

A Multitrait-Multimethod Validation of the Implicit Association Test

Implicit and Explicit Attitudes Are Related but Distinct Constructs

Brian A. Nosek and Frederick L. Smyth

University of Virginia, Charlottesville, VA, USA

Abstract. Recent theoretical and methodological innovations suggest a distinction between implicit and explicit evaluations. We applied Campbell and Fiske's (1959) classic multitrait-multimethod design precepts to test the construct validity of implicit attitudes as measured by the Implicit Association Test (IAT). Participants ($N = 287$) were measured on both self-report and IAT for up to seven attitude domains. Through a sequence of latent-variable structural models, systematic method variance was distinguished from attitude variance, and a correlated two-factors-per-attitude model (implicit and explicit factors) was superior to a single-factor-per-attitude specification. That is, despite sometimes strong relations between implicit and explicit attitude factors, collapsing their indicators into a single attitude factor resulted in relatively inferior model fit. We conclude that these implicit and explicit measures assess related but distinct attitude constructs. This provides a basis for, but does not distinguish between, dual-process and dual-representation theories that account for the distinctions between constructs.

Keywords: implicit social cognition, attitudes, individual differences, construct validity, structural equation modeling

Realizing that the human mind is more than the sum of its conscious processes, a number of theorists have proposed a conceptual distinction between evaluations that are the products of introspection, called explicit attitudes, and those that occur automatically and may exist outside of conscious awareness, called implicit attitudes (Greenwald & Banaji, 1995; Wilson, Lindsey, & Schooler, 2000). Greenwald and Banaji (1995, p. 8), for example, defined implicit attitudes as "introspectively unidentified (or inaccurately identified) traces of past experience that mediate favorable or unfavorable feelings toward an attitude object." This theory has developed in conjunction with the invention of measurement tools that assess automatic evaluative associations without introspection (e.g., Fazio, Sanbonmatsu, Powell, & Kardes, 1986; Greenwald, McGhee, & Schwartz, 1998; Nosek & Banaji, 2001; Wittenbrink, Judd, & Park, 1997).

Some experiences with these new measurement tools have spawned doubts about whether they measure attitudes at all (Karpinski & Hilton, 2001), and whether a conceptual distinction between implicit and explicit attitudes is worthwhile (Fazio & Olson, 2003). Fazio and Olson contend "it is more appropriate to view the *measures* as implicit or explicit, *not* the attitude (or whatever other construct)" (2003, p. 303). The purpose of the research we report was to test whether the structure of attitude variance derived from an implicit measure (the Implicit Association Test [IAT]; Greenwald et al., 1998) and from an explicit one

(semantic differentials) is best represented by one latent factor or by two correlated latent factors, when stripped of confounding method variance. If our hypothesis that the latter specification will fit the data better is sustained, it will support a view that substantively different attitude constructs, distinguishable from the techniques used to measure them, underlie data collected by explicit and implicit methods.

The finding would not, however, lend weight to one side or the other in the debate about origins of the distinction between implicit and explicit attitudes – i.e., are they derived from a single representation at different stages of processing (Fazio & Olson, 2003) or do they reflect distinct evaluative sources (Strack & Deutsch, 2004; Wilson et al., 2000). Also, this multitrait-multimethod (MTMM) approach does not identify the cognitive processes that distinguish the constructs.

The context of this analysis follows from Cronbach and Meehl's (1955) classic discussion of construct validation where a construct is an indeterminant function of representation and process. So, our reference to distinct implicit and explicit attitude constructs should be interpreted as referring to distinguishable attitudinal components, without implying a specific commitment to distinguishable formative processes, single versus dual mental representations, or single versus dual operative processes.

Any of these theoretical perspectives can explain dual constructs by postulating combinations of representations

and processes to account for the observed differences between constructs. Dual-construct validation justifies the need to have theoretical models account for the distinction without providing evidence for or against any particular explanation. Further illustration of the difference between construct validation versus commitments to dual-process or dual-representation theories appears in the discussion (see also Greenwald, Nosek, & Banaji, in press).

Construct Validation

The conceptual and empirical justification for a psychological construct requires development of a nomological net of facts, relationships, and validity evidence that clarifies the identity and utility of the construct (Cronbach & Meehl, 1955; McArdle & Prescott, 1992). The nomological net supporting the validity of implicit attitudes has been gaining strength (Greenwald & Banaji, 1995; Nosek, Greenwald, & Banaji, in press; Wilson et al., 2000). For example, Poehlman, Uhlmann, Greenwald, and Banaji (2004) conducted a meta-analysis of studies examining the predictive validity of the IAT, a measure thought to be influenced by automatic associations, and found that the IAT had robust predictive validity across domains, and outperformed self-report measures in some domains (stereotyping and prejudice), while self-report outperformed the IAT in other domains (e.g., political preferences). Also, recent social-neuroscience research finds evidence for a neurological distinction between automatic and controlled evaluative processes (Cunningham, Johnson, Gatenby, Gore, & Banaji, 2003; Cunningham, Johnson et al., 2004). Our research provides another avenue of evidence for this growing nomological net by examining the relationship between implicit and explicit attitude measures to determine whether they can be fairly interpreted as measuring a single construct, or whether they assess related, but distinct constructs.

Preliminary Evidence

Greenwald and Farnham (2000) observed that a model describing implicit and explicit self-esteem as distinct latent factors provided a better fit than a single self-esteem conceptualization. Likewise, Cunningham, Preacher, and Banaji (2001) found implicit and explicit measures of racial attitudes to reveal related, but distinct factors, as did Cunningham, Netlek, and Banaji (2004) for implicit and explicit ethnocentrism. Following this approach of comparing

single versus dual factors in structural equation modeling, we reanalyzed a large dataset reported by Nosek (2005). We found support for the generalizability of a model of distinct-but-related latent implicit and explicit attitude constructs across 56 of 57¹ widely varying attitude domains showing that this observation is quite general (see Table 1 and the supplement to this paper available at <http://brian-nosek.com/> for a full report). Even so, in this and the other previous structural modeling studies, specification of implicit and explicit attitude constructs is confounded with measurement method. As a consequence, a two-factor solution is an indeterminate function of both attitude and method variance.²

To transcend this inferential limitation, here, guided by principles articulated by Campbell and Fiske (1959), we use a MTMM design and comparative structural modeling analyses. This approach allowed us to distinguish attitude and method variance from IAT and thermometer ratings, the operationalizations of implicit and explicit attitudes, respectively. Our findings demonstrate (1) convergent and discriminant validity of the IAT, (2) that a model of distinct, but related implicit and explicit attitudes best fits the data, and (3) that this characterization is not attributable to attitude-irrelevant method variance of the IAT or of self-report.

MTMM and Confirmatory Factor Analysis

In their classic article on construct validation, Campbell and Fiske (1959) articulated a strategy for using MTMM matrices to evaluate convergent and discriminant validity. This strategy requires measurement of two or more ostensibly distinct trait constructs by two or more measurement methods. Convergent validity is demonstrated when indicators of a given trait (or, in our study, attitude) correlate highly across measurement method, while discriminant validity obtains when correlations between ostensibly different traits are low. Campbell and Fiske argued that “the clear-cut demonstration of the presence of method variance requires both several traits and several methods” (p. 85). They described ways to statistically evaluate the respective contributions of trait and method factors, but looked forward to continued progress in developing more rigorous validation methods.

Confirmatory factor analysis (CFA) has emerged as a tool well-suited to the partitioning of MTMM data envisioned by Campbell and Fiske (Jöreskog, 1974; Widaman, 1985). According to Marsh and Grayson (1995, p. 181),

1 The two-factor model for the males-females attitude domain failed to converge, leaving the hypothesis untested for this domain.

2 Cunningham, Nezlek et al. (2004) footnoted this limitation, but argued that, since a control IAT, birds vs. trees, did not load with five ingroup-outgroup IATs on an implicit ethnocentrism factor, IAT method variance was not a strong driver of the two-factor solution. They further suggested that if “systematic measurement error alone” was responsible for the implicit ethnocentrism factor, then the substantial correlation between implicit and explicit ethnocentrism ($r = .47$) would be unlikely. We do not disagree, and we test this supposition systematically.

Table 1. Results of one- and two-factor (oblique) attitude models fit to measures of 57 topics, a structural equation model reanalysis of Nosek (2005)

Attitude topic comparisons			One-factor model			Two-factor model			Factors $r(t)$
A	B	N	χ^2	ϵ_a	90%CI ϵ_a	χ^2	ϵ_a	90%CI ϵ_a	
Whites	Asians	279	44	.28	.21-.35	1.4	.04	.00-.17	.01 (0.1)
Cold	Hot	211	262	.79	.71-.87	0.5	.00	.00-.16	.13 (2.1)
Skirts	Pants	255	73	.37	.30-.45	0.1	.00	.00-.10	.15 (3.1)
Future	Past	235	76	.40	.32-.48	3	.10	.00-.23	.26 (1.9)
Thin people	Fat people	275	50	.30	.23-.37	0.4	.00	.00-.14	.27 (2.0)
Approaching	Avoiding	180	28	.27	.19-.36	0.3	.00	.00-.16	.29 (1.3)
Simple	Difficult	210	56	.36	.28-.44	0.0	.00	.00-.01	.29 (3.4)
Public	Private	196	40	.31	.23-.40	0.0	.00	.00-.05	.30 (3.7)
Freedom	Security	220	67	.38	.31-.47	1	.00	.00-.16	.31 (2.5)
Short people	Tall people	226	53	.34	.26-.42	3	.09	.00-.22	.31 (2.9)
Family	Career	238	42	.29	.22-.37	0.1	.00	.00-.12	.33 (2.7)
Married	Single	261	169	.57	.50-.64	0.5	.00	.00-.14	.33 (3.4)
Rich people	Poor people	222	89	.44	.37-.52	1	.00	.00-.17	.34 (3.2)
Education	Defense	214	50	.33	.26-.42	2	.07	.00-.21	.36 (3.7)
Letters	Numbers	228	75	.40	.33-.48	0.1	.00	.00-.12	.37 (3.4)
Nerds	Jocks	239	80	.40	.33-.48	0.0	.00	.00-.00	.38 (4.0)
Young people	Old people	250	42	.28	.21-.36	0.0	.00	.00-.00	.40 (3.1)
Imprisonment	Capital punishment	260	66	.35	.28-.43	0.2	.00	.00-.13	.41 (3.7)
Yankees	Diamondbacks	200	42	.32	.24-.40	1	.04	.00-.20	.41 (4.5)
Flexible	Stable	201	35	.29	.21-.37	0.0	.00	.00-.10	.42 (4.4)
Meg Ryan	Julia Roberts	255	41	.28	.21-.35	1	.03	.00-.17	.43 (4.7)
Emotion	Reason	175	84	.49	.40-.58	0.0	.00	.00-.09	.44 (3.9)
Conforming	Rebellious	208	132	.56	.48-.64	0.4	.00	.00-.16	.44 (4.2)
Summer	Winter	260	133	.50	.43-.58	0.4	.00	.00-.14	.44 (5.5)
Leaders	Helpers	265	60	.33	.26-.41	2	.07	.00-.19	.45 (5.4)
Tom Cruise	Denzel Washington	242	52	.32	.25-.40	0.3	.00	.00-.14	.46 (4.8)
Management	Labor	194	47	.34	.26-.43	0.1	.00	.00-.14	.47 (3.9)
Exercising	Relaxing	258	185	.60	.53-.67	0.3	.00	.00-.14	.48 (5.8)
Jay Leno	David Letterman	196	49	.35	.27-.43	0.0	.00	.00-.04	.48 (4.5)
American places	Foreign places	205	44	.32	.24-.40	0.0	.00	.00-.11	.49 (5.3)
Microsoft	Apple	205	77	.43	.35-.51	3	.09	.00-.23	.49 (4.1)
California	New York	253	66	.36	.29-.43	2	.05	.00-.18	.49 (4.3)
Tea	Coffee	250	90	.42	.35-.50	0.0	.00	.00-.00	.50 (5.6)
Abstaining	Drinking	249	100	.44	.37-.52	0.2	.00	.00-.13	.50 (5.9)
Christian	Jewish	253	77	.39	.31-.46	0.0	.00	.00-.00	.50 (5.2)
Classical	Hip Hop	243	106	.47	.39-.54	1	.00	.00-.16	.51 (6.0)
Northerners	Southerners	207	62	.39	.30-.47	1	.00	.00-.17	.51 (5.8)
Jews	Muslims	243	52	.32	.25-.40	2	.07	.00-.20	.52 (5.5)
Books	Television	233	77	.40	.33-.48	0.1	.00	.00-.11	.54 (5.2)
Cats	Dogs	258	77	.38	.31-.46	1	.00	.00-.15	.54 (5.6)

Table 1 continued.

Attitude topic comparisons			One-factor model			Two-factor model			Factors r (t)
A	B	N	χ^2	ϵ_a	90%CI ϵ_a	χ^2	ϵ_a	90%CI ϵ_a	
European Americans	African Americans	254	39	.27	.20–.35	0.0	.00	.00–.09	.55 (4.7)
American	Canadian	290	44	.27	.20–.34	0.5	.00	.00–.14	.55 (4.7)
Teen pop	Jazz	239	81	.41	.33–.48	1	.00	.00–.15	.56 (6.1)
Vegetables	Meat	234	76	.40	.32–.48	0.3	.00	.00–.14	.56 (5.4)
Social programs	Tax reductions	188	77	.45	.36–.53	5	.15	.04–.28	.57 (5.6)
USA	Japan	246	49	.31	.24–.39	0.0	.00	.00–.00	.57 (5.2)
Gun rights	Gun control	216	85	.44	.36–.52	1	.04	.00–.19	.59 (6.8)
Straight people	Gay people	175	35	.31	.22–.40	0.1	.00	.00–.14	.60 (5.3)
Religion	Atheism	211	160	.61	.53–.69	2	.05	.00–.20	.61 (7.1)
Coke	Pepsi	250	71	.37	.30–.45	1	.00	.00–.16	.66 (7.1)
Liberals	Conservatives	215	124	.53	.46–.61	1	.00	.00–.17	.67 (6.9)
Creationism	Evolution	231	99	.46	.39–.54	3	.08	.00–.21	.68 (7.2)
Feminism	Traditional values	226	75	.40	.33–.48	0.4	.00	.00–.15	.71 (7.4)
Gore	Bush	211	61	.37	.30–.46	0.0	.00	.00–.09	.74 (7.2)
Democrats	Republicans	195	42	.32	.24–.41	2	.08	.00–.23	.75 (6.7)
Prochoice	Prolife	242	52	.32	.25–.40	0.4	.00	.00–.10	.79 (8.6)
Females	Males	289	112	.44	.37–.51	*	*	*	*

Note. Attitude object A was implicitly preferred on average. All one-factor models have $df = 2$, and two-factor models have $df = 1$. ϵ_a = root-mean-square error of approximation. CI = confidence interval. $t = r/se$. Factor correlations in boldface have $t \geq 2.0$. * model did not converge.

“the operationalizations of convergent validity, discriminant validity, and method effects in the CFA approach apparently better reflect Campbell and Fiske’s (1959) original intentions than do their own guidelines.” By comparing the fits of nested structural models, the relative merits of alternative hypotheses concerning the structure of trait (attitude) and method variance can be systematically tested (Jöreskog & Sörbom, 1979; Loehlin, 2004; McDonald, 1985). We used this approach to distinguish method variance from trait variance for seven attitude object pairs. And, within this framework, we tested whether a single attitude construct or distinct implicit and explicit attitude constructs provides a better fit for the data.

This MTMM design, because it enables direct modeling of method factors, increases confidence that a finding of distinct implicit and explicit attitude factors indicates a substantive distinction and not one that is driven by confounding influences in the measurement requirements. Our implicit measurement instrument, the IAT, measures associations between concepts (e.g., thin–fat) and attributes (e.g., good–bad) by comparing the average response times for sorting exemplars of those concepts and attributes in two distinct response conditions – one in which sorting *thin* and *good* exemplars requires a single response and sorting *fat* and *bad* exemplars requires an alternate response, and a second in which sorting *fat* and *good* exemplars requires a single response and sorting *thin* and *bad* exemplars requires an alternate response.

This method is distinct from attitude self-report in which a participant self-assesses attitudes by reporting the magnitude of good or bad feelings on a response scale. Because of their radically different measurement properties, it is possible that the unique components of the better two-factor models observed in earlier research are the result of sources of method variance such as cognitive fluency or task-switching ability, two known influences on IAT effects (Mierke & Klauer, 2003; see Nosek et al., 2006, for a review).

Study Overview

Data are from four laboratory studies in which attitudes toward seven different attitude-object pairs were measured: flowers–insects, Democrats–Republicans, humanities–science, straight–gay, Whites–Blacks, creationism–evolution, and thin people–fat people. These domains were selected because on their face they cover a broad range of attitudes. This was important so that our goal of accounting for common method variance would not be confounded with common substantive variance. For example, Cunningham, Nezlek, and Banaji’s (2004) examination of implicit and explicit ethnocentrism specifically hypothesized that attitudes for a variety of ingroup-outgroup domains (e.g., poor–rich, Blacks–Whites, Jews–Christians) would share a

common ethnocentrism factor. They reported support for this idea and also found that implicit and explicit ethnocentrism factors were related, but distinct.

Our goal, in one sense, was the opposite of Cunningham et al.'s: they sought to demonstrate convergent validity between conceptually related attitude domains revealing an ethnocentrism factor. We sought to demonstrate *discriminant* validity *between* attitude domains – hypothesizing that the attitude domains would form distinct factors; and *convergent* validity *across* measurement types (IAT and self-report) – hypothesizing that the implicit and explicit attitude constructs would be related, but retain distinctiveness not accounted for by method factors. This simultaneous examination of discriminant and convergent validity is the core value of the MTMM approach.

Campbell and Fiske (1959) stressed that “One cannot define [a construct] without implying distinctions, and the verification of these distinctions is an important part of the validation processes” (p. 84). Self-report and IAT measures were obtained for each attitude object pair (e.g., Whites–Blacks) and participants were measured on multiple pairs. We fitted a sequence of nested covariance structure models, beginning with one in which method variance was partitioned into latent factors, and we predicted that a model specifying distinct implicit and explicit attitude factors would provide the best data fit, whether or not the partitioning of method variance proved important.

Method

Participants

A total of 287 Yale University undergraduates from four data collections in 2000 and 2001 ($n = 81, 86, 60, 60$) comprise the study sample.

Materials

Implicit Association Test (IAT)

One of the four samples received IATs for all seven object pairs, while the others received subsets of at least four pairs, including the flowers–insects and Democrats–Republicans pairs (patterns of measured variables are listed in the Appendix). All IAT category headings and exemplars are listed in the Appendix. IAT D scores were computed based on the scoring algorithm suggested by Greenwald, Nosek, and Banaji (2003), that is, by taking the dif-

ference in mean response latency between the two critical block conditions and scaling it by the participant's average latency standard deviation for both blocks. Most of the IATs administered across the four samples consisted of 56 trials in each of the critical double-discrimination blocks.³ IAT scores were removed from analysis if more than 10% of trials were unreasonably fast (< 300 ms) or if the error rate for any block of trials was greater than 39%. This cleaning resulted in elimination of less than 1% of IAT scores (14 of 1475).

In all data collections, participants first completed a flowers–insects IAT with the order of blocks conforming to that suggested by Greenwald et al. (2003), i.e., (1) a single-discrimination block of trials for bad–good exemplars, followed by (2) a single-discrimination block for flower–insect exemplars, then (3) a double-discrimination block of (counterbalanced) either flowers+good/insects+bad or flowers+bad/insects+good, followed by (4) another single-discrimination block to practice the switched flower–insect key assignments for the (5) final, reversed, double-discrimination block. For the remaining IATs, response blocks for all tasks were randomized and single-discrimination practice blocks were eliminated. For example, after the flowers–insects IAT, a participant could receive the race attitude compatible block of trials (i.e., compatible with the dominant prejudice) in which Black faces and bad words are to be categorized with one response key and White faces and good words are categorized with another key, without any opportunity to practice the simple, single-discrimination task of sorting Black from White faces; then the participant could receive the incompatible block for the gay–straight IAT (gay+good with one key, straight+bad with the other); then the Democrat+bad/Republican+good block; then the incompatible block for race attitude, etc. This atypical approach provides a tough test for identifying attitude factors because the component performance tasks are intermixed with performance tasks for the other attitude domains. In this way, we allow substantial opportunity for method factors to influence IAT performance and challenge the hypothesis that distinguishable attitude factors can be identified despite intermixed performance blocks.

To facilitate latent variable analyses, four IAT D scores were calculated for each attitude domain based on the difference between the means of each fourth of the trials, in turn, across the critical blocks. That is, for IATs with 56 critical trials in each double-discrimination block, the mean latency for trials 1–14 in one double-discrimination condition was compared with that of trials 1–14 in the other, and so on for sets of trials 15–28, 29–42, and 43–56.⁴

³ Exceptions to the 56-trial critical block format were the flowers–insects IAT for two of the samples (40 trials in critical blocks) and the creation–evolution and thin–fat IATs for 13 participants in one of the samples, also 40 critical trials.

⁴ We also conducted the analysis with just two IAT indicators, first half vs. second half of trials from each block, and found comparable results in our comparative structural modeling. Using four indicators was preferable, however, in terms of attaining stable estimates in the more complex models.

Self-Report Measures

Participants reported attitudes toward each of the target objects per pair independently using two 9-point semantic differentials. Anchors for these differentials varied across data collections, including, for a given study, two of the following four pairs: cold–warm, unpleasant–pleasant, bad–good, or unfavorable–favorable. A difference score was calculated between ratings of each object of a pair so as to be conceptually parallel to the relative character of IAT scores. Positive values indicate greater liking for the object that was implicitly preferred on average (–8 to +8). For example, a participant who rates gay people as a 5 on the bad–good scale and straight people as a 7 would have a difference score of 2, indicating relative preference for straight people.

Procedure

Similar procedures were used across all four data collections. After informed consent, participants completed a selection of IATs and self-report measures. The order of implicit and explicit measures was counterbalanced between subjects. The correlation matrices are substantively similar when each data collection is considered independently.

Analyses

Following guidelines suggested by Campbell and Fiske (1959), we first describe and offer interpretations of the reliability and validity relations in the MTMM matrix. We quickly proceed, however, to the CFA approach so that specific, rejectable hypotheses about the structure of the data can be tested. Following Widaman (1985), we compare model fits across a progression of five nested structural equation models designed to account for any common method variance and to test a primary hypothesis that our two measurement approaches, self-reports and IATs, assess distinct – but related – attitude constructs.

All five models can be understood with reference to the three abridged diagrams in Figure 1, each of which depicts just two of the seven attitude objects actually modeled. Our primary criterion for judging whether one model significantly improved on the fit of another is based on the root mean square error of approximation (RMSEA) of the change ($\epsilon_a\Delta$) in fit (Browne, 1991): The model fits are considered significantly different if the 95% confidence interval (CI) around $\epsilon_a\Delta$ does not encompass .05. This criterion is based on guidelines proposed by MacCallum, Browne, and Sugawara (1996), who suggest that models fitting with $RMSEA < .05$ should be considered close fits. Thus, if the 95% CI for the $\epsilon_a\Delta$ includes .05, the models being compared are considered close to one another in fit, and one would be preferred only to the extent that it involves fewer parameter estimates (i.e., is more parsimonious).

In Model 1 (see Figure 1) we specify two method factors to account for the covariances among the 42 observed indicators (14 explicit and 28 implicit) across all seven general attitude domains. Though the fit of this model would certainly improve on a null model, i.e., one positing no relations between the indicators, we did not expect it to be a good fit for the data. It was a useful first model, however, since it partitions any variance that is common to the measurement instruments. In Model 2, we built on Model 1 by specifying, in addition to the common method factors, one factor for each of the six indicators of a given attitude domain. Comparing the fit of this model with that of Model 1 allowed us to examine whether accounting for common variance within each of the seven attitude domains makes an important improvement over the simple two-factor specification of Model 1.

In Campbell and Fiske's terms, this comparison can serve to formally, i.e., statistically, test both the discriminant and convergent validity of our measures. Discriminant validity is demonstrated to the extent that the fit of Model 2 improves on that of Model 1. That is, specifying a factor for each of the seven attitude domains, which discriminates among them, is superior to modeling them as the same thing within a measurement technique. At the same time, convergent validity between the explicit and implicit approaches is demonstrated to the extent that Model 2 fits the data well. A good fit, in other words, would indicate that the explicitly and implicitly measured indicators have something in common, i.e., are converging on the same construct.

Finally, in Model 3, we represent the dual-construct hypothesis by specifying distinct but related (correlated) implicit and explicit attitude factors. If the fit of this model is significantly better than that of Model 2, whatever the convergence between the explicit and implicit measurement approaches, then it is practically and theoretically useful to treat them as two constructs rather than as one. Once again, in Campbell and Fiske's terms, this illustrates discriminant validity, but now between implicit and explicit attitude constructs, since variance common to measurement method is partitioned in the model.

In Models 4 and 5 (not depicted in Figure 1) we evaluate the utility of specifying the respective common implicit and explicit method factors. In other words, we answer the question of whether we really need to account for common method variance to appropriately model the patterns of covariation in these data. This is accomplished by specifying models that are identical to Model 3 except that, in turn, a common implicit method factor is not specified (Model 4) and a common explicit method factor is not specified (Model 5). By comparing the fits of each with that of Model 3, we may discern to what extent accounting for common method variance is important to understanding the structure of relations among these measures.

We used the full information maximum likelihood (FIML) estimation approach in our analyses (Enders & Bandalos, 2001; McArdle, 1994) so that model fit would

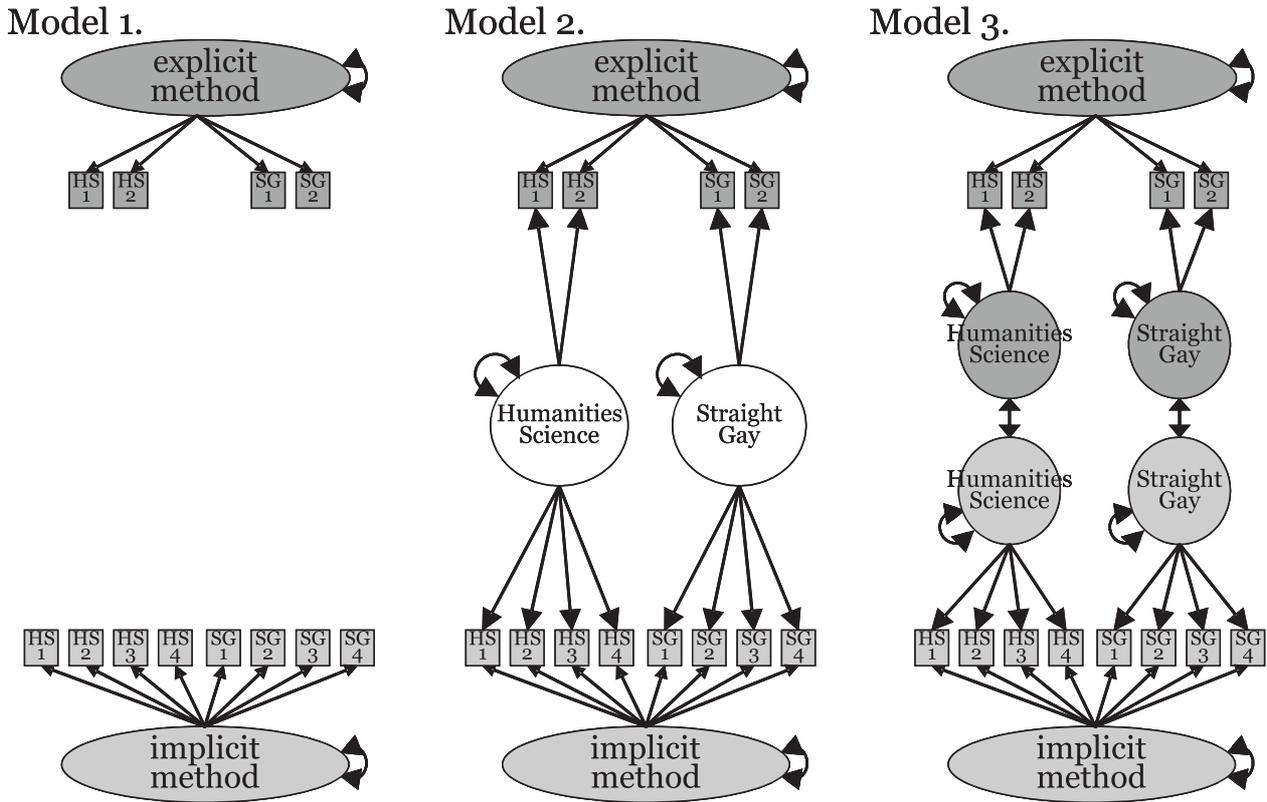


Figure 1. Abridged representation of structural model hypotheses (only two of seven attitude pairs depicted per model). Squares depict measured variables; circles depict latent factors. Darker shading indicates explicit measures/constructs; lighter indicates implicit measures/constructs.

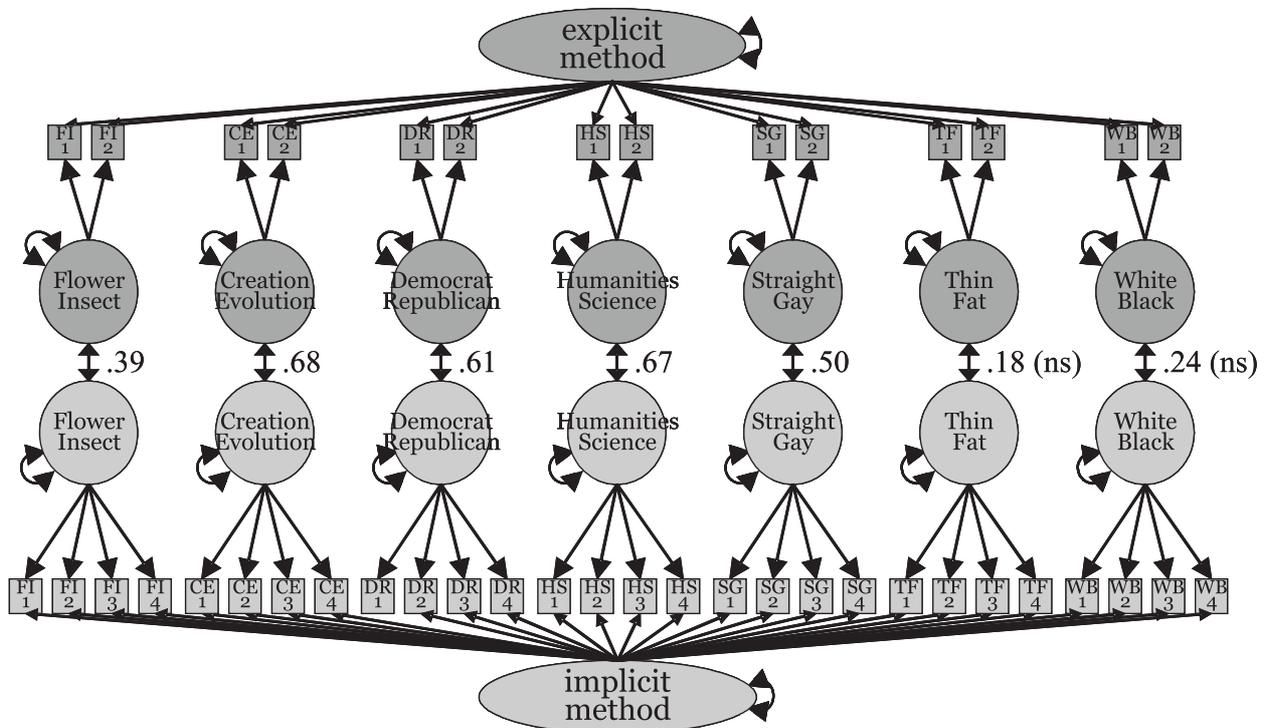


Figure 2. Diagram of Model 3. Squares depict measured variables; circles depict latent factors. Darker shading indicates explicit measures/constructs; lighter indicates implicit measures/constructs. Estimated implicit-explicit factor r is shown for each attitude domain.

Table 2. Multitrait-multimethod matrix for seven IAT and seven self-report attitude measures

	IAT							Self-Report						
	FI	DR	HS	SG	WB	CE	TF	FI	DR	HS	SG	WB	CE	TF
Mean	.60	.19	.33	.26	.19	.04	.25	4.09	2.19	.87	1.20	.10	-2.25	1.47
SD	.38	.44	.43	.38	.40	.44	.42	2.23	2.76	2.68	1.93	1.17	3.61	1.79
N	283	281	223	222	166	141	146	86	287	227	227	167	146	146
								^{aa} .227			^{aa} .167			^{aa} .86
IAT														
FI	(.80)													
DR	.10	(.87)												
HS	.02	.17	(.85)											
SG	.08	-.03	-.11	(.80)										
WB	.11	-.06	.10	.23	(.84)									
CE	-.03	-.14	-.07	.07	-.14	(.81)								
TF	.13	-.11	-.14	.10	.10	-.15	(.81)							
Self-Report														
FI	.27	-.01	.09	-.18	.02	.01	.12	(.96)						
DR	.01	.51	.17	-.11	-.07	-.02	-.12	.17	(.94)					
HS	.05	.09	.52	.00	.04	.06	-.11	.13	.09	(.91)				
SG	.04	-.16	-.28	.39	-.03	.18	.00	-.05	-.16	-.14	(.90)			
WB	.09	-.20	-.20	.08	.12	-.31	.16	.04	-.27	-.06	.25	(.88)		
CE	-.09	-.11	-.16	.05	-.16	.56	.01	.06	-.05	.00	.19	-.22	(.97)	
TF	-.02	.00	-.01	.05	.10	-.18	.13	.06	.04	.11	.22	.18	-.13	(.70)

Note: FI = flowers–insects, DR = Democrats–Republicans, HS = humanities–science, SG = straight–gay, WB = Whites–Blacks, CE = creationism–evolution, TF = thin–fat. Positive means indicate preference for first object of an attitude pair. IAT metric is D (Greenwald et al., 2003); Self-report metric is -8 to $+8$, a difference score derived from two 9-pt ratings. **Bold** = $p < .05$, Underline $p < .01$, *Italic* $p < .001$. Values in parentheses are split-half reliabilities (Cronbach's α): For IAT, split-halves were formed from alternating couplets of trials (couplet = one object stimulus and one evaluative stimulus); for self-report, reliabilities are for the two difference scores; α_n = the n upon which α was calculated if different from that used to calculate the mean (i.e., for three domains, only one explicit difference score was obtained). Gray shading indicates validity correlations. Data are pooled from four studies of Yale undergraduates that varied in terms of which attitude domains were measured: Study 1 ($n = 81$, FI, DR, HS, SG, WB), Study 2 ($n = 86$, FI, DR, HS, SG, WB, CE, TF), Study 3 ($n = 60$, FI, DR, HS, CE), Study 4 ($n = 60$, FI, DR, GS, TF).

be based on all available data. That is, even though our four samples varied in terms of the collection of attitude domains on which they were measured (see Table 2), the FIML technique allows variable interrelations and parameter standard errors to be estimated from the combined data of all samples. Our design included the key characteristic necessary for this planned missing data approach to work, i.e., overlapping measures across the samples. All participants received the flowers–insects and Democrat–Republicans IATs, one sample received implicit and explicit measures in all seven domains, and every domain was included in the test battery for at least two of the samples.

Results

Description of the MTMM Correlations

Table 2 is a MTMM matrix for the two methods and seven attitude object pairs. The first three rows of the table provide descriptive statistics for each of the fourteen measurements

(i.e., full IAT D scores and averages of the self-reports). Reliability estimates (Cronbach's α) are shown in parentheses along the main diagonal in the top-left (IAT) and bottom-right (self-report) panels, and intramethod (Campbell and Fiske's monomethod-heterotrait) correlations are listed in the other cells of these respective panels. For the IAT, reliabilities are based on D scores for split-halves formed from alternating couplets of trials, since a couplet consisting of an object stimulus (e.g., a science word) and an evaluative stimulus (e.g., a "good" word) occurred every two trials. Reliabilities for self-reports are based on the two difference scores derived for each set of object ratings. The median reliability for the IATs was $\alpha = .81$, and was $\alpha = .91$ for the self-reports. The generally low intramethod correlations (IAT median $r = -.03$ and self-report median $r = .04$), are evidence of within-method discriminant validity. That is, little evidence of common method variance is apparent across attitude domains. Even a more conservative approach of taking the absolute values of the intramethod correlations reveals weak relations across the IATs (median $|r| = .10$) and across self-reports (median $|r| = .13$).

Table 3. Goodness-of-fit indices for structural models representing multitrait-multimethod data

Model	χ^2	df	ϵ_a	$\Delta\chi^2/\Delta df$	95%CI $\epsilon_a\Delta$
[0] Null (no relations among 42 indicators)	3869	861	.110		
[1] Two uncorrelated method factors	3175	820	.100	684/41	.216–.253
[2] Seven uncorrelated attitudes + Model 1	1323	778	.049	1852/42	.370–.406
[3] Seven implicit & seven explicit attitudes + Model 1	959	771	.029	364/ 7	.378–.466
[4] Drop implicit method from Model 3	1023	799	.031	64/28	.065–.125
[5] Drop explicit method from Model 3	1012	785	.032	53/14 ^a	.066–.133 ^a

Note. ϵ_a = root-mean-square error of approximation (RMSEA) for the model. Unless noted, $\Delta\chi^2/\Delta df$ = change in χ^2 and degrees of freedom relative to the previous model. 95%CI $\epsilon_a\Delta$ = confidence interval around RMSEA of the change in fit between models; If .050 falls within the CI, then model fits are not considered significantly different. ^aChange is relative to Model 3. We tested an alternative model (2b) in which the correlation between explicit and implicit method factors was allowed, but the correlation was n.s. Similarly, we fit alternative Models 2a and 3a in which crossdomain attitude factor correlations were estimated, but neither resulted in a significant improvement in fit.

Correlations in the gray diagonal of the bottom-left panel can be used to assess convergent validity (monotrait-heteromethod), and discriminant validity (heterotrait-heteromethod) can be assessed by those off the diagonal in this panel. Supporting the interpretation of the IAT as a measure of attitudes, five of the seven convergent validity correlations, i.e., between implicit and explicit measures of the same attitude targets, were significantly positive ($p < .05$, r s ranging from .27 to .56) while the other two (thin-fat $r = .13$, Whites-Blacks $r = .12$) were not statistically significant. At the same time, discriminant validity correlations between attitude objects, across methods were weak (median $r = .04$, $|r| = .13$).⁵ These data are consistent with our hypotheses for the convergent and discriminant validity of the IAT and self-report across attitude domains. However, the power of the MTMM design is not fully harnessed by scrutinizing correlation matrices. The relative merits of competing hypotheses about the structure of the data can be tested formally by comparing the fits of nested, but differentially specified, structural equation models.

MTMM Structural Equation Models

Summary statistics from the confirmatory structural models are listed in Table 3 (details of each model's specifications and parameter estimates – all fit with *Mplus* statistical software (Muthén & Muthén, 1998–2004) – are available in the supplement to this paper at <http://briannosek.com/>).

Fit indices for a null model (0), in which means and variances are estimated for the 42 manifest variables, but with no interrelations between them, are provided as a baseline of (mis)fit ($\chi^2 = 3869$ on 861 df ; $\epsilon_a = .110$). In Model 1, the first of substantive interest (see Figure 1), two method factors are specified, one loading on the explicit indicators and one on the implicit.⁶ Relative to Model 0, the ratio of the change (improvement) in fit ($\Delta\chi^2 = 684$) to the change in df ($\Delta df = 41$) was statistically significant ($\Delta\chi^2/\Delta df = 17$), suggesting that the fit of Model 1 is superior, but χ^2 comparison does not take into account the sample N or model complexity. We focused instead on the 95% confidence interval around the RMSEA of this change (95%CI $\epsilon_a\Delta = .22$ –.25). This range does not include the $\epsilon_a < .05$ benchmark of close fit, and so we concluded that the fits of these two models are not close to one another. It is not surprising that this model improved on the null model, but, as expected, with $\epsilon_a = .10$, it is not a close representation of the data. In sum, this first comparison established that a statistically significant, but small, portion of the covariances is attributable to measurement method.

In Model 2, we added the specification of seven single attitude factors, one for each of the attitude object pairs, each defined by six indicators (four measured by IAT and two by self-report difference scores), and specified that these factors be uncorrelated with one another.⁷ The increased complexity of this model (42 additional parameters estimated) proved worthwhile in terms of improved fit; the RMSEA was .049 and this model was clearly superior to

5 The significant heterotrait correlations that emerged were fairly consistent across implicit and explicit measures, which implies a conceptual rather than methodological relationship among traits. The correlations pattern is consistent in direction with conservatism being associated with negativity toward outgroups (Jost, Banaji, & Nosek, 2004), and ethnocentrism (Cunningham, Nezlek, & Banaji, 2004).

6 Because relatively few participants ($n = 86$) were measured on both of the explicit thin-fat indicators, the uniquenesses (residual variances) for these two indicators were constrained to be equal (see parameter estimates in the *Mplus* output files available in the supplement to this paper at <http://briannosek.com/>). This was done in order to facilitate estimation of their parameters in the most complex models (3–5). The same pattern of comparative model results was observed when we eliminated all thin-fat indicators, implicit as well as explicit, from the models.

7 Though correlations of zero between these attitude factors is not a substantive part of the hypotheses of interest, we constrained them to zero so as not to spuriously improve the closeness of model fit. Even so, when we allowed the correlations between the attitude factors to be freely estimated (as we did in alternative Models 2a and 3a, specifications and estimates for which can be found in the supplement), the changes in fit from Models 2 and 3, respectively, were nonsignificant. We also tested an alternative model (2b) in which the correlation between explicit and implicit method factors was allowed, but the correlation was nonsignificant.

Model 1 ($\Delta\chi^2/\Delta df = 1852/42 = 44$; 95% CI $\epsilon_a\Delta = .37-.41$). The good fit of Model 2 is evidence for both the convergent validity of the IAT with the explicit measures and the discriminant validity of the attitude object pairs across the measures – i.e., the utility of representing the attitude domains as distinct constructs.

Our dual-construct hypothesis is represented by Model 3. Comparing this model's fit with that of Model 2 tests whether specifying distinct implicit and explicit attitude factors is superior to a single-attitude factor per domain model for these data. Model 3 fit with $\epsilon_a = .029$, compared with $.049$ for Model 2. The change in fit meets our criterion for constituting a significant improvement ($\Delta\chi^2/\Delta df = 364/7 = 52$; 95% CI $\epsilon_a\Delta = .38-.47$). A diagram representing this model, annotated with the estimated implicit–explicit attitude factor correlations, is shown in Figure 2. Five of the seven implicit–explicit latent attitude factor correlations were significantly positive, indicating that treating the factors as orthogonal is unjustified. Yet the comparison of this model's fit with that of Model 2 indicates that ignoring the distinctiveness of implicit and explicit attitudes, i.e., collapsing these indicators onto a single attitude factor, yields a relatively inferior model. Having partitioned the variance common to measurement method, this analysis supports a view that related but distinct implicit and explicit attitude constructs have been measured in these domains.

To test the importance of accounting for each of the two kinds of method variation, in Model 4 we eliminated the specification of an implicit method factor, while in Model 5 we eliminated the explicit method factor specification. When the common implicit method variance was no longer modeled (Model 4), the RMSEA was $.031$, only slightly different from that of Model 3 ($\epsilon_a = .029$). The indices of change in fit relative to Model 3 corroborate this small difference ($\Delta\chi^2/\Delta df = 64/28 = 2.3$; 95% CI $\epsilon_a\Delta = .07-.13$). The 95% confidence interval for the RMSEA of the change did not overlap the $.05$ benchmark for close fit, but it is clear that accounting for the implicit method variance is of little consequence in representing the relations between these variables. A similarly modest difference in fit resulted when explicit method variance was dropped (Model 5). Model 5's RMSEA was $.032$, and the change indices, though again exceeding our criterion for judging a significantly different fit, were of relatively modest magnitude ($\Delta\chi^2/\Delta df = 53/14 = 3.8$; 95% CI $\epsilon_a\Delta = .07-.13$). To summarize, there was relatively little common method variance to account for in these data; statistically significant, but small amounts were found for both the explicit and implicit measurements. This suggests that Model 3 would be a useful specification against which to judge models in future studies, but that little may be lost if method variance is not partitioned. Still, other research has demonstrated method influences on the IAT (e.g., Mierke & Klauer, 2003; Nosek et al., 2006) suggesting that a more conservative approach for future research would be to work toward specifications like that of Model 3 when possible.

Follow-Up Analyses

Absolute Values

We conducted additional analyses to evaluate the possibility that method variance is underestimated in these models because both the IAT and the explicit measures have rational zero points; both indicate a preference for one category (Democrats) compared to another (Republicans). Some factors may primarily influence the extremity of the score away from neutrality (0). For example, people who are more skilled at task-switching will achieve less extreme scores regardless of whether, for example, they are pro-Democrat or pro-Republican. To determine whether influences of this sort contributed significantly to the pattern of results we have reported, we used the absolute values of all scores in a reanalysis. This provides a liberal test for the method factor influences because it reduces the construct-variant variance by treating positive and negative score values as the same, and enhances the opportunity to see influences of extremity (distance from 0) as indicating common influence on the implicit or explicit measures.

Refitting the sequence of models summarized in Table 3 with the absolute values for each indicator yielded the same pattern of results (a table in which these results are listed is part of the online supplement, <http://brian-nosek.com/>). That is, both implicit and explicit method factors still made small but statistically noticeable contributions to accounting for relations among the absolute values of the indicators, and specifying correlated implicit and explicit attitude factors was again clearly superior to a single attitude specification.

Highly Correlated Domains

To more rigorously test the generality of our two-attitude specification, we fit the sequence of models only to the data for the three domains with relatively strong implicit–explicit correlations: creationism–evolution, Democrats–Republicans, and science–humanities. The raw implicit–explicit correlations (Table 2) for these domains were $.56$, $.52$, and $.51$, respectively, while the latent factor correlations estimated in Model 3 (see Figure 2) were $.68$, $.61$, and $.67$. Thus, by using the most highly correlated domains, and again partitioning common method variance, we increased the likelihood that a single attitude specification would suffice to account for variable interrelations. However, the two-attitude specification was again superior (a complete table of results is available in the online supplement). The RMSEA for the one-attitude factor specification (i.e., Model 2) was $.052$, but was $.015$ for the two-attitude specification (Model 3). The change in χ^2 between the two models was $\Delta\chi^2 = 86$ on $\Delta df = 3$, and the 95% CI around the RMSEA of this change was $.25-.38$.

Discussion

The results of this study add to construct validation evidence for the IAT as a measure of attitudes (Greenwald & Nosek, 2001). We applied the classic construct validation principles articulated by Campbell and Fiske (1959) through a MTMM structural modeling analysis of seven attitude object pairs measured by IAT and self-report. Because participants were measured across multiple trait (i.e., attitude) domains, we were able to model the variance that was common to a measurement technique. With common method factors thus specified (one implicit and one explicit), our subsequent models, first, a single-construct structure per domain, then a dual-construct alternative, could be more confidently interpreted as tests of competing hypotheses about attitudes. In other words, the leverage of the MTMM design allowed for separating attitudes from the peculiarities of the instruments necessary to measure them, a limitation common to previous structural modeling research that involved only one domain per analysis (e.g., Cunningham et al., 2001; Greenwald & Farnham, 2000).

Consistent with the results of the earlier single-domain research, but now with method variance partitioned, we found that specifying distinct but related attitude constructs was superior to a single attitude construct formulation. The convergent validity of the IAT was evidenced by significant factor correlations between the implicit and explicit attitude constructs in five of the seven attitude domains, while its discriminant validity was simultaneously evidenced by the statistical superiority of the two-attitude model to the single-attitude model. This correlated two-attitude factor specification was superior to the single-attitude specification even when we limited analyses to the three attitude domains in which the implicit–explicit correlations were highest. There was a modest amount of common method variance to account for, with statistically significant but small portions isolated from both the explicit and implicit measures, and this was also observed when absolute values of the indicators were used. This relatively small systematic method variance associated with the IAT bolsters conclusions that the *D* scoring algorithm, used here, reduces the impact of extraneous method variance compared to alternative algorithms (Greenwald et al., 2003; Mierke & Klauer, 2003). In sum, by isolating variance common to measurement technique, our findings allow for a stronger inference that implicit and explicit attitudes are distinct, though often related constructs.

The MTMM approach seems especially important to validation of implicit attitude constructs since identifying them relies on specially designed measurement techniques. As Campbell and Fiske (1959, p. 84) observed, “In any given psychological measuring device, there are certain features or stimuli introduced specifically to represent the trait that it is intended to measure. There are other features which are characteristic of the method being employed, features which could also be present in efforts to measure

other quite different traits.” The MTMM design, coupled with comparative structural modeling, allowed common implicit method variance to be distinguished from implicit attitude variance.

Other Components of the Nomological Net for the Implicit Attitude Construct

This research provides a basis for some key components of validation of the IAT and implicit attitudes. Yet the MTMM approach is not a panacea. It does not, for example, identify the processes that differentiate the implicit and explicit attitude constructs. Nor does it clarify whether the identified constructs, derived from measurement, correspond with their definitions stated at the beginning, derived from theory. There are a number of additional elements of the nomological net that will enhance understanding of the implicit attitude construct.

First, researchers in many laboratories are working to develop process models of the IAT (e.g., Conrey, Sherman, Gawronski, Hugenberg, & Groom, 2005; Mierke & Klauer, 2003; Rothenmund & Wentura, 2004). Also, our finding that common method variance did not account for the implicit–explicit distinction should not be mistaken for implying that the methodological distinctions between implicit and explicit measures are all relevant for distinguishing between implicit and explicit attitude constructs. For example, the IAT requires participants to categorize individual stimulus items into superordinate categories, whereas participants’ self-reports concern evaluations of the category labels only. If the IAT reflected evaluations of the stimulus items and not the categories, then the implicit–explicit distinction might be explained by this difference. This example is not likely to account for the differences because the IAT effect is largely category driven, with the exemplars serving to affect the construal of the category (DeHouwer, 2001; Nosek et al., 2004; Mitchell, Nosek, & Banaji, 2003; Nosek, Greenwald, & Banaji, 2005). However, it illustrates the point that the distinction between constructs confirmed here does not explain why those constructs are distinct. A model of how the IAT gives rise to its effects will clarify interpretation of the measured construct, such as which dimensions of automaticity are engaged in the assessment, and whether the identified implicit attitude construct conforms to the proposed theoretical definitions (see, for example, Conrey et al., 2005).

Second, the relationship strength between implicit and explicit attitudes varies as a function of features of the attitude objects (Nosek, 2005; and in our reanalysis of that data in Table 1). Weak implicit–explicit correspondence was found for some attitude object pairs (e.g., hot–cold), while strong correspondence was found for other pairs (e.g., Democrats–Republicans). Understanding the relationship between implicit and explicit attitudes will foster theoretical developments concerning the structure and

function of each. Nosek (2005) found evidence for four moderators of implicit and explicit attitude relations: self-presentation, attitude strength, attitude dimensionality, and attitude distinctiveness (see also Hofmann, Gschwendner, Nosek, & Schmitt, 2005, for a review).

Third, predictive validity is important for any method and construct, and attitudes are presumed to guide perception, judgment, and action. Since the IAT was first described (Greenwald et al., 1998), its predictive utility has been demonstrated in a variety of domains. A meta-analysis of 86 studies (Poehlman et al., 2004) corroborates the IAT's predictive validity, and also reinforces the dual-construct interpretation, with explicit measures showing better predictive validity in some domains (e.g., consumer attitudes), and the IAT showing better predictive validity in others (e.g., stereotyping and prejudice).

Fourth, our MTMM analysis involved a small subset of domains that are of interest to behavioral scientists. This can lead to questions of the generality of the conclusion that implicit and explicit attitudes are related, but distinct constructs. The fact that positive and significant implicit–explicit relations are observed is an existence proof that these attitude measures are related, at least under some conditions. On the other end of the spectrum, even though we examined a wide range of attitude domains in the preliminary analysis (Table 1) and used some domains that showed strong positive relations in the study, it is possible that some as yet unexamined domains could suggest that a single attitude factor structure is sufficient for that domain. Would this threaten the generality of our MTMM conclusions? The emerging nomological net described above suggests not. The evidence suggests that implicit and explicit attitudes are related, but distinct (this article), that the variation in implicit–explicit correlation between objects can be explained by multiple features of evaluation such as attitude strength or perceived distinctiveness (Nosek, 2005), and that both implicit and explicit attitudes have predictive validity (Poehlmann et al., 2004). So, even if we find a subset of domains that effectively form a single attitude factor those domains will stand in contrast to the many domains that do not. The theoretical questions, then, would still make reference to dual constructs, e.g., what causes implicit and explicit attitudes to be indistinguishable?

Fifth, the IAT is one example of measurement methods that are referred to as “implicit measures.” However, some of these measures are only weakly related (Bosson, Swann, & Pennebaker, 2000), and no measure is a process-pure assessment of a construct (e.g., Conrey et al., 2005). MTMM construct validation methods will be useful for clarifying the relations and identities of the variety of implicit methods, and for testing the utility of the two-construct view compared to alternatives.

A natural extension of this study would be to add another dimension of variation – type of implicit and explicit measure – beyond the single example of each used here (the IAT and semantic differentials). While collecting data with multiple implicit and explicit measures for multiple attitude

domains can be quite laborious (for both participant and researcher), such efforts may yield dividends in clarifying whether the measures that are collectively referred to as “implicit” are reasonably interpreted as assessing a single construct. Also, additional methodological innovations of multiple sessions and a planned missing data design in which a subset of measures is administered to any given participant may facilitate a more comprehensive investigation of the commonalities and distinctions among the growing variety of implicit and explicit measurement methods.

Finally, dual-process models of evaluation demonstrate a wide variety of perspectives on the interaction between processes or representations (Chaiken & Trope, 1999; Gilbert, 1999). Some models suggest that the distinction between implicit and explicit attitude measures reflects when in time they assess an evaluation along a single processing dimension, i.e., that explicit evaluations are “farther ‘downstream’ than automatically activated attitudes” (Fazio & Olson, 2003, p. 305). From this perspective, implicit and explicit attitudes are of the same evaluative “stuff” with the explicit attitude measures reflecting the evaluation plus any alterations caused by motivation and opportunity. Other models suggest that the implicit and explicit distinction is truly dual-process in that implicit and explicit evaluative processes or representations are distinct and both can influence perception, judgment, and action simultaneously, interactively, or in turn (e.g., Strack & Deutsch, 2004; Wilson et al., 2000).

Our MTMM results cannot be interpreted as contradicting the single processing stream metaphor or as supporting a dual-process or dual-representation model. We examined the relationship between implicit and explicit attitudes from an individual differences perspective, without directly addressing the cognitive processes or representations giving rise to the evaluations. It is not contradictory to propose distinct constructs for content that is produced by operations on a unitary mental representation. Constructs are a function of both representation and process, and the same information can be specified in multiple forms, without inconsistency. For example, to a chemist H_2O is H_2O , and the snowboarder's insistence that slush, snow, and ice are importantly different can elicit the nerdy retort “it's still all water.” The problem, of course, is that the chemist and the snowboarder are talking past each other. The chemist is concerned with the common core property regardless of how operations (heating, condensation) affect its expression. The snowboarder is concerned with the constructs that result from the operations: snow, slush, and ice and what they will predict for a day on the slopes.

This example also illustrates a hazard of reductionism: The effort required for the chemist to specify the snowboarder's constructs in terms of a single H_2O construct would produce ungainly process theories that would likely “miss the point” of the snow and ice constructs. In short, the present evidence for distinct implicit and explicit attitude constructs does not rule out the possibility that the two constructs derive from common evaluative content. Nor

does it rule in this possibility. Implicit and explicit attitudes could be like snow and ice (same base content, different forms following operations), they could be like oil and vinegar (a separable and mixable admixture of different content), they could be like flour/eggs/milk and a baked cake (same base content but fundamentally transformed through processing), or any of an infinite variety of alternative conceptions.

The most useful metaphor for the interaction of implicit and explicit evaluations is still a matter of debate (Gilbert, 1999). Our findings support a view that, whatever their source, implicit and explicit attitudes have substantive, independent, properties and that neither of two extreme conceptions is a useful description of the data: (1) Implicit and explicit attitude measures assess exclusive constructs that have nothing in common, and (2) implicit and explicit attitude measures assess the same thing, varying only as a function of common sources of methodological variation. Future research will clarify the moderators of their relationship (Nosek, 2005), illuminate the nature of the interaction between automatic and controlled processes (Chaiken & Trope, 1999; Strack & Deutsch, 2004), and identify better metaphors for understanding the nature of evaluation.

Conclusion

Convergent evidence across a variety of research programs suggests that the IAT is a valid measure of attitudes (see Nosek et al., 2006, for a review). Like other methods such as semantic differentials, Likert scales, sequential priming, and the Stroop task, the IAT can be adapted to measure evaluations of many types of social categories. The cumulative evidence identifies design factors that will influence the method's validity, and provides a nomological net of knowledge to accelerate validation of novel applications of the IAT.

The emergence of the implicit attitude construct has spurred investigations to test the strength and limitations of this concept and its measurement tools, like the IAT. We sought to strengthen its developing nomological net by harnessing the power of a MTMM design. We found simultaneous evidence of convergent and discriminant validity of the IAT and self-report as measures of related but distinct attitude constructs, and as distinct from methodological variation.

Acknowledgments

This research was supported by the National Institute of Mental Health (MH-R01 MH68447-01). We thank Aki Hamagami, Jack McArdle, John Nesselroade, Mahazarin Banaji, and Tony Greenwald for comments. Supplementary materials are available at <http://briannosek.com/>.

References

- Bosson, J.K., Swann, W.B., & Pennebaker, J.W. (2000). Stalking the perfect measure of self-esteem: The blind men and the elephant revisited? *Journal of Personality and Social Psychology*, *79*, 631–643.
- Browne, M.W. (1991). [FITMOD: A computer program for calculating point and interval estimates of fit measures.] Unpublished manuscript.
- Browne, M.W., & Cudeck, R. (1993). Alternative ways of assessing model fit. In K.A. Bollen & J.S. Long (Eds.), *Testing structural equation models* (pp. 136–162). Newbury Park, CA: Sage.
- Campbell, D.T., & Fiske, D.W. (1959). Convergent and discriminant validation by the multitrait-multimethod matrix. *Psychological Bulletin*, *56*, 81–105.
- Chaiken, S., & Trope, Y. (1999). *Dual-process theories in social psychology*. New York: Guilford.
- Conroy, F.R., Sherman, J.W., Gawronski, B., Hugenberg, K., & Groom, C. (2005). Beyond automaticity and control: The quad-model of behavioral response. *Journal of Personality and Social Psychology*, *89*, 469–487.
- Cronbach, L.J., & Meehl, P.E. (1955). Construct validity in psychological tests. *Psychological Bulletin*, *52*, 281–302.
- Cunningham, W.A., Johnson, M.K., Gatenby, J.C., Gore, J.C., & Banaji, M.R. (2003). Neural components of social evaluation. *Journal of Personality and Social Psychology*, *85*, 639–649.
- Cunningham, W.A., Johnson, M.K., Raye, C.L., Gatenby, J.C., Gore, J.C., & Banaji, M.R. (2004). Separable neural components in the processing of Black and White faces. *Psychological Science*, *15*, 806–813.
- Cunningham, W.A., Nezlek, J.B., & Banaji, M.R. (2004). Implicit and explicit ethnocentrism: Revisiting the ideologies of prejudice. *Personality and Social Psychology Bulletin*, *30*, 1332–1346.
- Cunningham, W.A., Preacher, K.J., & Banaji, M.R. (2001). Implicit attitude measures: Consistency, stability, and convergent validity. *Psychological Science*, *12*, 163–170.
- De Houwer, J. (2001). A structural and process analysis of the Implicit Association Test. *Journal of Experimental Social Psychology*, *37*, 443–451.
- Enders, C.K., & Bandalos, D.L. (2001). The relative performance of full information maximum likelihood estimation for missing data in structural equation models. *Structural Equation Modeling: A Multidisciplinary Journal*, *8*, 430–457.
- Fazio, R.H., & Olson, M.A. (2003). Implicit measures in social cognition research: Their meaning and use. *Annual Review of Psychology*, *54*, 297–327.
- Fazio, R.H., & Towles-Schwen, T. (1999). The MODE model of attitude-behavior processes. In S. Chaiken & Y. Trope (Eds.), *Dual process theories in social psychology* (pp. 97–116). New York: Guilford.
- Fazio, R.H., Sanbonmatsu, D.M., Powell, M.C., & Kardes, F.R. (1986). On the automatic activation of attitudes. *Journal of Personality and Social Psychology*, *50*, 229–238.
- Gilbert, D.T. (1999). What the mind's not. In S. Chaiken & Y. Trope (Eds.), *Dual process theories in social psychology* (pp. 3–11). New York: Guilford.
- Greenwald, A.G., & Banaji, M.R. (1995). Implicit social cognition: Attitudes, self-esteem, and stereotypes. *Psychological Review*, *102*(1), 4–27.

- Greenwald, A.G., & Farnham, S.D. (2000). Using the Implicit Association Test to measure self-esteem and self-concept. *Journal of Personality and Social Psychology*, 79, 1022–1038.
- 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. (2001). Health of the Implicit Association Test at age 3. *Zeitschrift für Experimentelle Psychologie*, 48, 85–93.
- 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, 197–216.
- Greenwald, A.G., & Nosek, B.A. (in press). Attitudinal dissociation: What does it mean? In R.E. Petty, R.H. Fazio, & P. Briñol (Eds.), *Attitudes: Insights from the new implicit measures*. Hillsdale, NJ: Erlbaum.
- Hofmann, W., Gschwendner, T., Nosek, B.A., & Schmitt, M. (2005). What moderates implicit–explicit consistency? *European Review of Social Psychology*, 16, 335–390.
- Jöreskog, K.G. (1974). Analyzing psychological data by structural analysis of covariance matrices. In R.C. Atkinson, D.H. Krantz, R.D. Luce, & P. Suppes (Eds.), *Contemporary developments in mathematical psychology: Measurement, psychophysics, and neural information processing* (Vol. 2, pp. 1–56). San Francisco: Freeman.
- Jöreskog, K.G., & Sörbom, D. (1979). In J. Magidson (Ed.), *Advances in factor analysis and structural equation models*. Cambridge, MA: Abt Books.
- Jost, J.T., Banaji, M.R., & Nosek, B.A. (2004). A decade of system justification theory: Accumulated evidence for the unconscious bolstering of the status quo. *Political Psychology*, 25, 881–919.
- Karpinski, A., & Hilton, J.L. (2001). Attitudes and the Implicit Association Test. *Journal of Personality and Social Psychology*, 81, 774–788.
- Loehlin, J.C. (2004). *Latent variable models: An introduction to factor, path, and structural analysis*. Hillsdale, NJ: Erlbaum.
- MacCallum, R.C., Browne, M.W., & Sugawara, H.M. (1996). Power analysis and determination of sample size for covariance structure modeling. *Psychological Methods*, 1, 130–149.
- Marsh, H.W., & Grayson, D. (1995). Latent variable models of multitrait-multimethod data. In R.H. Hoyle (Ed.), *Structural equation modeling: Concepts, issues and applications* (pp. 177–198). Thousand Oaks, CA: Sage.
- McArdle, J.J. (1994). Structural factor analysis experiments with incomplete data. *Multivariate Behavioral Research*, 29, 409–454.
- McArdle, J.J., & Prescott, C.A. (1992). Age-based construct validation using structural equation modeling. *Experimental Aging Research*, 18, 87–115.
- McDonald, R.P. (1985). *Factor analysis and related methods*. Hillsdale, NJ: Erlbaum.
- Mierke, J., & Klauer, K.C. (2003). Method-specific variance in the Implicit Association Test. *Journal of Personality and Social Psychology*, 85, 1180–1192.
- Mitchell, J.P., Nosek, B.A., & Banaji, M.R. (2003). Contextual variations in implicit evaluation. *Journal of Experimental Psychology: General*, 132(3), 455–469.
- Muthén, L.K., & Muthén, B.O. (1998–2004). *Mplus user's guide* (3rd ed.). Los Angeles: Muthén & Muthén.
- Nosek, B.A. (2005). Moderators of the relationship between implicit and explicit evaluation. *Journal of Experimental Psychology: General*, 135, 565–584.
- Nosek, B.A., & Banaji, M.R. (2001). The go/no-go association task. *Social Cognition*, 19, 625–666.
- Nosek, B.A., Greenwald, A.G., & Banaji, M.R. (2005). Understanding and using the Implicit Association Test: II. Method variables and construct validity. *Personality and Social Psychology Bulletin*, 31(2), 166–180.
- Nosek, B.A., Greenwald, A.G., & Banaji, M.R. (2006). The implicit association test at age 7: A methodological and conceptual review. In J.A. Bargh (Ed.), *Automatic Processes in Social Thinking and Behavior* (pp. 265–292). New York: Psychology Press.
- Poehlman, T.A., Uhlmann, E., Greenwald, A.G., & Banaji, M.R. (2004). *Understanding and using the Implicit Association Test: III. Meta-analysis of predictive validity*. Unpublished manuscript, Yale University, New Haven, CT.
- Project Implicit. (2002). *A virtual laboratory for research in implicit social cognition. Version 1.5*. <http://implicit.harvard.edu/>
- Rothermund, K., & Wentura, D. (2004). Underlying processes in the Implicit Association Test (IAT): Dissociating salience from associations. *Journal of Experimental Psychology: General*, 133, 139–165.
- Steiger, J.H. (2000). Point estimation, hypothesis testing, and interval estimation using the RMSEA: Some comments and a reply to Hayduk and Glaser. *Structural Equation Modeling*, 7, 149–162.
- Steiger, J.H., & Lind, J.C. (1980, May). *Statistically based tests for the number of factors*. Paper presented at the annual spring meeting of the Psychometric Society, Iowa City, IA.
- Strack, F., & Deutsch, R. (2004). Reflective and impulsive determinants of social behavior. *Personality and Social Psychology Review*, 3, 220–247.
- Widaman, K.F. (1985). Hierarchically nested covariance structure models for multitrait-multimethod data. *Applied Psychological Measurement*, 9, 1–26.
- Wilson, T.D., Lindsey, S., & Schooler, T.Y. (2000). A model of dual attitudes. *Psychological Review*, 107(1), 101–126.
- Wittenbrink, B., Judd, C.M., & Park, B. (1997). Evidence for racial prejudice at the implicit level and its relationship with questionnaire measures. *Journal of Personality and Social Psychology*, 72, 262–274.

Received October 12, 2005

Revision received January 19, 2006

Accepted January 20, 2006

Brian Nosek

Department of Psychology
Box 400400
102 Gilmer Hall
University of Virginia
Charlottesville, VA 22904-4400
USA
Tel. +1 434 924-0666
Fax +1 434 982-4766
E-mail nosek@virginia.edu

Appendix

IAT Category Labels and Exemplars

<i>Flowers</i> Daffodil Daisy Rose Sunflower	<i>Insects</i> Bugs Caterpillar Cricket Fly	<i>Democrats</i> Al Gore Democrat Liberal Left-wing Lieberman	<i>Republicans</i> George Bush Cheney Conservative Republican Right-wing
<i>Gay people</i> Gay Homosexual	<i>Straight People</i> Straight Heterosexual	<i>African-American</i> 6 face photos	half female
<i>Fat people</i> Chubby Fat Large Obese Overweight	<i>Thin people</i> Skinny Slender Slim Thin Underweight	<i>Creation</i> Adam and Eve Bible Creator God Religion Six Days	<i>Evolution</i> Darwin Eons Evolved Natural Selection Origin of Species Science
<i>Humanities</i> Arts English History Humanities Latin Music Philosophy Spanish	<i>Science</i> Astronomy Biochemistry Biology Chemistry Engineering Neuroscience Physics Science		

Additional IAT Procedure and Cleaning Details

Error Responses

When participants made errant responses, they were required to correct them by pressing the other response key; the latency for such trials was calculated from initial stimulus presentation until the corrected response.

Extreme Latencies

Beyond the disqualification of IAT scores when more than 10% of critical responses were faster than 300 ms, individual trial latencies faster than 400 ms or slower than 10,000 ms were not included in score calculation.

Structural Model Parameter Estimates

All models were run with Mplus software (Muthén & Muthén, 1998–2004) and results are available in a supplement that may be downloaded from <http://briannosek.com/>. For Models 1–5 listed in Table 3, implicit and explicit method factors were identified by fixing the factor variance to 1.0, and the covariance between the method factors was fixed at zero. Attitude factor covariances across attitude domains were also fixed to zero. Residual variances of the two explicit thin–fat indicators were constrained to be equal.

Patterns of Measured (x) Data

	Variable															
	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16
FIIATa	x	x	x	x	x	x	x	x	x	x	x	x	x			
FIIATb	x	x	x	x	x	x	x	x	x	x	x	x	x			
FIIATc	x	x	x	x	x	x	x	x	x	x	x	x	x	x		
FIIATd	x	x	x	x	x	x	x	x	x	x	x	x	x	x		
FIEXP1	x				x										x	
FIEXP2	x				x										x	
DRIATa	x	x	x	x	x	x	x	x	x	x	x			x		x
DRIATb	x	x	x	x	x	x	x	x	x	x	x			x		x
DRIATc	x	x	x	x	x	x	x	x	x	x	x			x		x
DRIATd	x	x	x	x	x	x	x	x	x	x	x			x		x
DREXP1	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x
DREXP2	x	x	x	x	x	x	x	x	x			x		x		x
SGIATa	x	x	x	x						x		x		x		x
SGIATb	x	x	x	x						x		x		x		x
SGIATc	x	x	x	x						x		x		x		x
SGIATd	x	x	x	x						x		x		x		x
SGEXP1	x	x	x	x	x	x				x	x	x	x	x	x	
SGEXP2	x	x	x	x	x	x						x		x		
HSIATa	x	x	x	x	x	x	x	x	x					x		
HSIATb	x	x	x	x	x	x	x	x						x		
HSIATc	x	x	x	x	x	x	x	x							x	
HSIATd	x	x	x	x	x	x	x	x						x		
HSEXP1	x	x	x	x	x	x	x	x	x			x		x		x
HSEXP2	x	x		x	x	x	x	x	x			x		x		x
WBIATa	x	x	x		x	x						x		x		
WBIATb	x	x	x		x	x						x		x		
WBIATc	x	x	x		x	x						x		x		
WBIATd	x	x	x		x	x						x		x		
WBEXP1	x	x	x	x	x	x						x		x		
WBEXP2	x	x	x	x	x	x						x		x		
CEIATa	x				x		x	x	x					x		
CEIATb	x				x		x	x	x					x		
CEIATc	x				x		x					x		x		
CEIATd	x				x		x					x		x		
CEEXP1	x				x		x	x	x					x		x
CEEXP2	x				x		x	x	x					x		x
TFIATa	x				x					x	x		x	x	x	
TFIATb	x				x					x	x		x	x	x	
TFIATc	x				x					x	x		x	x	x	
TFIATd	x				x					x	x		x	x	x	
TFEXP1	x				x					x	x		x	x	x	
TFEXP2	x				x									x		
<i>n</i>	84	76	2	1	1	1	54	2	1	56	1	1	2	1	1	3

Note: FI = flowers–insects, DR = Democrats–Republicans, HS = humanities–science, SG = straight–gay, WB = White–Black, CE = creation–evolution, TF = thin–fat.