Health of the Implicit Association Test at Age 3

Anthony G. Greenwald, University of Washington

Brian A. Nosek, Yale University

Abstract. Since its first publication in 1998, the Implicit Association Test (IAT) has been used repeatedly to measure implicit attitudes and other automatic associations. Although there have also been a few studies critical of the IAT, there now exists substantial evidence for the IAT's convergent and discriminant validity, including new evidence reported in several of the articles in this special issue. IAT attitude measures have often correlated only weakly with explicit (self-report) measures of the same associations. It therefore seems appropriate to conclude that the IAT assesses constructs that are often (but not always) distinct from the corresponding constructs measured by self-report.

The Implicit Association Test (IAT: Greenwald, McGhee, & Schwartz, 1998) offers a method for indirectly measuring strengths of associations between concepts. In taking the IAT, subjects sort stimuli representing four concepts using just two responses, each assigned to two of the four concepts. Usefulness of the IAT as a measure of association strength depends on the assumption that, if two concepts are highly associated, the IAT's sorting tasks will be easier when the two associated concepts share the same response than when they require different responses.

The example IAT illustrated in Figure 1 requires subjects to sort the four concepts of male, female, mathematics, and arts. This IAT measures associations involved in a 'gender-math' stereotype that associates mathematics more with male than female (Nosek, Banaji, & Greenwald, 2000b). Subjects who have these gender-stereotypic associations should find the IAT's task to be easier when they are asked to give the same response to *male* and *math* and a different response to *female* and *arts* (as in the 5th step of Figure 1) than when they must give one response to *female* and *math* and the other response to *male* and *arts* (3rd step of Figure 1).

In its typical uses, the IAT measures *relative* strengths of *pairs* of associations. The task schematized in Figure 1 provides an IAT effect measure that assesses the strengths of a gender-stereotypic pair of associations (*male-math* and *female-arts*, which should facilitate performance in Step 5 of Figure 1), relative to the strengths of a counterstereotypic pair of associations (*female-math* and *male-arts*, which should facilitate performance in Step 3). The faster Step 5's performance is compared to Step 3's, the more the IAT effect measure (computed as Step 3 latency minus Step 5 latency) indicates strength of the gender stereotype.

The first research using the IAT (Greenwald et al., 1998) established that the IAT could detect differences in association of concepts with valence (i.e., differences in attitudes) that were (a) almost universal in the population (e.g., preference for flowers over insects), (b) expected as differences between subject populations (between Korean Americans and Japanese Americans in preferences for their respective ethnicities), or (c) unrecognized mental residues of a racist culture (more positive associations with White than Black). Now, three years later, it is not too soon to be seeking answers to basic questions about validity of measures constructed using the IAT procedure, making use of all available evidence, including that reported in articles in this special issue.

Validity of the IAT

Internal Validity

Greenwald et al. (1998) demonstrated that the IAT was free of a few sources of procedural artifact. In particular, the IAT effect measure was uninfluenced by whether the pleasant category was assigned to the left hand or right hand, by whether categories used in the IAT were composed of 5 or 25 items, nor was it affected by variations (ranging from 150 ms to 750 ms) in the response-stimulus (intertrial) interval.[1] Further, Greenwald et al. found that effects obtained with the IAT were relatively impervious to variations in the manner of treating data from incorrect responses and from non-normal response latency distributions.[2]

Subsequent research extended evidence for the IAT's internal validity by establishing that IAT effect measures are not influenced by wide variations in familiarity of items used to represent contrasted attitude-object concepts (Dasgupta, McGhee, Greenwald, & Banaji, in press; Ottaway, Hayden, & Oakes, in press; Rudman, Greenwald, Mellott, & McGhee, in press). However, it has also been found that using items at or near the low extreme of familiarity results in problematic IAT data. This is known from studies by Rothermund and Wentura (this issue) and Brendl, Markman, and Messner (in press), both of which used pseudo-categories composed of entirely unfamiliar nonsense words. Results obtained from this procedure appear to indicate that the pseudo-categories are associated with negative evaluation. Possibly, this is a valid conclusion – categories with meaningless members may indeed be affectively negative. Alternatively, however, it may be that the IAT does not function as intended when one or more of its four concepts is composed of nonsense items. Pending further clarification, it will be prudent to interpret results cautiously when using the IAT with categories composed of items that are entirely unfamiliar to subjects. The question of effects of items selected to represent categories is considered further under the heading, "Understanding how the IAT works," toward the end of this article.

In the first experiments with the IAT, Greenwald et al. (1998) reported an influence on IAT effects of the order of administering the two critical IAT tasks (i.e., the order of Steps 3 and 5 in Figure 1). Performance on either of these tasks tends to be faster when it is performed as Step 3 rather than as Step 5 in Figure 1's sequence. This procedural effect, which is the most noticeable internal validity problem of the IAT, has been accommodated in most subsequent research by counterbalancing the order of these two tasks.

---

[1]Mierke and Klauer (this issue) found that the difference between intertrial intervals of 100 ms and 1000 ms influenced IAT effects, with larger effects obtained for the shorter interval. However, their results came from a procedure that randomly ordered the two 2-choice discrimination tasks in Step 3 and Step 5 of the IAT procedure. It is possible that this result will not occur with the strict alternation of 2-choice discriminations that was used by Greenwald et al. (1998) and also in most subsequent IAT research.

[2]The data-handling procedure that was introduced by Greenwald et al. (1998), and since used in many other studies, is to (a) drop the first two trials of each block of critical trials (in Steps 3 and 5 of Figure 1), (b) deal with outlier trials by recoding latencies less than 300 ms to 300 ms and those greater than 3000 ms to 3000 ms, (c) log transform the resulting values prior to averaging, and (d) subtract the mean for Step 5 from that for Step 3 (or vice versa, depending on desired direction of scoring). Data may be discarded for subjects who appear to be outliers in response speed or error rate, although error trials are otherwise included in latency computations (and these latencies are scored as latency to the correct response, which is required on each trial, rather than to the initial response). This combination of procedures has worked quite satisfactorily, but should be recognized as a conventional scoring algorithm, rather than as a theoretically mandated scoring procedure.

Internal validity of IAT measures was demonstrated in another way by Kim and Greenwald (1998), who found that subjects were unable to comply with an instruction to respond to an IAT measure of implicit race attitudes so as to indicate lack of automatic preference for White relative to Black. Banse et al. (this issue) similarly found an inability of their heterosexual subjects to fake positive implicit attitudes toward homosexuality on their IAT measure.

In the course of operating a demonstration IAT website (http://www.yale.edu/implicit/) Nosek, Banaji, and Greenwald (2000a) obtained evidence concerning effects of several procedural and individual difference variables on IAT measures. Their results (see Table 1) indicated (a) a consistent order effect, such that associations appear stronger when tested in Step 3 of the Figure 1 procedure rather than in Step 5 ($\bar{r}$ = .20), (b) a correlation of the IAT effect with an IAT error measure, meaning that the combined task (Step 3 or Step 5) that produces faster responding also has a lower error rate ($\bar{r}$ = .38), (c) smaller IAT effects for website visitors who had taken more previous IATs ($\bar{r}$ = –.15), and (d) IAT effects increasing slightly in magnitude with subject age ($\bar{r}$ = .08). Of the procedural variables summarized in Table 1, only overall speed of responding was not consistently related to IAT effect magnitude ($\bar{r}$ = –.04). Not shown in Table 1 is an analysis that was possible for a portion of the website test takers who, before taking one of the race attitude IATs, were asked to describe themselves as left-handed, right-handed, or ambidextrous. Subject handedness was found to have essentially zero relation to magnitude of the race IAT effect.

Psychometric Properties

In order to be useful in correlational studies, the IAT must be sensitive to individual differences. The primary test for sensitivity to individual differences is test-retest reliability. Across several studies, test-retest reliabilities of IAT measures have averaged above r = .6 (e.g., Bosson, Swann, & Pennebaker, 2000; Dasgupta & Greenwald, in press; Greenwald & Farnham, 2000). In this issue, Banse et al. reported quite satisfactory internal consistencies of IAT measures (average α > .80), although they also mentioned that test-retest reliabilities were lower than anticipated.

In July 2000, the demonstration IAT website was modified to allow recording of individual trial data. For the four IAT measures that were available on the website between July 2000 and November 2000, split-half reliabilities of IAT effect measures ranged from r = .89 to r = .92 (*N*s ranged from 9,491 to 22,648).

Convergent validity

The case for any measure's construct validity can be advanced by establishing that it correlates with other measures of the same construct. For the IAT's measure of implicit attitudes the most suitable potentially converging measure is the measure of association strengths obtained from the affective priming procedure introduced by Fazio, Sanbonmatsu, Kardes, and Powell (1986). Cunningham, Preacher, and Banaji (in press) have provided the strongest test yet reported of the relationship between parallel IAT attitude and affective priming measures. Cunningham et al. obtained data from the same subjects at four weekly sessions for both priming and IAT measures of implicit racial attitudes. Their multiple measurement strategy allowed use of a confirmatory factor analysis, which revealed a correlation of r = .55 between latent factors representing priming and IAT implicit attitude measures. They also found satisfactory fit for a second-order model in

which the first-order priming and IAT factors loaded on a single higher-order factor. Rudman and Kilianski (2000) also found convergence between related IAT and priming measures in a study of attitudes toward female authority. Mellott, Cunningham, Rudman, Banaji, and Greenwald (2001) summarized data from these and three other studies that used both priming and IAT implicit attitude measures. A problem in portions of these data sets was that the priming measures occasionally had vanishingly low internal consistency, making it unlikely that they would correlate with any other measures. However, in all cases for which priming measures had at least moderate internal consistency, statistically significant correlations with IAT effects were observed. Related to this observation, Bosson et al. (2000) reported a low correlation between IAT and priming measures of implicit self-esteem, while also noting that their priming measures had very low in test-retest reliability. In sum, the evidence from studies that have included both affective priming and IAT measures of implicit attitudes has been varied due to occasional low reliability of priming measures, but nevertheless supportive of convergent validity for these two measures.[3]

Discriminant validity

Greenwald et al. (1998) found that correlations between IAT and self-report measures of implicit attitudes tended to be weakly positive, reporting 16 such correlations with an average *r* of .25. Bosson et al. (2000) and Greenwald and Farnham (2000) similarly reported that correlations between IAT and self-report measures of self-esteem were weakly positive, and several of the articles in this issue similarly reported low correlations between parallel IAT and self-report measures.[4] The most extensive data on relations between nominally similar IAT and self-report measures comes from the demonstration website mentioned previously (Nosek et al., 2000a). Summarized in Table 1, the relationship between implicit and self-report measures of attitude and stereotype can be described as consistently positive, while also being quite variable in magnitude, with observed *r*s ranging from *r* = .08 to *r* = .47. The data in Table 1 indicate that the relationship between implicit and explicit preferences varies widely across attitude domains. Consequently, the available data on correlations of IAT measures with explicit measures provide a mixture of evidence for discriminant and convergent validity. There is not yet an established interpretation of the cause of variability in correlations between implicit and explicit attitude measures. Three interpretations that have been suggested are that implicit-explicit correlations are reduced (a) when self-reports (i.e., explicit measures) are inaccurate due to impression management by subjects in responding to politically sensitive questions, (b) when such inaccuracy results from subjects' poor introspective access to their attitudes (cf. Greenwald et al., in press) or (c) when attitudes are so homogeneous in the population that low correlations result from restricted range (cf. Greenwald et al., 1998, p. 1470).

---

[3]Marsh, Johnson, and Scott-Sheldon (this issue) reported a lack of correlation between priming and IAT measures in their study of attitudes toward condoms. Because their IAT measures differed in several respects from those used in most other studies, the interpretation of these low correlations is uncertain.

[4]An exception is the report by Banse et al. (this issue) that "No evidence was found for the hypothesis that implicit and explicit attitudes should be regarded as distinct constructs" (p. xxx). Their data did indeed show that for the domain of homosexuality, correlations between their IAT and self-report measures were high, in the vicinity of *r* = .60. Nevertheless, it should be borne in mind that implicit-explicit correlations in most other studies have been substantially lower.

Predictive Validity

This paragraph very briefly summarizes studies in which IAT measures have been shown to correlate in expected fashion either with group membership classifications or with measures of individual differences within groups. Correlations of IAT measures with group memberships include observations of (a) differences between Japanese Americans and Korean Americans in implicit attitudes toward their respective ethnic groups (Greenwald et al., 1998), (b) differences between East and West Germans in implicit attitudes toward their geographical regions (Kuehnen, Schiessl, Bauer, Paulig, Poehlmann, & Schmidthals, this issue), (c) differences between men and women in their implicit gender self-concepts (women showing stronger association of self with expressive relative to instrumental traits; Greenwald & Farnham, in press), and (d) differences between men and women in their implicit attitudes and implicit self-concepts relating to mathematics (women showing stronger association of both positive valence and of self to arts relative to math; Nosek, Banaji, & Greenwald, 2000b). Studies demonstrating predicted correlations of IAT measures with individual differences within groups include (a) Greenwald et al.'s (1998) finding that strength of implicit ingroup preference for Japanese and Korean Americans was predicted by measures of their immersion in their respective Asian cultures, (b) Phelps et al.'s (2000) observation that IAT-assessed implicit race prejudice correlated with fMRI-assessed activation of the amygdala (a subcortical structure associated with emotional learning and evaluation) in White subjects exposed to unfamiliar Black faces, and (c) Greenwald, Banaji, Rudman, Farnham, Nosek, and Mellott's (in press) multiple findings of evidence for cognitive balance among triads of implicit (but not explicit) measures.

Correlations with Behavior

Through most of its history in the 20th century, the attitude construct faced difficulties in establishing evidence for its predicted relations with behavior. It is therefore of great interest to investigate correlations of IAT measures with measures of social and other behavior. Several studies have demonstrated expected differences in IAT measures between groups that are defined by behavioral differences. Swanson, Rudman, and Greenwald (in press) showed expected differences in IAT-measured self-concept between smokers and nonsmokers (in association of self with smoking) and also between vegetarians and omnivores (in association of self with type of diet). Banse, Seise, and Zerbes (this issue) showed expected differences between homosexuals and heterosexuals in implicit attitudes toward homosexuality. In data obtained from the IAT website, Nosek et al. (2000a) showed that self-described supporters of each of the two major candidates in the 2000 U. S. Presidential election (George W. Bush and Al Gore) differed as expected in IAT-measured implicit attitudes toward the two candidates. Correlations between IAT-measured attitudes or self-concepts and consumers' brand or product choices have been demonstrated by Maison, Greenwald, and Bruin (in press) and by Brunel, Collins, Greenwald, and Tietje (1999). In the first published clinical application of the IAT, Teachman, Gregg, and Woody (in press) demonstrated that IAT measures of implicit attitude, fear, disgust, and of implicit associations with danger discriminated high-snake-fear from high-spider-fear subjects. A combination of Teachman et al.'s four implicit measures correctly classified 92% of their subjects into the appropriate fear group (i.e., spider or snake) as assessed by self-report. In a follow-up in a treatment setting, Teachman and Woody (2000) found that the IAT was sensitive to the effects of treatment for spider phobia. In addition, their pre-treatment IAT was predictive of behavioral

approach to a spider at pre-treatment, and their post-treatment IAT predicted approach to a spider at a follow-up session two months later.

Other studies have demonstrated IAT-behavior correlations within college student samples. Rudman and Glick (in press) reported that prejudice against female job applicants was associated with IAT-assessed (but not explicit) gender stereotypes. Nosek et al. (2000b) found that IAT-measured implicit attitudes toward math correlated with SAT math performance. McConnell and Leibold (in press) showed 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.

Understanding How the IAT Works

Mierke and Klauer (this issue) have provided an overview of several suggested theoretical interpretations of cognitive mechanisms underlying the IAT effect. The various suggested interpretations are *figure-ground asymmetry* (Rothermund & Wentura, this issue), *criterion shift* (Brendl, Markman, & Messner, in press), *environmental associations* (Karpinski & Hilton, in press), *task-set switching* (Mierke & Klauer, this issue), and *acquired meaning of response keys* (De Houwer, 2000). In addition, a few studies have addressed the theoretical question of whether the IAT draws on associations at the level of *category features* or on ones at the level of *stimulus features* (De Houwer, in press; Steffens & Plewe, this issue).[5] Research on alternative theoretical interpretations has not yet progressed enough to establish any theoretical interpretation of the IAT effect. However, some of the available results are certain to figure importantly in the IAT effect's explanation. One of these is Mierke and Klauer's (this issue) demonstration that task-set switching is an important ingredient of the IAT effect, a result confirmed in an unpublished study by Sriram and Lee (1999). Another is De Houwer's (in press) demonstration that the IAT effect is determined more by the identity of categories used in the IAT than by the features of the stimuli used as category instances.

De Houwer's (2000) demonstration notwithstanding, several articles have shown that properties of instances can be important. As described earlier, both Rothermund and Wentura (this issue) and Brendl et al. (in press) have shown that pseudo-categories composed of nonsense items function in the IAT as if they had negative valence, even though it seemed plausible, a priori, that they should be neutral. Steffens and Plewe (this issue) additionally showed that IAT attitude measures functioned inappropriately when an irrelevant attribute (in their case, gender) was confounded with the valence attribute. Mitchell, Nosek, and Banaji (2000) have also shown that IAT attitude measures are affected when category contrasts contain dimensional confounds – in their research liked Black athletes and disliked White politicians could be classified either by race or by occupation.

---

[5]Because of the availability of the overview by Mierke and Klauer in this issue, these theories and the data bearing on them are not described here. One significant theory-relevant result not included in Mierke and Klauer's overview is the finding of Chee, Sriram, Soon, and Lee (2000) from a study that obtained fMRI data from subjects engaged in performance of an IAT measure of implicit attitudes toward flowers and insects. After finding high levels of activity in left dorsolateral prefrontal cortex during portions of the IAT that were opposed by strong associations (e.g., a combined task requiring the same response to insects and pleasant words), Chee et al. concluded that inhibitory processes were importantly involved in responding to the IAT.

In summary of the available evidence on effects of items selected as category instances in the IAT: The IAT appears relatively insensitive to choice of stimuli used to represent categories (best demonstrated by De Houwer, in press), so long as (a) the items well represent the categories, and (b) the items do not permit alternative construal of the discrimination the subject is to make (e.g., male-female rather than good-bad as in Steffens & Plewe, this issue; athlete-politician rather than black-white as in Mitchell et al., 2000).

Comments on Articles in This Issue

Wherever appropriate in earlier parts of this article, comments on other articles in this issue have been integrated with review of other published studies that provide evidence concerning the validity of the IAT. This section provides a few additional comments on the present issue's articles.

Kuehnen et al. (this issue) reported an interesting effect of priming the East-West (German) contrast and that contrast's relevance to self. Their priming procedure appeared to make pro-West thoughts more salient, as evidenced by both an increase in ingroup-favoring IAT effects for West Germans and a decrease in ingroup-favoring effects for East Germans. This indication of sensitivity of the IAT to a context manipulation is supported by the results of several other studies that will appear in a forthcoming special Attitudes and Social Cognition section of the *Journal of Personality and Social Psychology*. There is, accordingly, no doubt that IAT measures are susceptible to influence by a variety of independent variable manipulations. Establishing theoretical interpretations of these priming-like effects will be a challenge for the near future.

Marsh et al.'s (this issue) finding of correlations between IAT measures and reports of condom use with casual partners is, in part, a very welcome demonstration of potential practical usefulness of the IAT. At the same time, their IAT procedures and scoring methods differed enough from those used in most other IAT research so that enthusiasm about their findings should be limited to cautious optimism. The findings in the domain of sexual orientation by Banse et al. (this issue), which were obtained with procedures closely resembling IAT procedures in many other studies, are likewise very encouraging in regard to prospects for interesting applications.

Mierke and Klauer's (this issue) article elegantly fills a void in theorization and experimentation on cognitive mechanisms underlying the IAT effect. Their demonstration of the contribution of task-set switching to IAT effect measures justifies a decision that had been arrived at on a much more informal basis by Greenwald et al. (1998). That is, after obtaining results similar to those of Mierke and Klauer in pilot testing of potential IAT procedures, Greenwald et al. settled on a procedure for Steps 3 and 5 of the IAT that, in effect, capitalized on the task-set switching effect by alternating the two 2-choice discriminations on successive trials.

Conclusion

In three years of existence, the IAT has been exposed to moderate stress, as researchers (including several in this issue) have sought to identify limits on its usefulness and, on occasion, to suggest that the IAT may not be at all useful for its intended purposes of measuring association strengths (Bosson et al., 2000; Brendl et al., in press; Karpinski & Hilton, in press; Rothermund & Wentura, this issue). At the same time, a rapidly increasing body of published and in-press studies

(about 30 reviewed in this article) now indicates that the IAT often assesses individual difference dimensions for which self-report measures are insensitive. Consequently, there now exists a sizeable body of construct validity evidence for the IAT.

Results that establish limits on usefulness of the IAT can themselves be very useful. The limits so far identified concern the selection of stimuli that represent target or attribute concepts in IAT measures. As a consequence, it seems appropriate to suggest that the IAT either should not be used with categories composed of nonsense materials, or should be used in this fashion only with considerable caution in interpreting results (Brendl et al., in press; Rothermund & Wentura, this issue). Additionally, concept exemplars used in the IAT should not be selected so that they make it easy for respondents to generate alternative construals of the concept's identity (Mitchell et al., 2000; Steffens & Plewe, this issue).

Empirical stress-testing is very healthy for a procedure of the IAT's youth, and the IAT has obviously benefited from the research attention it has received. If the future holds similar good fortune, additional evidence for validity should be on the horizon, along with further understanding of the conditions and limits of successful use of the IAT. Three of the topics considered in the present overview are ones for which new developments will be especially welcome: (a) identification of cognitive mechanisms underlying the IAT effect, (b) identification of variables that moderate the strength of correlation between IAT and parallel self-report measures, and (c) identification of additional behaviors that are predicted effectively by IAT measures.

References

Banse, R., Seise, J., & Zerbes, N. (2001). Implicit attitudes towards homosexuality: Reliability, validity, and controllability of the IAT. *Zeitschrift für Experimentelle Psychologie*, 48 (this issue), xx-xx.

Blair, I., & Ma, J., & Lenton, A. (2000). *Imagining Stereotypes Away: The Moderation of Automatic Stereotypes Through Mental Imagery*. Manuscript under editorial review.

Bosson, J. K., Swann, W. B., & Pennebaker, J. W. (2000). Stalking the perfect measure of self-esteem: The blind men and the elephant revisted? *Journal of Personality and Social Psychology*, *79*, 631-643.

Brendl, C. M., Markman, A. B., & Messner, C. (in press). How do indirect measures of evaluation work? Evaluating the inference of prejudice in the Implicit Association Test. *Journal of Personality and Social Psychology*, in press.

Brunel, F. F., Collins, C. M., Greenwald, A. G., & Tietje, B. C. (1999, October). Making the private public, accessing the inaccessible: Marketing applications of the Implicit Association Test. Paper presented at meetings of the Association for Consumer Research, Columbus, Ohio.

Chee, M. W. L., Sriram, N., Soon, C. S., & Lee, K. M. (2000). Doroslateral prefrontal cortex and the implicit association of concepts and attributes. *NeuroReport*, *11*, 135-140.

Cunningham, W. A., Preacher, K. J., & Banaji, M. R. (in press). Implicit attitude measures: Consistency, stability, and convergent validity. *Psychological Science.*

Dasgupta, N., & Greenwald, A. G. (in press). Exposure to admired group members reduces automatic intergroup bias. *Journal of Personality and Social Psychology.*

Dasgupta, N., McGhee, D. E., Greenwald, A. G., & Banaji, M. R. (in press). Automatic preference for White Americans: Eliminating the familiarity explanation. *Journal of Experimental Social Psychology.*

De Houwer, J. (in press). A structural and process analysis of the Implicit Association Test. *Journal of Experimental Social Psychology*, in press.

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.

Greenwald, A. G., Banaji, M. R., Rudman, L. A., Farnham, S. D., & Mellott, D. S. (2001 [in press]). A unified theory of implicit attitudes, stereotypes, self-esteem, and self-concept. *Psychological Review.*

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.

Haines, E. L. (1999). *Elements of a Social Power Schema: Gender Standpoint, Self-Concept, and Experience*. Unpublished doctoral dissertation, City University of New York.

Karpinski, A., & Hilton, J. L. (in press). Attitudes and the Implicit Association Test. *Journal of Personality and Social Psychology.*

Kim, D-Y., & Greenwald, A. G. (1998, May). Voluntary controllability of implicit cognition: can implicit attitudes be faked?   Paper presented at meetings of the Midwestern Psychological Association, Chicago.

Kuehnen, U., Schiessl, M., Bauer, N., Paulig, N., Poehlmann, C., & Schmidthals, K. (2001). How robust is the IAT? Measuring and manipulating implicit attitudes of East- and West-Germans. *Zeitschrift für Experimentelle Psychologie*, 48 (this issue), xx-xx.

Maison, D., Greenwald, A. G., & Bruin, R. (in press). The Implicit Association Test as a measure of implicit consumer attitudes. *Polish Psychological Bulletin, 2*, xxx-xxx.

Marsh, K. L., Johnson, B. T., & Scott-Sheldon, L. A. J. (2001). Heart versus reason in condom use: Implicit and explicit attitudinal predictors of sexual behavior. *Zeitschrift für Experimentelle Psychologie*, 48 (this issue), xx-xx.

Mellott, D. S., Cunningham, W. A., Rudman, L. A., Banaji, M. R., & Greenwald, A. G. (in preparation). Do the IAT and priming measure the same construct? Evidence for the convergence of implicit measures. University of Washington.

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

Mierke, J., & Klauer, K. C. (2001). Implicit Association Measurement with the IAT: Evidence for an Effect of Supervisory Processes. *Zeitschrift für Experimentelle Psychologie*, 48 (this issue), xx-xx.

Mitchell, J. P., Nosek, B. A., & Banaji, M. R. (2000). Multitudes of attitudes: Automatic evaluation of multiply categorizable targets. Manuscript submitted for publication.

Nosek, B., Banaji, M. R., & Greenwald, A. G. (in press). Harvesting implicit group attitudes and beliefs from a demonstration website. *Group Dynamics.*

Nosek, B., Banaji, M. R., & Greenwald, A. G. (2000). Math = Male, Me = Female, therefore Math ≠ Me. Manuscript submitted for publication. Yale University, New Haven, CT.

Ottaway, S. A., Hayden, D. C. & Oakes, M. A. (in press). Implicit attitudes and racism: effect of word familiarity and frequency on the Implicit Association Test. *Social Cognition.*

Phelps, E. A., O'Connor, K. J., Cunningham, W. A., Gatenby, J. C., Funayama, E. S., Gore, J. C., & Banaji, M. R. (in press). Amygdala activation predicts performance on indirect measures of racial bias. *Journal of Cognitive Neuroscience*, *12*, 729-738.

Rothermund, K., & Wentura, D. (2001). Figure-ground asymmetries in the Implicit Association Test. *Zeitschrift für Experimentelle Psychologie*, 48 (this issue), xx-xx.

Rudman, L. A., Ashmore, R. D., & Gary, M. (1999). *Implicit and Explicit Prejudice and Stereotypes: a Continuum Model of Intergroup Orientation Assessment.* Manuscript submitted for publication.

Rudman, L. A., & Glick, P. (in press). Prescriptive gender stereotypes and backlash toward agentic women. *Journal of Social Issues.*

Rudman, L. A., Greenwald, A. G., & McGhee, D. E. (in press). Self-esteem and gender identity are manifest in implicit gender stereotypes. *Personality and Social Psychology Bulletin.*

Rudman, L. A., Greenwald, A. G., Mellott, D. S., & McGhee, D. E. (1999). Measuring the automatic components of prejudice: Flexibility and generality of the Implicit Association Test. *Social Cognition*, 17, 437-465.

Rudman, L. A., & Kilianski, S. E. (2000). Implicit and explicit attitudes toward female authority. *Personality and Social Psychology Bulletin*, *26*, 1315-1328.

Sriram, N., & Lee, I. (1999). The multiplicative effects of context switching, congruity, predictability, and dominance in speeded semantic classification. Unpublished manuscript, National University of Singapore.

Steffens, M. C., & Plewe, I. (2001). Items' cross-category associations as a confounding factor in the Implicit Association Test. *Zeitschrift für Experimentelle Psychologie*, 48 (this issue), xx-xx.

Swanson, J. E., Rudman, L. A., & Greenwald, A. G. (in press). Using the Implicit Association Test to investigate attitude-behavior consistency for stigmatized behavior. *Cognition and Emotion.*

Teachman, B. A., Gregg, A. P., & Woody, S. R. (in press). Implicit associations of fear-relevant stimuli among individuals with snake and spider fears. *Journal of Abnormal Psychology.*

Teachman, B. A., & Woody, S. R. (2000). Unpublished manuscript.

Concept Contrasts for the Gender Self-Concept IAT

	Math versus Arts		Male versus Female	
concepts	Math	Arts	Male	Female
items	math algebra geometry calculus equations	poetry art dance literature novel	male man boy he sir	female woman girl she lady

Sequence of Tasks for Gender Self-Concept IAT

Steps	Concepts for left response	Concepts for right response
1	Male	Female
2	Arts	Math
3	Male or Arts	Female or Math
4	Math	Arts
5	Male or Math	Female or Arts

Figure 1. **Illustration of the Implicit Association Test to measure gender self-concept**. The IAT starts by introducing subjects to the four concepts used in a series of 5 tasks. In this illustration, one pair of concepts is introduced in the first task by asking subjects to respond with left key to words representing *male* and with right key to words representing *female*. In the second task, the second pair of concepts is introduced, with subjects asked to respond left to words representing *arts* and right to words representing *math*. The third step introduces a combined task, in which words representing either *male* or *arts* get the left response and words representing either *female* or *math* get the right response. The fourth task reverses the first, and the fifth task combines the tasks of the 2nd and 4th steps. The IAT effect measure is constructed by comparing performance in the 3rd and 5th steps. If the subject responds more rapidly in the *male*-or-*math* vs. *female*-or-*arts task* than in the *female*-or-*math* vs. *male*-or-*arts* task, this indicates that, in combination, the *male-math* and *female-arts* associations are stronger than the *female-math* and *male-arts* associations.

Table1. Correlations of IAT Effects with Method Variables and Explicit Measures

Task	Order of critical tasks[a]	IAT error effect	Speed of responding	Prior IAT experience	Subject age	Explicit measure
Math/arts attitude	.21	.38	-.10	-.08	.11	.47
Race attitude (name stimuli)	.19	.31	-.07	-.14	.09	.24
Gender/science stereotype	.29	.49	-.03	-.09	.09	.21
Race attitude (face stimuli)	.14	.51	-.07	-.11	-.01	.17
Self-esteem	.20	.28	.05	-.19	.15	.17
Gender/career stereotype	.15	.28	-.04	-.13	.09	.17
Age attitude (face stimuli)	.15	.41	.01	-.21	.03	.13
Age attitude (name stimuli)	.24	.19	-.05	-.25	.11	.08
Average *r* (unweighted)	**.20**	**.38**	**-.04**	**-.15**	**.08**	**.21**

Note. Entries are product-moment correlations based on data provided by drop-in visitors to the IAT demonstration website (Nosek, Banaji, & Greenwald, 2000a; http://www.yale.edu/implicit/) between October 1999 and April 2000. Sample sizes range from $N = 28{,}108$ for Math/arts attitude to $N = 160{,}857$ for the race attitude measure with face stimuli. For the smallest sample size, any correlation greater than $r = .012$ is statistically significant. [a]Positive correlations indicate faster performance for a given pair of associations when it is tested in Step 3 of Figure 1's procedure rather than Step 5.