

Running head: ATTITUDES AND THE IAT

Attitudes and the Implicit Association Test

Andrew Karpinski

James L. Hilton

University of Michigan

Journal of Personality and Social Psychology, In Press

Abstract

To date, few studies have closely examined the relationship between the Implicit Association Test (IAT) and explicit attitudes. Across three studies, we examined the correlations between the IAT and explicitly measured attitudes where the measures were collected at multiple levels of specificity, across several different attitude domains, and across a wide range of attitude distributions. In all cases, the patterns of correlation that emerged support the view that the IAT is independent from explicit attitudes. In Study 2, we examined the relationships between the IAT, explicit attitudes, and behavior and again found evidence for the independence of the IAT and explicit attitudes. Although explicit attitudes predicted choice behavior in the second study, the IAT did not. Finally, in Study 3, we found that the IAT was affected by exposing participants to new associations between attitude objects, while the explicit attitudes remained unchanged. Taken together, these results support an environmental association interpretation of the IAT in which IAT scores are interpreted to reflect the extent to which various attitudinal objects are associated in a person's environment rather than the extent to which the person endorses the attitude objects.

Attitudes and the Implicit Association Test

Traditional models of attitudes assume that attitudes consist of three components: a cognitive component, an affective component, and a behavioral component. Moreover, these attitudes are thought to be open to conscious inspection, though their expression often depends upon their social desirability. In other words, traditional models assume that if you want to know someone's beliefs, feelings, and behavioral tendencies toward an object, all you need to do is measure his or her attitude—provided that you are not investigating a socially sensitive area. More recent models of attitudes, however, suggest that attitudes often exist outside of conscious awareness and control (Greenwald & Banaji, 1995). These “implicit attitudes” are thought to shape our automatic reactions to attitude objects and to thereby shape our subsequent interactions with them.

But how should we conceive of the relationship between implicit and explicit attitudes? One possibility is that implicit and explicit attitudes reflect a single attitudinal construct. According to this view, attitudes are similar to icebergs, with explicit attitudes residing above the surface of conscious control and implicit attitudes residing below. Because implicit and explicit measures tap a single attitudinal construct, albeit in different places, one implication of this view is that, given the right conditions, implicit attitudes, explicit attitudes, and attitude related behaviors should all correlate

A second possibility, however, is that implicit and explicit attitudes are independent from each other. For example, Wilson and colleagues argue that we may have dual attitudes toward objects, one implicit and one explicit (Wilson, Lindsey, & Schooler, in press). From this perspective, the IAT and explicit attitude measures tap different underlying constructs. As such, one implication of this view is that correlations between the IAT and explicit attitude measures

should be low to non-existent, and these two measures may predict completely different aspects of behavior.

The available evidence on this relationship is mixed (cf. Wittenbrink, Judd, & Park, 1997; von Hippel, Sekaquaptewa, & Vargas, 1997; Dovidio, Kawakami, Johnson, Johnson, & Howard, 1997). Some evidence suggests that implicit and explicit attitude measures tap the same underlying attitude, while other research and theorizing indicate that implicit and explicit attitude tests measure different underlying constructs. A primary goal of the research reported here was to investigate the relationship between implicit and explicit attitudes, focusing especially on the IAT.

The Implicit Association Test

Recently, Greenwald, Banaji, and their colleagues have proposed in a series of articles, conference presentations, and interviews that the implicit association test (IAT) is a measure that taps these implicit attitudes (Greenwald, McGhee, & Schwartz, 1998; Greenwald, 1998; Chamberlain, 1998; Schwartz, 1998, 1999). The IAT is thought to measure implicit attitudes by examining the automatic associations between various attitude objects and various evaluative attributes (see Greenwald et al., 1998). Specifically, the IAT measures how closely associated any given attitude object (e.g., a flower or an insect) is with an evaluative attribute (e.g., pleasant or unpleasant words) and assumes that the more closely related the objects and attributes are, the stronger the implicit attitude is.

Consider, for example, an IAT experiment designed to measure attitudes toward insects and flowers. The IAT involves five stages of activity. In the first stage, participants categorize target words that are relevant to the attitude objects. In the case of insects and flowers, this means that participants first categorize words (e.g., cockroach and rose) as either insects or flowers by

pressing a key with their left hand if it is an insect and pressing a different key with their right hand if it is a flower. In the second stage, participants categorize a different set of words (e.g., happy and rotten) as either pleasant or unpleasant by pressing a key with their left hand if it is an unpleasant word and pressing a different key with their right hand if it is a pleasant word. These first two stages are learning stages where participants become familiar with the categorization tasks. In the third stage, the previously learned categorizations are combined. Participants are instructed to press a key with their left hand if any given word is either an insect or an unpleasant word and to press a different key with their right hand if any given word is either a flower or a pleasant word. In the fourth stage, the response keys are reversed. Participants now must press a key with their right hand if the word is an insect and press a different key with their left hand if the word is a flower. In the final stage, the new attitude object categorization practiced in stage four is combined with the categorization of the evaluative attributes learned in stage two. In the case of insects and flowers, participants are instructed to press a key with their left hand if any given word is either a flower or an unpleasant word and to press a different key with their right hand if any given word is either an insect or a pleasant word.

An overall IAT score is obtained by taking the difference in response times between the two combined stages. Individuals who respond more quickly when pleasant and flower are paired together on the same response key than when pleasant and insect are paired together are said to have more positive associations toward flowers than insects. Conversely, individuals who respond more quickly when pleasant and insect are paired on the same response key than when pleasant and flower are paired together are said to have more positive associations toward insects than flowers.

What consistently emerges in the IAT is that people are faster to respond when generally liked items are paired with positive words than when generally disliked items are paired with positive words. The IAT reveals that people have more positive associations with flowers than with insects and with musical instruments than with weapons (Greenwald et al., 1998). These results appear to provide evidence for a non-reactive measure of people's attitudes that can generalize across attitude objects. Simply substitute gender-related terms and you have a measure of gender bias (Rudman, Greenwald, & McGhee, 1996). Substitute age-related terms and you have a measure of ageism (Nosek, Greenwald, & Banaji, 1998). Substitute race-related terms and you have an implicit measure of racial prejudice (Greenwald et al., 1998).

To date, the IAT has generated enormous interest in both the general public and the scientific community. An IAT website (Nosek et al., 1998) was unveiled shortly after the first IAT paper appeared in press. At this website anyone can take several IATs that reportedly measure unconscious levels of racism, ageism, gender bias, and self-esteem. By January of 2000, over 500,000 IATs had been completed on the IAT website (Nosek, Cunningham, Banaji, & Greenwald, 2000) and both ABCNews.com (Chamberlain, 1998) and the Associated Press (Tibbets, 1998) had published articles on the IAT. Among the scientific community the IAT is of interest because it is easy to administer, robust, and produces large effect sizes, particularly in comparison to other implicit measures (Greenwald et al., 1998). Indeed, although the IAT has appeared in relatively few peer-reviewed articles, the website lists over forty researchers who are currently using the IAT.

The Relationship between the IAT and Explicit Attitude Measures

From Greenwald et al.'s (1998) initial paper, it is possible to find evidence supporting both the unitary and independent models of the relationship between implicit and explicit attitudes.

On the one hand, collapsing across individuals the findings that typically emerge in IAT studies converge with explicit attitude measures. Traditional attitudinal measures indicate that people have more favorable attitudes toward flowers than insects and toward musical instruments than weapons, and the IAT reveals these same patterns of results. On average, there is a closer association between flowers or musical instruments and pleasant words than between insects or weapons and pleasant words (Greenwald et al., 1998, experiment 1). Similarly, explicit attitude measures reveal that White people, on average, express more favorable explicit attitudes toward White targets than toward Black targets, and the IAT reveals that, on average, there is a closer association between White names and pleasant words than for Black names and pleasant words (Greenwald et al., experiment 3).

On the other hand, other IAT results suggest that the IAT and explicitly measured attitudes are independent constructs. Although the group-based averages on the IAT and explicit attitude measures typically correspond, the correlational data suggest that the two constructs are independent at the level of the individual. Across three studies and multiple IATs, Greenwald et al. found no consistent correlations between the IAT and explicitly measured attitudes (Greenwald et al., 1998).

If the IAT and explicit attitudes are measuring independent constructs, then what is the IAT measuring? Again, we can find some preliminary cues from Greenwald et al.'s (1998) initial studies. In Study 2, Korean-American participants showed an IAT bias for Korean names over Japanese names while Japanese-American participants showed a bias for Japanese names over Korean names. Interestingly, this effect was moderated by the degree to which participants were immersed in Asian culture. The more a Korean participant was immersed in Asian culture, the greater bias he or she showed for Korean names. Similarly, the more a Japanese participant was

immersed in Asian culture, the greater bias he or she showed for Japanese names. People who are more immersed in Asian culture are presumably more likely to be exposed to the knowledge of that culture, including its stereotypes. Perhaps it is this immersion that the IAT detects, rather than the extent to which a person consciously or unconsciously endorses cultural stereotypes.

These results suggest an environmental association model of IAT effects. Whereas explicit attitudes assess an individual's level of endorsement toward an attitude object, the IAT may tap the associations a person has been exposed to in his/her environment. According to the environmental-associations model of the IAT, a high score on a White-Black IAT, for example, should not be seen as indicating that the individual has more favorable evaluations of Whites compared to Blacks. Instead, the score may simply indicate that the individual has been exposed to a larger number of positive-White and negative-Black associations than negative-White and positive-Black associations. Given the high levels of exposure to and awareness of cultural stereotypes (Devine, 1998), perhaps it is not surprising that most of the participants in Greenwald et al.'s (1998) Study 3 showed an IAT bias against Blacks.

What this analysis suggests is that the relationship between the IAT and explicitly measured attitudes remains open to investigation. The goal of the three studies reported here was to examine this relationship more thoroughly. In these studies we first looked at the correlations between the IAT and explicitly measured attitudes in a fine-grained analysis. Next, we examined the relationship between the IAT, explicitly measured attitudes, and behavior. Finally, we specifically investigated the environmental associations model of the IAT by examining changes in the IAT and explicitly measured attitudes in response to contextual changes in the environment.

Studies 1a and 1b

As previously mentioned, Greenwald et al. (1998) found no reliable correlation between the IAT and explicit attitude measures. There are, however, several reasons to question the lack of correlations that emerged. Several of the studies measured attitudes toward racial or ethnic groups (Experiments 2 & 3). Failed correlations in these studies are not surprising because of social desirability concerns. A great deal of social psychological literature indicates that people frequently respond in a socially acceptable manner when explicitly reporting attitudes toward various ethnic groups (Devine & Elliot, 1995). If participants censored their true attitudes toward the ethnic groups on the explicit measures, then we would not expect the IAT and the explicitly measured attitudes to correlate. The IAT could well reflect the participants' "true attitudes" while the explicit attitude measures respond primarily to social desirability concerns.

In addition, two methodological issues may have prevented the correlations from emerging even in domains where social desirability concerns were absent (Greenwald et al., 1998, study 1). First, the IAT and the explicit attitudes were measured at different levels of specificity and this mismatch may have masked the underlying correlations (Fishbein & Ajzen, 1974). Specifically, for the explicit measures participants rated the terms “flower” and “insect” on semantic differentials and feeling thermometers. The ratings for “insect” were then subtracted from the ratings of “flower” to obtain an explicit measure of attitudes toward flowers versus insects. In other words, these explicit attitude measures were obtained by having participants make category-level responses. In contrast, the IAT required participants to respond at the level of the individual item. Specific target words (e.g., rose or wasp) were categorized as either being a flower or an insect and an IAT score was obtained by averaging across the responses to the insect target words and subtracting that time from the average response time to the flower target

words. Thus, the IAT measure was obtained by examining responses to specific category exemplars. As a result of this mismatch in the level of responding, the expected correlations between the IAT and explicitly measured attitudes may have failed to emerge simply because the explicit attitude measures were taken at too broad of a level.

Second, Greenwald et al. (1998) may not have found a correlation between the IAT and explicit measures because the explicit ratings of flowers and insects may have been restricted in their range. Most participants liked flowers much more than insects, and there was little variability in the reports of these attitudes. As Greenwald et al. noted, this restricted variability of the explicitly measured attitudes may have prevented the expected correlations from emerging.

In order to begin examining the relationship between the IAT and explicit attitudes more closely, we first conducted two studies. In both studies, participants first completed a flower-insect IAT and then rated flowers and insects on an explicit rating scale. The IAT and explicit ratings that we used were designed to capture data at the level of the individual item (e.g., rose) and at the level of the category (e.g., flower). In the first study (1a), the explicit measures consisted of valence ratings of the categories (e.g., flower and insect) and of each of the target words (e.g., rose and wasp) used in the IAT. These valence ratings were made on a scale from "extremely positive" to "extremely negative." In the second study (1b), participants rated the categories and target words on feeling thermometers, ranging from "cold or unfavorable" to "hot or favorable". We substituted the feeling thermometer ratings for the valence ratings used in the first study in order to obtain a second, and perhaps more personal, judgement of the attitude objects. In addition, we changed one of the insect target words from "spider" to "fly". Several participants in Study 1a correctly pointed out that a spider is not actually an insect. By obtaining both IAT scores and explicit attitude measures at several levels, our primary goal in Studies 1a

and 1b was to replicate the results of Greenwald et al. (1998) and to extend the analyses in ways that rule out the methodological and statistical concerns that were raised.

Method

Participants

Forty-three and 28 students enrolled in an introductory psychology course at the University of Michigan participated in Study 1a and 1b, respectively. All participants received course credit for their participation.

Materials & Apparatus

The materials and procedures of this study closely resembled those used by Greenwald et al. (1998). In particular, we borrowed five pleasant words (cheer, pleasure, happy, love, and peace), five unpleasant words (death, filth, jail, murder, and ugly), five flower names (carnation, daisy, lily, rose, and tulip), and five insect names (ant, cockroach, maggot, spider [1a only], fly [1b only], and wasp).

The participants completed the IAT on a 7200 Power Macintosh computer with 15-inch color monitors.

Procedure

The participants were tested in groups of up to six individuals at a time. An experimenter greeted the participants and asked them to first read and sign written consent forms. Each participant was seated in a small cubicle containing only a computer. All instructions were presented both verbally and in writing.

The IAT. In presenting the IAT to the participants, we followed the methodology outlined by Greenwald et al. (1998), with a few exceptions. First, the IAT computer program used by Greenwald et al. computes average IAT scores across target words within a category. Because

we wanted to examine the relationships between IAT scores and the explicit attitude measures at the category level and on an item-by-item basis, we reprogrammed the task using Psyscope (Cohen, MacWhinney, Flatt, & Provost, 1993) in order to gain access to response times for individual target words.. Second, because Greenwald et al. found that response keys assigned to pleasant items, category set size, and the interval between the response and the next item did not make a difference in the IAT scores, we fixed these procedural variables at one level for all participants. Third, Greenwald et al. found that the order of compatibility conditions with the IAT (i.e., flowers paired with pleasant words first or flowers paired with unpleasant words first) did have an effect on IAT scores. Smaller IAT effects were obtained when flowers were first paired with unpleasant words (the incompatible combination). Because the direction of the effect did not change, however, we did not manipulate this factor. Instead, all participants completed the compatible condition first. Fourth, to reduce any order effects resulting from the order of presentation of the combined tasks and to reduce fatigue, we reduced the total number of target word presentations. Greenwald et al. presented two blocks of 50 trials for the single-categorization blocks, and four blocks of 50 trials for the critical combined trials. Participants in our task responded to one block of 40 trials for the single-categorization blocks, and one block of 80 trials for the critical combined trials. Fifth, Greenwald et al. provided feedback to their participants when the participants made incorrect categorizations. Because error rates in the IAT are typically quite low, and because participants report knowing when they have made an error even without feedback, we did not provide feedback for incorrect responses. Sixth, participants in our study completed the evaluative attribute discrimination (i.e., pleasant versus unpleasant discrimination) before they completed the target-concept discrimination (i.e., flower versus insect discrimination). Because the first two stages are essentially practice sessions, this change,

like the others, should not make any difference.

In all other respects, we followed the methods outlined by Greenwald et al. (1998). Specifically, participants completed five stages in the following order: (1) initial evaluative attribute discrimination (pleasant versus unpleasant), (2) initial target-concept discrimination (flower versus insect), (3) initial combined task (pleasant + flower versus unpleasant + insect), (4) reversed target-concept discrimination (insect versus flower), and (5) reversed combined task (pleasant + insect versus unpleasant + flower). In each stage, participants saw each target word four times, and the target words were selected randomly, without replacement. Thus, there were 40 stimulus word presentations for stages one, two, and four, and 80 stimulus word presentations for stages three and five. The first two presentations of each target word in a stage were considered practice trials and were not included in the analyses.

Each stage was preceded by a set of instructions concerning the dimension(s) of the categorization task and the appropriate key response. If the target word was a member of the category listed on the left side of the screen, the participants were to respond with the A key. If the target word was a member of the category listed on the right side of the screen, the participants were to respond with the 5 key on the numeric keypad. Once participants read the instructions, they were instructed to proceed with the task. Each target word appeared centered on the screen with category reminder labels appropriately positioned on the left or right sides of the screen. The target word remained on the screen until the participants responded. The interval between the response and the next item was 30 milliseconds.

Explicit attitude measures. After completing the IAT, the participants completed an explicit measure of their attitudes toward the categories (e.g., flower and insect) and target words (e.g., rose and wasp) used in the IAT. Participants in Study 1a were asked to rate how positive or

negative they found each target word on a scale from -100 (Extremely negative) to 100 (Extremely positive). Participants first rated eight practice words (lucky, abuse, violet, fly, sunrise, pollute, daffodil, and termite). The four category words and 20 target words from the IAT were then presented in a random order. Participants in Study 1b rated each word on a feeling thermometer rather than on a valence scale. Specifically, participants were asked to rate their general level of warmth or coolness toward each word on a feeling thermometer ranging from 0 (cold or unfavorable) to 100 (hot or favorable). The midpoint, 50, was labeled neutral. Finally, the participants were thoroughly debriefed about the experiment and thanked for their participation.

Results

Analysis of the Error Rate

Recall that our participants did not receive an error message when they gave an incorrect response. Despite this fact, across both studies the error rate was nearly identical to the error rate of Greenwald et al. (1998) -- about 4%. For Study 1a, the data from one participant were excluded from all analyses due to an error rate of 26%. With this participant removed, the error rate averaged 4% percent and ranged from 0% to 11%. For Study 1b, all data were used in the analyses because no participant showed an excessively high error rate. In this study, the error rate on the IAT averaged 4% and ranged from 0% to 10%.

The IAT Measure

In each stage, the relevant target words were each presented four times. However, two of these presentations were practice. Using the non-practice responses, we computed an IAT score for each participant. To compute this score, we examined the ten target words for each category within each stage (five words, each presented twice) and took the median of those ten words. We

then averaged the median reaction time responses within each stage and subtracted the stage three times from the stage five times.

Alternative methods of calculating IAT scores include the use of truncated means and log-transformed means (see Greenwald et al., 1998). However, a median-based approach has several advantages. First, median-based measures are not sensitive to the skewness of a participant's time distribution at each stage, and they are robust to outliers. As a result, it is not necessary to create arbitrary cut-off points for the exclusion of outlying observations. When medians are used, all original data points may be included in the analysis. Second, because transforming means is a relatively extreme data-analytical technique, it should only be used when alternate methods of data-analysis break down. Although the stage three and stage five response times in our studies were skewed, the IAT scores (the difference in the response times at stage three and stage five) were not skewed, indicating that a transformation was not necessary. Median-based measures also have the virtue of being easy to interpret. In contrast to transformed means, which must be re-transformed for interpretation, median based scores can be interpreted in their raw form. For all of these reasons, all of the IAT results that are reported in this paper were based on medians. However, we also analyzed the data using both log-transformed and truncated means (Greenwald, et al., 1998). Except where noted, all three types of analyses yielded similarly significant results.

The IAT scores for each stage are listed in Table 1. In both Studies 1a and 1b, participants' IAT scores exhibited a bias in favor of flowers, all F 's > 124 , p 's $< .001$, d 's > 1.72 (see Table 2). These findings replicate those of Greenwald et al. (1998). Participants were faster when responding to the congruent pairings (flower + pleasant and insect + unpleasant) than when responding to the incongruent pairings (insect + pleasant and flower + unpleasant).

Explicit Attitude Measures

The primary purpose of the present study was to examine the relationship between the IAT and explicit attitudes. Recall that the main difference between Studies 1a and 1b was how explicit attitudes were accessed. In Study 1a, participants rated the valence of all the categories and target words used in the IAT, while in Study 1b participants rated each word on a feeling thermometer.

Study 1a

We combined the valence ratings in two ways to create different levels of explicit attitude scores. First, following Greenwald et al. (1998), we created an explicit attitude measure based on category level judgements (i.e., insect and flower) by subtracting ratings of the category-word “insect” from ratings of the category-word “flower” (valence [category]). Second, we created an explicit attitude measure based on the ratings of the individual target words within each category (valence [item]) by subtracting the average valence ratings of the five insect words (e.g., fly and maggot) from valence ratings of the five flower words (e.g., rose and lily). This item-based explicit attitude measure should have a better chance of correlating with the IAT because participants make responses for both valence (item) and the IAT at the level of the target word.

The data for both explicit measures are contained in Table 2. Like the IAT, these explicit measures of attitudes are comparative scores examining the difference between ratings or responses to flowers and insects, with positive numbers indicating a bias for flowers over insects. Consistent with Greenwald et al. (1998), the IAT and explicit attitude measures both reveal a bias in the direction of flowers.

Study 1b

The feeling thermometer ratings obtained in Study 1b reveal a similar picture. Again, two measures of explicit attitudes were created from the feeling thermometer ratings. First, we created an explicit attitude measure based on category level judgements (i.e., insect and flower) by subtracting the thermometer ratings of the category insect from ratings of the category flower, FT (category). FT (category) is identical to the feeling thermometer measure used by Greenwald et al. (1998, Study 1). Second, we averaged the feeling thermometer ratings of the five target insect words and subtracted it from the average of the feeling thermometer ratings of the five target flower words to create a feeling thermometer score based on the ratings of the individual items, FT (item). The means are presented in Table 2, where it can be seen that participants had more favorable explicit attitudes toward flowers than insects. Again, the means indicate that, on average, participants showed both an IAT “preference” for flowers and an explicit preference for flowers as measured by the feeling thermometer.

Correlations between the IAT and Explicit Attitude Measures

In examining the direction and magnitude of the IAT and explicitly measured attitudes across Studies 1a and 1b, it would be tempting to conclude that the IAT and explicit measures correspond. However, in order to test whether those with higher IAT scores are the ones with higher explicit attitude scores, it is necessary to conduct a correlational analysis of the relationship between the IAT and the explicit attitude measures.

If the IAT and the explicit attitude measures tap the same attitudinal construct, then we would expect to find a positive correlation between the two measures. The correlations between the IAT scores and the explicitly measured attitudes are presented in Table 3. Although there was a strong positive correlation between the two explicitly measured attitudes ($r = .82$ and $r = .62$,

for Studies 1a and 1b respectively), there was no relationship between the IAT and either explicit measure of attitude. The correlations ranged between $r = -.11$ and $r = -.02$ in Study 1a and between $r = -.31$ and $r = -.19$ in Study 1b.

One possible explanation for the lack of correlation between the IAT and the explicit attitude measures is that the attitude measures were not specific enough. When the IAT and explicit attitude scores were calculated, we averaged across the individual target words. Perhaps the process of averaging across all the target words masked the relationship between the IAT and the explicit measures. To examine this possibility, we analyzed the relationship between the IAT for each target word separately. Specifically, for each of the 20 target words used in the IAT, we calculated an item-based IAT score. For example, by subtracting the average of the two non-practice stage three (consistent pairing) responses to the target word “rose” from the two non-practice stage five (inconsistent pairing) responses to the word “rose,” we obtained an item-based IAT score for rose ($IAT[rose]$). This process was repeated for all 20 target words. IAT scores obtained in this manner are not comparative scores like the standard IAT. Instead, each item-based score measures the degree to which a person has associations between the pleasant/unpleasant evaluative dimension and the target word. For flower and pleasant target words, higher scores indicate that a person has more positive than negative associations with the target word, while for insect and unpleasant target words lower scores indicate that a person has more positive than negative associations with the target word. These item-based scores were then correlated with the participants’ valence ratings of each of the target words (e.g., $IAT[rose]$ was correlated with the participants’ valence ratings of rose).

If the IAT is related to our explicit attitude measures, then we would expect to find correlations between the item-based IAT scores and the valence ratings of the individual items.

For flower words, we would expect these correlations to be positive. Positive numbers on the item-based flower IAT indicate a greater number of pleasant than unpleasant associations with that particular flower and a positive valence rating indicates a more favorable attitude toward that flower. For insect words, we would expect the correlations to be negative because positive numbers on the item-based insect IAT indicate a greater number of unpleasant than pleasant associations with that particular insect, while positive numbers on the valence ratings again reflect a favorable attitude toward that particular insect. By the same logic, we would expect a positive correlation between item-based pleasant IAT scores and valence rating of those pleasant words, but a negative correlation between item-based unpleasant IAT scores and valence rating of those unpleasant words.

The correlations are presented in Table 4. Despite the fine-grained level of this analysis, we find no evidence for a correlation between the IAT and explicit attitude measures. Of the 40 hypothesized correlations, fewer than half are in the predicted directions, one reached significance in the predicted direction, and two correlations (for tulip and cheer in Study 1b) were significant in the opposite direction.

Was There a Restriction of Range Problem?

As mentioned in the introduction, one reason why Greenwald et al. (1998) might have failed to find a correlation between the IAT and explicit attitude measures may have been statistical in nature. If all participants have similar scores on a variable that is being correlated with another variable, this restriction of range of the variable can reduce the absolute value of the correlation.

In Study 1a, there was no evidence for a restriction of range problem. First, we examined the range of the item-based IAT scores. These scores displayed a minimum range of 1500.5 ms for flowers (from -272ms to 1228.5ms, IAT [daisy]) and 928ms for insects (from -216ms to 712ms,

IAT[cockroach]). In addition, the overall IAT scores ranged from -19.75ms to 441.50ms. For the valence ratings, a similar picture emerged. The item-based valence ratings of all flowers and insects covered a minimum of 50% of the possible range. In addition, 47.5% of the possible range was covered by the valence (category) measure.

Analysis of Study 1b revealed the same findings. In this case, the item-based IAT scores displayed a minimum range of 722.75ms for flowers (IAT[rose]) and 1006.50ms for insects (IAT[maggot]). Furthermore, the overall IAT scores ranged from 22.5ms to 424.08ms. For the feeling thermometer ratings, the minimum range covered for the item-based flower ratings was 50% of the scale (IAT[lily]) and 35% for the insect item-based ratings (IAT[maggot]). Similarly, the category-based feeling thermometer covered 65% and 41% of the rating scale for flowers and insects, respectively.

These analyses reveal a significant amount of variability in the IAT scores and the explicit attitude ratings at both the item and category level. Despite this variability, correlations between the IAT and the explicit measures failed to emerge.

Discussion

The results from Studies 1a and 1b suggest that explicit attitudes and the IAT are independent and they highlight a problem with relying on means to infer relationships among variables. Looking only at the means for the IAT and explicit attitudes, it would be easy to conclude that they are related. After all, both measures appear to reveal the same evaluative tendencies. Participants indicate that they like flowers more than they like insects on both the IAT and on the explicit measures. But when we examined the correlations, we found that the IAT did not correlate with the participants' explicit attitudes.

Importantly, these correlations failed to emerge under fairly ideal conditions. First, the study was conducted in a domain where there is little reason to believe that participants would have been concerned with controlling their explicit attitudes. Unlike studies involving prejudice, where it would be easy to imagine the correlations failing to emerge because participants monitor the expressions of their explicit attitudes, this study involved attitudes toward flowers and insects. Second, the correlations failed to emerge despite the fact that we calculated them at multiple levels of specificity (i.e., at the level of the category, averaged across items, and at the level of the individual item). It is unlikely that the correlations failed to emerge because of a mismatch between the level of responding tapped by the IAT and the level of responding tapped by the explicit measures. Third, the correlations failed to emerge despite ample variability in the participants' responses. Finally, we replicated the results in two independent samples.

Although the correlational data from the first pair of studies suggest that explicitly measured attitudes and the IAT are independent, two findings would make this case stronger. First, despite the considerable variability that emerged in the IAT and explicit ratings of flowers and insects, it is still true that a majority of our participants in the first studies preferred flowers over insects. This general preference for flowers over insects still leaves non-significant correlational findings open to interpretations based on range restrictions. A greater range of attitude scores would eliminate this alternative interpretation. Second, the case for independence would be stronger if it could be shown that explicit attitudes and the IAT differ in their ability to predict behavior.

Study 2

With these concerns in mind, the goals of the second study were twofold. The first goal was to replicate the correlational findings of the first study in an attitude domain where we expected to find a greater range of attitudes. The second goal was to examine the relationships among the

IAT, explicit attitude measures, and behavior. To accomplish these goals, we moved away from the domains of flowers and insects and into the domain of consumer choice behavior.

Specifically, we were interested in how well the IAT and explicit attitude measures would predict participants' behavior when they were given a choice between two foods: an apple and a candy bar. To these ends, participants in the second study took an apple/candy bar IAT, completed a set of explicit attitude measures concerning apples and candy bars, and then chose either an apple or a candy bar to eat.

There are three important aspects of this design. First, our goal was to examine a consequential behavior that would be amenable to implicit influence. In line with this goal, Dovidio and colleagues (Dovidio et al., 1997) found that attitudes measured implicitly tend to predict spontaneous or non-verbal behaviors, while those measured explicitly tend to predict deliberative behaviors. In the realm of consummatory behaviors, researchers have found that people only engage in a thoughtful, deliberative consideration of their attitudes for high-involvement purchases (Herr, 1995). Because the choice of apple versus candy bar is a relatively spontaneous choice that does not involve a great deal of personal involvement, it seemed reasonable to assume that it could be influenced by implicit attitudes.

Second, an interesting aspect of this study is that both attitude objects (i.e., apples and candy bars) are likely to be positively valenced for most participants. With oppositely valenced attitude objects there is a clear consistent pairing (e.g., flower + pleasant and insect + unpleasant) and a clear inconsistent pairing (e.g., insect + pleasant and flower + unpleasant) that the IAT measures. With an IAT designed to measure associations toward apples and candy bars, there is no obvious consistent or inconsistent pairing. As a result, the second study provides a test of whether the IAT can be used in situations where both of the attitude objects are similarly valenced.

Third, there is a potential problem with too simple a design. If participants complete the explicit attitude measures before they make their choice, they may feel that because they just reported their attitudes toward apples and candy bars, they should behave in a manner consistent with those attitudes. If the IAT were to fail to predict behavior in this situation, it could be due to the fact that the participants first reported their attitudes explicitly. To investigate this possibility, we included two conditions. In one condition (IAT + explicit measures), participants completed an apple/candy bar IAT and explicit attitude measures before choosing between the apple and candy bar. In the other condition (IAT only), participants only completed the IAT before choosing between the apple and candy bar. The design allowed us to conduct separate analyses for each condition to examine if implicit and explicit attitudes predict choice behavior.

Method

Participants

Eighty-five students enrolled in an introductory psychology course at the University of Michigan participated in this experiment. All participants received course credit for their participation.

Procedure

Participants were run in groups of up to six at a time. Each group was randomly assigned to be in either the IAT + explicit measures or the IAT only condition. The participants first completed an apple/candy bar IAT. The categories pleasant and unpleasant and their associated target words were the same as those used in Studies 1a and 1b. Five target words were selected to be representative of the concept apple: red, cider, pie, Red Delicious, and Macintosh. Five target words were also selected to be representative of the concept candy bar: chocolate, peanuts,

wrapper, Snickers, and Hershey's. Otherwise, the procedure of the IAT was identical to that of Studies 1a and 1b.

An independent sample of 30 individuals rated how well each of these target words was representative of the concepts of apples and candy bars. Participants rated each target word on a scale from -3 (Strongly associated with apples) to +3 (Strongly associated with candy bars). The neutral point was labeled "Associated with neither apples nor candy bars." The results of the pretest indicate that all five apple target words were strongly associated with apples ($M_{\text{apple}} = -2.57$, $F(1,29) = 883$, $p < .001$) and all five candy bar words were strongly associated with candy bars ($M_{\text{candy bar}} = 2.47$, $F(1,29) = 740$, $p < .001$).

After completing the IAT, those participants in the IAT + explicit measures condition completed the explicit measures. Participants completed a semantic differential and a feeling thermometer measure for both apples and candy bars. The feeling thermometer measure was identical to the one described in Study 1b, with the substitution of the categories “apple” and “candy bar”. Participants only rated the category words, and not each of the individual target words. For the semantic differential, participants rated “apple” and “candy bar” on five bipolar dimensions: ugly-beautiful, bad-good, unpleasant-pleasant, foolish-wise, and awful-nice. Each dimension was rated on a 7-point scale ranging from -3 (the negative pole) to +3 (the positive pole) and participants were instructed to circle zero if the anchoring adjectives were irrelevant to the concept. Afterward, participants completed questions concerning how much they liked eating apples and candy bars, how often they ate apples and candy bars, and if given a choice between an apple and a candy bar, which they would choose. Each of these questions was answered on an 8-point scale ranging from 1 (strongly disagree) to 8 (strongly agree).

All participants were then presented with a “Fun-size Snickers”® candy bar and a “Red Delicious” apple. They were informed that they could choose one and only one of these objects to eat or to take home with them. The participants were then thoroughly debriefed about the experiment and thanked for their participation

Results

Analysis of the Error Rate

The data from four participants who had an IAT error rate of 18% or greater were removed from all of the analyses. Once these participants were removed, the average error rate was 3% and ranged from 0% to 10%.

Did Completing Explicit Attitude Measures Change Behavior?

Participants in the IAT + explicit measures condition completed the explicit attitude measures before choosing their apple or candy bar. As noted above, a potential problem with the design is that the act of explicitly reporting one’s attitudes toward apples and candy bars may change one’s behavior. If it does, then the study would provide a very weak test of the relationship between the IAT and behavior. In order to examine this possibility, we analyzed the participants’ choices of apples versus candy bars in the two conditions. Their choices suggest that they were not affected by the manipulation, $\chi^2(1) = .0191, p = .92$. Participants who completed the explicit measures selected an apple 47.5% of the time, and participants in the IAT only condition selected an apple 46.3% of the time. However, because the relationship between attitudes and behavior is complex, we examined if the IAT predicted behavior separately for each condition.

The IAT Measure

As in the first two studies, we computed IAT scores by subtracting Stage 3 response times from Stage 5 response times. As can be seen from Table 5, the IAT scores indicate that participants had more positive associations with apples than with candy bars, $F(1,80) = 123.32$, $p < .001$.

Explicit Attitude Measures

How do the IAT results correspond to the preferences revealed by the explicit measures? To examine this question, we computed feeling thermometer scores by subtracting the participants' feeling thermometer ratings of the word "candy bar" from their ratings of the word "apple." For the semantic differential, liking, and intent measures, we summed the participants' responses to the candy bar items and then subtracted them from their responses to the apple items. Only the semantic differential revealed a significant preference, and it was in line with the IAT results (see Table 5), with participants "preferring" apples over candy bars. The means that emerged in this study were similar to the means in the first two studies, in that both the IAT and the explicit attitude measures seem to indicate similar preferences.

The Relationship Between the IAT and Explicit Attitude Measures

Although the mean differences appear to tell similar stories, a critical issue, given that explicit reports of attitudes toward apples and candy bars should be relatively free of social desirability pressures, is whether the IAT correlates with the explicit measures. As can be seen in Table 6, it did not. Paralleling the findings of Studies 1a and 1b, the correlations between the IAT measures and the explicit attitude measures ranged from $r = -.10$ to $r = .16$, with none approaching significance. Although the various explicit attitude measures correlated with each other, there was no relationship between the IAT and any of the explicit attitude measures.

It is worth noting that the restrictions of range concerns identified in the earlier studies do not emerge here. In this study, we find great variety in the participants' preferences. In fact, on the feeling thermometer ratings, 32.5% of the participants indicated a preference for apples, 32.5% indicated a preference for candy bars, and 35% indicated equal preference for the two. This variability in responding makes it difficult to argue that the lack of correlations between the IAT and the explicit attitude measures was due to restricted range.

Relationships to Behavior

Up to this point, the results from this study closely resemble the results from the first two studies. Recall, however, that the primary purpose of this study was to examine whether the IAT and the explicit attitude measures predict behavior. Because our behavior was a dichotomous choice between an apple and a candy bar, we tested this hypothesis using a logistic regression with the IAT as the predictor and behavioral choice as the outcome. This analysis revealed that the IAT failed to predict the participants' behavior, $p = .847$ (see top of Table 7).

How did the explicit measures do at predicting behavior? To examine this question, we conducted two additional logistic regressions using the feeling thermometer ratings and the semantic differentials as predictors of behavioral choice. In contrast to the analysis based on the IAT, these analyses revealed that both explicit measures were significant predictors of behavior, $p's < .015$ (see top of Table 7).

Next, we conducted a logistic regression on behavior in which the IAT and the two explicit measures were entered simultaneously as predictors to see which account for unique variance in the prediction of the behavior. Consistent with the univariate results, when all three predictors were entered simultaneously, the two explicit measures were marginally significant predictors of behavior and the IAT was unrelated to behavior (see middle of Table 7).

Because of the possibility that completing the explicit attitude measures may have inadvertently affected the IAT-behavior relationship, we examined whether the IAT predicted the choice behavior in the condition where no explicit attitude measures were given. Consistent with the preceding analyses, we found no evidence that the IAT predicted the behavior (see bottom of Table 7).

Finally, although the global IAT measure did not predict the choice behavior, it is possible that a more specifically measured implicit attitude would predict the behavior. Thus, we tested the hypothesis that the IAT would predict behavior using a more specific IAT measure. We computed an item-based IAT based on participants' responses to the target words Red Delicious (IAT [Red Delicious]) and Snickers (IAT [Snickers]) in the IAT task. We then conducted two logistic regressions. First we attempted to predict the selection of the Red Delicious apple with IAT (Red Delicious). Although IAT (Red Delicious) is a marginally significant predictor of the behavior, the prediction is in the wrong direction. Participants whose IAT (Red Delicious) score indicated a preference for apples were marginally more likely to select a Snickers over a Red Delicious apple, $\beta = -.0006$, $(\beta/SE(\beta))^2 = 2.834$, $p = .0923$. Next, we attempted to predict the selection of a Snickers candy bar with IAT (Snickers). This analysis revealed that IAT (Snickers) scores failed to predict the behavior, $\beta = .0003$, $(\beta/SE(\beta))^2 = 0.970$, $p = .3248^1$.

Discussion

The primary goal of the second study was to examine the links between explicit attitudes, the IAT, and behavior. The results showed that although explicitly measured attitudes predicted whether a participant chose an apple or a candy bar, the IAT failed to predict the same choice behavior. These findings provide additional evidence for the independence of the IAT and explicitly measured attitudes.

Although the IAT did not predict behavior in this study, it may well predict behavior in other situations. After all, this was only one study and much of the research on the links between attitudes and behavior indicates that attitudes and behavior must be measured appropriately and in the right situations in order to find a relationship between the two (see Ajzen, 1996, 1985; Fishbein & Ajzen, 1975). It is certainly possible that future studies will find that the IAT predicts some behaviors. For example, others have found that different implicit attitude measures predict behaviors that tend to be out of one's conscious control, such as eyeblinking (Dovidio et al., 1997). Although we used a relatively spontaneous behavior in this study, it is possible that the predictive power of the IAT is limited to more non-conscious behaviors. But the important point for the current thesis is that the explicit attitude measures and the IAT had different predictive power with regard to behavior.

The correlational results of this study also support the independence model of explicit attitudes and the IAT and conceptually replicate the correlations found in Studies 1a and 1b. Once again, we failed to find a correlation between the IAT and explicitly measured attitudes. One question that can be raised, however, is whether social desirability obscured the correlations between the IAT and the explicit attitude measures. After all, candy bars may be seen as unhealthy and people may be unwilling to admit that they prefer an unhealthy snack over an apple. If this were the case, however, we might expect the IAT to reveal more positive association with candy bars than would the explicit attitude measures. In fact, the opposite pattern emerged. The IAT revealed that over 90% of the participants had more positive associations with apples than with candy bars, while the explicit attitude measures revealed that some participants preferred apples, while others preferred candy bars. In addition, half of the participants selected a candy bar instead of an apple. If social desirability influenced responses,

then we should not have found as many participants selecting a candy bar. Thus, it is unlikely that social desirability affected participants' answers in ways that prevented us from finding a correlation between the IAT and explicitly measured attitudes.

Finally, it is worth noting that the potential concerns about the range restrictions that were raised in the flower and insect studies do not apply here. We failed to find correlations between the IAT and explicit measures despite relatively equal numbers of pro-apple and pro-candy bar participants in the sample.

Study 3

So far, the results suggest that the IAT and explicitly measured attitudes tap independent constructs. However, we do not yet know how they are independent. Traditional conceptualizations of attitudes assume that attitudes tap endorsement (Fishbein & Ajzen, 1975; Ajzen, 1985). Ask a person to complete a feeling thermometer about various political candidates and you have a fairly good idea of the extent to which the person endorses each candidate. According to an endorsement model, the IAT taps endorsement like explicit attitude measures, but the endorsement domain it taps is independent from the endorsement domain that is tapped by explicit attitude measures. Another possibility, however, is that the IAT measures a construct unrelated to endorsement. Perhaps, for example, it measures the extent to which various attitudinal objects are associated in the person's environment. According to an environmental association model, the IAT may tell us what associations the person has been exposed to in his or her environment, rather than the extent to which the person endorses the attitude object. In other words, the environmental-associations model posits a dissociation between explicitly measured attitudes and the IAT, consistent with Devine's (1989) dissociation of exposure to stereotypic

knowledge (which may be measured by the IAT) and personal beliefs (which may be measured by explicit attitude scales).

Studies 1 and 2 provide some indirect evidence for the environmental association model of the IAT. In our society, people are exposed to many positive associations with flowers and many negative associations with insects. Conversely, there are relatively few negative associations with flowers and positive associations with insects. If the IAT measures the extent to which various attitude objects are associated in a person's environment, then we would expect Americans in general to show an IAT preference for flowers over insects, which is exactly what we found in Study 1. In this case, endorsement and environmental associations lead to the same prediction -- preference for flowers over insects.

Interestingly, an environmental association perspective can also explain the relatively surprising finding from Study 2 that most participants displayed an IAT preference for apples over candy bars. In our society, there is an abundance of positive associations and virtually no negative associations with apples. For candy bars, however, the messages are much more mixed. Although there are positive associations with candy bars (e.g. they certainly taste good), there are also many negative associations (e.g. they are high in fat and can lead to tooth decay). Thus, if one were simply to compare the number of positive and negative messages associated with apples and candy bars, apples would be clear winners. In Study 2, the IAT results are consistent with this logic.

If the IAT taps associations while explicit attitudes measure evaluation, then manipulating the frequency with which certain concepts are paired in the environment should result in a change in IAT scores. Consider for example, the standard "youth bias" IAT (see Nosek, Greenwald, & Banaji, 1998) where people are faster on the IAT to respond to youth + pleasant

and elderly + unpleasant than they are to respond to youth + unpleasant and elderly + pleasant. According to an endorsement model of IAT effects, the youth bias occurs because people have more favorable evaluations of the young than of the elderly. But what would happen to the IAT scores and the explicit attitudes of a person who is repeatedly exposed to youth paired with unpleasant words and with elderly paired with pleasant words? According to the environmental-associations model, these pairings should strengthen the connection between the concepts "youth" and "unpleasant" and between "elderly" and "pleasant." When the person subsequently takes a youth-elderly IAT, he or she should be relatively quicker to respond to the now more familiar pairing of youth with negative words and also elderly with positive words. In other words, changes in the environment (i.e., which concepts co-occur) may affect IAT responses and attenuate the standard "youth bias" on the IAT.

Notice that it is not clear how these same manipulations would affect explicitly measured attitudes. To the extent that explicitly measured attitudes are affected by an individual's explicit evaluation of the attitude object and any social norms that are relevant to the attitude in that specific situation (Fishbein & Ajzen, 1974, 1975), the repeated pairing of concepts may have no effect on a person's explicit attitudes. A person who sees the concepts "youth" and "unpleasant" and "elderly" and "pleasant" repeatedly paired may not consider the information relevant to his or her explicit evaluation of the young and old.

The purpose of the last experiment was to examine support for the environmental associations model by investigating if changes in the frequency of associations between concepts would affect IAT scores, but not explicitly reported attitudes. To this end, participants in this study were randomly assigned to one of two groups: the youth=good group and the elderly=good group. All participants first completed a youth-elderly IAT in order to obtain a baseline measure

of associations between these categories. Next, participants were exposed to 200 word pairings under the guise of a memory experiment. This task was presented as a completely separate task from the previously taken IAT. Participants' reaction times were not recorded and they did not make specific responses to specific target words. Participants simply tried to remember which words were paired together, and how often those words were paired. Participants in the youth=good condition read 200 pairings of the word youth with positive words and elderly with negative words, while those in the elderly=good condition read 200 pairings of the word elderly with positive words and youth with negative words. Subsequently, all participants completed a second youth-elderly IAT to measure any change in IAT scores. Participants also completed explicit measures of attitudes toward youth and the elderly. While we expected all participants to show a bias in favor of youth over the elderly, we hypothesized that those in the youth=good condition would show an increased IAT bias in favor of youth, compared to baseline IAT times. On the other hand, those in the elderly=good condition would show an attenuation of the IAT bias in favor of youth, compared to IAT baseline times.

Method

Participants

Fifty students enrolled in an introductory psychology course at the University of Michigan participated in this experiment. All participants received course credit for their participation.

Procedure

Participants were run in groups of up to four at a time. Each participant was randomly assigned to be in either the youth=good or elderly=good condition. The participants first completed a baseline youth/elderly IAT. The categories pleasant and unpleasant and their associated target words were the same as those used in Studies 1 and 2. Five target words were

selected to be representative of the concept young: teenager, newlywed, college, twenty-one, and adolescent. Five target words were also selected to be representative of the concept elderly: gray hair, seventy, grandparent, retired, and mature. Otherwise, the procedure of the IAT was identical to that of Experiments 1 and 2.

After completing the baseline IAT, all participants were led to believe they were participating in an "association detection task." Participants were informed that they would be viewing 200 word pairings. Their task was to remember what words were paired together and how frequently each word pairing occurred. Participants believed there would be a recall task later in the experimental session.

We used this association detection task to manipulate the associations a participant had with youth and with elderly. Those in the youth=good condition were exposed to 10 different word pairs: youth - cheer, youth - pleasure, youth - happy, youth - love, youth - peace, elderly - death, elderly - filth, elderly - jail, elderly - murder, and elderly - ugly. Those in the elderly=good condition were exposed to the following 10 word pairings: elderly - cheer, elderly - pleasure, elderly - happy, elderly - love, elderly - peace, youth - death, youth - filth, youth - jail, youth - murder, and youth - ugly. For both conditions, each word pair was presented 20 times, for a total of 200 word pairings.

On each trial, participants saw the word pair appear near the center of the computer screen, with one word appearing just to the left of center and one appearing just to the right of center. Within each word pair, the ordering of the words was counterbalanced. Once the participant read the word pair AND one second elapsed, the participant could advance to the next word pairing by pressing the space bar. At the conclusion of the word-pairing task, participants completed a second youth/elderly IAT that was identical to the baseline IAT task.

Next, participants completed a series of explicit attitude measures. Participants first completed semantic differentials and feeling thermometers regarding their attitudes toward youth and elderly. The attitude scales were identical to those used in Studies 1 and 2, except that the targets of the scales were changed to "Elderly" and "Youth." Second, participants completed the Attitudes Toward Old People Scale, ATOP (Kogan, 1961), on a scale from 1 (Disagree Strongly) to 7 (Agree Strongly) with high reliability ($\alpha = .86$). This scale was recoded so that negative numbers indicate positive attitudes toward the elderly and positive numbers indicate more unfavorable attitudes. The participants were then thoroughly debriefed about the experiment and thanked for their participation.

Results

Analysis of the Error Rate

The data from five participants who had an IAT error rate of 10% or greater on either IAT were removed from all of the analyses. Once these participants were removed, the average error rate was 4% and ranged from 0% to 9%. One participant failed to complete the explicit attitude measures, and was excluded from those analyses.

The IAT Measure

As in the first two studies, we computed IAT scores based on the difference between responses to target words at Stage three and Stage five. Positive scores indicate faster times associating youth with positive and elderly with negative than associating elderly with positive and youth with negative. IAT scores at both the baseline and follow-up stages indicate that participants showed a bias for youth over elderly, all F 's (1,44) > 58, p 's < .001, d 's > 1.14 (see top of Table 8).

Did Word Co-Occurrences Affect IAT Scores?

A 2 (condition: Youth=good vs. Elderly=good) x 2 (time: Baseline vs. Follow-up) ANOVA with repeated measures on time was conducted on IAT scores (see Figure 1). This analysis revealed the predicted condition by time interaction, $F(1,43) = 4.75$, $p = .035^2$. Post hoc tests reveal that IAT scores for participants in the youth=good condition did not change after the presentation of the word pairings, $F(1,21) = 0.33$, $p = .571$, perhaps indicating that their scores were already at the ceiling. However, IAT scores for participants in the elderly=good condition significantly decreased after the presentation of the word pairings, $F(1,22) = 8.90$, $p = .007$.

Explicit Attitude Measures

Recall that for the explicit measures, we did not obtain baseline (pre-manipulation) measurements. Collapsing across conditions, the explicit attitude scores reveal a mixed picture (see bottom of Table 8). The results of the feeling thermometer indicate that participants had slightly unfavorable attitudes toward the elderly, $F(1,43) = 3.83$, $p = .057$. Conversely, the ATOP scale revealed that participants had slightly favorable attitudes toward the elderly, $F(1,43) = 17.80$, $p < .001$. Finally, the semantic differential indicated no difference in attitudes toward youth and elderly, $F(1,43) = 0.01$, $p = .979$.

From the results, the feeling thermometer and ATOP appear to show a conflicting pattern. However, it is informative to consider how these measures were constructed. The feeling thermometer, like the semantic differential, was calculated by subtracting ratings for elderly from ratings for youth. These two explicit attitude measures, like the IAT, are measures of comparative attitudes. A high score indicates a bias for youth, compared to the elderly. On the other hand, the ATOP measures attitudes toward the elderly only. A negative score indicates positive attitudes toward the elderly relative to the midpoint of the scale. Taken together, the

results from the feeling thermometer and the ATOP indicate that our participants had positive attitudes toward both youth and elderly, but that their attitudes toward youth were more positive than their attitudes toward the elderly.

Did Word Co-Occurrences Affect the Explicit Measures?

Although the explicit attitude measures show a "youth bias" just like the IAT scores, the explicit attitudes, in contrast to the IAT, showed no effect from the manipulations. That is to say, the feeling thermometer, the semantic differential, and the ATOP scale were unaffected by the association manipulation, all F 's (1,42) < 1 (see Table 9).

One could argue that the design for this study favored our hypothesis. By obtaining baseline and follow-up measures for the IAT, but only follow-up measures for the explicit attitudes, there is greater power to detect a difference on the IAT than for the explicit attitude measures. Indeed, if we only examine the follow-up IAT scores by condition, the results are in the predicted direction, but they do not reach conventional levels of significance, $F(1,43) = 2.50$, $p = .121$. However, a comparison of effect sizes for the follow-up analysis of IAT and explicit attitude scores reveals a much larger effect size for the IAT scores ($d = .46$) than for the explicit attitude measures (d 's < .082). These differences in effect sizes suggest that the observed differences between the IAT and explicit attitude ratings are real and not simply the result of differences in statistical power.

These results add additional support to the environmental association explanation of the independence between the IAT and the explicit attitude measures. While personal evaluations play a large role in the expression of explicit attitudes, the IAT appears to be driven by simple associations between concepts.

The Relationship Between the IAT and Explicit Attitude Measures

The correlational data also support the notion of the independence between the IAT and explicit attitude measures. Although the IAT and the explicit attitude measures reveal the same pattern of means, we found little evidence for any correlation between the two. As can be seen in Table 10, although the explicit attitude measures all correlate positively with each other, but they did not correlate with the IAT.

Discussion

Taken together the results of the final study add additional support to the claim of independence between the IAT and explicit attitude measures. As in the previous two studies, correlations between the IAT and explicit attitude measures again failed to emerge. More importantly the fact that IAT scores moved in the direction of our association manipulation, while explicit attitudes were unchanged, provides preliminary support for the environmental association model of the IAT.

An interesting methodological concern arose from this study. We used the Attitudes Toward Old People scale as an explicit attitude measure and attempted to compare this scale to IAT scores. Similarly, others have used the Modern Racism Scale (McConahay, 1986) as an explicit attitude measure in IAT studies examining racial biases (see Greenwald et al., 1998). However, it is unclear how to interpret these scales relative to the IAT. The IAT measures a comparative bias. For example, the youth/elderly IAT measures the tendency to have positive associations with the elderly as compared to the tendency to have positive associations with youth. On the other hand, most rating scales, such as the ATOP scale, measure attitudes toward a single attitude object -- in this case the elderly. This dimensional difference argues for caution when comparing means and effect sizes concerning the IAT and uni-dimensional attitudinal scales. Ideally, bi-

dimensional explicit attitude scales, such as the feeling thermometer and semantic differential, should be used when making comparisons between the IAT and explicit attitude measures to disambiguate any relationships that are found.

General Discussion

Across all three studies, we found converging evidence for the independence of the IAT and explicit attitude measures. First, across all three studies, we consistently failed to find any correlations between the IAT and explicit attitude measures, even when social desirability pressures were minimized. Correlations failed to emerge despite examining the correlations at multiple levels of analysis, across several different attitude domains, and across a wide range of attitude distributions. Second, we found that explicit attitudes and the IAT were differentially predictive of behavior. Third, we found that the IAT was affected by exposing participants to new associations between attitude objects, while the explicit attitude ratings remained unchanged. Taken together, these three studies provide strong evidence for the independence of the IAT and explicit attitude measures. In addition, the results of these studies provide preliminary support for the environmental association model of the IAT. According to this model, the IAT taps the associations a person has been exposed to in his or her environment, and not that individual's level of endorsement regarding the attitude object. Consistent with the environmental-associations model, changes in environmental frequency led to changes in the IAT, but not in explicit attitudes.

Support for the association model can also be found in the stereotyping literature. To the extent that the IAT measures associations, it should be possible to manipulate IAT scores by making a set of previously learned associations salient. If prominent, well-to-do Blacks (such as Michael Jordan or Oprah Winfrey) were made temporarily accessible, for example, then a

Black/White IAT should reveal less of a bias against Blacks. On the other hand, if several criminal Blacks were made temporarily accessible, then a Black/White IAT should reveal more of a bias against Blacks. Consistent with this analysis, Blair & Ma (1998) found that thinking of a “typical strong woman” reduced the strength of the gender bias that frequently emerges on the IAT.

If the environmental association model correctly describes the underlying process of the IAT, then we need to reevaluate interpretation of the IAT as an implicit attitude and as a predictor of behavior. IAT scores may reveal little about a person's beliefs, and much about his or her environment or culture. By this interpretation, showing a White bias on a Black/White IAT does not necessarily indicate that a person holds deep-rooted prejudices against Blacks or that the person will discriminate against Blacks. A more appropriate interpretation would be that the IAT reflects the fact that an individual lives in an environment or culture where Blacks are devalued relative to Whites. Given that we live in such a culture, it is not surprising that most Americans, White or Black, show a White IAT bias (Greenwald et al., 1998; Nozok et al., 1998).

Relationship Between the IAT and Other Implicit Attitude Measures

Although the current studies suggest that the IAT does not correspond to explicit attitudes or behavior, other studies have found evidence that other implicit measures of attitudes do predict behavior. For example, measures of prejudice based on linguistic biases predict how threatening a person will rate a Blacks target (von Hippel et al., 1997). Using a priming task, Dovidio and colleagues (Dovidio et al., 1997) found that implicitly measured racial attitudes predict nonverbal behavior, such as blinking and amount of eye contact, during an interaction with a Black partner. Similarly, Fazio and his colleagues have developed an implicit measure of prejudice based on automatic evaluations that predicts how smoothly Whites will interact with a

Black partner (Fazio, Jackson, Dunton, & Williams, 1995). Studies like these suggest that implicitly measured racial attitudes hold considerable promise for furthering our knowledge of prejudice and discrimination.

To date, however, there is no evidence that the IAT measures the same constructs as any of these implicit attitude measures. Cameron, Alvarez, & Bargh (2000), for example, examined the inter-relationships between numerous implicit "attitude" measures (Fazio et al., 1995; Wittenbrink et al., 1997; von Hippel et al., 1997), and their relationship to behavior. The IAT failed to correlate reliably and in the predicted direction with any of the other implicit attitude measures. In addition, while a feeling thermometer measure of racial attitudes predicted friendliness in an interaction with a Black confederate, scores on a Black-White IAT were unrelated to friendliness in the interaction.

Indeed, if it turns out that the IAT measures environmental associations, rather than a person's level of endorsement, it may be misleading to refer to the IAT as an implicit attitude measure. After all, an implicit attitude is defined as an introspectively unidentified (or inaccurately identified) trace of past experience that mediates favorable or unfavorable thought, feeling, or action toward an object (Greenwald & Banaji, 1995). According to the environmental association model, although the IAT does measure a trace of past experience, there is no evidence that it mediates evaluative thought, feeling, or action.

Coda

The quest for an attitude pipeline is a recurrent theme in the attitude literature. Indeed, it could be argued that the desire to find a measure that taps attitudes in ways that are impervious to self-presentation represents a kind of "Holy Grail" for attitude researchers. With such a measure we could explore confidently all those attitudes that people privately hold, but seldom reveal. It

would be open season on attitudes toward race, sex, money, and all the domains of private life. With such a measure we could explore attitudes that exist outside of conscious awareness. And with such a measure we might gain an important tool for educating the public about our less socially desirable attitudes. Racial prejudice, for example, often fails to emerge on traditional attitude measures. Yet discrimination clearly permeates American culture, and a valid attitude pipeline might provide one way of revealing the unpleasant truth of prejudice.

Despite the appeal of an attitude pipeline, the search has been filled with many false starts. At various points, galvanic skin response (Podlesny & Raskin, 1978), non-verbal behavior (DePaulo, 1994), facial expressions (Ekman, O'Sullivan, Friesen, & Scherer, 1991), and response latency (Larson, 1932) have all been put forward as ways to tap attitudinal "truth." Upon closer analysis, however, each of these measures has failed to emerge as a valid predictor of a person's true thoughts or attitudes (for example, see Iacono & Lykken, 1997; Bashore & Rapp, 1993; Steinbrook, 1992; Furedy & Heslegrave, 1988; Lykken, 1979). In light of these failures, the results of the current research suggest that we should also be cautious about the pipeline potential of the IAT, specifically, and implicit measures more generally.

References

- Ajzen, I. (1985). From intentions to actions: A theory of planned behavior. In J. Kuhl & J. Beckman (Eds.), Action-control: From cognition to behavior (pp. 11-39). Heidelberg: Springer.
- Ajzen, I. (1996). The social psychology of decision making. In E. T. Higgins & A. W. Kruglanski (Eds.), Social psychology: Handbook of basic principles (pp. 297-325). New York: Guilford.
- Bashore, T. R., & Rapp, P. E. (1993). Are there alternatives to traditional polygraph procedures? Psychological Bulletin, 113, 3-22.
- Blair, I., & Ma, J. (1998). Imagining stereotypes away: The moderation of implicit stereotypes through counterstereotype imagery. Unpublished research: University of Colorado, Boulder, CO.
- Cameron, J. A., Alvarez, J. M., & Bargh, J. A. (2000). Examining the Validity of Implicit Measures of Prejudice. Poster presented at the first annual meeting for the Society of Personality and Social Psychology, Nashville, TN.
- Chamberlain, C. (1998, October 14). Roots of racism revealed. ABCNews.com, Health and Living. Available URL:
http://www.abcnews.go.com/sections/living/InYourHead/allinyourhead_11.html
- Cohen, J. D., MacWhinney, B., Flatt, M., & Provost, J. (1993). PsyScope: A new graphic interactive environment for designing psychology experiments. Behavioral Research Methods, Instruments & Computers, 25, 257-271.
- DePaulo, B. M. (1994). Spotting lies: Can humans learn to do better? Current Directions in Psychological Science, 3, 83-86.

- Devine, P. G. (1989). Stereotypes and prejudice: Their automatic and controlled components. Journal of Personality and Social Psychology, 56, 5-18.
- Devine, P. G., & Elliot, A. J. (1995). Are racial stereotypes really fading? The Princeton trilogy revisited. Personality & Social Psychology Bulletin, 21, 1139-1150.
- Dovidio, J. F., Kawakami, K., Johnson, C., Johnson, B., & Howard, A. (1997). On the nature of prejudice: Automatic and controlled processes. Journal of Experimental Social Psychology, 33, 510-540.
- Ekman, P., O'Sullivan, M., Friesen, W. V., & Scherer, K. R. (1991). Face, voice, and body in detecting deceit. Journal of Nonverbal Behavior, 15, 125-135.
- Fazio, R., Jackson, J., Dunton, B., & Williams, C. (1995). Variability in automatic activation as an unobtrusive measure of racial attitudes: A bona fide pipeline. Journal of Personality and Social Psychology, 69, 1013-1027.
- Fishbein, M., & Ajzen, I. (1974). Attitudes towards objects as predictors of single and multiple behavioral criteria. Psychological Review, 81, 59-74.
- Fishbein, M., & Ajzen, I. (1975). Belief, attitude, intention, and behavior: An introduction to theory and research. Reading, MA: Addison-Wesley.
- Furedy, J. J., & Heslegrave, R. J. (1988). Validity of the lie detector: A psychophysiological perspective. Criminal Justice & Behavior, 15, 219-246.
- Greenwald, A. G. (1998). The Implicit Association Test: Demonstration and Overview of Findings. Paper presented at the Annual Meeting of the Society of Experimental Social Psychology, Lexington, KY.
- Greenwald, A. G., & Banaji, M. R. (1995). Implicit social cognition: Attitudes, self-esteem, and stereotypes. Psychological Review, 102, 4-27.

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.

Herr, P. A. (1995). Whither fact, artifact, and attitude: Reflections on the theory of reasoned action. Journal of Consumer Psychology, 4, 371-80.

Iacono, W. G., & Lykken, D. T. (1997). The validity of the lie detector: Two surveys of scientific opinion. Journal of Applied Psychology, 82, 426-433.

Kogan. 1961. Attitudes toward old people: The development of a scale and an examination of correlates. Journal of Abnormal & Social Psychology, 62, 44-54.

Larson, J. A. (1932). Lying and its detection. Chicago: University of Chicago Press.

Lykken, D. T. (1979). The detection of deception. Psychological Bulletin, 86, 47-53.

McConahay, J. B. (1986). Modern racism, ambivalence, and the modern racism scale. In J. F. Dovidio & S. L. Gaertner (Eds.), Prejudice, Discrimination, and Racism. Orlando, FL: Academic Press.

Nosek, B. A., Cunningham, W., Banaji, M. R., & Greenwald, A. G. (2000). Measuring implicit attitudes on the internet. Poster presented at the first annual meeting for the Society of Personality and Social Psychology, Nashville, TN.

Nosek, B. A., Greenwald, A. G., & Banaji, M. R., (1998). Implicit association test [Website]. Available URL: <http://www.yale.edu/implicit/>.

Podlesny, J. A., & Raskin, D. C. (1977). Physiological measures and the detection of deception. Psychological Bulletin, 84, 782-791

Rudman, L., Greenwald, A. G., & McGhee, D. E. (1996, October). Powerful women, warm men? Implicit associations among gender, competency, and nurturance. Paper presented at the meeting of the Society of Experimental Social Psychology, Sturbridge, MA.

Schwartz, J. (1999, January 22). 200,000 Web Tests Show Unconscious Roots Of Racism, Ageism. (University of Washington Press Release) Available URL:
<http://140.142.3.7/newsroom/news/1999archive/01-99archive/k012299.html>.

Schwartz, J. (1998, September 23). Pervasiveness of prejudice to be demonstrated at press conference; Americans can go on line to test themselves . (University of Washington Press Release) Available URL: <http://140.142.3.7/newsroom/news/1998archive/09-98archive/k092398a.html>.

Steinbrook, R. (1992). The polygraph test: A flawed diagnostic method. New England Journal of Medicine, 327, 122-123.

Tibbets, G. (1998, October 14) Test measures unconscious bias, dismays takers. The Standard Times.

von Hippel, W., Sekaquaptewa, D., & Vargas, P. (1997). The Linguistic Intergroup Bias as an implicit indicator of prejudice. Journal of Experimental Social Psychology, 33, 490-509.

Wilson, T. D., Lindsey, S., & Schooler, T. Y. (in press). A Model of Dual Attitudes. Psychological Review.

Wittenbrink, B., Judd, C. M., & Park, B. (1997). Evidence for racial prejudice at the implicit level and its relationship with questionnaire measures. Journal of Personality and Social Psychology, 72, 262-274.

Author Note

Andrew Karpinski and James L. Hilton, Department of Psychology, University of Michigan.

Correspondence concerning this article should be addressed to Andrew Karpinski or James L. Hilton, Department of Psychology, 525 East University, Ann Arbor, MI 48109-1109.

Electronic mail may be sent to rtkarpin@umich.edu or hilton@umich.edu.

The authors wish to thank Laura Citrin, Steve Fein, Rich Gonzalez, Shelley Keiper, Markus Kemmelmeier, Charles Perdue, Chris Peterson, Denise Sekaquaptewa, and Oscar Ybarra for their helpful comments on an earlier version of this manuscript.

Footnotes

¹ We did not obtain explicit attitude ratings of each individual item used in this study. Thus, it was not possible to perform a similarly fine-grained analysis for the explicit measures.

² The condition by time interaction is marginally significant when the IAT scores are analyzed using truncated means and transformed means, $F(1,43) = 2.49$, $p = .09$ and $F(1,43) = 3.19$, $p = .08$, respectively.

Table 1

Study 1: Measures of Central Tendency for Each Stage of the IAT.

	<u>Study 1a</u>		<u>Study 1b</u>	
	<u>M</u>	<u>SD</u>	<u>M</u>	<u>SD</u>
Stage 1	546	85	557	72
Stage 2	586	54	617	117
Stage 3	547	64	604	76
Stage 4	540	56	630	105
Stage 5	720	119	771	111

Note: All times are reported in milliseconds.

Table 2

Study 1: Summary Statistics for the IAT and Explicit Attitude Measures

Attitude Measure	<u>M</u>	<u>SD</u>	<u>d</u>	<u>F</u>	<u>p</u>
Study 1a					
IAT	173.12	100.60	1.72	124.37	< .001
Valence (Category)	84.07	55.35	1.51	96.91	< .001
Valence (Item)	106.17	47.45	2.23	210.24	< .001
Study 1b					
IAT	167.22	77.78	2.15	129.41	< .001
FT (Category)	54.89	30.82	1.78	88.82	< .001
FT (Item)	59.36	18.52	3.20	287.77	< .001

Note: Positive numbers indicate a preference for flowers over insects. For Study 1a, the range for both valence measures was -200 to 200. For Study 1b, the range for both valence measures was -100 to 100.

Table 3

Study 1: Correlations between the IAT and Explicit Attitude Measures

Measure	IAT	Valence (Category)	Valence (Item)
IAT	1.0	-.11	-.02
Valence (Category)	-.31	1.0	.82**
Valence (Item)	-.19	.62**	1.0

Note: Correlations above the diagonal are from Study 1a (N=42) ; Correlations below the

diagonal are from Study 1b (N=28). + $p < .10$. * $p < .05$. ** $p < .01$.

Table 4

Study 1: Item by item analysis of the correlation between the IAT and explicit valence ratings

Flowers	Study 1a	Study 1b	Insects	Study 1a	Study 1b
Daisy	-.087	.076	Ant	.075	.044
Lily	-.058	-.044	Cockroach	-.401**	-.179
Rose	.119	-.252	Maggot	.087	.035
Tulip	.055	-.377*	Spider/Fly	.152	-.146
Carnation	.204	-.198	Wasp	-.073	-.177
Flowers	Study 1a	Study 1b	Insects	Study 1a	Study 1b
Cheer	-.047	-.384*	Death	.188	.005
Pleasure	.068	-.173	Filth	.101	.238
Happy	.147	.065	Jail	-.210	-.261
Love	-.275+	.002	Murder	.148	-.177
Peace	-.189	-.368+	Ugly	.075	-.214

Note: + $p < .10$. * $p < .05$. ** $p < .01$.

Table 5

Study 2: Summary Statistics for the IAT and Explicit Attitude Measures

Attitude Measure	<u>M</u>	<u>SD</u>	<u>d</u>	<u>F</u>	<u>p</u>
IAT	138.08	111.91	1.23	F(1,80) = 123.32	< .001
Feeling Thermometer	2.25	20.32	0.11	F(1,39) = 0.49	.488
Semantic Differential	4.25	5.81	0.73	F(1,39) = 21.42	< .001
Like Eating	-0.07	4.91	0.01	F(1,39) = 0.01	.924
Intent	0.47	2.98	0.16	F(1,38) = 0.99	.326

Note: Higher numbers indicate a preference for apples over candy bars. The IAT includes all 81 participants. The feeling thermometer, semantic differential, like eating, and intent measures include only those participants who completed the explicit measures. The intent measure was centered so that positive numbers indicate an intent to choose an apple while negative numbers indicate an intent to choose a candy bar.

Table 6

Study 2: Correlations Between the IAT and Explicit Attitudes

Measure	IAT	Explicit Measures			
	1	2	3	4	5
1. IAT	1.0				
2. Feeling Thermometer	.16	1.0			
3. Semantic Differential	-.09	.46**	1.0		
4. Like eating	-.10	.66**	.32*	1.0	
5. Intent	-.03	.52**	.14	.67**	1.0

Note: + $p < .10$. * $p < .05$. ** $p < .01$. $N = 40$.

Table 7

Study 2: Logistic Regression Predicting Behavioral Choice

Measure	β	$(\beta/SE(\beta))^2$	p-value	Odds Ratio
IAT + Explicit Measures Condition :				
Each variable entered in a separate regression				
IAT	.001	0.04	.847	1.000
Feeling Thermometer	.060	5.87	.015	1.062
Semantic Differential	.184	6.03	.014	1.202
IAT + Explicit Measures Condition :				
All variables entered simultaneously				
IAT	.001	0.02	.901	1.000
Feeling Thermometer	.054	3.54	.060	1.056
Semantic Differential	.163	3.44	.064	1.178
IAT + Only Condition				
IAT	-.002	0.67	.411	0.998

Table 8

Study 3: Summary Statistics for the IAT and Explicit Attitude Measures

Attitude Measure	<u>M</u>	<u>SD</u>	<u>d</u>	<u>F</u>	<u>p</u>
Baseline IAT	198.95	171.63	1.16	F(1,44) = 60.46	< .001
Follow-up IAT	181.12	158.34	1.14	F(1,44) = 58.88	< .001
Feeling Thermometer	6.82	23.11	0.29	F(1,43) = 3.83	.057
Semantic Differential	-.02	5.66	0.01	F(1,43) = 0.01	.979
Attitudes Toward Old	- 0.33	0.52	0.65	F(1,43) = 17.80	< .001

People Scale

Note: Higher IAT scores indicate a bias for youth over elderly. Higher feeling thermometer and semantic differential scores indicate more favorable attitudes toward youth than elderly. The Attitudes Toward Old People scale has been centered so that positive numbers indicates unfavorable attitudes toward the elderly.

Table 9

Study 3: Explicit Attitude Measures Across Conditions

Attitude Measure	<u>Youth = Good</u>		<u>Elderly = Good</u>	
	<u>M</u>	<u>SD</u>	<u>M</u>	<u>SD</u>
Feeling Thermometer	5.91	24.23	7.73	22.45
Semantic Differential	0.13	5.12	-0.18	6.27
Attitudes Toward Old People Scale	0.15	0.60	0.19	0.43

Note: N=44. Higher scores indicate unfavorable attitudes toward the elderly.

Table 10

Study 3: Correlations Between the IAT and Explicit Attitudes

Measure	IAT	Explicit Attitude Measures		
	1	2	3	4
1. IAT	1.0			
2. Feeling Thermometer	.03	1.0		
3. Semantic Differential	.12	.64**	1.0	
4. Attitudes Toward Old People Scale	-.06	.38*	.32*	1.0

Note: * $p < .05$. ** $p < .01$. $N = 44$.

Figure Caption

Figure 1: Changes in youth-elderly IAT scores as a function of viewing associations.

