Validity of the Salience Asymmetry Interpretation of the IAT:

Comment on Rothermund and Wentura (2004)

Anthony G. Greenwald, University of Washington

Brian A. Nosek, University of Virginia

Mahzarin R. Banaji, Harvard University

K. Christoph Klauer, University of Freiburg

Draft of September 5, 2004

Address correspondence to:
Anthony G. Greenwald
Department of Psychology
University of Washington
Box 351525
Seattle, WA 98195-1525

email: agg@u.washington.edu
phone: (206) 543-7227
fax: (206) 685-3157

Length (present / previous draft):

Abstract (146 / 146)
Text (4478 / 4851)
References (578 / 644)
Author Note (91 / 91)
Appendixes (0 / 1087)
Footnotes (147 / 235)
TOTAL (5440 / 7054)

Abstract

The Implicit Association Test (IAT) requires responding to category contrasts such as young vs. old, male vs. female, and pleasant vs. unpleasant. In introducing the IAT, A. G. Greenwald, D. E. McGhee, and J. L. K. Schwartz (1998) proposed that IAT measures reflect mental structures involving the nominal features of the IAT's categories (e.g., age features for young vs. old; gender features for male vs. female; valence features for pleasant vs. unpleasant). In contrast, K. Rothermund and D. Wentura (2004) proposed that IAT performance is dominated by salience asymmetries of the IAT's pairs of contrasted categories. To assess relative contributions of nominal feature contrasts vs. salience asymmetries we (a) briefly summarize the extensive evidence now available to support construct validity of the IAT as a measure based on nominal category features and (b) present two new experiments that yielded results problematic for the salience asymmetry interpretation.

Validity of the Salience Asymmetry Interpretation of the IAT:

Comment on Rothermund and Wentura (2004)

Rothermund and Wentura (2004; R&W hereafter) presented multiple experiments to support their conclusion that the Implicit Association Test (IAT; Greenwald, McGhee, & Schwartz, 1998; GM&S hereafter) is strongly influenced by *salience asymmetries* of the categories used in IAT measures.¹ R&W introduced their salience asymmetry hypothesis as follows:

“Assume for a moment that the two categories of both the target and the attribute dimension indeed differ in salience. . . . In this case, participants will find it easier to respond if the salient categories of both dimensions (the “figures”) are mapped onto one response and the non-salient categories (the “background”) are mapped onto the other response” (R&W, p. 140).

Overview

We agree with R&W that salience asymmetries have the potential to contribute to IAT effects, much as do any other features that afford a basis for distinguishing among categories. We nevertheless disagree with R&W about the importance of salience asymmetries relative to variations in the *nominal features* that distinguish categories used in the IAT. Nominal features are ones indicated by the names commonly used to identify categories. To give a few examples: age is the nominal feature when the category contrast is young vs. old; gender is the nominal feature when the contrast is male vs. female; and valence is the nominal feature when the contrast is pleasant vs. unpleasant. For these contrasts and others, R&W proposed that asymmetries in salience may be more significant contributors to IAT measures than are variations in the nominal features of the contrasted categories. In this comment we consider the possibilities for distinguishing the nominal feature and salience asymmetry interpretations and present some relevant data.

In preparing this comment, we discovered that what initially appeared to be our strongest disagreement with R&W was inconsequential. Correspondence with Klaus Rothermund

(personal communication, May 5, 2004) established that R&W had assumed a definition of *association* different from the one assumed by GM&S. This was not initially apparent because neither GM&S nor R&W had explicitly stated a definition of association. As will be seen, the different conceptions of association allowed (a) GM&S to declare that the IAT measured association strengths, (b) R&W to declare that the IAT did not measure association strengths, and (c) these two assertions not to involve an empirical disagreement.

We start by describing the different conceptions of association used by GM&S and R&W. This allows us to observe that the disagreement between GM&S and R&W is confined to determining the relative contributions of salience asymmetries and nominal category features to IAT measures. To anticipate our conclusion: Although we agree with R&W that salience asymmetries have the potential to influence IAT performance, it remains for further research to establish that this possible influence threatens uses of the IAT to provide implicit measures of constructs in the domains of attitudes, stereotypes, self-concepts, and self-esteem.

Theory-uncommitted and Theory-committed Conceptions of Association

Association has a long history in psychology, including a wide variety of theory-based conceptions derived from the learning theories of Thorndike, Pavlov, Hull, and Tolman, among others (for overviews, see learning theory texts such as that of Bower & Hilgard, 1981). Before the 20th century, the topic of *association of ideas* had a very long history in philosophy. Among the prior philosophical conceptions of the association of ideas are (a) Aristotle's (ca. 350 B.C./1930) view that there are four bases for association: contiguity, frequency, similarity, and contrast, and (b) David Hume's (1740/1939) view that the bases for association can be reduced to three principles: resemblance (i.e., similarity), contiguity in time or place, and causality (effect associated with its cause). Many other philosophers, as well as many psychologists, have endorsed similarly broad conceptions of association, which are not tied to any specific theory of the structure of associative mental representations. This widely used *theory-uncommitted* conception of association was the one implicitly used in the GM&S article that introduced the

IAT.

Like GM&S, R&W also provided no explicit definition of association. However, correspondence with Klaus Rothermund (personal communication, May 5, 2004) established that R&W assumed a conception based on modern theories of semantic networks (e.g., Quillian, 1967) and spreading semantic activation (e.g., Collins & Loftus, 1975). As will be explained in the next few paragraphs, identification of this difference — between GM&S's theory-uncommitted conception and R&W's theory-committed conception — removes much of the apparent disagreement between their respective interpretations of the IAT.

A critical passage in R&W appears in their section, "Pitting the Two Accounts Against Each Other" (p. 158). In that passage R&W described a *strategic recoding* interpretation of the IAT in which they used as an illustration the face–name thought experiment that was used by GM&S (p. 1464) to introduce the IAT. In the face–name example one task combination (male names and male faces getting one response; female names and female faces getting the other) makes it easy to use a shared feature (gender) to group the two categories assigned to each response. The other task combination (male names and female faces getting one response; female names and male faces getting the other) provides no shared feature that can be used to simplify the task. Because R&W's semantic-network conception of association did not include similarity due to shared features as a basis for association, they did not consider GM&S's face–name example to involve a role of associations in IAT performance. In contrast, GM&S — regarding similarity due to the shared gender feature as a valid basis for association in their theory-uncommitted conception — understood their thought experiment as being consistent with an associative basis for IAT performance.

Although GM&S used no theory of the structure of associative mental representations in presenting their interpretation of the IAT as a measure of association strengths, it is nevertheless possible to explain the IAT with a theory-committed interpretation of association. That was done recently in this journal by Hall, Mitchell, Graham, and Lavis (2003), who based

their interpretation of the IAT on the theory of *acquired equivalence* (Dollard & Miller, 1950, p. 101). The theory of acquired (or learned) equivalence is that pairs of stimuli that have a common associate (e.g., faces and names both associated with specific genders) become associated to each other.

In summary, what appeared to be a central disagreement in interpretation of the IAT between GM&S and R&W proved to be no more than different preferences for defining the concept of *association*. This definitional disagreement has implications for choice of language to describe results that are expected to occur in similar empirical form by both GM&S and R&W. Although the definitional disagreement can therefore be set aside for the remainder of this article, we can recommend to ourselves and to others that discussions of associative interpretations of the IAT should hereafter be explicit about their assumed conceptions of association.

Disagreements That Have Empirical Implications

When we go beyond the definitional disagreement that has just been set aside as non-central, we find two empirically addressable questions that stem from R&W's theorization about salience asymmetries. The first, and more important, question concerns the relative contributions of nominal features and salience asymmetries to observed IAT effects. If past IAT results are due primarily to salience asymmetries (as supposed by R&W), then published interpretations in terms of nominal features (as supposed by GM&S) are largely in error. The second question concerns whether R&W's empirical results are validly interpreted as demonstrating salience asymmetry effects.

As we seek to establish in the remainder of this article, for the question about relative importance of salience asymmetries and nominal features we find it implausible that salience asymmetries are more important than nominal features, chiefly because of the strength of published construct validity evidence for interpretations of IAT measures in terms of nominal features. The second question, concerning construct validity of the salience asymmetry

interpretation of R&W's findings, is more difficult to address because research on salience asymmetries in the IAT is so new that there do not yet exist well established research operations for manipulating and measuring salience asymmetry. To add to the available evidence, we report relevant data from two new experiments. These experiments did not support what we presumed to be predictions from R&W's account of salience asymmetry.

Construct Validity of Nominal-Feature Interpretations of IAT Measures

The main evidence for the role of nominal features of the IAT's categories in IAT measures is the extensive existing evidence for construct validity of IAT measures. That evidence can be described in four categories: (a) known groups differences in IAT measures, (b) correlations of IAT measures with self-report measures, (c) predictive validity of IAT measures in studies of prejudice and stereotyping, and (d) use of IAT measures to confirm consistency theory predictions. Because there are multiple published studies in each of these categories, there is no need for a detailed review of evidence here. We limit the summary of these categories to brief descriptions, accompanied by mention of selected relevant studies.

Known groups differences in IAT measures. Some of the earliest evidence for construct validity of IAT measures came from findings that known groups differed in expected ways on IAT measures. Nosek and Greenwald (2001, pp. 88–89) summarized 12 studies showing expected IAT differences between groups such as vegetarians and omnivores, smokers and nonsmokers, homosexuals and heterosexuals, snake phobics and spider phobics, and East Germans and West Germans. Subsequently, more studies of this type have accumulated (reviewed by Poehlman, Uhlmann, Greenwald, & Banaji, 2004).

Correlations of IAT measures with self-report measures. Numerous studies have demonstrated that the IAT's implicit measures of attitudes are positively correlated with self-report measures of attitudes toward the nominal categories used in the IAT. Hofmann, Gawronski, Gschwendner, Le, and Schmitt (2004) reported a meta-analytic summary of such correlational results for IAT and corresponding self-report measures. For 81 studies, they

reported an average effect size r of .24 (also noting moderation by several variables). In a multi-level analysis of correlations between IAT and self-report attitude measures for 57 attitude objects (average N per attitude object = 202), Nosek (2004) reported an average correlation with self-report of $r = .36$, with individual correlations as high as $r = .70$. Greenwald, Nosek, and Banaji (2003) reported disattenuated correlations of IAT with self-report for four attitude objects, using very large internet samples (N s between 6,811 and 10,537). The disattenuated correlations ranged from a low of $r = .23$ for age attitudes to a high of $r = .86$ for attitudes toward candidates in the 2000 United States Presidential Election. These consistently positive correlations provide strong presumptive evidence that nominal features play an important role in IAT attitude measures.

Predictive validity of IAT measures in studies of prejudice and stereotyping. The first study showing prediction of behavioral measures of prejudice by an IAT attitude measure was McConnell and Leibold's (2001) finding that white subjects' IAT-measured implicit racial attitudes predicted several nonverbal indicators of apparent discomfort displayed during a videotaped laboratory interaction with an African American. A recent meta-analysis of predictive validity correlations of IAT measures involving prejudices or stereotypes showed a mean effect size of $r = .25$ (32 independent effect sizes) with a 95% confidence interval of $\pm .06$ (Poehlman et al., 2004).

Use of IAT measures to confirm consistency theory predictions. Greenwald, Banaji, Rudman, Farnham, Nosek, and Mellott (2002) reported a series of studies testing consistency-theoretical predictions of relationships among individual difference measures of attitude, stereotype, self-concept, and self-esteem. These predictions were repeatedly confirmed in designs that used IAT measures of the various constructs. For additional reports that include such findings, see Aidman and Carroll (2003), Nosek, Banaji, and Greenwald (2002), Hummert, Garstka, O'Brien, Greenwald, and Mellott (2002), Jost, Pelham, and Carvallo (2002), and Rudman, Greenwald, and McGhee (2001).

The extensive evidence for construct validity of IAT measures might fit with R&W's salience asymmetry interpretation if salience asymmetries are generally confounded with nominal-feature contrasts in such ways as to yield the same patterns expected for nominal-feature interpretations. R&W indeed did propose such a correlation between salience asymmetries and nominal feature contrasts in interpreting their Experiments 2a and 2b. Of course, if this correspondence is postulated to be a general state of affairs, the salience asymmetry interpretation becomes entirely indistinguishable from a nominal-features interpretation. However, because R&W did not assert or imply that salience asymmetries and nominal-feature differences are generally confounded, their position certainly permits the possibility of empirically distinguishing the two interpretations.

Construct Validity of Salience-Asymmetry Interpretations of IAT Measures

The salience asymmetry interpretation implies correlations of IAT measures with measures of salience asymmetries. In particular, if there are salience asymmetries in both of an IAT's category contrasts, performance in the IAT should be faster when the higher salience categories of each contrast are assigned to the same response than when these two high-salience categories must receive different responses. An important contribution of R&W was to report evidence of this type.

The accumulated evidence for validity of the salience asymmetry interpretation is at present modest. One desirable respect in which R&W's evidence could be extended is through the development of multiple additional measures of salience asymmetry, to supplement the one that they developed and reported. We report some data below using additional salience asymmetry measures that we developed as variations of the one reported by R&W.

A second respect in which R&W's evidence could usefully be extended is to incorporate procedures more closely resembling the procedures of the published studies on which the main evidence for construct validity of the IAT rests. In their experiments, R&W extensively used two procedures that had previously been described as creating threats to the construct validity of

IAT measures: First, in five of their nine IAT experiments, R&W used non-categories (nonsense strings, unrelated neutral words, unknown names) in place of one category of a contrasted pair. On the basis of their review, Greenwald and Nosek (2001) concluded that the IAT does not function properly when non-categories are used in this fashion. Second, R&W used millisecond-unit IAT measures rather than using either the log-transformed latency measure used in most of the existing published IAT literature or the improved *D* measure introduced by Greenwald et al. (2003). Proneness of the millisecond-unit measure to cognitive skill artifact was demonstrated by Cai, Sriram, Greenwald, and McFarland (in press), Greenwald et al. (2003), and Mierke and Klauer (2003). The two new experiments reported here avoided these problematic procedures.

Additional Evidence Concerning Construct Validity of the

Salience Asymmetry Interpretation: Two New Experiments

In order to add to the available evidence on construct validity of R&W's salience asymmetry account of the IAT, we conducted two new experiments that are described briefly here.²

Experiment 1: No Effect of a Strong Salience Asymmetry Manipulation

This experiment applied a salience manipulation to an IAT closely modeled after one of the two IATs used in GM&S's Experiment 1. The target concept contrast was flowers vs. insects and the attribute contrast was pleasant vs. unpleasant valence. This IAT reliably produces an effect that GM&S interpreted as showing a stronger association of flowers (than insects) with positive valence. R&W should interpret this IAT either in terms of (a) an effect of salience asymmetries — assuming that insect items are more salient than flower items and that unpleasant items are more salient than pleasant items, or (b) a strategic recoding of the two contrasts (of the type described on p. 158 of their article) due to the subject using shared valence features to improve performance when the flower and pleasant categories are assigned to one key (with insect and unpleasant to the other). As previously noted, the latter interpretation is empirically equivalent to GM&S's association-strength interpretation, albeit with

different language.

Experiment 1's plan was to overpower any natural salience asymmetry between the flower and insect categories by presenting items for one of these two categories in a bright red font, with all items for the remaining three categories of the IAT in blue font. This manipulation should produce an unambiguous salience asymmetry such that when the flower category is in red it should be more perceptually salient than the insect category; likewise, when the insect category is in red it should be more salient than the flower category. To simplify description of IAT tasks in describing procedures, we identify each of the IAT's two combined tasks by naming two categories that are assigned to the same response. Thus the *flower+pleasant* (equally, the *insect+unpleasant*) task was performed with the categories *flower* and *pleasant* assigned to one response and the categories *insect* and *unpleasant* to the other response.

Subjects for Experiment 1 were 30 University of Washington undergraduate students, half assigned to the flower-salient condition (i.e., flower items in red) and half to the insect-salient condition (insect items in red). The stimulus items for all categories were a subset of those used in the original GM&S (Experiment 1) flower–insect IAT. Three measures of salience asymmetry were used in Experiment 1, administered after all other procedures. These salience asymmetry measures assessed natural salience asymmetries of the flower–insect and pleasant–unpleasant category contrasts, rather than testing salience asymmetries due to use of red font (which needed no testing to confirm). To save space, we omit presentation of the details of these three measures (see Footnote 3).

If, as assumed by R&W (p. 140) *unpleasant* items are more salient than *pleasant* items, the straightforward prediction of the salience asymmetry hypothesis is that Experiment 1 should reveal (a) faster performance for *insect+unpleasant* in the insect-salient condition and (b) faster performance for *flower+unpleasant* in the flower-salient condition. (I.e., if *unpleasant* is more salient than *pleasant*, then the conditions mentioned as expected to have faster performance are those in which the two more salient categories are assigned to the same response.)

The data were analyzed by computing differences between the two combined tasks such that higher values indicated faster performance for *insect+unpleasant* than for *insect+pleasant*. In terms of the salience asymmetry hypothesis, this difference was expected to be numerically higher for the insect-salient condition than for the flower-salient condition. The observed mean differences were 384 ms for the insect-salient condition and 440 ms for the flower-salient condition. Although these two values were not significantly different, the direction of their difference was actually opposite to that predicted by the salience-asymmetry hypothesis, $t(28) = -0.69$, $p = .49$.

A related finding was recently reported by Mierke and Klauer (2003, Experiment 1a). Mierke and Klauer employed an IAT with geometric forms and the two contrasts, red–blue and small–large. All red objects were small and all blue objects were large. This contingency between color and size of stimulus objects (i.e., shared features) led to an IAT effect of faster performance with *red+small* than with *blue+small*. When the contingency was reversed for another group of subjects — that is, when all blue objects were small and red objects large — the IAT effect was reversed even though any salience asymmetries between red and blue and between small and large objects should have remained constant. Mierke and Klauer's finding was entirely consistent with a nominal feature interpretation.

If the finding of Mierke and Klauer's (2003) Experiment 1a and present Experiment 1 are taken together, salience asymmetries appear neither necessary (Mierke and Klauer's Experiment 1a) nor sufficient (present Experiment 1) to induce IAT effects. Nevertheless, it remains plausible that salience asymmetries might cause IAT effects in the absence of stronger cues to association (such as shared meaning or size–color contingency). This could explain R&W's finding, in their Experiment 3b, in which a salience manipulation by color distinctiveness caused a reversal of IAT-like effects in the absence of other bases of association. In R&W's Experiment 3b, an old–young contrast was combined with a yellow–green color discrimination. There is little reason to expect an association of old or young more with one than the other of

these two colors.

Experiment 2: Lack of Correlation Between Salience Asymmetries and Individual Differences in Implicit Gender Identity

In their Experiment 2a R&W reported several findings that supported the salience asymmetry hypothesis. Present Experiment 2 sought to replicate their main finding, that individual differences in a measure of salience asymmetry of male and female categories were correlated with individual differences in a gender self-concept IAT.

Procedures of present Experiment 2 differed from R&W's Experiment 2a in three ways, all of which were expected to strengthen tests of possible salience asymmetry effects: (a) we used Greenwald et al.'s (2003) *D* measure for the IAT (while also replicating R&W's analysis using a millisecond-unit IAT measure); (b) we used three salience asymmetry measures modeled on (but not identical to) R&W's visual search measure of salience asymmetry; and (c) in place of R&W's use of masculine and feminine trait words to represent the male–female contrast connotatively, we used nouns (boy, man; girl, woman) and pronouns (he, him, his; she, her, hers) that represented the gender contrast denotatively.

Subjects were 54 new volunteers (25 male, 29 female) from the same undergraduate population as for Experiment 1. Two female subjects who provided incomplete data were excluded from all analyses. Analyses that included the salience asymmetry measures are described omitting data from ten other subjects (5 male, 5 female) who had very slow performance on one or more of the salience asymmetry measures. Exclusion of these subjects mildly improved power of statistical tests.

The only variations from Experiment 1's procedure were (a) use of a gender self-concept IAT in place of the flower–insect IAT of Experiment 1, (b) absence of any manipulation of salience in the IAT, and (c) inclusion, prior to the IAT, of two self-report gender self-concept measures (results from which are tangential to present purposes and are not described here). As in Experiment 1, the three salience asymmetry measures were administered after all other

procedures.

Results confirmed the nominal-feature expectation that performance on the gender self-concept IAT would be sharply different for male and female subjects. The IAT measure was scored so that numerically higher scores indicated faster performance for *self+female* than *self+male*. For the *D* measure (Greenwald et al., 2003) women showed the expected higher score ($M = 0.51$, $SD = 0.38$) than men ($M = -0.48$, $SD = 0.38$), $t(50) = 9.44$, $p = 10^{-12}$. This result indicated, as desired, that Experiment 2's pronoun-item gender self-concept IAT produced more clearly defined sex differences than did the trait-item gender self-concept IAT of R&W's Experiment 2a. The male–female difference of 1.00 on the *D* measure in the present experiment was 61% greater than the difference (0.62) observed in a parallel analysis conducted on R&W's Experiment 2a.³

Correlations between salience asymmetries and IAT measures. From the principle that “The less familiar of two categories is . . . more salient and constitutes a figure against the background of the familiar category” (R&W, p. 140), both women and men should show greater salience of *other* (the less familiar category) than of *self*, while women should show greater salience of *male* than *female* and men the reverse. The salience measure was scored so that faster identification of male words (the salience asymmetry hypothesis's expected result for women) yielded the numerically higher score. The IAT was scored so that higher numbers indicated faster performance in the *female+self* combined task (the expected result for women for both the nominal feature and salience asymmetry hypotheses). With these directions of scoring, the salience asymmetry hypothesis predicted a positive correlation between the IAT and salience asymmetry measures. (The nominal feature interpretation has no prediction for this correlation.) The correlation of the gender self-concept *D* measure with the average of our three salience asymmetry measures was $r = -.15$, $p = .35$. For the millisecond-unit measure, the correlation was $r = -.19$, $p = .23$. These correlations did not confirm the salience asymmetry hypothesis's prediction and were actually non-significantly opposite in direction from that

prediction. The results from our salience asymmetry measures also did not confirm R&W's expectations that items in the less familiar category in each contrast should be more salient than items from the more familiar category (R&W, p. 140).⁴

Discussion

The first and most important conclusion of this article is that disagreements between the conclusions of R&W, on the one hand, and those of GM&S and other IAT researchers, on the other, are less than might initially appear. The respective authors' different (but unstated) conceptions of *association* led to the appearance of greater disagreement than actually exists. With this difference in use of language identified, R&W's strategic recoding interpretation becomes empirically interchangeable with GM&S's association-strength interpretation. Both of these interpretations credit variation in IAT measures to relations among the nominal features of the categories used in the IAT. To the extent that the IAT measures either association strength (in GM&S's usage) or strategic recoding (as used by R&W), the implications for construct validity of IAT measures are the same.⁵ Notwithstanding this agreement on possible validity of IAT interpretations in terms of nominal features, there remains a substantial disagreement between GM&S and R&W. GM&S, along with most other researchers who have published IAT studies, regarded the nominal-feature interpretation as sufficiently potent to be useful in the analysis of group and individual differences in IAT measures of social cognitive constructs such as attitudes, stereotypes, self-concepts, and self-esteem. R&W disagreed, proposing that the salience asymmetry interpretation is sufficiently potent to call into question claims for construct validity of IAT measures in terms of nominal features.

In support of the nominal-feature interpretation of the IAT we made three points: (a) published empirical studies provide extensive evidence for validity of the nominal-feature interpretation, (b) evidence for construct validity of R&W's salience asymmetry interpretation is, at present, confined to the findings reported in their article, and (c) findings from the present article's two new experiments were generally consistent with expectations of nominal feature

interpretations, while deviating substantially from theoretical expectations of the salience asymmetry hypothesis.

It is apparent that our conclusions differ noticeably from those of R&W. One might respond to this difference of opinion by conducting studies to identify possible crucial differences between procedures of R&W's studies and those used in the present two experiments. However, the goal of identifying such procedures is of minor importance in comparison with the value of conducting further research that more directly assesses the construct validity of the nominal-feature and salience asymmetry interpretations of the IAT.

References

- Aidman, E. V., & Carroll, S. M. (2003). Implicit individual differences: Relationships between implicit self-esteem, gender identity, and gender attitudes. *European Journal of Personality*, 17, 19–36.
- Aristotle (1930). On memory and reminiscence. In Ross, W. D. (Ed.), *The works of Aristotle* (vol. 3). Oxford: Clarendon Press. (Original publication ca. 350 B.C.)
- Bower, G. H., & Hilgard, E. R. (1981). *Theories of Learning*. New York: Prentice-Hall.
- Cai, H., Sriram, N., Greenwald, A. G., & McFarland, S. G. (in press). The Implicit Association Test's *D* measure can minimize a cognitive skill confound: Comment on McFarland and Crouch (2002). *Social Cognition*.
- Collins, A. M., & Loftus, E. F. (1975). A spreading-activation theory of semantic processing. *Psychological Review*, 82, 407--428.
- Dollard, J., & Miller, N. E. (1950). *Personality and psychotherapy*. New York: McGraw-Hill.
- Greenwald, A. G., Banaji, M. R., Rudman, L. A., Farnham, S. D., Nosek, B. A., & Mellott, D. S. (2002). A unified theory of implicit attitudes, stereotypes, self-esteem, and self-concept. *Psychological Review*, 109, 3-25.
- 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.
- Hall, G., Mitchell, C., Graham, S., & Lavis, Y. (2003). Acquired equivalence and distinctiveness in human discrimination learning: Evidence for associative mediation.

Journal of Experimental Psychology: General, 132, 266–276.

Hofmann, W., Gawronski, B., Gschwendner, T., Le, H., & Schmitt, M. (2004). *A Meta-Analysis on the Correlation between the Implicit Association Test and Explicit Self-Report Measures*. Manuscript submitted for publication. Northwestern University.

Hume, D. (1939). *A Treatise of Human Nature*. London: J. M. Dent and Sons. (Original publication 1740)

Hummert, M. L., Garstka, T. A., O'Brien, L. T., Greenwald, A. G., Mellott, D. S. (2002). Using the Implicit Association Test to measure age differences in implicit social cognitions. *Psychology and Aging*, 17, 482-495.

Jost, J. T., Pelham, B. W., & Carvallo, M. R. (2002). Non-conscious forms of system justification: Implicit and behavioral preferences for higher status groups. *Journal of Experimental Social Psychology*, 38, 586–602.

McConnell, A. R., & Leibold, J. M. (2001). Relations among the Implicit Association Test, discriminatory behavior, and explicit measures of racial attitudes. *Journal of Experimental Social Psychology*, 37, 435–442.

Mierke, J., & Klauer, K.C. (2003). Method-specific variance in the Implicit Association Test. *Journal of Personality and Social Psychology*, 85, 1180–1192.

Nosek, B. A. (2004). *Moderators of the Relationship between Implicit and Explicit Attitudes*. Manuscript submitted for publication. University of Virginia.

Nosek, B.A., Banaji, M.R., & Greenwald, A.G. (2002). Math = male, me = female, therefore math ...me. *Journal of Personality and Social Psychology*, 83, 44-59.

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. [This is soon to be submitted for publication.]

Quillian, M. R. (1967). Word concepts: A theory and simulation of some basic semantic capabilities. *Behavioral Sciences*, 12, 410-430.

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.

Rudman, L. A., Greenwald, A. G., & McGhee, D. E. (2001). Implicit self-concept and evaluative implicit gender stereotypes: Self and ingroup share desirable traits. *Personality and Social Psychology Bulletin*, 27, 1164–1178.

Author Note

Anthony G. Greenwald, Department of Psychology, University of Washington, Seattle, Washington; Brian A. Nosek, Department of Psychology, University of Virginia, Charlottesville, Virginia; Mahzarin R. Banaji, Department of Psychology, Harvard University, Cambridge, Massachusetts; K. Christoph Klauer, Institut für Psychologie, Albert-Ludwigs-Universität Freiburg, Freiburg, Germany.

This research was supported by grants from National Institute of Mental Health, MH-41328, MH-01533, and MH-57672.

Correspondence concerning this article should be addressed to Anthony G. Greenwald, University of Washington, Department of Psychology, Box 351525, Seattle, WA, USA, 98195-1525. Electronic mail should be sent to agg@u.washington.edu

Footnotes

1. Because the IAT procedure has been described in several publications (including R&W), we dispense with an additional description here.
2. More detailed descriptions of the two experiments can be obtained from the first author.
3. The authors thank Klaus Rothermund for providing the data of R&W's Experiment 2a for reanalysis.
4. However, none of our three salience asymmetry measures was identical to the R&W's single measure and, of course, our items were in English language rather than German. Consequently, the findings involving our salience asymmetry measures are not directly comparable to R&W's.
5. R&W and GM&S do differ slightly in describing the automatic versus controlled nature of processes involving use of nominal features. R&W assume that strategic recoding is an exclusively controlled (conscious) processes, whereas GM&S did not take a position on the involvement of automatic versus controlled processes in the use of nominal features in the IAT, allowing both possibilities.