it is already widely (and implicitly) used to estimate reliability in implicit cognition research (e.g., split-half). Here we highlight the assumptions that underpin these various approximations and provides concrete methods for testing them. We then draw upon an existing dataset to provide a concrete example of our ideas (Part IV). We close by discussing the implications and consequences of our framework for reliability in particular and implicit cognition research more generally.

Part I: Clarifying the Difference between Procedures and Tests

A major impediment to communication and scientific progress is the tendency for researchers to use different terms when referring to the same phenomenon while at other times using similar terms to refer to different phenomena. This is particularly evident in the domain of implicit cognition, where tasks such as the IAT, AMP, and evaluative priming are variously labeled as implicit measures (Nosek et al., 2011), indirect measurement procedures (De Houwer, 2005), or tests (Greenwald et al., 1998) whereas the term “implicit measures” is sometimes reserved for scores obtained from those tasks (De Houwer & Moors, 2010). This terminological inconsistency can lead to conceptual confusion and theoretical misunderstanding and requires researchers to be clear in their respective terminology. In what

follows we outline precisely what we mean by procedure, test, sample, context, and measure before applying these terms to the concept of reliability.

Measurement Procedure

We conceptualize a measurement procedure as that core property of a task that does not vary from one measurement occasion to the next. Others may refer to this as a ‘paradigm’. For instance, when it comes to the IAT, researchers can systematically vary task parameters in the pursuit of their scientific goal, from the number of trials or blocks, nature of stimuli employed, or type of responses registered. In this way, the very same measurement procedure can be instantiated in a near infinite number of ways given the potential number of stimuli, responses, and other variable task parameters. What remains the same and unifies different variants of the IAT is a certain procedural consistency, namely way in which stimulus categories are mapped onto responses. More specifically, in the IAT, four stimulus categories are mapped on two responses, with two categories assigned to the first response and two to the second response, in a way that varies across the two main phases of the task.² Likewise, whereas the stimulus onset asynchrony (SOA), type of trials, and other parameters may vary in an AMP, the presentations of primes and the requirement to assign a target to one of two categories remains constant across (and thus unifies) different instantiations of the AMP. Such a consistent procedural core can be identified for any procedure whether it is an IAT, AMP, or another type of indirect procedure (also see De Houwer, 2003). If this core is changed in any way then the researcher has constructed a novel measurement procedure.

² Different descriptions of a measurement procedure are certainly possible. Good descriptions are those that are precise (i.e., allow one to determine whether a specific test is a variant of a specific procedure), broad (i.e., allow one to encompass many different variants of the procedure), and reveal how different procedures relate to each other and other procedures that are used in psychology.

Test

A test refers to the specific ways in which the measurement procedure has been arranged in order to produce an outcome. Whereas the measurement procedure is defined in terms of a core (and unchanging) property, tests represent the constellation of task parameters that are varied by the researcher in concrete studies. Amongst other things, IAT researchers can choose the number, order, and sequence of trials and blocks while AMP researchers can select the content and nature of the target and prime stimuli. Tests can be derived from the same or different measurement procedures with the intention of capturing the same, different, or related phenomena. For example, two tests can be constructed from the same measurement procedure (IAT) in order to measure two different domains (racial versus political attitudes). Yet it is also possible to construct a number of tests from the same procedure in order to examine similar domains (e.g., race IAT involving words vs. pictures). Although the latter tests involve the same procedure being used to achieve a similar goal, different task parameters were selected (and thus technically different tests were constructed). It is important to realize that each change – how trivial it might seem – results in the construction of an entirely new test, and like direct procedures such as questionnaires, whenever a new test is constructed its psychometric qualities need to be evaluated. This practice is widely accepted in psychometric research because even small changes can, in principle, have big effects on reliability and validity. We see no reason why this practice should not be upheld for tests derived from indirect measurement procedures as well. Just as it is inappropriate to equate the psychometric qualities of tests from one measurement procedure (AMP) with another (IAT), it is also inappropriate to equate the psychometric qualities of two tests from the same measurement procedure (e.g., two IATs). This is true regardless of whether the two tests are designed to assess the same phenomenon (e.g., race IAT using words versus pictures) or different phenomena (e.g., race versus self-esteem IAT).

Population and Sample

Besides identifying a measurement procedure and defining the parameters of one's test it is also necessary to define the population to which that test will be administered. In psychometrics, a sample is typically defined as that subset of the population that has been selected for testing. When it comes to research on implicit cognition, samples will typically be comprised of individuals, who are not static, unchangeable entities but rather organisms that constantly adapt and evolve. Individuals represent the constituent elements of a sample and can be matched in terms of, or vary along, single or multiple dimensions. In the previous section we argued that psychometric qualities are unique to a given test and that it is inappropriate to equate tests derived from the same or different measurement procedures. The exact same reasoning also holds for populations. Once a researcher has specified the population of interest and selected a representative sample from that population, she is then able to administer her test and infer the psychometric qualities of that test. These qualities only hold for the test given that population and cannot be generalized to other populations. Imagine, for instance, that an IAT designed to assess automatic attitudes to mathematics is administered to a group of students who study either psychology or mathematics. Whereas the psychology students will likely be *heterogeneous* in terms of their IAT scores, such that some individual's detest mathematics while others love it, their counterparts in the mathematics department may be a relatively more *homogenous* insofar as they consistently evaluate mathematics in a positive light. As soon as a test is administered to a sample of a previously untested population, its psychometric qualities need to be re-evaluated. Thus tests are constructed in order to examine a particular phenomenon for a given population. It is the *interaction* between a test and sample derived from that population (within a broader context) that determines the psychometric qualities of that test.

General Context

Even when researchers set out to construct a specific measurement environment and expose a specific sample of individuals to that environment, it is important to appreciate that this interaction will always take place within a wider context. For instance, race IATs may be taken during times of racial tension or peace, AMPs assessing religious attitudes may be administered inside a church or at a bar whereas an IAT investigating social conformity could be delivered by a person wearing informal clothing or a white lab coat. Therefore, whereas a measurement procedure is defined in terms of a core unchanging property of a task, and a test in terms of the constellation of task parameters that the researcher can vary, the general context refers to that constellation of contextual factors that are either present or manipulated during testing, but that do not comprise the test-specific context itself. In much the same way that changes in the parameters of a test or individuals that comprise a sample will influence the psychometric qualities of a test, so too will manipulations that alter the general context. Put another way, the test and population can remain the same and yet the psychometric qualities of that test can change due to variations in the wider context.

Measure

When a test is constructed based on a certain measurement procedure and administered to a specific sample within a broader context, some form of behavior will be observed and registered. In the context of research on implicit cognition this behavior will usually take certain forms, from the speed and accuracy of a response (Greenwald et al., 2009) to neural or physiological activity (Stanley, Phelps, & Banaji, 2008). Likewise, a wide variety of data-analytic techniques can be applied to these responses, from difference (D) score calculations to log and Gaussian transformations or even probabilistic indices (e.g., De Schryver, Hussey, De Neve, Cartwright, & Barnes-Holmes, 2018). These techniques

represent the set of consecutive steps that one has to follow in order to summarize the observed behavior. For instance, IAT scoring algorithms (e.g., Greenwald, Nosek, & Banaji, 2003) specify the trials to be excluded, how the tails of the distribution and error-trials should be treated, what kind of transformation should be performed, and what mathematical operations should be applied.

Critically, during the construction of a test the researcher will specify the particular responses that are registered and the algorithmic parameters that are used to derive a score on the basis of those responses. As noted by De Houwer, Teige-Mocigemba, Spruyt, and Moors (2009), the selection of the response and the algorithm is an integral part of the core procedure. In-line with De Houwer et al., we reserve the term “measure” for the outcome obtained whenever some data-analytic technique is applied to observed behavior. Depending on the technique involved the measure may reflect various aspects of the observed behavior (e.g. the maximum observed reaction time, the percentage of errors, the difference in median reaction time and so on). ³ Figure 1 visualizes the nature of implicit measures and provides an overview of the relationship between procedures, tests, samples, contexts, and measures.

³ It's worth noting that the application of a scoring algorithm to observed behavior does not imply or require that any (psychological) meaning be given to the measure. That is, a measure, defined as a summary of observed behavior, is merely the result of some set of analytic operations applied to behavior. Tests, however, are constructed with a specific theoretical or scientific goal in mind - namely - to infer knowledge about properties or attributes of individuals or groups of individuals. The relation between a to-be-measured attribute and the behavior observed on a test therefore requires some set of theoretical inferences that extend beyond the test and measure (De Houwer et al., 2009b).

Figure 1. A visual representation of a typical measurement environment in implicit cognition research. A test is derived from a certain indirect measurement procedure and administered to a specific sample within a general context and some form of behavior is observed and registered. Based on the data-analytic techniques employed, a measure is obtained.

Part II: Defining Reliability

What is Reliability?

Regardless of the procedure employed, or the specific test constructed, researchers are typically interested in the relationship between the observed score (measure) and the theoretical variable of interest. Unfortunately, any such relationship is immediately complicated by the fact that no test is capable of perfectly capturing the theoretical variable of interest and that *error* is always going to be an unavoidable companion to any observations that we make. This error typically comes in two forms: either random or systematic (see Meredith & Teresi, 2006). Random error is often viewed as trial-by-trial variability that results from unknown but non-systematic sources, such as inadvertently turning one's head away from the screen during an IAT trial or accidentally pressing a certain button during an

AMP trial. Systematic error reflects trial-by-trial variability that stems from systematic sources such as situational or individual differences (e.g., the failure to counterbalance participant gender during a priming study on gender or systematic differences in stimulus familiarity during an IAT). Researchers are concerned with these different sources of error given that error is present in many, if not most situations, and will potentially undermine the relationship between an observed outcome and a theoretical variable. Therefore the degree of *measurement precision* that can be achieved when a test is administered to a given population will be determined by the extent to which error contributes to the scores obtained from that test. It is this measurement precision that researchers are referring to when they talk about the *reliability* of a test (see Borsboom, 2005; Mellenbergh, 1996).

Classical Test Theory and Reliability

Historically, the concept of reliability has nearly always been interpreted through the lens of Classic Test Theory (CTT) (Lord & Novick, 1968). At its core, this account makes several assumptions. The first is that each of the observed scores (Y) obtained from a test for a given population (P) can be decomposed into two independent elements: the underlying true score (T) and an error score (E) (i.e., $Y = T + E$ for P). Based on the assumptions that T and E are independent, it can be shown that observed score variance (σ_Y^2) in the population is equal to the sum of true-score variance (σ_T^2) and error score variance (σ_E^2) (i.e., that $\sigma_Y^2 = \sigma_T^2 + \sigma_E^2$). This decompositional perspective of observed, true, and error scores reveals a key point: as the contribution of error-score to observed score variance increases the relative contribution of true-score variance to observed score variance will (by definition) decrease.

Put simply, reductions in error-score variance will lead to increases in measurement precision (for a fixed amount of true-score variance) whereas the opposite is true whenever the contribution of error-score variance grows (i.e., our measurement precision will decrease

for a fixed amount of true-score variance). Therefore, from a Classical Test Theory perspective, when we speak of reliability we are referring to the ratio of true-score variance (σ_T^2) to observed score variance (σ_Y^2) and this ratio is usually written in the following way: $\rho_{YY'} = \sigma_T^2 / \sigma_Y^2$, with Y' referring to observed scores obtained from a parallel test (*see below*).

Theoretical implications versus practical applications. Although CTT has proven highly influential, its limitations are widely known. These stem from the fact that the theory does not make any assumptions about an underlying theoretical construct that is responsible for the behavior obtained from a test (Borsboom & Mellenbergh, 2002; Sijtsma, 2009). As we just mentioned, CTT defines reliability as the ratio of true-score to observed score variance. Calculating this ratio is complicated by the fact that the true score is an unobservable quantity, as is the notion of true-score variance. We therefore need a way to make the unobservable observable. Lord and Novick (1968) demonstrated that by first constructing two parallel tests (obtaining Y and Y'), and then examining the covariance between the measures derived from those tests, true-score variance could be obtained if one assumes that participants are (a) identical with respect to each of the test administrations and that (b) the two tests are administered under the same contextual or environmental conditions. Yet, by definition, even minor modifications to the test, general context, and/or sample will yield new (true) test scores. Therefore the only parallel test that can be constructed from this perspective is the test itself, administered under identical environmental conditions, which seems only possible when the participant is somehow “brain-washed” and exposed to the test once again (see Borsboom, 2005, for an in-depth discussion).

Lord and Novick’s solution may be theoretically possible, but in reality, it is practically untenable. Reliability will always remain a theoretical construct and can never be observed. This is because a precise estimation of reliability requires the general context, test-specific context, and sample to somehow remain constant across repeated test administrations.

Yet as we have seen in Part I there are many different ways in which the context and sample vary from one moment to the next and the probability that all these parameters will remain fixed across time and space is practically zero. Therefore, in reality, researchers can only ever *approximate* the reliability of a test and tend to do so using one of three “proxies” known as internal consistency reliability (consistency), parallel test reliability (equivalence), and test-retest reliability (stability).

When we combine the arguments from Parts I and II an important point emerges: whenever researchers set out to approximate reliability they are (implicitly) deciding which set of (test, context, and sample) parameters they will manipulate in order to examine measurement precision. In Table 1 we provide a taxonomy that highlights which aspects of the test, sample, and context will likely vary or remain constant when these three proxies (consistency, equivalence, and stability) are approximated.

Table 1. An overview of different approximations of reliability (consistency, equivalence and stability) and the likely elements of the general context, test-specific context and sample that will vary (v) or remain constant (c).

Reliability				
	Theoretical	Approximations		
Parameter		Internal Consistency (<i>Consistency</i>)	Parallel Tests (<i>Equivalence</i>)	Test-Retest (<i>Stability</i>)
General Context	C	C	C/V	C/V
Test-Specific Context	C	V	V	C
<i>Stimuli</i>	C	V	C/V	C
<i>Responses</i>	C	C	C/V	C
<i>Algorithm</i>	C	C	C/V	C
Sample	C	C	C/V	C/V
<i>Individual</i>	C	V	V	V

Table 1 clarifies that the three main methods of approximating reliability (consistency, stability, and equivalence) each differ in terms of the environmental parameters that vary or remain constant. Take consistency: researchers interested in implicit cognition typically rely on consistency as an approximation of reliability whenever they assess data from a single test administration. Given that most indirect procedures are completed in a relatively short period,

the general context will (likely) remain constant from one trial to another. In contrast, certain elements of the test-specific context will vary (e.g., one trial on a race IAT might present an image of a black individual while another might present a White individual) while others will remain fixed (e.g., the same class of responses are emitted during the AMP whereas the same algorithms are applied to the data from different AMP trials). Although the sample as a whole will tend to remain constant (providing attrition is low) the individuals that comprise the sample may vary in how they respond to the elements that comprise the test.

Now compare this to stability and equivalence. Implicit cognition researchers typically rely on stability whenever they want to assess the consistency of their data from one test administration to the next. They rely on equivalence whenever they want to assess the consistency of their data across tests. Although the general context may be similar across repeated test administrations the probability that it can vary is higher than when consistency is estimated within a single test administration. The test specific context will vary in a similar way to that seen when internal consistency is assessed while the individuals that comprise the sample will likely change across repeated test administrations (e.g., people may alter how they behave at Time 2 based on their experiences at Time 1).

Part III: A General Framework for Approximating Reliability

Latent Variable Models

Unsurprisingly, the aforementioned conceptual problems of CTT have led to alternative psychometric perspectives on reliability. One of these perspectives is based on a class of Latent Variable Models (LVM) that seek to explain different patterns of behavior in terms of some underlying theoretical construct. Before we showcase how these models address the issues of reliability we will first provide a brief introduction to LVM and related concepts.

The main aim of LVM is to explore the relationship between some unobservable or latent construct (η) and one or more observable indicators ($Y_1, Y_2, \dots, Y_k$). Confirmatory Factor Analysis (CFA) is a well-known example of LVM and is often used to explore how the latent construct is linearly related to its various indicators. In CFA the relationship between a latent construct and its indicators is often modeled using a regression equation and is defined by three (to-be-estimated) sets of parameters: the intercept (a_i), slope or factor 'loading' (b_i) and error scores (e_i) (see Figure 2). Thus CFA can be mathematically defined as $Y_i = a_i + b_i\eta + e_i$, for indicator (Y_i) (Jöreskog, 1971; Raykov, 2004; Brown, 2006; Graham, 2006).

Figure 2. A visual illustration of a CFA model where the relationship between an unobservable or latent variable (η) and its observable indicators (Y_i) is determined by three to-be-estimated parameters: the intercept (a_i), slope (b_i), and error term (e_i).

Three different types of CFA models are typically constructed and each differs in the specific assumptions it makes about the above parameters (a_i , b_i and e_i). First, *congeneric* models allow all of the parameters to vary under the assumption that the indicators measure the same latent construct (Brown, 2006). Second, (*essentially*) *tau-equivalent* models take a different approach, constraining the factor loadings (b_i) to be equal, while allowing the other

model parameters to vary.⁴ Third, *parallel* models take the most restrictive route by forcing the factor loadings and the error variances to be equal. In this way a parallel model is nested in a tau-equivalent model and a tau-equivalent model is nested in a congeneric model (i.e., all three models have the same structure but differ in the parameters that they fix or let vary). To test if the model actually meets its criteria (i.e., that certain parameters are equal to one another) model comparison tests are conducted (e.g. a χ^2 - difference test; see Brown, 2006).

Indicators in Implicit Cognition Research. The relationship between the latent construct and its indicators is pivotal for CFA. Hence, it is important to clearly specify what an indicator is. By definition, an *indicator* has to fulfill two basic conditions: it must vary systematically with changes in the latent construct and increase or decrease monotonically with the latent construct (Lord & Novick, 1968). In other words, when higher (or lower) scores are observed on the indicator this must be caused by increases (or decreases) in the latent construct's values. This relatively simple story is immediately complicated when it comes to indirect procedures like the IAT and certain types of priming. Due to the ways in which these tasks are constructed, the responses observed in one condition (e.g. those obtained from congruent or incongruent blocks of trials) cannot be considered as indicators of the construct of interest. Although these responses are certainly related to the latent construct they often violate the aforementioned conditions (e.g., they do not always vary monotonically with the construct of interest). Several authors have acknowledged this finding in the literature. For instance, Nosek and Sriram (2007) argued that the responses obtained from congruent or incongruent blocks on an IAT cannot be treated as direct indicators of the underlying construct and that, just like the Stroop task, it is only by contrasting these two

⁴ An essentially tau equivalent model differs from a tau equivalent model by allowing the intercept a_i to vary across components. While both models assume that the component measures relate in a similar way to the latent construct η , essentially tau equivalent models allow for different means across the measures. The component measure and the (essentially) tau-equivalent model make no restriction with respect to the error variances. In both models, different measurement precision can be obtained for the different components.

conditions that one can obtain an indicator in indirect procedures such as the IAT. Therefore when we speak of indicators we are referring to the application of some mathematical operation to the behavior obtained from a test. This is precisely the same way in which we defined a measure in Part I.

The process of selecting an indicator, that is (certain) elements of a test to which mathematical operations will be applied in order to generate a measure, can be defined as *parceling*. It is through this process of parceling that one obtains parcels. We define a parcel as *an indicator for which the same mathematical operations that would be applied to the entire test have been applied to a subset of elements of the test.*⁵ There appear to be two main ways of parceling one's data. On the one hand, a distributed parceling strategy can be used wherein items are either randomly or sequentially assigned to a parcel. The simplest version of this involves dividing one's data into two halves based on an odd/even split (LeBel & Paunonen, 2011), although researchers have also carved their data into three (Cunningham et al., 2001), four (Mierke, & Klauer, 2003) or even more parcels using a similar strategy (Schmitz, Teige-Mocigemba, Voss, & Klauer, 2013). On the other hand, a homogenous parceling strategy can be applied wherein the items that comprise the parcel are in some way related to one another, and thus form a homogenous group (Marsh, Lüdtke, Nagengast, Morin, & Von Davier, 2013).

Two points are worth noting here. On the one hand, there are many researcher degrees of freedom when it comes to parceling: one not only has to decide how many parcels should be created but also what parceling strategy should be used (homogenous vs. distributed). Specifying *a priori* how and why one will parcel their data places constraints on this process

⁵ Note that in the LVM literature parcels are typically defined as “indicators comprised of the sum (or average) of two or more items, responses, or behaviors” (Little, Cunningham, Shahar, & Wideman, 2002, p.152). Yet many mathematical operations can be applied to elements of a test (above and beyond simple summation or averaging) when constructing parcels.

and thus limits those degrees of freedom. On the other hand, certain parceling strategies have the potential to camouflage unwanted sources of variance. A popular idea in the implicit cognition literature is that the trials that make up a certain phase of a test (e.g., those in the congruent block) are all equivalent and that the responses to any one trial within this phase will be similar to the next (i.e., there is a high exchangeability of trials within a block). Therefore the variation in responding to these trials (at least within participants) is thought to represent random error variance (Blanton, Jaccard, & Burrows, 2015). If this idea is correct, then parceling one's data in different ways (e.g., either randomly, sequentially, or even homogeneously) should yield similar outcomes.

Imagine, for instance, that a researcher decides to assign trials from a race IAT to different parcels based on their content. She assigns Black exemplars to a 'Black' parcel, White exemplars to a 'White parcel', positive items to a 'positive parcel' and negative items to a 'negative' parcel. If the above exchangeability assumption holds then these content-based parcels should all be equivalent - or at least tau equivalent. Yet there is no guarantee that this assumption will hold. For instance, there could be some systematic confound that applies to certain test-elements but not others. For instance, participants may be more familiar (thus able to respond quicker) to in-group exemplars and be less familiar with (and thus slower to respond to) out-group exemplars. They may also be better able to process certain types of stimuli (e.g., images of faces) compared to others (e.g., names). The test-specific context might also influence responding in other ways (e.g., images presented on black vs. white backgrounds might facilitate or undermine how readily people identify Black and White faces). Now when a distributed parceling strategy is applied the above systematic variance is potentially 'hidden' or 'spread' across the various parcels - and as such - the latent variable might capture that systematic variance and thus overestimate consistency (see Little, Rhemtulla, Gibson, & Schoemann, 2013; Maul, 2013). Yet a content-based strategy avoids

this problem by forcing the systematic variance to remain unique to each parcel. We therefore recommend that researchers stay vigilant to such unwanted sources of variance that may contaminate their consistency estimates. One way to do this is to use a homogenous (content-based) parceling strategy that allow for these systematic sources of variance to be identified.

Latent Variable Models and Reliability

So, taking a step back, what does all of this have to do with reliability? As we outlined above, CTT makes certain theoretical assumptions about reliability that seem rather difficult (if impossible) to instantiate in reality. Faced with the need to estimate the precision of their measurements, researchers have often ignored these theoretical assumptions and instead tried to approximate the degree to which their data was internally consistent, equivalent, or stable from one measurement occasion to the next. We also argued that each of these approximations is basically interested in answering a different question and will therefore proceed from the position that certain aspects of the general context, test, and sample should be fixed or allowed to vary. Therefore we need a way of conceptualizing these different approximations and of putting their respective assumptions to the test.

Latent Variable Models can do just this. They not only address many of the theoretical problems outlined above but, perhaps more importantly, they provide an overarching method for conceptualizing the different approximations and for testing their respective assumptions. Whereas CTT (strictly speaking) considers different parts of a test, or different tests to measure different psychological concepts, LVM allows us to assume that performance on different implicit measures can actually be underpinned by the same latent construct. In this way LVM go “beyond classical test theory in that they attempt to construct a hypothesis about the data-generating mechanism in which the attribute is explicitly represented as a latent variable” (Borsboom, 2005, p.49). LVM allows us to explore how parcels are related to a latent variable by comparing different models (e.g., does a parallel or tau-equivalent model fit

the data adequately?) One of the main advantages of thinking in this way is that it allows researchers to examine the assumptions that underpin the various approximations of reliability. For example, by using CFA the specific assumptions that different approximations of reliability make, can be tested by examining the covariation between the different parts of a test (Yang & Green, 2011). In what follows we will unpack these points in greater detail. Specifically, we will describe the three reliability approximations (consistency, stability, and equivalence), indicate what their underlying assumptions are, as well as the consequences that follow when their assumptions are violated. We will then show how CFA can be used to explore these assumptions and to approximate reliability.

Consistency as an Approximation of Reliability

We define *consistency* as the interrelatedness of *a priori* constructed parcels of a test.⁶ In the next section we will show how the parameter estimates obtained from a CFA model can be used to generate a consistency estimate known as coefficient ‘omega’ (ω , see McDonald, 1978, 1999; Raykov, 2001, 2004; Bentler, 2009). We then argue that two of the most common ways of approximating consistency in the implicit cognition literature (*split-half reliability* and *Cronbach’s alpha*) are just two instances of this coefficient omega (Raykov, 2004). The value of the (omega) approach lies in the fact that it cannot only generate split-half and Cronbach’s alpha but also yields a reliability estimate in situations where the assumptions underpinning the above do not hold (i.e., for congeneric models).

Coefficient Omega. Although little used in the implicit cognition literature, estimating coefficient omega via CFA is a “*widely applicable latent variable modeling approach to point and interval estimation*” of consistency (Raykov, 2004, p.324) that can be mathematically

⁶ Note that the concept of internal consistency can also be defined in other ways. For instance, one can view it as a measure of *homogeneity* (i.e. the degree of unidimensionality of the items; for a discussion see Sijtsma, 2006; Clark & Watson, 2009; Yang & Green, 2011).

defined as $\rho_{YY'} = \frac{(\sum_{i=1}^k b_i)^2}{(\sum_{i=1}^k b_i)^2 + \sum_{i=1}^k \theta_{ii}}$, where k is the number of parcels, b_i is the factor loading,

and $\theta_{ii} = \text{Var}(e_i)$ is the error variance of the i^{th} parcel.⁷ The advantage of this approach is that one can easily estimate consistency in many different ways using a single formula. For instance, imagine that you have two parcels ($k = 2$), and that your factor loadings ($b_1 = b_2$) and error terms ($e_1 = e_2$) are equal to one another. The solution yielded by the above formula will be the same as what one obtains from the split-half reliability formula. Now imagine that you have k parcels and that the factors loadings of those parcels ($b_1 = b_2 = \dots b_k$) are equal. This same formula will yield an outcome similar to that obtained from the Cronbach's alpha formula given the same number of parcels. Interestingly, in situations where factor loadings (or factor loadings and error terms) are not equal, the omega formula can still be used to generate an unbiased consistency estimate given a specific set of indicators. Although split-half reliability and Cronbach's alpha represent the most popular means of determining the consistency of one's data, the assumptions that underlie these estimates are rarely tested (or at least rarely reported). For instance, when one calculates split-half reliability using the Spearman-Brown correction (which can be mathematically defined as $\hat{r}_{sb} = \frac{2\hat{r}}{1+\hat{r}}$) they are assuming that the two parcels they have created are parallel in nature (Eisinga, Grotenhuis, & Pelzer, 2013). Recall that from a CFA perspective this means that the measures obtained from those parcels are thought to reflect the same underlying construct, that the factor loadings are identical, that the error variances are equal, and that the errors do not correlate. Likewise, when we use Cronbach's alpha, we assume that the parcels are (essentially) tau-equivalent. That is, we make a similar set of assumptions as above except that the condition of equal error variances is relaxed (Novick & Lewis, 1967; Lord & Novick, 1968; Yang & Green, 2011; Eisinga et al., 2013). Whenever these assumptions are violated we end up with problematic

⁷ Note that this version of the formula only holds for uncorrelated errors. For correlated errors we refer the reader to Raykov (2004, p.304).

consistency estimates that may be over- or underestimated. Put another way, even in situations where the underlying measurement model is not parallel (which is required for split-half reliability)⁸ or tau-equivalent (which is required for Cronbach's alpha) but simply congeneric, researchers can still use the above formula to estimate consistency.

The identification of the measurement model is thus an important first step before approximating reliability using the CFA methodology. Yang and Green (2011) argue that a congeneric model should first be fit to the data so that each parameter can be freely estimated. A tau-equivalent model should then be explored, and if it fits the data equally well, then it should be preferred over the congeneric model. If it does not then the congeneric model should be accepted. If the tau-equivalent model fits the data equally well then it should then be compared with a parallel model (and the later only accepted when it also fits the data).

Equivalence as an Approximation of Reliability

We define *equivalence* as the interrelatedness of measures obtained from two or more tests. When reliability is approximated in this way the researcher assumes that the scores obtained from two independent administrations of the same test reflects the same construct. If this is true then correlations between test scores can be considered as a direct estimate of reliability. This presents an immediate (practical) problem of constructing tests that are actually parallel to one another: in a strict sense, two tests can be considered *parallel*, if and only if, they are *exactly equivalent*. Both tests should not only measure the same (underlying) construct but also do so with the same degree of precision (i.e., observed score means, variances, and co-variances of different parts of the tests should be equal to one another).

As we mentioned earlier on, the conditions of parallelism can easily be met in the

⁸ Some authors have avoided choosing a specific parceling strategy by adopting a bootstrap procedure. This involves first splitting the data into two equal halves, calculating a correlation coefficient, and then repeating this process many times (e.g., Ravenzwaaij, van der Maas, & Wagenmakers, 2011). The mean correlation coefficient (Spearman-Brown corrected) they obtain is then used as a consistency estimate

context of a thought experiment. Yet those conditions are difficult, if not impossible, to realize in the real world. Consider studies that have set out to measure the same construct using “parallel tests”. Dasgupta, McGhee, Greenwald, and Banaji (2000) assessed implicit racial preferences using IATs with pictorial or name stimuli. Likewise, Asendorpf, Banse, & Mücke, (2002) examined implicit shyness while Greenwald et al. (1998) examined racial evaluations using two IATs with different target stimuli. According to our framework, the authors of these studies manipulated aspects of the test-specific context. For instance, they varied the nature of the stimuli while fixing the type of responses registered and algorithms applied to the data. Although their samples always remained constant across testing the individuals that comprise those samples likely differed from one test administration to the next. Whenever researchers manipulate aspects of the test-specific context, general-context, and/or the individual they produce tests that differ from one another in some way, and therefore, they need to determine if those tests are parallel as initially assumed.

Yet as is the case with consistency, the assumptions underpinning equivalence estimates are rarely - if ever - tested or reported in the literature (i.e., same latent construct and same degree of measurement precision). Again these assumptions can easily be assessed via CFA by identifying the underlying measurement model. For instance, one could construct four parcels (e.g., based on content) for one test and then do so again for the second test. If the conditions of parallelism hold then a model where both factor loadings and error variances are equal among related content parcels , should fit the data. In conditions where the model does not fit the data researchers cannot claim that the two tests are parallel - and by definition - the observed correlation coefficient between test scores cannot be considered as a good approximation of reliability.

Stability as an Approximation of Reliability

We define *stability* as the interrelatedness of measures obtained from the same test that has been administered two or more times. Stability and equivalence bear many similarities to one another: in each case a test is administered at least two times and the correlation between test scores are assumed to reflect the same latent construct. From our perspective, the key difference between equivalence and stability is the extent to which they involve manipulating aspects of the test-specific context. Whereas equivalence does involve manipulating the test-specific context (because different tests are used) stability does not (because the same test is used). Yet in both cases the general-context and the individual may vary given that both approximations involve repeatedly administering tests across time (see Table 1). The extent of this variation may depend on the time lag between test administrations.

The assumptions underpinning stability are the same also those underpinning equivalence and can therefore be tested in the same way. For instance, one could construct two parcels for one test and another two parcels for a second test. If the conditions of parallelism hold then a model where the factor loadings are equal, and error variances are equal, should fit the data. In conditions where the model does not fit the data researchers cannot claim that test scores are stable - and by definition - the observed correlation coefficient between test scores cannot be considered a good approximation of reliability.

Part IV: An Empirical Example

We have just introduced a LVM approach to the approximation of reliability. This approach may be new for many researchers interested in the study of implicit cognition. Even those who are acquainted with a LVM approach they might never have applied it to the topic of reliability. Others may have been aware of both LVM and its links to reliability but never approximated the latter using the former because access to relevant software was either unavailable or too costly. A short example of this approach in action will serve as a tutorial on the basic ideas outlined here. It will also show how a CFA model can be fit to data in order to

generate consistency estimates using the free and open-source software known as R (R Development Core Team, 2016) and the add-on package known as lavaan (Rosseel, 2012).⁹

This section will also showcase how the choices one makes when approximating consistency have real and tangible consequences for the types of estimates one obtains. This will be achieved by drawing on a data set that was originally reported by Bar-Anan and Nosek (2014). Although these authors administered several indirect procedures in different domains (i.e., the so-called “Attitudes 3.0” data set), we will only focus on the data of those who completed a race IAT. In this IAT the target labels were *White People* and *Black People* while the target stimuli were six pictures of white people and six pictures of black people. The attribute labels were *Good* and *Bad* while the attribute stimuli were six positive and six negative words. The authors calculated IAT scores using the D2 scoring algorithm. They then estimated the internal consistency of their data (using the Cronbach’s alpha formula) by constructing three parcels based on a distributed parceling strategy wherein items were sequentially assigned to three different subsets. In this way they obtained an internal consistency estimate of .86.

In what follows, we first approximate reliability via the “traditional” approaches used in implicit cognition research (i.e., split-half reliability and Cronbach’s alpha). We then apply the CFA-methodology using the two-step method recommend by Yang and Green (2011) (i.e., where the CFA model is first assessed and then coefficient omega estimated). Although Bar-Anan and Nosek’s sample comprised White and Black participants, we apply the CFA approach in three ways: to a sub-sample of White participants ($N = 161$), a sub-sample of Black participants ($N = 161$), and a combination of the two ($N = 322$). Doing so will show

⁹ For readers interested in learning more about CFA we recommend Brown (2006) whereas those looking to learn more about fitting CFA models using R can read Finch and French (2015). Although our example focuses on (internal) consistency it could easily be adapted for stability and equivalence as well.

how consistency estimates can fluctuate depending on the parceling strategies (and samples) involved, and thus why researchers need to be aware of these issues (see Appendix for R code that will allow the reader to reproduce the results of the combined sample).

Split-half reliability. When researchers have the aim to determine the consistency of an IAT based on a certain IAT data set, they typically rely on either split-half or Cronbach alpha. Take split-half reliability: by splitting the test into two parcels of equal length we obtain two measures. Given that the number of items within each parcel is half that of the original test the correlation coefficient will tend to underestimate reliability. In order to correct for this the coefficient is typically adjusted using the Spearman-Brown prophecy formula, leading to what is known as stepped-up reliability (Lord & Novick, 1968). In the implicit cognition literature split-half reliability is usually calculated by creating two different parcels based on the index number of the trial (i.e., whether it was odd or even numbered; see Glashouwer, Smulders, de Jong, Roefs, & Wiers, 2013; Krause, Back, Egloff, & Schmukle, 2011; Schmukle, & Egloff, 2006). The popularity of this (distributed) parceling strategy may stem from the idea that it leads to two different tests that are assumed to be parallel in nature. Applying this approach to Bar-Anan and Nosek's race IAT yields a reliability coefficient that is higher for the White participants (0.79) and combined sample (0.80) than the Black participants (0.68).

Cronbach's alpha. One could also estimate consistency by creating three or more parcels (using either a distributed or homogenous parceling strategy) and then calculating Cronbach's alpha. In practice there are often few limitations placed on researchers at this analytic stage (i.e., there are many researcher degrees of freedom). For instance, Bar-Anan and Nosek (2014) constructed three parcels based on a distributed parceling strategy (without any clear motivation for why they did so). But they could have equally constructed four, five...or any number of parcels and used either of the above strategies. That said, there may be logical reasons for choosing one strategy over another.

It is worth noting here that tests are comprised of different elements could share some systematic variance that is ‘reliable’ but unrelated to the latent construct of interest. This type of variance might lead to an overestimation of consistency especially when a distributed parceling strategy is employed. To illustrate, imagine that we adopt two strategies. The first is a homogenous parceling strategy where trials were assigned to four different parcels based on their respective content: with Black stimuli assigned to a Black parcel, white stimuli to a White parcel, positive adjectives to a Positive parcel, and negative adjectives to a Negative parcel. In the second case we constructed four parcels via a distributed parceling strategy, with trials assigned sequentially based on their index number. Calculating Cronbach’s alpha for the homogeneous parcels revealed that data from White participants (0.78) and the combined sample (0.80) was once again more internally consistent than from the Black participants (0.69). As expected, the consistency estimates for the distributed parcel data were generally higher, and the same pattern was observed (White participants (0.85); combined sample (0.86); Black participants (0.77)). However, there are no guarantees that the homogenous parcels will meet the assumptions underlying Cronbach’s alpha and might produce biased consistency estimates. We therefore need a way to test these assumptions, and as discussed above, this can be achieved using CFA.

Coefficient omega. Applying this logic to the homogenous data from above means that a congeneric model (one for each of our three samples) should be constructed with the parcels as indicators and one latent variable. To identify the scale of the latent variable its variance was set to 1.00. The congeneric model seemed to fit the data well in all three cases.¹⁰ When these congeneric were then compared with tau-equivalent models, the former were found to

¹⁰ As rules of thumb we use the following criteria for good model fits: p-value of χ^2 -test > .05; RMSEA < .06; CFI > .95 (for more see Brown, 2006). For the combined sample model fits were $\chi^2(2) = 0.15$, $p = .93$, CFI=1.00, RMSEA = 0.00. For the White participants model fits were $\chi^2(2) = 2.75$, $p = .25$, CFI=1.00, RMSEA = 0.05. For the Black participants model fits were $\chi^2(2) = 3.41$, $p = .18$, CFI=0.99, RMSEA = 0.07.

fit the data better than the latter, which is indicated by highly significant p -values. Moreover, bad fit indices were obtained for the tau-equivalent models (see Table Y for factor loadings and error variances of the congeneric models). The parameter estimates of the congeneric models were then used to calculate coefficient omega – one for each of our samples. This can be achieved by using the lavaan syntax to define the coefficient. One advantage of the lavaan syntax is that it also provides standard errors which can be used to calculate a 95% confidence interval ($\hat{\tau} \pm 1.96 \text{ SE}$). It is also possible to use a bootstrap-based confidence interval with the lavaan method. Calculating omega for the homogeneous parcels revealed that data from White participants (0.80) and the combined sample (0.82) was once again more internally consistent than from the Black participants (0.72) (see Table 2).

Table 2. Consistency estimates for the White participants, Black participants, and combined sample from Bar-Anan and Nosek (2014). Four estimates are provided for each sample: Split-half reliability (with Spearman-Brown corrections), Cronbach's alpha (distributed parceling strategy; Distr), Cronbach's alpha (homogenous parceling strategy; Hom), and coefficient omega.

	Combined Sample (N = 322)	White Sample (N = 161)	Black Sample (N = 161)
Split-half	.80	.79	.68
alpha (Distr)	.86	.85	.77
alpha (Hom)	.80	.78	.69
Omega	.82 (SE = 0.017) 95 CI={.79-.85}	.80 (SE = 0.031) 95 CI={.74-.86}	.72 (SE = 0.038) 95 CI={.65-.79}

Table 3. Factor loadings and error terms obtained from the congeneric model with four homogenous parcels (White parcel (WP); Black parcel (BP); Negative parcel (NP) and Positive parcel (PP) for the White participants, Black participants, and combined sample from Bar-Anan and Nosek (2014).

	WP	BP	NP	PP
Total sample (N=322)	0.60 (0.64)	0.50 (0.75)	0.88 (0.22)	0.84 (0.29)
White sample (N = 161)	0.55 (0.70)	0.49 (0.76)	0.89 (0.21)	0.85 (0.28)
Black sample (N = 161)	0.47 (0.78)	0.36 (0.87)	0.83 (0.31)	0.73 (0.46)

Given that only a congeneric model fits the homogenous data the necessary condition of tau-equivalence does not hold - a condition upon which the use of Cronbach's alpha rests -

and thus this consistency estimate will likely be biased. Indeed, as Table 2 shows, the use of Cronbach's alpha (for the homogeneous data) leads to an underestimate of internal consistency. However, consistency estimates are still lower compared to the estimates based on the distributed parcels. Table 3 also shows that there is a large difference in the factor loadings for picture (Black and White) compared to the word items (Positive and Negative) as well as for items from the Black parcel compared to Positive parcel (both for White participants and Black participants), questioning the assumption of item exchangeability.

The above demonstrates how consistency estimates can vary as a function of the sample and parceling strategies (distributed vs. homogenous) involved (see Table 2). Consistency estimates were higher for White participants and the combined sample relative to Black participants, and (as expected) a distributed parceling strategy produced larger estimates than a homogenous one. In this case, some systematic (error) variance – and thus variance unrelated to the construct of interest - seems to be captured by the latent variable. The results also demonstrate the importance of identifying the measurement model: by applying the Cronbach alpha formula (without testing the condition of tau-equivalence) researchers may obtain biased consistency estimates.

Part V: Conclusion

In the course of designing experiments as well as collecting and analyzing data, researchers face many decisions: they need to figure out how they should structure the test and general-context, what sample they should select for testing, what data they should collect and how they should analyze that data. They are afforded many degrees of freedom during this decision making process and even more whenever it is hidden from view or they fail to specify it on an *a priori* basis. This can lead to problematic situations in which “researchers explore various analytic alternatives to search for a combination that yields ‘statistical significance’ and to then report only what ‘worked’” (Simmons et al., 2011, p.1359).

This is also true for reliability. Every researcher is faced with multiple decisions when attempting to determine the precision of their measurement outcomes; from the reliability approximations they select (consistency, stability, equivalence), to how the data should be parceled, how many parcels they should create, and what should be assigned to those parcels (homogenous vs. heterogeneous). In the absence of clear guidelines and standards for approximating reliability the degrees of freedom afforded to researchers are high and their reliability-related decisions they make can be hidden or simply ignored. We argue that this state of affairs reflects current practice within the implicit cognition literature where the motivation for, and assumptions underpinning reliability, are rarely if ever explicated or tested. This leads to a situation where there are little, if any, constraints that prevent researchers from ‘cherry-picking’ those reliability estimates that presents the most favorable impression of their implicit measures.

In the current paper we attempted to address the above problems by providing a framework for approximating reliability. We first argued that any discussion about reliability has to begin from a position of conceptual clarity. We highlighted the difference between a test and procedure and showed how implicit measures represent *outcomes* that are obtained when a constellation of test-specific elements are administered to a specific sample of individuals who complete that test within a general context. Thus when we speak about the reliability of a test we always have to consider the interaction between these three elements (sample, test, context). Second, we offered a framework that highlights the likely elements of the general context, test-specific context, and sample that can remain constant or vary. Third, we provided a LVM perspective on reliability providing a clear way of approximating reliability (consistency, equivalence, stability) and testing the assumptions that underpin those approximations. By comparing conventional (split-half and Cronbach’s alpha) and more recent ways of approximating reliability (omega using the CFA approach), by highlighting the

concept of parceling, and by applying the above to real-world data, we then showed that the way in which one manipulate these elements can have an impact on the consistency, stability, or equivalence estimates obtained, demonstrating the need for a clear set of standards.

To conclude, when it comes to reliability we recommend that the following practices be adopted: authors should (a) motivate why they decided to approximate reliability in a given way before data collection and report this rule in their article, (b) test the assumptions that underpin the reliability approximation they selected, and if these assumptions fail to hold, then (c) apply an appropriate correction or alternative approximation. We also encourage reviewers and editors to ensure that authors follow these recommendations. Doing so requires relatively little from all parties but will ensure the further improvement of transparent and good scientific practices (and thus cumulative scientific progress). It seems likely that high researcher degrees of freedom are partially responsible for the so-called replication crisis in psychological science as well as the historical issues surrounding the reliability of implicit measures. We hope our framework and suggestions help address both.

References

- Aiken, L. S., West, S. G., Sechrest, L., Reno, R. R., Roediger III., H. L., Scarr, S., ... & Sherman, S. J. (1990). Graduate training in statistics, methodology, and measurement in psychology: A survey of PhD programs in North America. *American Psychologist*, 45(6), 721.
- Asendorpf, J. B., Banse, R., & Mücke, D. (2002). Double dissociation between implicit and explicit personality self-concept: the case of shy behavior. *Journal of personality and social psychology*, 83(2), 380.
- Bar-Anan, Y., & Nosek, B. A. (2014). A comparative investigation of seven indirect attitude measures. *Behavior research methods*, 46(3), 668-688.
- Barnes-Holmes, D., Barnes-Holmes, Y., Power, P., Hayden, E., Milne, R., & Stewart, I. (2006). Do you really know what you believe? Developing the Implicit Relational Assessment Procedure (IRAP) as a direct measure of implicit beliefs. *The Irish Psychologist*, 32(7), 169-177.
- Bentler, P. A. (2009). Alpha, dimension-free, and model-based internal consistency reliability. *Psychometrika*. doi:10.1007/s11336-008-9100-1
- Blanton, H., Jaccard, J., & Burrows, C. N. (2015). Implications of the implicit association test D-transformation for psychological assessment. *Assessment*, 22(4), 429-440.
- Borsboom, D. (2005). *Measuring the mind: Conceptual issues in contemporary psychometrics*. Cambridge University Press.
- Borsboom, D., & Mellenbergh, G. J. (2002). True scores, latent variables, and constructs: A comment on Schmidt and Hunter. *Intelligence*, 30(6), 505-514.

- Bosson, J. K., Swann Jr., W. B., & 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.
- Brown, T. A. (2006). *Confirmatory factor analysis for applied research*. Guilford Publications.
- Cunningham, W. A., Preacher, K. J., & Banaji, M. R. (2001). Implicit attitude measures: Consistency, stability, and convergent validity. *Psychological science*, 12(2), 163-170.
- Dasgupta, N., McGhee, D. E., Greenwald, A. G., & Banaji, M. R. (2000). Automatic preference for White Americans: Eliminating the familiarity explanation. *Journal of Experimental Social Psychology*, 36(3), 316-328.
- De Houwer, J. (2005). What are implicit measures and indirect measures of attitude. *Social Psychology Review*, 7(1), 18-20.
- De Houwer, J., & Moors, A. (2010). Implicit measures: Similarities and differences. In B. Gawronski & K. B. Payne (Eds.), *Handbook of implicit social cognition: Measurement, theory, and applications* (pp. 176-193). New York: Guilford Press.
- De Houwer, J., Teige-Mocigemba, S., Spruyt, A., & Moors, A. (2009). Implicit measures: A normative analysis and review. *Psychological bulletin*, 135(3), 347.
- De Schryver, M., Hussey, I., De Neve, J., Cartwright, A., & Barnes-Holmes, D. (2018). The PI-IRAP: An alternative scoring algorithm for the IRAP using a probabilistic semiparametric effect size measure. *Journal of Contextual Behavioral Science*, 7, 97-103.

- De Schryver, M., Hughes, S., Rosseel, Y., & De Houwer, J. (2015). Unreliable Yet Still Replicable: A Comment on LeBel and Paunonen (2011). *Frontiers in psychology*, 6.
- Eisinga, R., Grotenhuis, M. T., & Pelzer, B. (2013). The reliability of a two-item scale: Pearson, Cronbach, or Spearman-Brown?. *International journal of public health*, 1-6.
- Fazio, R. H. (2001). On the automatic activation of associated evaluations: An overview. *Cognition & Emotion*, 15, 115-141
- Fazio, R. H., & Olson, M. A. (2003). Implicit measures in social cognition research: Their meaning and use. *Annual review of psychology*, 54(1), 297-327.
- Gawronski, B., & De Houwer, J. (2014). Implicit measures in social and personality psychology. *Handbook of research methods in social and personality psychology*, 2.
- Gawronski, B., LeBel, E. P., & Peters, K. R. (2007). What do implicit measures tell us?: Scrutinizing the validity of three common assumptions. *Perspectives on Psychological Science*, 2(2), 181-193.
- Gawronski, B., & Payne, B. K. (Eds.). (2011). *Handbook of implicit social cognition: Measurement, theory, and applications*. Guilford Press.
- Glashouwer, K. A., Smulders, F. T., de Jong, P. J., Roefs, A., & Wiers, R. W. (2013). Measuring automatic associations: Validation of algorithms for the Implicit Association Test (IAT) in a laboratory setting. *Journal of Behavior Therapy and Experimental Psychiatry*, 44(1), 105-113.
- Graham, J. M. (2006). Congeneric and (essentially) tau-equivalent estimates of score reliability what they are and how to use them. *Educational and Psychological Measurement*, 66(6), 930-944.

- Greenwald, A. G., McGhee, D. E., & Schwartz, J. L. K. (1998). Measuring individual differences in implicit cognition: The Implicit Association Test. *Journal of Personality and Social Psychology*, 74, 1464-1480.
- Greenwald, A. G., Nosek, B. A., & Banaji, M. R. (2003). Understanding and using the implicit association test: I. An improved scoring algorithm. *Journal of personality and social psychology*, 85(2), 197.
- Greenwald, A. G., Poehlman, T. A., Uhlmann, E. L., & Banaji, M. R. (2009). Understanding and using the Implicit Association Test: III. Meta-analysis of predictive validity. *Journal of personality and social psychology*, 97(1), 17.
- Jöreskog, K. G. (1971). Statistical analysis of sets of congeneric tests. *Psychometrika*, 36(2), 109-133.
- Kawakami, K., & Dovidio, J. F. (2001). The reliability of implicit stereotyping. *Personality and Social Psychology Bulletin*, 27(2), 212-225.
- Krause, S., Back, M. D., Egloff, B., & Schmukle, S. C. (2011). Reliability of implicit self-esteem measures revisited. *European Journal of Personality*, 25(3), 239-251.
- LeBel, E. P., & Paunonen, S. V. (2011). Sexy but often unreliable: The impact of unreliability on the replicability of experimental findings with implicit measures. *Personality and Social Psychology Bulletin*, 37(4), 570-583.
- Little, T. D., Cunningham, W. A., Shahar, G., & Widaman, K. F. (2002). To parcel or not to parcel: Exploring the question, weighing the merits. *Structural equation modeling*, 9(2), 151-173.

- Little, T. D., Rhemtulla, M., Gibson, K., & Schoemann, A. M. (2013). Why the items versus parcels controversy needn't be one. *Psychological Methods*, 18(3), 285.
- Lord, F. M., Novick, M. R. (1968). *Statistical Theories of Mental Test Scores*. Reading, MA: Addison-Wesley.
- Mierke, J., & Klauer, K. C. (2003). Method-specific variance in the implicit association test. *Journal of personality and social psychology*, 85(6), 1180.
- Marsh, H., Lüdtke, O., Nagengast, B., Morin, A. J. S., & Von Davier, M. (2013) Why item parcels are (almost) never appropriate: two wrongs do not make a right-camouflaging misspecification with item parcels in CFA models. *Psychological Methods*, 18, 257-284. DOI: 10.1037/a0032773
- Maul, A. (2013). Method effects and the meaning of measurement. *Frontiers in Psychology*, V4, Article 169, 1-13.
- McDonald, R. P. (1978). Generalizability in factorable domains: "domain validity and generalizability": 1. *Educational and Psychological Measurement*, 38(1), 75–79.
- McDonald, R. P. (1999). *Test theory: A unified treatment*. Hillsdale: Erlbaum
- Mellenbergh, G. J. (1996). Measurement precision in test score and item response models. *Psychological Methods*, 1(3), 293.
- Meredith, W., & Teresi, J. A. (2006). An essay on measurement and factorial invariance. *Medical care*, 44(11), S69-S77.
- Nosek, B. A., Greenwald, A. G., & Banaji, M. R. (2007). The Implicit Association Test at age 7: A methodological and conceptual review. *Automatic processes in social thinking and behavior*, 265-292.

- Nosek, B. A., Hawkins, C. B., & Frazier, R. S. (2011). Implicit social cognition: From measures to mechanisms. *Trends in cognitive sciences*, 15(4), 152-159.
- Nosek, B. A., & Sriram, N. (2007). Faulty assumptions: a comment on Blanton, Jaccard, Gonzales, and Christie (2006). *Journal of Experimental Social Psychology*, 43(3), 393-398.
- Novick, M. R., & Lewis, C. (1967). Coefficient alpha and the reliability of composite measurements. *Psychometrika*, 32(1), 1-13.
- Payne, B. K., Cheng, C. M., Govorun, O., & Stewart, B. (2005). An inkblot for attitudes: Affect misattribution as implicit measurement. *Journal of Personality and Social Psychology*, 89, 277-293
- R Core Team (2016). *R: A language and environment for statistical computing*. R Foundation for Statistical Computing, Vienna, Austria. URL <https://www.R-project.org/>.
- Raykov, T. (2001). Bias of coefficient α for fixed congeneric measures with correlated errors. *Applied Psychological Measurement*, 25, 69–76.
- Raykov, T. (2004). Behavioral scale reliability and measurement invariance evaluation using latent variable modeling. *Behavior Therapy*, 35, 299-331.
- Rosseel, Y. (2012). lavaan: An R package for structural equation modeling. *Journal of Statistical Software*, 48(2), 1-36.
- Schmitz, F., Teige-Mocigemba, S., Voss, A., & Klauer, K. C. (2013). When scoring algorithms matter: Effects of working memory load on different IAT scores. *British Journal of Social Psychology*, 52(1), 103-121.

- Schmukle, S. C., & Egloff, B. (2006). Assessing anxiety with extrinsic Simon tasks. *Experimental Psychology*, 53(2), 149-160.
- Sijtsma, K. (2009). Reliability beyond theory and into practice. *Psychometrika*, 74(1), 169-173.
- Simmons, J. P., Nelson, L. D., & Simonsohn, U. (2011). False-positive psychology undisclosed flexibility in data collection and analysis allows presenting anything as significant. *Psychological science*, 0956797611417632.
- Stanley, D., & Phelps, E., & Banaji, M. R. (2008). The Neutral Basis of Implicit Attitudes. *Current Directions in Psychological Science*, 17(2), 164-70.
- Uhlmann, E. L., Pizarro, D. A., & Bloom, P. (2008). Varieties of social cognition. *Journal for the Theory of Social Behaviour*, 38(3), 293-322.
- Yang, Y., & Green, S. B. (2011). Coefficient alpha: A reliability coefficient for the 21st century? *Journal of Psychoeducational Assessment*, 0734282911406668.

Appendix

```
# libraries
require(lavaan)

# Cronbach alpha formula
alpha <- function(data){
  S <- cov(data)
  p <- ncol(S)
  p/(p-1) * (sum(S) - sum(diag(S)))/sum(S)
}

# Read Data (N=322)
# S1 - S2 = Split Half Parcel 1 (2): odd (even) trials
# A1 - A4 = Sequential Parcel 1 (4)
# WP, BP, NP, PP = White, Black, Negative, Positive Parcel

Data <- read.table("http://users.ugent.be/~mldschr/omega/IAT_Data_Sample_322.dat",
  header = TRUE)

# Split-half reliability:
rsh <- cor(Data$S1,Data$S2)
rsh

# Spearman-Brown correction
2*(rsh) / (1+rsh)

# Cronbach alpha
alpha(Data[,c(7:10)])

# 2 Omega using sem
# Standard errors based on bootstrap

# Congeneric Model
M1 <- '
  L =~ lam1*WP + lam2*BP + lam3*NP + lam4*PP

  WP~~e1*WP
  BP~~e2*BP
  NP~~e3*NP
  PP~~e4*PP

  #Reliability omega
  omega := (lam1 + lam2 + lam3 + lam4)^2 /
    ((lam1 + lam2+ lam3 + lam4)^2 + (e1 + e2 + e3 + e4))
  '

fit1 <- cfa(M1, data=Data, std.lv=TRUE, se="boot")
summary(fit1, fit.measures=TRUE, standardized=TRUE)

# Tau-equivalent Model
M2 <- '
  L =~ lam1*WP + lam1*BP + lam1*NP + lam1*PP
  WP~~e1*WP
  BP~~e2*BP
  NP~~e3*NP
  PP~~e4*PP
  '

fit2<- cfa(M2, data=Data, std.lv=TRUE)

summary(fit2, fit.measures=TRUE, standardized=TRUE)

# Compare Tau-equivalent Model with Congeneric Model
anova(fit1,fit2)
```