Psychology of Women Quarterly, **25** (2001), 155–165. Blackwell Publishers. Printed in the USA.

A COMPARATIVE STUDY OF THE IMPLICIT AND EXPLICIT GENDER ATTITUDES OF CHILDREN AND COLLEGE STUDENTS

John J. Skowronski
Northern Illinois University

Melissa A. Lawrence
Illinois State University

Implicit attitudes and explicit attitudes toward men and women and toward male soldiers and female soldiers were assessed in fifth-graders (28 male, 31 female) and college students (43 male, 42 female). Women were rated more positively than men on an explicit attitude measure. Similarly, female soldiers were rated more positively than male soldiers, except among college men, who were pro–male soldier. Different results emerged from an Implicit Association Test using names of men and women (general gender condition) or of male soldiers and female soldiers (soldier name condition). Latencies indicated pro-female attitudes in the soldier name condition and among women and college students. Error rates also indicated pro-female attitudes, except for a pro-male preference among men in the general gender condition. Reasons that implicit and explicit attitude measures may produce such divergent results are discussed.

One important element of gender-based thought is the extent to which an evaluative response to another person is based on the other's gender. One hypothesis that can be derived from cultural myth is that men should be better liked than women, and some research indirectly supports this conjecture. For example, traits associated with the male stereotype are often positive (e.g., independent, objective, competitive, and logical), whereas the stereotype for women includes such negative traits as unstable, emotional, and illogical (Broverman, Vogel, Broverman, Clarkson, & Rosenkrantz, 1972). Additional indirect support for a pro-male preference is provided by extensive workplace discrimination against women (see Deaux & LaFrance, 1998). However, although at least one direct assessment of gender attitudes supports the hypothesis of a pro-male preference (McKee & Sherriffs, 1957), the results of several other studies contradict this pro-male hypothesis. These studies show that when gender attitudes are directly assessed people have more positive attitudes toward women than men (Eagly & Mladinic, 1989; Eagly, Mladinic, & Otto, 1991; Haddock & Zanna, 1994).

Theorists have tried to account for the inconsistencies in this literature in various ways. For example, the concept of benevolent sexism (Glick & Fiske, 1997) suggests that women are viewed in a subjectively positive way but are also viewed in stereotypically restricted roles. Hence, the fact that women are the objects of discrimination is not necessarily inconsistent with a positive view of women. Instead, discrimination (and a negative affective response) may occur when women attempt to move outside of their socially prescribed roles.

A second approach to reconciling the inconsistencies in the literature (Haddock & Zanna, 1994) argues that category construals may determine evaluative responses. That is, the term "woman" might provoke different reactions depending on whether "woman" is thought of as synonymous with "housewife" (which may produce a positive evaluative response in some individuals) or as synonymous with "feminist" (which may provoke a negative response).

A third interpretation of these discrepancies focuses on methodology. The exclusive use of explicit measures in much of the previous gender attitude research may not provide a complete picture of people's gender attitudes. For example, one approach to attitudes suggests that at least some attitudes reflect traces of prior experiences, many of which might not be easily recalled or reported. In support of this view, Cacioppo, Marshall-Goodell, Tassinary, and Petty (1992) have shown that attitudes can be classically conditioned without awareness. Similarly, Murphy, Zajonc, and Monahan (1995) have shown that attitudes might develop without awareness via mere exposure. Explicit attitude measures might not be adequate measures for attitudes developed in such ways. In

Address correspondence and reprint requests to: John Skowronski at Northern Illinois University, Department of Psychology, DeKalb, IL 60115-2892. E-mail: jskowron@niu.edu

addition, conscious processes accompanying explicit attitude responses might distort the attitudes reflected in those responses. This can occur for many reasons, including people's attempts at deception (Arkin & Lake, 1983) and the use of overly thoughtful attitude decision rules (Wilson, Hodges, & Lafleur, 1995).

The pitfalls associated with explicit attitude assessments might be avoided by measuring participants' *implicit attitudes,* the "introspectively unidentified (or inaccurately identified) traces of past experience that mediate favorable or unfavorable feeling, thought, or action toward social objects" (Greenwald & Banaji, 1995, p.8). One technique used to measure implicit attitudes is the Implicit Association Test (IAT). This test purports to measure implicit attitudes toward social categories such as gender "by measuring their automatic evaluation," the affect that is automatically and unintentionally activated in response to a category. This technique has shown promise in the measurement of social categories such as race and gender (Cunningham, Preacher, & Banaji, in press; Dasgupta, McGhee, Greenwald, & Banaji, 2000; Greenwald, McGhee, & Schwartz, 1998; Ottaway, Hayden, & Oakes, in press; Rudman, Greenwald, Mellott, & Schwartz, 1999). In this article we report the results of a study that used the IAT to explore implicit attitudes toward gender categories.

One of the issues explored in this study was the relation between age and implicit gender attitudes. Given a consistent gender attitude bias in the environment, if processes such as classical conditioning or mere exposure are responsible for the development of implicit gender attitudes, then one might expect those attitudes to become stronger with increasing life experience (i.e., age). Hence, any implicit gender preferences that emerge in the IAT (regardless of whether they are pro-female or pro-male) should be stronger in adults than in children. Furthermore, such age-related strengthening may not emerge on an explicit attitude measure. Concerns about political correctness and self-presentation might emerge later in life, so that adults, but not children, may mask their attitudes by responding in politically correct ways. To explore such possibilities we used an explicit measure of gender attitudes similar to one used previously (see Eagly et al., 1991) and compared the results obtained from our implicit and explicit attitude measures.

A second issue explored in our study were the effects of benevolent sexism on implicitly and explicitly measured attitudes. Some theorists would argue that there is a shift in people's conceptions of gender roles and the perceived gender appropriateness of occupations across the age range (see Huston, 1985). For example, acculturation processes might cause adults (especially adult men) to view women's occupation of stereotypically male roles (such as soldier) as inappropriate. Consequently, some adults might hold explicit negative attitudes toward women in such occupations. However, these culturally determined notions should not affect implicit attitudes, so they should not appear on implicit attitude tests. In fact, one can argue that people are now frequently exposed to positive examples of female soldiers (e.g., Xena the Warrior Princess, the Pink Power Ranger), so people may have developed positive implicit attitudes toward this gender category. One implication of this line of reasoning is that some adults who have negative explicit attitudes toward female soldiers might have positive implicit attitudes toward female soldiers. Another implication is that a manipulation that alters the categories used in attitude assessment tasks (e.g., women/men versus female soldiers/male soldiers) may have little effect when attitudes are measured implicitly but should alter the pattern of attitude responses obtained on explicit attitude measures. However, such shifts should only be observed among adults. Altering the categories used in the attitude assessment test should have little effect on the responses of children.

One other possible source of divergence between implicit attitudes and explicit attitudes is suggested by work on in-group preferences. People prefer in-group members to out-group members, an outcome that extends to gender groups (Lorenzi-Cioldi, 1991). Such preferences might be partially developed through repeated positive experiences with in-group members (Stangor, Sullivan, & Ford, 1991). In-group preferences formed via direct repeated experience ought to be amenable to measurement via tests of implicit attitudes, such as the IAT. However, because of egalitarian social norms (Monteith, 1993), such in-group preferences might not readily emerge on explicit attitude measures. Egalitarian norms may induce people to alter the explicit attitudes that they report toward in-groups and out-groups. This reasoning suggests that the apparent gender attitudes of men and women toward in-group and out-group members might depend on whether those attitudes are measured implicitly or explicitly.

To explore these ideas we conducted an experiment that assessed some people's attitudes toward men and women and other people's attitudes toward male and female soldiers. These attitudes were assessed in two ways: using an explicit measure (semantic differential attitude ratings) and using an implicit measure (the IAT). We expected that explicit gender attitudes would favor women (as in Eagly et al., 1991), but that the strength of this pro-female preference would be related to the age and gender of the participants and the exact categories used in the attitude rating task (men vs. women, or male soldiers vs. female soldiers). We also hypothesized that there would be a pro-female preference in implicit attitudes and were interested in the extent to which this pro-female attitude preference was related to the age and gender of the participants and the categories used in the attitude assessment task. Finally, we explored the extent to which people's implicit gender attitudes and their explicit gender

attitudes were similar or whether they sometimes diverged and whether the divergences that emerged were related to participant gender, participant age, and the categories used in the attitude tasks.

METHOD

Participants

There were two groups of participants. The first group, recruited from a public school system in a small, Midwestern town populated predominantly by middle-class families, was composed of 59 fifth-grade students (28 boys and 31 girls). The median age of the fifth-graders was 11 years of age. The sample was a convenience sample: Teachers indicated a willingness to allow us access to the fifth-grade students in the school. Prior to the study, each fifth-grader was given a contact letter describing the study and obtained written parental consent for the child's participation. Each child provided written consent at the time of participation. The experiment was oversubscribed: There were more volunteers than necessary. Volunteers were selected on a "first-come, first-served" basis.

The second group of participants was composed of 85 college undergraduates (43 men and 42 women) who were enrolled in introductory psychology courses at a large, Midwestern university. These participants received credit toward fulfillment of a course requirement in return for participation. The median age of the college-age participants was 18 years of age.

All participants were native English-speakers. Because of the composition of the populations at both schools, the student samples from both locations were almost entirely composed of individuals of Caucasian ethnic heritages.

Attitude Measures

Explicit attitude measures. Each participant completed two types of attitude measures. The first was an explicit attitude measure derived from Eagly et al. (1991; also see Eagly & Mladinic, 1989). The measure asked participants to rate the strength of their feelings about each of two gender-related categories. One group of participants rated the gender categories *Males and Females*[1] (general name condition) a second group rated the categories *Male Soldiers* and *Female Soldiers* (soldier name condition).

Participants made their ratings on semantic differential scales. The dimensions rated were *good-bad, positive-negative, valuable-useless, pleasant-unpleasant,* and *nice-awful.* Each response scale ranged from +3 to –3. The positive descriptor was on the high end and the negative descriptor was on the low end of each scale. Participants were instructed to mark the middle of the range if they considered both anchoring adjectives to be irrelevant to the judgment.

Implicit attitude measures. The second type of attitude measure was derived from a computerized version of the *Implicit Association Test* (IAT). The version that we used was designed by Shelly Farnham and was downloaded from the Web (currently available at hive-mind.com/shelly/IAT). Our version of the IAT procedure asked participants to repeatedly categorize target stimuli. The target stimuli were seven male names, seven female names, seven pleasant words, and seven unpleasant words. The words that we used had been pretested along several dimensions to rule out possible word-related confounds (see Greenwald et al., 1998). The pleasant words were gold, joy, smile, peace, paradise, sunshine, and warmth. The unpleasant words were abuse, corpse, death, filth, poison, slime, and pain. The names used as exemplars of the male category were Patrick, Zachary, Michael, Adam, Andrew, William, and David. The names used as exemplars of the female category were Melissa, Jennifer, Sarah, Julie, Laurie, Katie, and Maria.[2]

Each participant first engaged in two separate classification judgment tasks using these stimuli. In one classification task, participants were presented with a series of words, one at a time, and were asked to classify the words as pleasant or unpleasant. In the second classification task, participants were presented with a series of names, one at a time, and were asked to classify the names as male or female. Participants who used the general gender categories (male, female) in the explicit attitude task were told nothing about the male and female names. However, those who used the soldier categories in the explicit attitude task were told in the implicit task that the names were male soldier names and female soldier names. The names used in the implicit task were identical across conditions, regardless of the categories used in the explicit task.

The words and names were next combined into a mixed stimulus list and were presented one at a time. In one condition, participants were told to press one of the response keys if the stimulus was either a female name (or a female soldier name in the soldier classification condition) or a pleasant word and to press the other key if the stimulus was either a male name (or a male soldier name in the soldier classification condition) or an unpleasant word. In a second condition, participants were told to press one of the response keys if the stimulus was either a female name (or a female soldier name in the soldier classification condition) or an unpleasant word and to press the other key if the stimulus was either a male name (or a male soldier name in the soldier classification condition) or a pleasant word. Participants performed each of these classification tasks twice with the second repetition of the task coming right after the first. The order of the names and words was varied across each of the two repetitions. The first repetition contained practice trials; the second repetition contained the critical trials.

The order in which participants completed these two classification tasks was varied. Some participants complet-

ed the task in which the female names shared a response key with the positive words first; other participants completed the task in which female names shared a response key with negative words first. This order was yoked to the order of presentation of the explicit attitude scales. Those who received the male explicit scale first also received the implicit task pairing the male judgment with the pleasantness judgment on the same response key first; those who received the female explicit task first received the implicit task pairing the female judgment with the pleasantness judgment on the same response key first.

The IAT was administered on IBM-compatible desktop computers. Participants viewed the display from a distance of about 65 cm and made response choices with the left forefinger (using the A key) and the right forefinger (using the 5 key on the right side numeric keypad). Each separate classification task consisted of a block of 20 trials. Each block started with instructions that described the category discrimination for the block and the assignment of response keys (left or right) to the categories. Reminder labels in the form of category labels appropriately positioned to the left or right remained on the screen during each block. Each block contained words selected randomly and with replacement (independently for each participant) until 20 trials had been achieved for the block. In the complex classification tasks that combined name and word categorizations, gender names and words appeared on alternating trials. We note that this procedure deviates slightly from the usual IAT procedure in which the words and names are used without replacement. The fact that our results are similar to the results obtained by others who have used the IAT with general gender categories (e.g., Mitchell et al., 1999) suggests that our results cannot be attributed to these minor alterations to the IAT procedure.

Procedure

When registering for the experiment, participants were informed that they would be performing one task that assessed their attitudes toward men and toward women and a second task in which they would make classification judgments about names and words. These instructions were reiterated immediately prior to the start of the experiment. The participants completed the experiment at their respective schools. The fifth-graders used a classroom containing school computers. The college students used computer-equipped rooms in the psychology laboratories.

Participants always completed the explicit attitude rating task first. This task lasted approximately 5 minutes. This explicit-first procedure was adopted to reinforce the use of the "soldier" category on the implicit task. We were concerned that participants in this condition might not use the soldier label in the implicit task and simply focus on the gender implications of the names. We thought that always forcing explicit ratings to first be made of male and female soldiers would help to avoid this potential problem.

The participants were instructed to begin the computerized IAT task when they finished the explicit measure. The instructions for the second task were on the computer screen; no other instructions were given. The computerized task lasted approximately 15 minutes. After completing the computerized IAT task participants were debriefed and dismissed.

RESULTS

Explicit Measures

Two explicit attitude scores were calculated for each participant. The first score was calculated by averaging each participant's responses to all of the items assessing the participant's attitude toward women (Cronbach's α = .85). The second score was calculated by averaging each participant's responses to all of the items assessing the participant's attitude toward men (Cronbach's α = .79). These scores were entered into a mixed ANOVA. The between-participants variables in the ANOVA were participant gender (male or female), participant age (fifth-grader or college student), gender categories rated (men/women or male soldiers/female soldiers), and scale order (male explicit task first or female explicit task first). The within-participant variable was explicit gender attitude (males or females).

We expected that the results from these explicit attitude measures would replicate the results of Eagly et al. (1991), who found a pro-female preference in explicit attitudes. However, we were also interested in whether this pro-female preference would be qualified by the gender of the participants, the age of the participants, and the gender categories rated. A gender attitude main effect, F (1,128) = 23.35, MSE = .31, p < .001, did emerge in the data such that explicit attitudes were more positive toward women (M = 1.86, SD = .90) than toward men (M = 1.57, SD = .90). However, interpretation of this main effect was, indeed, qualified by two interactions. One interaction was between explicit gender attitude and participant age, F (1, 128) = 8.14, MSE = .31, p < .005. However, more important was a significant interaction between explicit gender attitude, participant age, participant gender, and the gender categories rated, F (1, 128) = 3.96, MSE = .31, p < .05. The means for this interaction (see Table 1) indicate that the pro-female preference was present in college-age women regardless of whether those women rated women in general or rated female soldiers. However, in college-age men the pro-female preference was present only when those men rated the general gender categories. College-age men showed a pro-male preference, not a pro-female preference, when rating male soldiers and female soldiers. This reversal is undoubtedly responsible for the significant three-way interaction, F (1, 77) = 4.58, MSE = .36, p < .05, between explicit gender attitudes, participant gender, and the specificity of the gender labels that emerged from a

Table 1

Interaction Between Explicit Gender Attitude, Participant Age, Participant Gender, and Classification Categories for Explicit Gender Attitude Ratings (Standard Deviations in Parentheses)

		College Participants	
		Attitude Toward	
		Males	*Females*
Male Participants	General Gender	1.42 (1.08)	1.98 (1.11)
	Soldiers	1.47 (.88)	1.16 (1.08)
Female Participants	General Gender	2.30 (.57)	2.48 (.54)
	Soldiers	1.59 (1.01)	1.73 (1.10)
		Fifth-Grade Participants	
		Attitude Toward	
		Males	*Females*
Male Participants	General Gender	1.37 (.92)	1.80 (.72)
	Soldiers	1.63 (.59)	2.13 (.54)
Female Participants	General Gender	1.24 (.92)	1.92 (.50)
	Soldiers	1.41 (.73)	1.88 (.64)

follow-up ANOVA conducted only on the data from the college-age participants. In comparison, in a second follow-up ANOVA this interaction was not significant in the fifth-grader data, F (1, 51) = .60, MSE = .24, p > .44. The fifth-graders consistently evinced a pro-female preference, regardless of participant gender or the types of gender categories used in the rating task, F (1, 51) = 32.90, MSE = .24, p < .0001.[3]

Implicit Attitudes

Easy IAT classification trials should be characterized by low latencies (faster responding) and low error rates; hard trials should be characterized by high latencies (slower responding) and high error rates. Thus, a pro-female preference in implicit attitudes would be indicated if the classification task was easier (faster and lower error rate) when the response keys paired female names with positive words and male names with negative words; a pro-male preference would be indicated if the classification task was easier (faster and lower error rate) when the response keys paired male names with positive words and female names with negative words.

Latency data. Two average response latencies were calculated for each participant. The first was calculated from each participant's response latencies (in msec) to the trials on which female names shared a response key with pleasant words and male names shared a response key with unpleasant words. The second average was calculated from each participant's response latencies to the trials on which female names shared a response key with unpleasant words and male names shared a response key with pleasant words. These two averages constitute the response key pairing within-participant variable in the experimental design. The other four variables in the design were all between-participant variables. These were: (1) participant age (college-age or fifth-grade), (2) participant gender (male or female), (3) classification categories used (general gender name condition or soldier name condition), and (4) task order (task pairing female names and pleasant words on one response key first, or task pairing female names and unpleasant words on one response key first). Manipulation of task order is a necessary part of the IAT procedure and ensures that response key pairing effects that emerge on the IAT are not an artifact of task order. However, because in this design interactions involving both task order and response key pairing are confounded, we make no attempt to interpret interactions involving both of these variables.[4]

The analysis yielded a significant response key pairing effect, F (1, 128) = 22.31, MSE = 22412.47, p < .001, and the means for this effect were suggestive of a pro-female implicit attitude preference. Participants responded more quickly (indicating easier responding) when female names shared a key with pleasant words, and male names shared a key with unpleasant words (M = 905.70 ms, SD = 310.42) than when female names shared a key with unpleasant words, and male names shared a key with pleasant words (M = 1001.15 ms, SD = 270.78). However, interpretation of this effect is qualified by three interactions.

The first of these interactions, depicted in Figure 1, was between response key pairing and participant age, F (1,128) = 6.54, MSE = 22412.47, p < .02. The means for this interaction indicate that the implicit pro-female preference was present only in college-age individuals. Simple effects tests indicate that the pro-female preference in implicit attitudes was significant for college students, F (1, 77) = 64.93, MSE = 11166.70, p < .001, but not for fifth-graders, F (1, 51) = 1.13, MSE = 39391.38, p > .29. A second significant interaction, depicted in Figure 2, involved the response key pairing and participant gender variables,

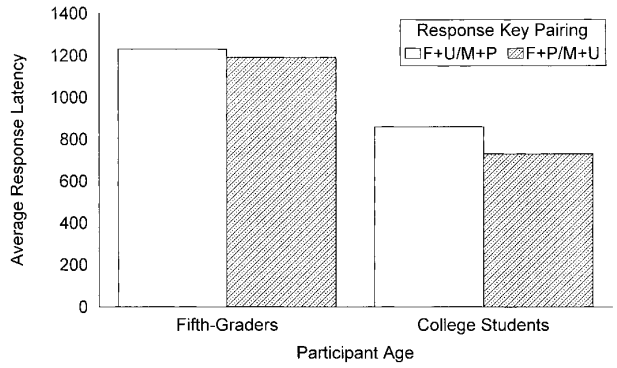

Fig. 1. IAT latency data: Interaction indicating that the pro-female preference in implicit attitudes emerged for college students and not for fifth-graders.

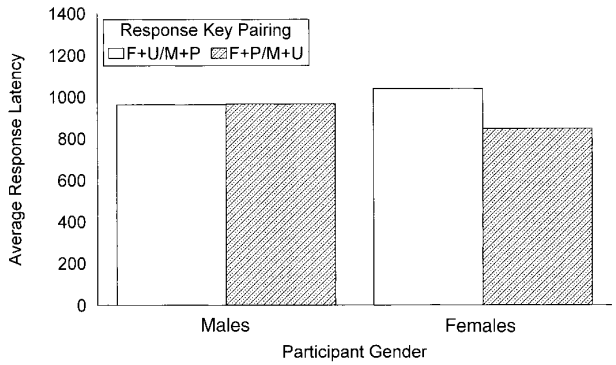

Fig. 2. IAT latency data: Interaction indicating that the pro-female preference in implicit attitudes emerged in female participants and not in male participants.

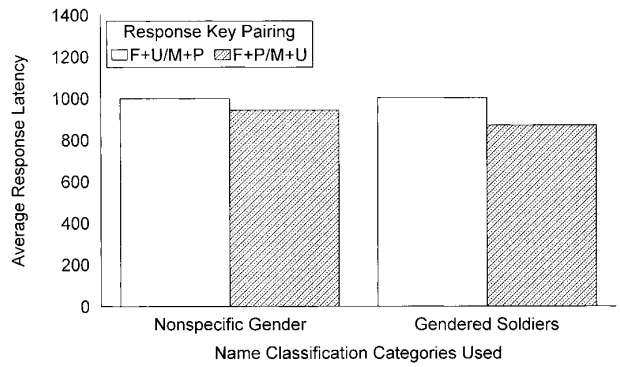

Fig. 3. IAT latency data: Interaction indicating that the pro-female preference in implicit attitudes emerged only when the names in the IAT were described as the names of male or female soldiers and not when the names were simply described as the names of males or females.

F (1, 128) = 29.12, MSE = 22412.47, $p < .001$. The means for this interaction suggest that the pro-female preference in implicit responses was present only for female participants, and a simple effects test conducted on the data collected from female participants confirmed this conclusion, F (1, 65) = 67.54, MSE = 17397.29, $p < .001$. This pro-female preference was not indicated in the data collected from male participants, F (1, 63) = .180, MSE = 27586.86, $p > .67$. A third significant interaction, depicted in Figure 3, involved the response key pairing and classification categories variables, F (1, 128) = 4.28, MSE = 22412.47, $p < .05$. The pattern of means for this interaction indicated that the pro-female preference existed across both types of categories. However, the effect was significant only in the soldier name condition. Simple effects tests decomposing this interaction indicate that the pro-female preference was significant when the categorization task involved soldier names, F (1, 66) = 34.15, MSE = 15498.73, $p < .001$, but was not significant when the general gender name categories were used, F (1, 62) = 2.59, MSE = 29772.26, $p > .11$.

Error rate data. Two average error rate scores were created for each participant. The first was tallied from each participant's incorrect responses to trials on which female names shared a response key with positive words, and male names shared a response key with unpleasant words. The second average was tallied from each participant's incorrect responses to trials on which female names shared a key with unpleasant words, and male names shared a response key with pleasant words. These two averages were entered into an ANOVA using the same five variables described in the analysis of the response latency measure.

A significant three-way interaction between response key pairing, participant gender, and the classification labels used, F (1, 128) = 7.56, MSE = .32, $p < .008$, indicated that different implicit gender attitude patterns emerged in different conditions of the study. The means for this interaction, depicted in Table 2, indicated that a pro-female preference was present in women regardless of the name condition (general gender or soldier). In contrast, men exhibited a pro-female preference only when they classified names by soldier gender. Men showed an in-group (pro-male) preference, not a pro-female preference, when classifying names by general gender. Follow-up analyses confirm this interpretation. A separate ANOVA conducted only on the data from female participants yielded a response key pairing main effect, F (1, 65) = 18.44, MSE = .27, $p < .001$. This effect was not significant in the ANOVA conducted on the data from the male participants, F (1, 63) = 1.22, MSE = .36, $p > .27$. Instead, the data for male participants revealed a significant interaction between response key pairing and classification categories used, F (1, 63) = 5.02, MSE = .36, $p < .03$.

The omnibus analysis also yielded a significant interaction between response key pairing and participant gender, F (1, 128) = 13.04, MSE = .32, $p < .001$. Interpretation of

Table 2

Error Rate Means (in %) for Interaction Between Response Key Pairing, Participant Gender, and Classification Categories for IAT Responses (Standard Deviations in Parentheses)

	Male Participants	
	Response Key Pairings	
	Female+pleas./ Male+unpleas.	*Female+unpleas./ Male+pleas.*
General Gender	8.29 (.09)	4.89 (.05)
Soldiers	7.50 (.09)	9.06 (.08)
	Female Participants	
	Response Key Pairings	
	Female+pleas./ Male+unpleas.	*Female+unpleas./ Male+pleas.*
General Gender	6.14 (.06)	11.29 (.09)
Soldiers	5.58 (.05)	8.37 (.06)

this effect is superseded by the significant three-way interaction involving these variables, so it will not be discussed further.[5]

Correlations Between Implicit and Explicit Measures

The differing result patterns for the implicit and explicit measures of gender attitudes suggest that responses to these different measures do not reflect the same underlying attitude construct. To more formally assess this possibility, we explored correlations among dependent measures. Some correlations used the mean latencies, error rates, and explicit attitude responses derived from each participant's responses. Other analyses correlated scores reflecting the difference between a participant's mean response in one condition of the experiment and their mean response in the other condition (e.g., the mean latency in the female+negative/male+positive condition minus the mean latency in the female+positive/male+negative condition correlated with the mean female explicit rating minus the mean explicit male rating).

Only one correlation between a variable derived from the IAT measures and a variable derived from the explicit attitude scale was significant. This correlation suggests that longer latencies in the female+positive/male+negative IAT condition were correlated with lower ratings of men on the explicit attitude scale, r (142) = –.18, p < .03. However, if these two variables were being driven by a common attitude construct, then the correlation between them should be positive (shorter latencies to the negative+male combination on the IAT ought to be related to lower male ratings). Hence, overall there is little evidence in these correlations to suggest that the implicit and explicit attitude measures assess the same underlying attitude construct.

Comparing the Latency and Error Measures

In comparison, the data do suggest that the two different variables derived from the IAT (error and latency) did, at least to some extent, reflect the same underlying construct. For example, using the difference scores described above, the correlation between the size of the IAT effect in the response latencies and the size of the IAT effect in the error rates was significant and positive, r (142) = .29, p < .001.

However, we note that despite this positive correlation, the latency and error data were only moderately correlated. The fact that these variables were not identical is also evident when one compares the results of the ANOVAs conducted on each of the variables. For example, participant age played a significant role in the latency analyses, but not in the error analyses. Interestingly, the correlation between the latency and error measures of the IAT effect also varied by age. There was a strong correlation between the two implicit measures among college students, r (83) = .49, p < .001, but not among the fifth-graders, r (57) = .06, p > .64.

We can only speculate about the cause of these differences. However, one possibility is that relative to the college students, fifth-graders might be more distractible across IAT trials. If this were the case, then one might expect the variability of the mean response latencies in the fifth-graders to be higher than the variability of the mean response latencies for the college students. This notion is supported by separate examination of the MSEs in the separate ANOVAs conducted on the latency data from the college students and the fifth-graders. The MSE for fifth-graders (39391.37) was much larger than the MSE for college students (11166.70). This disparity was significant when tested by means of an F_{max} test, F (1, 77) = 3.53, p < .01. This disparity did not characterize the error rate data, suggesting that the error rate data might be a more sensitive measure of implicit attitudes among fifth-graders than response latency.

DISCUSSION

The results of several previous studies (Eagly & Mladinic, 1989; Eagly et al., 1991; Haddock & Zanna, 1994) have suggested that people's attitudes toward women are more positive than their attitudes toward men. The data reported in the present paper replicated these results under conditions similar to those used in the previous studies, namely, with explicit measures using general gender categories with college students as participants. However, the results of our experiment also qualify this finding of a pro-female preference in several ways.

The first qualification is that, on the explicit attitude measure, a pro-female preference did not emerge for college men rating soldiers. Instead, college men had more positive attitudes toward male soldiers than they did toward female soldiers. This outcome may reflect societal changes. That is, the fact that this pro-male effect is limited to college-age male students may reflect the fact that young people are frequently exposed to women in traditional male roles (e.g., Xena, The Warrior Princess and the Pink Power Ranger) and may see nothing amiss with women in these roles. In addition, our younger participants may not, as yet, have fully assimilated the sex-typed gender roles specified by society. Furthermore, an in-group bias may induce women to favor women, regardless of occupational setting. Finally, as implied by the concept of benevolent sexism, college-age men may focus on notions such as the relative inappropriateness of women in male-dominated occupations, extant regulations that forbid women from engaging in combat, and women's stereotypical physical limitations (in comparison to men). Such ideas could cause men to lower their explicit attitudes toward female soldiers.

A second qualification on the generality of the pro-female preference appears on the implicit attitude measure. For female participants the pro-female preference was robustly present in both the latency and error measures

of implicit attitudes. In comparison, male participants showed no evidence of gender preference on the latency measure and actually provided some evidence of a pro-male preference on the error measure when rating the general gender categories. One possible explanation for the elimination of the pro-female preference in the implicit responses of men is that there are actually multiple influences at work. One may be a general pro-female preference; a second may be an in-group bias. For women, both influences work together, causing a large pro-female preference in the IAT responses. For men, these influences work at cross-purposes, canceling each other out. However, this does not capture the whole story: It does not account for the pro-female preference in men's responses to soldier names. Hence, additional factors affecting implicit responses may be at work in this condition.

The fact that the gender category manipulation had an impact on implicit attitude responses is consistent with the notion that the responses to the IAT are governed by the categories used in the task, not the individual names or words themselves (see Mitchell et al., 1999 for a similar demonstration using racial categories). However, the strengthening of the pro-female preference for names classified as soldiers in our study was unexpected. In this regard, it is useful to recall that the IAT measures the strength of implicit attitudes toward one category *relative to* another. Hence, the pro-female preference could be larger in the soldier condition because implicit attitudes toward male soldiers are more negative than implicit attitudes toward men in general. Of course, the effect could be larger because of a more positive implicit attitude toward female soldiers than toward women in general. Although the former strikes us as more plausible than the latter, we cannot empirically favor either interpretation at this time.

Another interesting outcome was that the implicit pro-female preference varied with participant age (at least on the latency measure). If this outcome is not an artifact of other variables (such as the increased distractibility of fifth-graders), then it might be explained by implicit learning mechanisms. Nonconscious implicit attitudes may result from repeated and consistent experience with the environment (through mechanisms such as classical conditioning and mere exposure effects). Hence, increased exposure to the environment (via aging) should produce stronger implicit attitudes. From this perspective, then, it makes sense that the pattern that emerges in the implicit data should be stronger in college students than in fifth-graders. In contrast, for at least some individuals the pro-female preference in explicit attitudes decreased with age. This outcome is consistent with the notion that many consciously controlled ideas (such as egalitarian norms and political correctness concerns) are used in the explicit attitude constructions of older individuals. Younger individuals are no doubt less likely to be concerned with such social niceties.

The different patterns of results for the implicit and explicit attitude measures suggest that the measures reflect different aspects of attitudes. This is further suggested by the absence of a correlation between the implicit and explicit attitude measures. However, as Eagly (personal communication, 1999) has noted, such outcomes might be an artifact of the different stimuli that are used in the implicit and explicit tasks. Other IAT results (e.g., Mitchell et al., 1999) suggest that responses to the IAT involve the categories instantiated by the stimuli and not the stimuli themselves. This suggests that the stimuli used in the implicit task are functionally equivalent to the stimuli used in the explicit task. Nonetheless, in view of Eagly's different stimuli argument and in lieu of further evidence, it is probably wise to exercise caution when interpreting the different results obtained on the implicit and explicit tasks.

One issue raised by this lack of correlation between implicit and explicit measures is the question "Which is the real attitude?" A simplistic answer would be to claim that the implicit attitudes are the real pipelines to the soul. After all, such assessments bypass the duplicitous strategies that might be used to keep one's true attitude from being exposed to public scrutiny. However, our own answer to this question would be both—and neither. In our view, an attitude is a construction derived from both conscious and nonconscious sources, and different attitude tasks differentially tap into these sources. We believe that in some cases, the "true" attitude might more appropriately be derived from people's attempts at modifying their initial, automatic evaluative responses. For example, Monteith (1993) suggested that nonprejudiced people are able to consciously use their explicit egalitarian ideals to overcome prejudiced attitudes that may have been implicitly learned from repeated exposure to a prejudiced environment. In what sense is this conscious correction not reflective of the individual's "true" attitude? After all, it is the explicit goal of the individual to be fair, and it is the automatically activated evaluative information that is "false" in the sense that it is not reflective of this goal. However, rather than become entangled in the false pursuit of the "true" attitude, we instead believe that a more fruitful approach to the problem would be to stick to identifying when, why, and how attitude responses will be governed by automatically activated evaluations that emerge in response to attitude objects, and when, why, and how those evaluations are altered by conscious thought.

Finally, in evaluating the results obtained in the research reported in this paper, it is useful to keep several limitations in mind. The first derives from the structure of the IAT. As currently construed, the IAT provides only a relative index, not absolute indices, of implicit attitudes toward social concepts. The enhancement of the pro-female preference for the soldier names that emerged in our experiment is a good example of this problem. Because of the nature of the IAT, it is unclear if this enhanced pro-female preference was caused by heightened positive attitudes toward the female soldier category or heightened

negative attitudes toward the male soldier category. Similarly, the implicit pro-female preference exhibited by some participants in our study simply reflects the ordering of the implicit attitudes toward men and females, but not the absolute level of those attitudes. Hence, we do not know if implicit attitudes toward women were positive, negative, or neutral; we know only that they were more positive than implicit attitudes toward men. The IAT could be more useful if it could be structured so that the outcome was more reflective of the absolute strength of an attitude toward a concept. Alternatively, given the utility of converging operations in research, alternative methods such as the techniques used by Fazio and his collaborators (e.g., Fazio & Dunton, 1997), may help to further clarify the exact quality and strength of implicit gender attitudes.

A second limitation derives from the cross-sectional design of our experiment. An obvious direction for further research would be a longitudinal study tracking the formation and change of gender attitudes across the life span. The data provided by such a study would be uncontaminated by extraneous differences (such as cross-generational differences in social norms) that contaminate the cross-sectional design that we used in the present study.

A third limitation is related to characteristics of our samples. In considering the implications of our results, it would be wise to remember that the age ranges of our groups were quite restricted, and that our samples overwhelmingly consisted of individuals of Caucasian ethnic heritages. It would clearly be desirable to extend these results to samples that include substantial proportions of individuals whose ethnic heritage is not Caucasian. Moreover, it would be desirable to extend the age range of the samples, certainly to middle-aged and older adults on the older end, as well as to children who are even younger than fifth-graders (although this latter desire may be restricted by technical issues, such as vocabulary). In addition, some might consider our sample sizes to be on the small side. However, clearly our sample size was sufficiently large to produce numerous effects, many of which were expected or predicted. However, it is possible that there are a few Type II errors that may have resulted from our design. This might be particularly true of our simple effects tests, where the statistical power of the tests was substantially reduced by the requisite partialling of the data. Either a greater number of participants, or the collection of a larger number of observations from each participant, would provide more statistical power in the detection of these effects.

A fourth limitation on our results also is a sampling issue, but it involves limited sampling of names and affectively toned words. One possibility is that the words that we used might be semantically associated with the male and female constructs, and it might be the case that these semantic associations, and not implicit attitudes, are responsible for the effects that emerged. For example, from one perspective, some of the negative words (abuse, corpse, death, pain) are associated with war, traditionally a playground for men. However, it is unclear whether this is really a problem; in fact, the IAT may work because of such associations. That is, one way to conceptualize implicit attitudes is that they are some combination of the affective responses to the associates of the attitude object. Nonetheless, the possibility that such effects can be solely determined by the semantic characteristics of the words, and not by the emotional characteristics, is a real one, and is made more plausible by the limited set of words that we used. However, we also note that other studies (e.g., Mitchell et al., 1999) have used a broader set of words and names and have essentially come up with similar results. This convergent validity suggests that our results are not an artifact of stimulus selection.

A final limitation of the present research is our use of a limited number of gender categorizations. We used the categories of male soldiers and female soldiers in the present study because of our intuition that the female soldier category was perceived to be extremely gender-inappropriate for women, so that it provided a good chance of yielding alterations in the patterns of explicit attitudes that emerged in comparison to the general female category. Obviously, other classification schemes may be used toward similar ends. For example, one possibility may be to use a classification scheme that involves placing names in leadership categories (e.g., female CEOs or female presidents vs. male CEOs or male presidents; see Carpenter & Banaji, 1999). Such schemes have the advantage of being more typical in society, and they may produce different patterns of implicit associations than observed with the soldier classification scheme. For example, people may not have more negative associations toward male CEOs than toward men in general, as may have happened with our use of male soldiers relative to men in general. In addition, female CEOs may have acquired negative associations (e.g., from television shows such as Melrose Place and real-life examples such as Leona Helmsley). The implication is that the pro-female preference in implicit attitudes observed with the general gender categorization scheme and with the soldier classification scheme might be eliminated or even reversed when using a classification scheme based on leadership or power.

These limitations simply point to the fact that a single study is not definitive; many other studies will be necessary to obtain a full view of attitudes toward women. However, the data collected in studies such as these continue to pose several delightful puzzles for researchers. One of these puzzles is the apparent disconnect between implicit and explicit attitudes. A second of these puzzles is the apparent disconnect between the pro-female preference in attitudes and the frequent anti-female bias in behavior. It is our hope to collect additional data with respect to these puzzles in the future and to begin to provide insight into the answers to these questions.

Initial submission: April 24, 2000
Initial acceptance: August 2, 2000
Final acceptance: October 30, 2000

NOTES

1. One reviewer was concerned that the terms that we used differed from the terms used by Eagly and her collaborators, which were "men" and "women." Given that our data substantially replicated Eagly's results, these term changes are probably not of great concern.
2. One reviewer wondered about the characteristics of the names, noting that outcomes in other studies may have been affected by the ways in which names differ other than gender (Kasof, 1993). Because the IAT measures liking for the *social categories* instantiated by the names, not the *individual stimuli,* such differences should not be critical. Furthermore, the male names that we used were slightly more common than the female names, which should lead to a pro-male preference in our study (assuming that frequency and liking are correlated). This is obviously inconsistent with the pro-female preference that repeatedly emerged.
3. There were interactions between participant gender, participant age, gender categories rated, scale order, and explicit gender attitude, F (1, 128) = 6.52, MSE = .31, p < .02; between scale order and explicit gender attitude, F (1, 128) = 5.21, MSE = .31, p < .03; and between scale order, explicit gender attitude, and age, F (1, 128) = 4.03, MSE = .31, p < .05. Careful examination of the means suggests that the theoretical implications of the effects that emerged for the other variables in the analysis are not substantially altered by these interactions with the order variable.
4. There was a significant interaction between the order in which the tasks were performed and the response key pairing, F (1, 128) = 15.50, MSE = 22412.47, p < .001. This interaction has previously been observed in the IAT paradigm (Greenwald et al., 1998), and reflects the fact that the second task in the sequence is difficult because it involves a reversal of the response pattern established on the first task. In any case, careful examination of the means suggests that the theoretical implications of the response key pairing effect are not altered by this interaction.
5. There was a significant interaction between the order in which the tasks were performed and the response key pairing, F (1, 128) = 10.90, MSE = .32, p < .001. As noted in footnote 4 this effect typically emerges in the IAT paradigm.

REFERENCES

Arkin, R. M., & Lake, E. A. (1983). Plumbing the depths of the bogus pipeline: A reprise. *Journal of Research in Personality, 17,* 81–88.

Broverman, I., Vogel, S. R., Broverman, D. M., Clarkson, F. E., & Rosenkrantz, P. S. (1972). Sex role stereotypes: A current appraisal. *Journal of Social Issues, 28,* 59–78.

Cacioppo, J. T., Marshall-Goodell, B. S., Tassinary, L. G., & Petty, R. E. (1992). Rudimentary determinants of attitudes: Classical conditioning is more effective when prior knowledge about the attitude stimulus is low than high. *Journal of Experimental Social Psychology, 28,* 207–233.

Carpenter, S. J., & Banaji, M. R. (1999). Implicit attitudes toward female leaders. Paper presented at the 71st meeting of the Midwestern Psychological Association, Chicago, IL.

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

Dasgupta, N., McGhee, D. E., Greenwald, A. G., & Banaji, M. R. (2000). Automatic preference for White Americans: Eliminating the familiarity explanation. *Journal of Experimental Social Psychology, 36,* 316–328.

Deaux, K., & LaFrance, M. (1998). Gender. In D. T. Gilbert, S. T. Fiske, & G. Lindzey (Eds.), *The handbook of social psychology* (4th ed., Vol. 1, pp. 788–827). Boston: McGraw-Hill.

Eagly, A. H., & Mladinic, A. (1989). Gender stereotypes and attitudes toward women and men. *Personality and Social Psychology Bulletin, 15,* 543–558.

Eagly, A. H., Mladinic, A., & Otto, S. (1991). Are women evaluated more favorably than men? An analysis of attitudes, beliefs and emotions. *Psychology of Women Quarterly, 15,* 203–216.

Fazio, R. H., & Dunton, B. C. (1997). Categorization by race: The impact of automatic and controlled components of racial prejudice. *Journal of Experimental Social Psychology, 33,* 451–470.

Glick, P., & Fiske, S. T. (1997). Hostile & benevolent sexism: Measuring ambivalent sexist attitudes toward women. *Psychology of Women Quarterly, 21,* 119–135.

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,* 1465–1479.

Haddock, G., & Zanna, M. P. (1994). Preferring housewives to feminists: Categorization and the favorability of attitudes toward women. *Psychology of Women Quarterly, 18,* 25–52.

Huston, A. C. (1985). The development of sex typing: Themes from recent research. *Developmental Review, 5,* 1–17.

Kasof, J. (1993). Sex bias in the naming of stimulus persons. *Psychological Bulletin, 113,* 140–163.

McKee, J. P., & Sherriffs, A. C. (1957). The differential evaluation of males and females. *Journal of Personality, 25,* 356–371.

Mitchell, J. P., Nosek, B. A., & Banaji, M. R. (1999). Dissociated implicit attitudes: Examples from race, gender and profession. Paper presented at the 71st meeting of the Midwestern Psychological Association, Chicago, IL.

Monteith, M. J. (1993). Self-regulation of prejudiced responses: Implications for progress in prejudice-reduction efforts. *Journal of Personality and Social Psychology, 65,* 469–485.

Murphy, S. T., Zajonc, R. B., & Monahan, J. L. (1995). Additivity of nonconscious affect: Combined effects of priming and exposure. *Journal of Personality and Social Psychology, 69,* 589–602.

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

Rudman, L. A., Greenwald, A. G., Mellott, D. S., & Schwartz, J. L. K. (1999). Measuring the automatic components of

prejudice: Flexibility and generality of the Implicit Association Test. *Social Cognition, 17,* 437–465.

Stangor, C., Sullivan, L. A., & Ford, T. E. (1991). Affective and cognitive determinants of prejudice. *Social Cognition, 9,* 359–380.

Wilson, T. D., Hodges, S. D., & Lafleur, S. J. (1995). Effects of introspecting about reasons: inferring attitudes from accessible thoughts. *Journal of Personality and Social Psychology, 69,* 16–28.