

Draft of chapter prepared for: Petty, R. E., Fazio, R. H., & Briñol, P. (Eds.), *Attitudes: Insights from the New Implicit Measures*. Hillsdale, NJ: Erlbaum.

Attitudinal Dissociation: What Does it Mean?

Anthony G. Greenwald and Brian A. Nosek

Abstract

Many recent experiments have used parallel Implicit Association Test (IAT) and self-report measures of attitudes. These measures are sometimes strongly correlated. However, many of these studies find apparent dissociations in the form of (a) weak correlations between the two types of measures, (b) separation of their means on scales that should coincide if they assess the same construct, or (c) differing correlations with other variables. Interpretations of these empirical patterns are of three types: *single-representation* — the two types of measures assess a single attitude, but under the influence of different extra-attitudinal *process* influences; *dual-representation* — the two types of measures assess distinct forms of attitudes (e.g., conscious vs. unconscious; implicit vs. explicit); and *person vs. culture* — a variant of the dual-representation view in which self-report measures reflect personal attitudes, whereas IAT measures reflect non-attitudinal cultural or semantic knowledge. Proponents sometimes interpret evidence for single versus dual *constructs* as evidence for single versus dual structural *representations*. Behavioral evidence can establish the discriminant validity of implicit and explicit attitude phenomena (dual constructs), but cannot choose among single- vs. dual-representation interpretations because the distinct constructs remain susceptible to interpretation in terms of either one or two representations. Selecting among representational accounts must therefore be based on considerations of explanatory power or parsimony.

Attitudinal Dissociation: What Does it Mean?

A byproduct of increasing recent attention to implicit measures of attitudes is the controversial hypothesis of dissociated attitude representations. As applied to attitudes, this reference to *dissociation* implies the existence of distinct structural representations underlying distinguishable classes of attitude manifestations. In psychology, appeals to dissociation range from the mundane to the exotic. At the mundane end, the dissociation label may be attached to the simple absence or weakness of correlation between presumably related measures. At the exotic end, dissociation may be understood as a split in consciousness, such as mutually unaware person systems occupying the same brain. While recognizing this breadth of uses, we focus in this chapter on the specific usage in which “dissociation” refers to structurally separate and presumably independently functioning mental representations within the same brain. We shall keep this focus in sight by frequently referring to “structural dissociation”.

Empirical Data Patterns and Dissociation

Consider a research finding that might be observed in a person whose cerebral hemispheres have been surgically separated to control epileptic seizures. This hypothetical subject is asked to view words and then attempt immediately to recognize each word by pointing to it in a list containing additional distracter words. If the to-be-identified word is briefly flashed to the left of a visual fixation point (and is therefore transmitted by optic nerves to the right cerebral hemisphere) performance will be excellent if the left hand (under control of the right hemisphere) does the pointing, but will be at chance if the right hand does the pointing. The reverse pattern (excellent with right hand, but chance with left hand) will result for words flashed to the right half of the visual field. This result illustrates *double dissociation*, a pattern of directionally opposite effects of an independent variable under two levels of a second independent variable. Double dissociation data patterns are often taken to justify a conclusion that structurally separate mental systems are involved in the performances. In this case, the separate

systems would be ones operating independently within the left and right cerebral hemispheres.¹

More ordinary (i.e., single) dissociation data patterns also take the form of a statistical interaction effect lacking the juxtaposed opposite-direction effects that identify double dissociation. Two measures show an *empirical dissociation* pattern when they respond differently to procedural variations and/or when they have different observed relationships to other measured variables. At the level of data (rather than theory), dissociation corresponds approximately to the notion of *discriminant validity*. Discriminant validity refers to the distinctness of *empirical constructs* (Cronbach & Meehl, 1955), whereas structural dissociation refers to distinctness of hypothesized *mental representations*. In the split-brain illustration of double dissociation one can describe the left-hand and right-hand response measures not only as having discriminant validity as measures, but also as corresponding to structurally distinct (dissociated) right and left hemisphere operations.

Evidence for Implicit–Explicit Dissociation

Figure 1 shows the regression of an IAT measure of implicit age attitude onto a parallel explicit (self-report) measure. For both measures a score of zero is interpreted as indicating attitudinal indifference between the concepts *young* and *old*. For the IAT measure, the zero score indicates that the respondent is equally fast at classifying young-appearing faces together with pleasant-meaning words and old-appearing faces together with pleasant-meaning words. The data set is one for which methods and sample were described by Greenwald, Nosek, and Banaji (2003), and the IAT measure is Greenwald et al.'s *D* measure. The self-report measure was based on three items

¹ A conclusion of structurally distinct systems does not require that the distinction be identified as one between *conscious* vs. *unconscious* systems. The left hemisphere of the split brain subject may not know what the right hemisphere is doing, but this does not mean either that one hemisphere has an unconscious representation of the other's conscious knowledge or that one hemisphere is operating consciously and the other unconsciously. Implicit and explicit attitude measures may likewise show double dissociations (e.g., Perugini, 2005), which, likewise, does not oblige a conclusion that one attitudinal system is conscious and the other is not.

(Greenwald et al., p. 216). The first of the three items used a 5-point Likert format in which the middle alternative (scored zero) was “I like young people and old people equally”; the other two items were used 11-point thermometer rating scales for the concepts young and old, combined subtractively into a difference score. The explicit measure in Figure 1 averaged the Likert and difference scores, with each measure divided by its standard deviation (i.e., preserving the zero-point locations) before averaging.

Although a positive relationship between the two measures (a positive regression slope) is visible in Figure 1, it is a decidedly weak positive relationship, corresponding to a correlation of $r = .16$. A correlation this weak is sometimes taken to indicate implicit–explicit dissociation. A second possible indicator of dissociation in Figure 1 is that the regression function deviates substantially from passing through the origin. Alternately described, there was a substantial difference in means for the two measures. In standard deviation units on scales for which zero indicates evaluative indifference between young and old and positive scores indicate preference for young, the mean of the explicit measure was 0.39, whereas the mean of the implicit measure was 1.35. This is nearly a full standard deviation difference, with the implicit measure showing substantially greater relative positivity for young than the explicit measure, $t(10254) = 75.5$ (a value of t that leaves p too small to be computed by standard statistical software).

Figure 2 shows still a third possible indicator of dissociation, in the form of the finding that a demographic variable, chronological age, has a well defined relation with the explicit age attitude measure ($r = -.194$, $N = 10,266$, $p = 10^{-87}$), but no relation with the implicit age attitude measure ($r = -.012$, $N = 10,266$, $p = .23$). The data in Figure 2 can also be described as showing an interaction effect of age and the implicit–explicit attitude variation, $t(10,188) = 14.11$, $p = 10^{-44}$.

A weakness of the evidence for dissociation in Figure 2 is the lack of any sure indication that the explicit measure’s relation to age has something to do with attitudes. Perhaps older subjects, who may be more conservative than young subjects, are reluctant to use responses at the end points of self-report measures. An age difference

in response style could therefore explain Figure 2's data pattern without concluding that there is less explicit favorableness toward the young with increasing age. The explanation just offered is perhaps implausible because (a) the explicit measures are not extreme even for younger subjects and (b) subjects of greater age may have a good reason (approaching old age) for having genuinely increased explicit favorableness toward the concept *old*. For these reasons, even the small correlation between the two measures shown in Figure 1 suggests that the two measures have something in common.

The evidence for attitudinal dissociation would be stronger if Figure 2 showed not just a lack of relation between age and implicit attitude, but a relationship opposite in direction to that found for age and explicit attitude — i.e., a double dissociation. An opposite-direction relation could not readily be dismissed by supposing that it could be due to the implicit measure being a poor measure. Even with such a (hypothetical) double-dissociation pattern, however, it might be assumed that the opposite direction relationship with the implicit measure was due to some non-attitudinal process associated with age that affects the IAT measure.²

In summary, Figures 1 and 2 provide an implicit–explicit data pattern that includes three components: (a) low intercorrelation, (b) separation of means, and (c) different relationships to a third variable. How compelling is this collection of patterns as evidence for two structurally distinct attitude representations? The collection of three patterns is certainly more compelling than is the low correlation by itself but, nevertheless, it is less than fully compelling. As we shall now show, one also needs some assurance that the data for both measures are relevant to attitudes.

Discriminant and Convergent Validity

²The original scoring procedure for the IAT might well have contained such an undesired effect of age, due to the characteristic slower responding of elderly subjects. Slower responding on RT measures tends to produce artifactually large differences in RTs between experimental conditions. However, introduction of the *D* measure sharply reduced that obviously non-attitudinal influence on IAT measures (cf. Greenwald et al., 2003).

The foregoing hopefully establishes that considerations of construct validity are essential in interpreting empirical data patterns. To justify interpretation of empirically distinct “implicit” and “explicit” attitude constructs data must meet an unusual combination of two validity-related criteria. They must show both (a) *discriminant validity*, such as by having different patterns of relationship to other variables, thereby establishing that the two measures are not measures of identically the same construct, and (b) *convergent validity*, which establishes that the two measures also warrant interpretation as reflecting the same type of construct. This is an interesting paradox of dissociation — one must demonstrate that two measures assess the same *type* of construct while, simultaneously, demonstrating that they must represent different *forms* of that construct.

For the split-brain case that we have are treating as a prototype of structural dissociation, most observers will readily agree that both the discriminant and convergent empirical validity criteria are met. The data directly provide evidence for discriminant validity — the right-hand and left-hand recognition measures have opposite patterns of relation to the independent variable of left vs. right visual hemifield stimulus presentation. Also, the convergent validity criterion is satisfied intuitively, because the two measures are identical except for the right–left switch — there is no plausible alternative to viewing them both as measures of recognition memory.

What about the situation for implicit and explicit attitude measures? How can the discriminant and convergent validity criteria be met simultaneously? Demonstrating discriminant validity — which requires showing different patterns of relationship to other variables — is straightforward. Figure 2’s data illustrate this. Discriminant validity justifies the use of the distinct construct terminology, in this case “implicit” and “explicit,” though it does not establish difference in the process(es) and/or representation(s) that generate the data. More difficult is meeting the convergent validity criterion – i.e., what justifies a conclusion that the constructs legitimately share use of the term “attitude”? The weak positive correlation between implicit and explicit measures (shown in Figure 1) helps, but does not suffice both because of its weakness and because the correlation could be due to some shared non-attitudinal influence. Each measure must also correlate with other variables in a way that makes plausible that the measures are both

attitude measures. However, these correlations cannot be with the *same* other variable for each measure — if they were, then the discriminant validity requirement for dissociation would be undermined.

Not all data sets that include implicit and explicit measures show the dissociation-suggestive patterns of Figures 1 and 2. Figure 3 shows a regression of implicit on explicit attitude from a data set obtained with procedures very similar to those that obtained Figure 1's data, differing only in the attitude object. Unlike Figure 1, Figure 3 reveals a high implicit–explicit correlation ($r = .730$). Also unlike Figure 1, the difference between means of the implicit and explicit measures is very small — 0.04 SD units, quite unlike the 0.96 SD units for the data in Figure 1. The same data set of Figure 3 can be seen in Figure 4 to show patterns in which the implicit and explicit measures have virtually identical relations to another variable, education level. Quite clearly, the data in Figures 3 and 4 do not show even one of the three dissociation-suggestive patterns evident in Figures 1 and 2.

Three Interpretations

Interpretations that we here label *single-representation*, *dual-representation*, and *person vs. culture* have received greatest attention in discussions of published data that, like the prior examples, showing either relationship or lack of relationship between implicit and explicit attitude measures.

Single-representation interpretations treat all appearances of attitudinal dissociation as illusory. All attitude manifestations — implicit and explicit — are attributed to a single form of mental attitude representation. Appearances of dissociation such as weak correlation and differing relationships with other variables are interpreted in terms of processes that are assumed to be different in the implicit and explicit measurement situations. In the most fully developed analysis of the single-representation type, Fazio (this volume; 1990; Fazio & Olson, 2003) interprets explicit measures as subject to motivational and ability or opportunity influences that differ from the influences on implicit measures.

“[F]rom the perspective of the MODE model, [overt, explicit, expressions of attitude] are, for want of any better expression, farther “downstream” than automatically activated

attitudes [i.e., implicit measures]. Responding to an explicit measure is itself a verbal behavior that can be affected by motivation and opportunity, as well as whatever is automatically activated.” (Fazio & Olson, 2003, p. 305).

The second interpretation of empirical dissociation patterns identifies implicit and explicit measures of attitude with structurally distinct mental representations of attitudes. Several such two-attitude views have been offered (see Chaiken & Trope, 1999; Wilson, Lindsey, & Schooler, 2000). These views often characterize the representations underlying implicit measures as operating automatically and perhaps unconsciously, while treating representations underlying explicit measures as operating consciously and with deliberate thought (see also Strack & Deutsch, 2004).

The third interpretation conceives implicit and explicit measures as due to distinct categories of influences that are represented by the labels, *culture* and *person*. Implicit measures (and perhaps the IAT more than other implicit measures) are assumed to represent the influence of culture, while explicit measures capture influences operating within the person. This person vs. culture interpretation has sometimes been stated so as to suggest that influences from culture are in the category of semantic knowledge (like one’s knowledge of names of countries and meanings of words) rather than in the category of attitudes (cf. Karpinski & Hilton, 2001; Olson & Fazio, 2004). However, the person and culture labels can fit equally well with their being conceived as two varieties of attitudinal knowledge, making it a variant of the dual-representation position.

Evaluating the Three Interpretations

To what extent can behavioral evidence for dissociation resolve questions of how many attitude representations exist? Perhaps the most disconcerting conclusion of this chapter is that there is actually *no* possibility for using behavioral evidence to choose decisively among the single-representation, dual-representation, and person vs. culture interpretations of dissociation data patterns.³

³ Dunn and Kirsner (1988) are more sanguine about demonstrating structural dissociation with behavioral measures. They describe a “reversed association” data pattern that can justify concluding that different “processes” are involved in two performances. Their analysis does not consider the distinction between processes and structural representations. In our view (which is not developed formally here in parallel fashion to that of Dunn and Kirsner), this added layer of distinctions removes the possibility of using behavioral data to choose between single- and dual-representation structural views.

Although demonstrations of simultaneous convergent and discriminant validity contribute toward a conclusion in favor of structural dissociation, they do not oblige such a conclusion. It is possible to explain the empirically distinct constructs in terms of a single type of structure. Nosek and Smyth (in press) illustrated the possibility of having distinguishable empirical constructs based on a single structure with the physics of H₂O. Snow, ice, water, and steam are empirically distinct phenomena that share a single structural form – H₂O. The differences among the four phases of H₂O are explained, not as differences in molecular structure, but as the result of *processes* — triggered by environmental variations of temperature and pressure — operating on a single molecular structure. Even without a distinct structural representation, it is quite useful to treat them as distinct constructs for many applications.

The H₂O example illustrates that empirically distinct constructs can derive from a single representation. The reverse is also true. Behavioral evidence suggesting consistency between measures could obscure the existence of distinct underlying representations. The data in Figures 3 and 4 might be taken to reveal the operation of one and the same attitudinal representation underlying both the implicit (IAT) and explicit (self-report) measures. Nevertheless, nothing about those data patterns demands the conclusion that a single representation underlies both types of measure. The measures could reflect representations that, despite being structurally dissociated, have been shaped by the same experiences. For example, imagine that the explicit measure in Figures 3 and 4 was not self-rated attitude, but a sibling's estimate of the participant's IAT assessed attitude. A strong correlation would indicate that persons may have accurate knowledge of their siblings' attitudes as measured by the IAT. Despite the strong correlation, there is most certainly a structural dissociation in the underlying data – they reside in different brains (see Nosek, 2005).

Returning to our original example of structural dissociation with split-brain patients, why is it that we can be confident in interpreting a structural dissociation in that case, but not in the case of attitudes? There is an important feature that distinguishes brains and attitudes: brains are physical entities, attitudes are not. Attitudes, like other psychological constructs, are hypothetical and unobservable. This means that resolution

of “how many structural representations” is not possible because latent constructs do not (at least, not yet) correspond to a known physical structure.

In summary, impressive as the double dissociation data pattern is, there is nothing in that empirical pattern that, by itself, requires an interpretation in terms of structurally distinct underlying representations. The only meaningful inferences from behavioral data are discriminant and convergent validation of empirical constructs. The empirical constructs “implicit attitudes” and “explicit attitudes” can reasonably be interpreted as deriving from either a single-representation or a dual-representation structure. No behavioral evidence can demand a conclusion that one is right and the other is not.

Convergent and Discriminant Validity Evidence Supports Two Attitude Constructs

As already described, even an empirically clear double-dissociation finding does not demand theoretical interpretation as the product of structurally distinct representations. And, the lack of dissociation might, in isolation, decrease the plausibility of distinct structural representations, but it does not require such a conclusion (i.e., highly correlated sibling responses does not mean that they share a brain). A non-structural theoretical interpretation for double-dissociation empirical data requires only some plausible explanation of distinct influences operating on each type of measure, such as explanation in terms of differences in processes engaged by the measurement procedures. In the case of implicit and explicit attitude measures, there are generally numerous differences in measurement procedure. Also, as was previously explained, even quite clear non-dissociation data patterns (as in Figures 3 and 4) are open to interpretation as being produced by structurally distinct representations.

Although these issues have been regularly discussed in philosophy of science, they are still frequently misapplied. For example, on distinguishing implicit and explicit attitude measures, Fazio and Olson (2003; pp 302-303) conflate constructs with representations stating that:

A second troublesome aspect of the implicit-explicit distinction is that it implies preexisting dual attitudes (or whatever the construct of interest might be) in memory. That is, if the terms refer to the constructs themselves, then

both an implicit and an explicit attitude presumably exist in memory (see Wilson et al. 2000).

Fazio and Olson (2003) continue: “For these reasons, it is more appropriate to view the *measure* as implicit or explicit, *not* the attitude (or whatever other construct)” (page 303; italics in original).

Following the discussion above, the construct terms “implicit attitude” and “explicit attitude” do not, as Fazio and Olson worry, commit attitudes to originating from dual-representations. Their preference to limit the implicit-explicit (or indirect-direct) terminology to measures appeals to a distinction that is methodological, not theoretical. Psychological theories explain relations between constructs, not measures.⁴ Procedural differences between measures can be determined without conducting empirical research and have no direct implications for psychological theory or construct validation (De Houwer, in press). In other words, the description of *measures* as implicit-explicit (or indirect-direct), holds no matter what behavioral evidence is gathered.

With the ambiguity of behavioral data, is any purpose served by debating whether behavioral data patterns such as Figures 1 and 2 are usefully interpreted as evidence for structurally dissociated underlying attitudinal representations? Even though the behavioral evidence does not afford a conclusion that one of the three theoretical interpretations is *the* correct one, nevertheless it is reasonable to use behavioral evidence to compare the three interpretations in terms of *construct validity*.

As a broad methodological topic, validity deals with *justification* for descriptions of research findings. Construct validity refers to the justifications for statements about research conclusions offered in the language of theoretical constructs. Without being able to declare in any decisive way that any of the structural interpretations of dissociation data patterns can be dismissed as incorrect, it is still possible to talk about empirically distinguishing constructs. So, instead of resolving single versus dual-

⁴ For example, one does not discuss solid and liquid measures of H₂O. More useful are the constructs ice and water, and explaining their relationship involves a theory in which processes such as heat application or removal lead to transformation of one to the other.

representations, convergent and discriminant validity can help distinguish the value of single versus dual attitude *constructs*.

In appraising construct validity of the implicit–explicit relation, the most important construct validity evidence is provided by studies that have reported correlations of IAT and self-report measures with attitude-relevant behaviors. Poehlman, Uhlmann, Greenwald, and Banaji (2005) recently collected and meta-analyzed these studies, yielding four conclusions that bear on evaluating the three interpretations of dissociation, at the level of constructs, not representations:

1. IAT measures showed consistent positive correlations with behavioral indicators of attitude at moderate levels (average effect size $r = .27$). These relationships were not significantly influenced by any of several potential moderators that were examined.

2. Correlations of explicit measures with behavioral measures of attitude (average effect size $r = .35$) were on average slightly and significantly higher than those of IAT measures, but moderating effects were found. Correlations of explicit attitude measures with behavior significantly “weakened in socially sensitive outcome domains and for responses that are difficult to consciously control”.

3. IAT measures significantly outperformed self-report measures in predicting behavior in the domain of prejudice and stereotyping (average effect size r s were .25 and .13, respectively), which are domains often considered to be socially sensitive.

4. When self-report and IAT measures were highly correlated with each other — a circumstance occurring especially in domains of political and consumer attitudes — both types of measures were more strongly correlated with behavior than when implicit–explicit correlations were low.

These meta-analytic conclusions conform to the unusual combination of convergent and discriminant validity described previously. The convergent validity evidence that justifies interpreting both IAT and self-report as measures of *attitude* is that both types of measure display reliable positive correlations with measures of attitude-relevant behavior. The consistent finding of positive correlations between IAT and self-report measures that has been found in other meta-analyses (Hofmann, Gawronski, Gschwendner, Le, & Schmitt, 2005; Hofmann, Gschwendner, Nosek, & Schmitt, 2005; Nosek, 2005) further supports convergent validity.

Part of the discriminant validity evidence that justifies interpreting the measures as assessments of distinct constructs (implicit and explicit) is finding that correlations between IAT and self-report measures are only modestly positive on average. This type of finding (illustrated in Figure 1) is not by itself convincing evidence of discriminant validity of implicit and explicit measures, because it has the possibly uninteresting explanation that one or both of the measures is psychometrically weak. More important for discriminant validity therefore were Poehlman et al.'s (2005) meta-analytic findings that IAT and self-report attitude measures differed in their relations with other variables. Correlations involving explicit attitude measures, but not IAT measures, were moderated by judged social desirability pressures of the measurement situation and, likewise, correlations of explicit but not IAT attitude measures with behavior were moderated by judged conscious controllability of the measured behavior. A useful summary of the overall meta-analytic evidence is one that has previously been offered by several researchers — IAT measures appear especially useful in predicting attitude-relevant behavior that plausibly occurs without planning and deliberation, whereas self-report best predicts the complementary category of attitude-relevant behaviors that are deliberate or planned (Asendorpf, Banse, & Mucke, 2002; Perugini, 2005).

The summary statement just given fits well with a dual-construct conception in which IAT and self-report measures reflect different types of attitudes. However — to restate a point made a few times previously in this chapter — the behavioral meta-analytic findings cannot be taken as disproving one or another of the different *representation* interpretations. The interpretation of multiple representations is an arbitrary decision about the psychological taxonomy on which psychological processes operate. Psychological taxonomies are organizational schemes, not theories (Willingham & Goedert, 2001). Constructs are hypothetical and tentative at the same time that they are useful and powerful. Whether implicit and explicit attitudes are conceived as dual-representation or a single-representation might be based on explanatory power and parsimony of the resulting theory, rather than more directly on empirical findings. If one theory must postulate dozens of interacting processes in order to maintain a sensible single representation account of existing data, and another theory can account for the same data more directly by use of a dual representation conception, then the latter

theory might justifiably be preferred to the former. In both cases, however, the empirical data would support an interpretation of dual constructs.⁵

Poehlman et al.'s (2005) meta-analysis also sheds light on interpretation of the person vs. culture distinction of the difference between IAT and self-report measures. The meta-analytic finding that IAT attitude measures effectively predicted attitude-relevant behavior is difficult to reconcile with the interpretation that the IAT provides a measure of cultural knowledge that is distinct from the person's own evaluations (i.e., attitudes). Nevertheless, an advocate of the person vs. culture interpretation might explain the IAT's ability to predict attitude-relevant behavior by suggesting that non-attitudinal cultural knowledge can influence behavior outside of awareness. This stipulation would bring the person vs. culture interpretation into agreement with the meta-analytic findings. Although this variant of the person vs. culture interpretation cannot be faulted on logical grounds, it does render that interpretation empirically indistinguishable from one in which the culturally produced knowledge is regarded as affective or attitudinal in nature. Said another way, with the stipulation that non-attitudinal cultural knowledge can influence attitude-relevant behavior, the term 'cultural knowledge' serves only to describe the presumed origins of the knowledge, not its implications for behavior (Nosek & Hansen, 2005). This theoretical flexibility is just one more symptom of the difficulty of using behavioral data to choose between theoretical interpretations.

Conclusions

Two issues make the question of "how many attitude representations are there" unresolvable. First, psychological constructs are hypothetical, resisting definitive decisions about number or form. Theories can explain the same behavioral data as multiple processes operating on a single representation, one process operating on multiple representations, or any admixture of representations and processes. Selection among theories is based on explanatory power and parsimony, not clarification of how many representations actually exist. Second, even if psychological constructs were treated as physical entities, behavioral dissociation data is not sufficient to determine

⁵ Importantly, the meta-analytic evidence does not resolve what is responsible for the discriminant validity

whether one, two, or more representations are operating. Dissociation increases the potential utility of conceiving of multiple representations, and association decreases the potential utility. But, as described, convergent validity can mask underlying multiple representations (e.g., self-ratings and sibling judgments), and discriminant validity can mask underlying singular representations.

Although we are confident that the *single-representation vs. dual-representation* debate will not be resolved decisively by behavioral data, fortunately no such uncertainty attends the question of whether two *theoretical constructs* are needed to map the implicit–explicit attitude domain. It appears unequivocally established that two constructs are needed. The relevant data are those that establish discriminant validity of the implicit–explicit distinction for attitudes described above. Even staunch adherents of the *single-representation* view must concede that the implicit–explicit distinction has been established at the level of empirical constructs. Among such advocates, Fazio and Olson (for example) account theoretically for the contrast between implicit and explicit attitudes by appeal to distinct *processes* — ones involving motivation and ability or situational opportunity — that can be applied to a single type of structural attitude representation (Fazio & Olson, 2003). Others prefer to treat the two constructs not as process variations applied to a single type of mental structure, but as structurally distinct attitude representations. Among those taking the latter structural dissociation view, there are two camps — one that describes the two types of representations as being attitudinal in nature and another that describes the distinction in terms of the contrast between an attitudinal representation and a cultural or semantic-knowledge representation.

References

- Asendorpf, J. B., Banse, R. & Mucke, D. (2002). Double dissociation between implicit and explicit personality self-concept: The case of shy behavior. Journal of Personality and Social Psychology, 83, 380-393.
- Chaiken, S., & Trope, Y. (Eds.) (1999). *Dual process theories in social psychology*. New York: Guilford.
- De Houwer, J. (in press). What are implicit measures and indirect measures of attitudes? *Social Psychology Review*.
- Dunn, J. C., & Kirsner, K. (1988). Discovering functionally independent mental processes: The principle of reversed association. *Psychological Review*, 95, 91–101.
- Fazio, R. H. (1990). Multiple processes by which attitudes guide behavior: The MODE model as an integrative framework. *Advances in Experimental Social Psychology*, 23, 75-109.
- 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., & Sriram, N. (2006). Consequential validity of the Implicit Association Test: Comment on the article by Blanton and Jaccard. *American Psychologist*, in press.
- Hofmann, W., Gawronski, B., Gschwendner, T., Le, H., & Schmitt, M. (2005). A meta-analysis on the correlation between the Implicit Association Test and explicit self-report measures. *Personality and Social Psychology Bulletin*, 31, 1369–1385.
- Hofmann, W., Gschwendner, T., Nosek, B. A., & Schmitt, M. (2005). What moderates implicit–explicit consistency? *European Review of Social Psychology*, 16(10), 335-390.
- Karpinski, A., & Hilton, J. L. (2001). Attitudes and the Implicit Association Test. *Journal of Personality and Social Psychology*, 81, 774–788.
- Nosek, B. A. (2005). Moderators of the relationship between implicit and explicit evaluation. *Journal of Experimental Psychology: General*, 134, 565-584.

- Nosek, B. A., & Hansen, J. (2005). *The associations in our heads belong to us: Attitudes and knowledge in implicit cognition*. Unpublished manuscript, University of Virginia at Charlottesville.
- Nosek, B. A., & Smyth, F. L. (in press). Implicit and explicit attitudes are related, but distinct constructs. *Experimental Psychology*.
- Olson, M. A., & Fazio, R. H. (2004). Reducing the influence of extrapersonal associations on the Implicit Association Test: Personalizing the IAT. *Journal of Personality and Social Psychology*, 86, 653–667.
- Perugini, M. (2005). Predictive models of implicit and explicit attitudes. *British Journal of Social Psychology*, 44, 29–45.
- Poehlman, T. A., Uhlmann, E., Greenwald, A. G., & Banaji, M. R. (2004). *Understanding and using the Implicit Association Test: III. Metaanalysis of predictive validity*. Unpublished manuscript, Yale University, New Haven, CT.
- Strack, F., & Deutsch, R. (2004). Reflective and impulsive determinants of social behavior. *Personality and Social Psychology Review*, 8(3), 220–247.
- Willingham, D. B., & Goedert, K. (2001). The role of taxonomies in the study of human memory. *Cognitive, Affective, and Behavioral Neuroscience*, 1(3), 250-265.
- Wilson, T. D., Lindsey, S., & Schooler, T. Y. (2000). A model of dual attitudes. *Psychological Review*, 107(1), 101–126.

Chapters in this volume by Fazio, Wilson (and perhaps Petty).

Figure Captions

1. Regression of an IAT measure of implicit age attitude on a parallel self-report measure. This analysis is based on data from Greenwald, Nosek, and Banaji (2003). The regression reveals both a weak positive correlation between the IAT and self-report measures and a wide separation between their means on standardized scales for which the zero points of both indicate evaluative indifference between young and old. The self-report measure shows much weaker attitudinal preference for young relative to old. See text for further discussion.
2. IAT and self-report age attitude measures of Figure 1, plotted to reveal that the IAT measure is unrelated to variations in age of respondents, whereas the self-report attitude measure shows a regular reduction in relative preference for young as respondent age increases. Such distinctive patterns of correlation with other variables suggest dissociation of mental representations underlying the implicit and explicit attitude measures.
3. Regression of an IAT measure of implicit attitudinal preference for George W. Bush, relative to John F. Kerry, on a parallel self-report measure. Data from Greenwald, Nosek, and Sriram (2006). This regression illustrates both a strong positive correlation between IAT and self-report measures, and no separation between their means on standardized scales for which the zero points of both indicate evaluative indifference between the two presidential candidates. These observations suggest lack of implicit–explicit dissociation. See text for further discussion.
4. IAT and self-report political preference measures of Figure 3, plotted to reveal that both measures have the same relation to variations in education level of respondents. Such similar patterns of correlation with other variables suggest singleness (lack of dissociation) of mental representations underlying the implicit and explicit attitude measures.

Figure 1

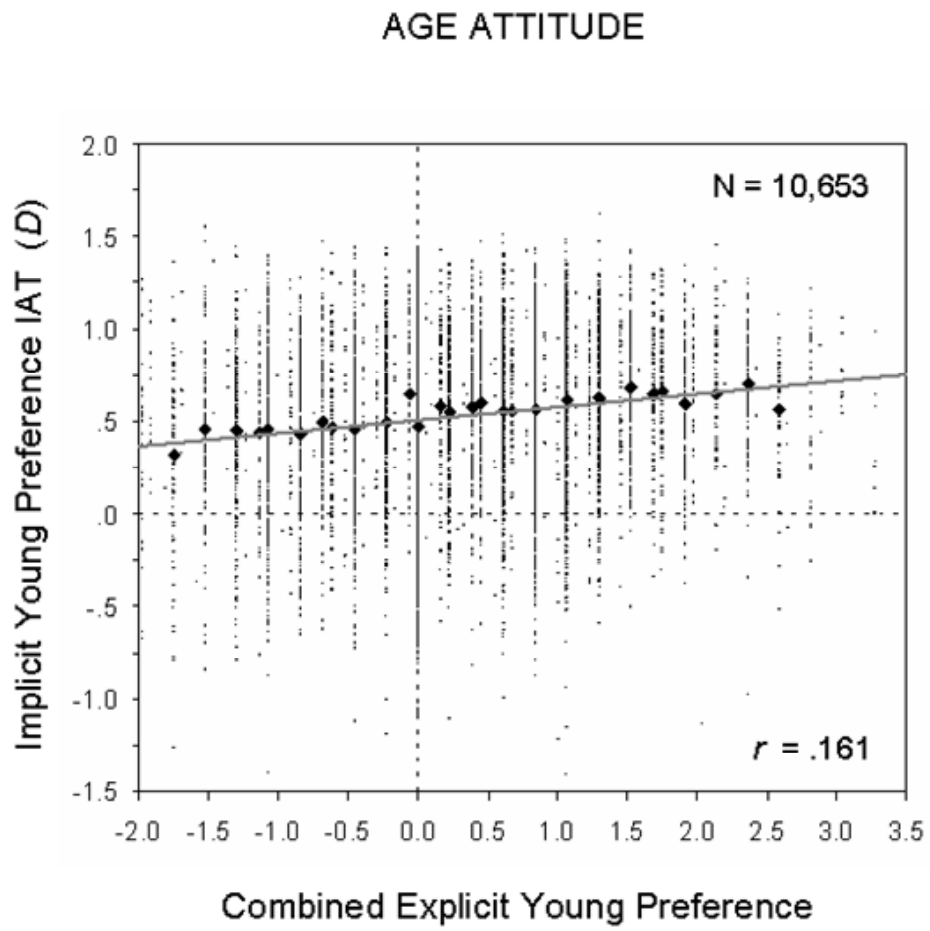


Figure 2

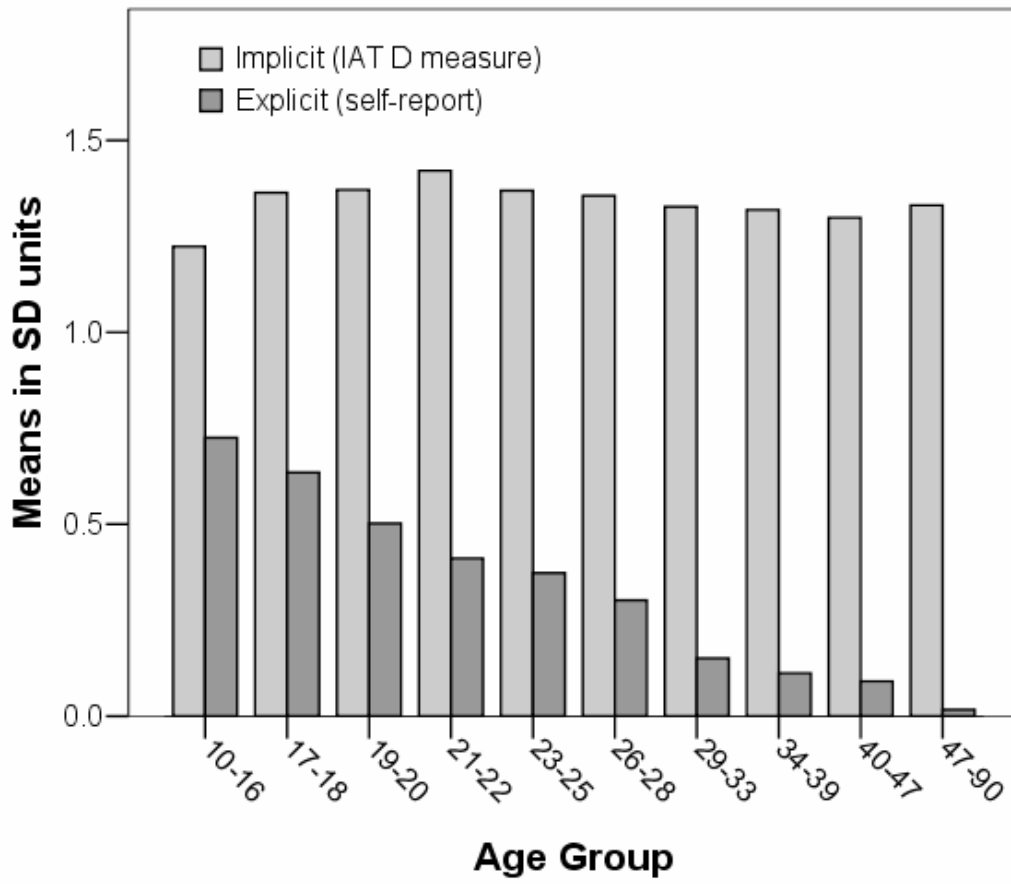


Figure 3

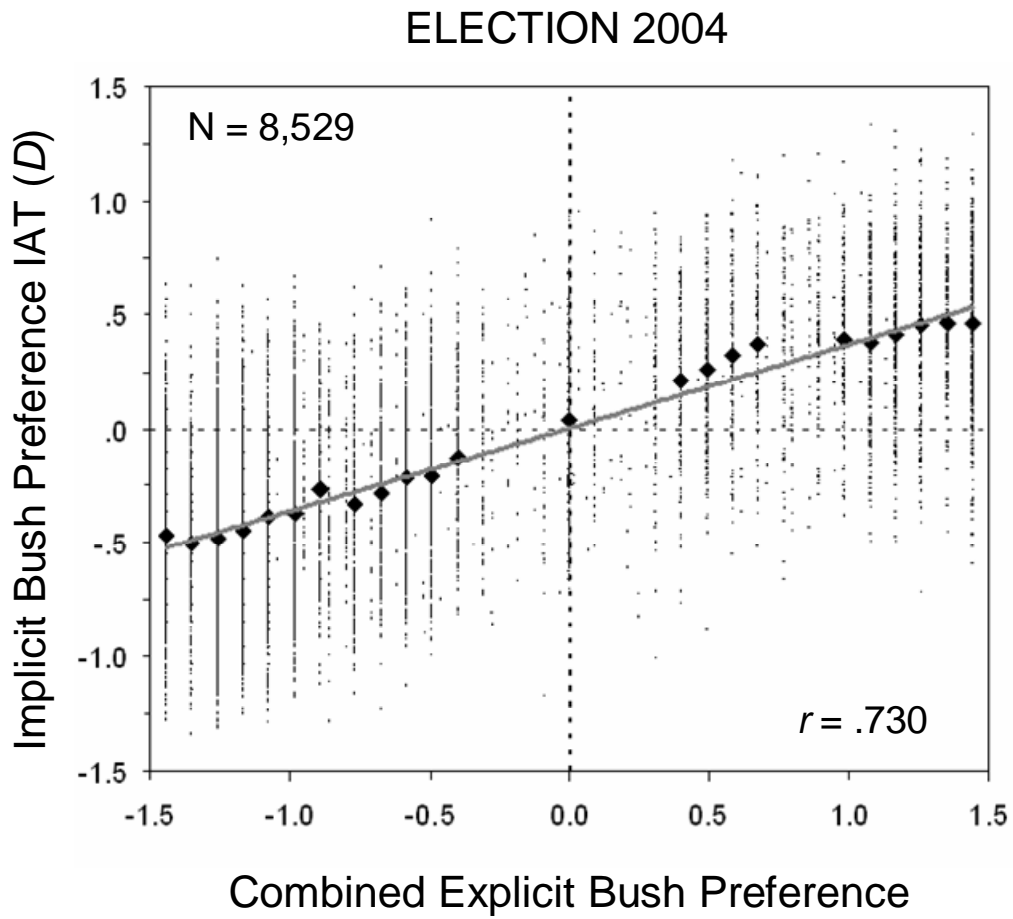


Figure 4

