Research Article

AUTOMATIC STEREOTYPING

Mahzarin R. Banaji[1] and Curtis D. Hardin[2]

[1]*Yale University and* [2]*University of California, Los Angeles*

Abstract—*Two experiments tested a form of automatic stereotyping. Subjects saw primes related to gender (e.g.,* mother, father, nurse, doctor*) or neutral with respect to gender (e.g.,* parent, student, person*) followed by target pronouns (stimulus onset asynchrony = 300 ms) that were gender related (e.g.,* she, he*) or neutral (*it, me*) or followed by nonpronouns (*do, all*; Experiment 2 only). In Experiment 1, subjects judged whether each pronoun was male or female. Automatic gender beliefs (stereotypes) were observed in faster responses to pronouns consistent than inconsistent with the gender component of the prime regardless of subjects' awareness of the prime–target relation, and independently of subjects' explicit beliefs about gender stereotypes and language reform. In Experiment 2, automatic stereotyping was obtained even though a gender-irrelevant judgment task (pronoun/not pronoun) was used. Together, these experiments demonstrate that gender information imparted by words can automatically influence judgment, although the strength of such effects may be moderated by judgment task and prime type.*

Based on recent theory and research on the role of unconscious processes in beliefs about social groups (Banaji & Greenwald, 1994; Bargh, 1994; Greenwald & Banaji, 1995), we report two experiments that provide stricter tests than previously conducted of a form of automatic stereotyping. Several recent experiments have demonstrated that stereotyping can occur implicitly, without subjects' conscious awareness of the source or use of stereotypic information in judgment (Banaji & Greenwald, 1995; Banaji, Hardin, & Rothman, 1993; Devine, 1989). In this article, we focus on a particular brand of stereotyping than can occur even when the perceiver retains awareness of the source of influence on judgment, but is unable to readily control the stereotypic response. Stereotyping, like other cognitive processes, consists of both automatic and controlled components, and the particular form of automaticity that is involved (e.g., awareness, intentionality, efficiency, and controllability) has been of recent interest (see Bargh, 1994). In the present experiments, we demonstrate that gender, as lexically coded in English, can operate automatically in judgment, even when the primary (denotative) meaning is not about gender.[1]

The semantic priming procedure is commonly used to examine automatic information processing and, in particular, to reveal the strength of association between two concepts that exists independently of conscious thought. Developed more than 20 years ago, this procedure has led to important discoveries about attention, signal processing, and semantic memory (Meyer & Schevaneveldt, 1971; Neely, 1977; Posner & Snyder, 1975). The first of these tests showed the now well-known effect that response latency to a target word is facilitated to the extent to which a prime word that appears prior to the target word is semantically related to it. In addition, the technique has recently been successfully adapted to demonstrate the operation of automatically activated attitudes or evaluations (Bargh, Chaiken, Govender, & Pratto, 1992; Fazio, Sanbonmatsu, Powell, & Kardes, 1986; Perdue & Gurtman, 1990).

Our primary interest lies in beliefs, and for the present research, we adapted the semantic priming procedure to provide a strict test of the extent to which beliefs about gender (i.e., gender stereotypes) operate automatically. Two words were presented in close succession, and the relationship between them was captured by reaction time (RT) to judge the second (target) word. In both experiments, the central empirical question of interest was, what is the influence of the gender code of a prime on speeded judgments of gender-consistent or gender-inconsistent targets? Faster judgments on targets that follow gender-congruent primes than on targets that follow gender-incongruent primes (i.e., gender-based priming) are taken as evidence for the automatic use of gender stereotypes.[2] We also examine related questions, such as (a) the relationship between automatic stereotyping and traditionally used explicit stereotyping measures, (b) the role of awareness of gender as a potential source of influence on performance, (c) the gender relevance of the judgment task, and (d) the gender strength of the primes.

Although variations of the semantic priming procedure have been used in previous research on stereotypes, the experimental procedures of these studies did not adhere to conventional standards for revealing automatic information use.[3] For example, Dovidio, Evans, and Tyler (1986) presented a prime (*black* or *white*) followed by a target (*intelligent* or *lazy*) and asked subjects if the target "could ever be true" of the prime category. They found that white subjects were reliably faster to respond to stereotype-related traits than stereotype-unrelated traits following the primes *white* and *black*. Such experiments

Address correspondence to Mahzarin Banaji, Department of Psychology, Yale University, P.O. Box 208205, New Haven, CT 06520-8205, e-mail: mbanaji@minerva.cis.yale.edu, or Curtis Hardin, Department of Psychology, UCLA, 1282A Franz Hall, Box 951563, Los Angeles, CA 90024-1563, e-mail: hardin@psych.ucla.edu. All materials may be obtained from the authors.

1. These demonstrations, although showing evidence for the automatic use of gender information, should not be taken to imply that seemingly automatic responses can never be controlled. The effects of automatically activated information are controllable under theoretically specified conditions (Bargh, 1994; Blair & Banaji, 1995).

2. The term *stereotype* has been a construct of changing meaning in social psychology, and our use of it is in keeping with recent definitions that reduce it to refer to beliefs about the attributes of social groups (Ashmore & Del Boca, 1981; Greenwald & Banaji, 1995).

3. In experiments in which the currently specified conditions of automaticity were met, the findings addressed the role of automatic evaluation rather than automatic belief (Gaertner & McLaughlin, 1983; Perdue & Gurtman, 1990, Experiment 2).

were critical in setting the stage for the present study, but did not strictly test automatic stereotyping. Most notably, the time between the onset of the prime and target (stimulus onset asynchrony, or SOA) in the previous studies was long enough to allow strategic processing, casting doubt about the automaticity of the process being measured. In the present experiments, we used a 300-ms SOA, a condition known to capture relatively automatic processes (Neely, 1977, 1991).

Further, Dovidio et al. (1986) required subjects to deliberately link the prime and target by asking whether the target "could ever be true" of the prime category, and Neely (1977) demonstrated that such explicit expectations do affect RT under long SOAs such as those used by Dovidio et al. (1986). Although these judgment tasks demonstrate important differences in judgment latencies for stereotyped traits following the prime *black* or *white*, these tasks do not index automatic processes that may occur outside conscious deliberation of prime–target relationships. In the present experiments, we used judgment tasks that required no attention to the relationship between prime and target. Indeed, subjects were instructed to ignore the prime word and classify the target word as a male or female pronoun (Experiment 1) or a pronoun or not a pronoun (Experiment 2).

In addition, the experiments we report differ from previous research in the number and type of stimuli that were used. Instead of the repeated presentation of two-category labels as primes (*black* or *white*, *young* or *old*), we used 150 primes signifying gender in a variety of ways, including those associated to gender by definition (e.g., *mother, father, waiter, waitress*) or by normative base rates (e.g., *doctor, nurse, mechanic, secretary*), or neutral with respect to gender (e.g., *humanity, citizen, people, cousin*). The larger set of primes more fully represents the social category, allows use of primes other than category labels alone, and permits a comparison of the strength of primes that denote gender and of primes that connote gender. Primes also included so-called generic masculine terms (e.g., *mankind*) to allow a test of whether such words automatically connote maleness or perform the more inclusive function that critics of nonsexist language assert is the case.

In the choice of target words, we departed from the almost exclusive reliance of past research on trait adjectives. In both experiments, we used pronouns because they inescapably mark gender (e.g., *she, he*). However, the judgment task itself differed across the two experiments in whether the decision focused on gender (male or female; Experiment 1) or grammatical form (pronoun or not pronoun; Experiment 2). To date, no studies of stereotyping have used a task that does not focus attention on the category of interest (e.g., gender, race). A finding that automatic stereotyping occurs even when a gender-irrelevant task is used would attest to the potency of automatic gender stereotypes.

EXPERIMENT 1

Experiment 1 tested whether gender information in words is automatically used in judgment as assessed by faster response times when the genders of the prime and target words match (e.g., *doctor-he, nurse-she*) than mismatch (e.g., *doctor-she, nurse-he*). In addition, measures of beliefs about explicit gender stereotypes, language reform, and the influence of gender in everyday life were included to test the relationship between automatic stereotyping and more traditional explicit measures of gender stereotyping.

Method

Subjects

Sixty-eight subjects (32 female, 36 male) from the introductory psychology pool at Yale University participated in partial fulfillment of a course requirement.

Materials and apparatus

The experimental task was administered on IBM-PS2 microcomputers running Micro-Experimental Laboratory software (Schneider, 1990). Subjects entered judgments on protruding keys, marked "M" and "F," affixed to the *f* and *j* keys. Key position was reversed for half the subjects.

Two hundred primes were divided evenly among four categories: male related, female related, neutral with respect to gender, and nonword letter string (*ZZZZZ*). Within each of the first three prime categories, words were chosen to appear virtually equally from two subcategories. The first subcategory contained words associated to gender by normative base rates. These words were chosen on the basis of 1990 census data indicating occupations that were heavily skewed (over 90%) toward the participation of either females (e.g., *nurse, secretary*) or males (e.g., *doctor, mechanic*) or that had equal participation (e.g., *reporter, postal clerk*). In addition, several other words having strong stereotypical associations to one gender or the other were included (e.g., *feminist, god*). The second subcategory contained words associated to gender by definition, that is, words that expressly refer to gender (e.g., *woman, man*), kinship terms (e.g., *mother, father*), and titles (e.g., *mr, mrs, king, queen*). Within this subcategory, words containing male morphemes (e.g., *salesman*), female morphemes (e.g., *salesgirl*), or neutral morphemes (e.g., *chairperson*) were also included. Targets were the six most common pronouns in English, half male (*he, him, his*) and half female (*she, her, hers*).

Three measures were designed to assess explicit beliefs regarding gender stereotypes, language reform, and the influence of gender in peoples' lives.[4]

Design and procedure

For each trial, events occurred in the following order: First, an orientation symbol (+) appeared for 500 ms. Then the prime word appeared for 200 ms, followed by a blank screen for 100 ms. Finally, the target pronoun appeared and remained on the screen until a response was entered. Subjects made 432 judgments (not including practice and buffer trials) divided equally among the eight prime-target categories (prime: male, female,

4. For a description of these measures, see Hardin and Banaji (in press).

neutral, nonword; target: male pronoun, female pronoun) within each of three blocks of trials that were counterbalanced across subjects. Prime and target stimuli were paired randomly for each subject.

The design was a 4 (prime gender: female, male, neutral, nonword) × 2 (target gender: female, male) × 2 (subject gender: female, male) mixed factorial, with subject gender the between-subjects factor. Subjects judged each pronoun as either male or female. They were instructed to ignore the primes and judge the targets as quickly and accurately as possible. Subjects then completed the three explicit measures of gender beliefs. Finally, they were probed for their awareness of the hypotheses and debriefed.

Results and Discussion

Reported results are based on correct judgments, excluding responses that were extreme outliers. Consistent with other studies employing this procedure, the error rate was low (1,117 of 29,502 judgments, or 3.8%). RTs greater than 3 *SD* above the mean (>1,300 ms) were identified as outliers and excluded (208 trials, or 0.7%). In sum, 95.6% (28,193 judgments) were retained in the reported analyses. The pattern of results is unchanged when these data are included. To achieve a better approximation to the normal distribution, analyses were performed on a log transformation of the raw RT latencies. Thirty-seven of 68 subjects were aware of the gender relationship between primes and targets. However, consistent with the assumption that this procedure reflects relatively automatic processing, the pattern of results was identical for both aware and unaware subjects, all $Fs < 1$.

As shown in Figure 1, the predicted gender priming effect was obtained, indicating that judgment was faster when target gender matched than mismatched prime gender. The omnibus Prime Gender (female, male, neutral, nonword) × Target Gender (female, male) × Subject Gender (female, male) three-way analysis of variance yielded the predicted Prime Gender × Target Gender interaction, $F(3, 198) = 72.25$, $p < .0001$. The specific Prime Gender × Target Gender interaction (excluding the

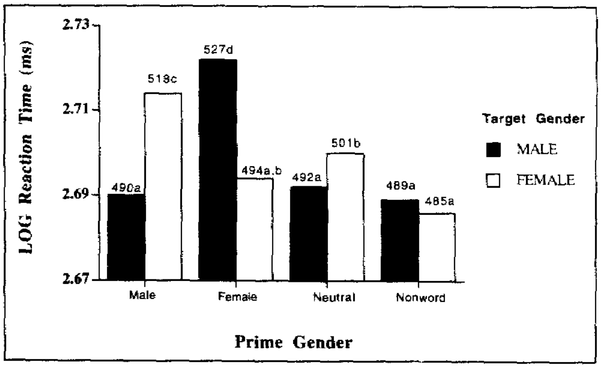

Fig. 1. Mean reaction time to judge words as male or female as a function of prime gender and target gender (Experiment 1, $n = 68$). Bars with shared subscripts are not significantly different from each other ($p > .05$).

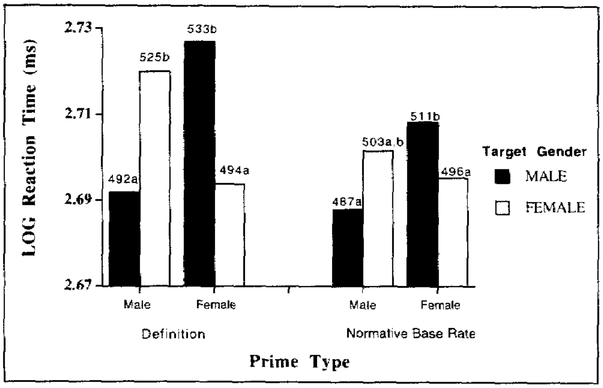

Fig. 2. Mean reaction time to judge words as male or female as a function of prime type, prime gender, and target gender (Experiment 1, $n = 68$). Bars with shared subscripts are not significantly different from each other ($p > .05$).

neutral conditions) was also reliable, $F(1, 66) = 117.56$, $p < .0001$. Subjects were faster to judge male pronouns after male than female primes, $t(67) = 11.59$, $p = .0001$, but faster to judge female pronouns after female than male primes, $t(67) = 6.90$, $p = .0001$. In addition, subjects were faster to respond to targets preceded by male ($M = 2.702$) than female primes ($M = 2.708$), $F(1, 66) = 7.52$, $p < .01$. No other reliable main effects or interactions were obtained as a function of either subject gender or target gender ($Fs < 1$).

The automatic gender priming effect was obtained for primes related to gender both by definition (e.g., *mother, father, man, woman*), $F(1, 67) = 103.97$, $p < .0001$, and by normative base rate, $F(1, 67) = 18.61$, $p < .0001$. However, as shown in Figure 2, the gender priming effect was significantly larger for primes related to gender by definition, as revealed by the three-way Prime Type (definition, normative base rate) × Prime Gender (female, male) × Target Gender (female, male) interaction, $F(1, 67) = 13.67$, $p < .0005$.

Generic masculine terms contributed to the automatic gender priming effect. After primes containing the morpheme *man* (e.g., *fireman, mankind, human, man*), judgments were faster for male pronouns ($M = 2.687$) than female pronouns ($M = 2.709$), $t(67) = 4.06$, $p < .0001$. The relationship also held under the most conservative analysis, in which terms that are sometimes used to refer only to men (e.g., *man, fireman*) were excluded. Primes considered to be generic masculine terms in virtually all contexts (e.g., *mankind, layman*) produced faster judgments for male pronouns ($M = 2.689$) than female pronouns ($M = 2.712$), $t(67) = 2.18$, $p < .05$.

Finally, we examined terms that differed in no way except for the gender of their suffix (e.g., *chairman, chairwoman, chairperson*). As expected, the gender of the suffix did influence response latencies as indicated by a Prime Gender (male, female, neutral) × Target Gender (male, female) interaction, $F(2, 112) = 11.59$, $p < .0001$. Judgments were faster for male pronouns after words with male ($M = 2.693$) than female ($M = 2.722$) suffixes, $t(67) = 3.29$, $p < .01$, whereas judgments were marginally faster for female pronouns after primes with female ($M = 2.694$) than male ($M = 2.716$) suffixes, $t(67) = 1.81$, $p =$

.07. In addition, judgments were faster when primes with female suffixes were followed by female pronouns (M = 2.694) than male pronouns (M = 2.722), $t(67) = 2.82$, $p < .01$. Interestingly, neutral *-person* suffixes after the identical words did not produce equivalent responses to female and male pronouns. Instead, after these primes, subjects were still faster to judge male targets (M = 2.704) than female targets (M = 2.722), $t(67) = 2.64$, $p < .05$.[5]

Relations between explicit beliefs and automatic gender stereotyping were examined by computing a correlation between each of the three explicit belief measures and a gender priming score, which was calculated by subtracting log RT for gender-congruent trials from log RT for gender-incongruent trials. None of the three correlations of explicit measures with the priming score was significant (language reform: $r[67] = -.003$, $p = .978$; role of gender in everyday life: $r[68] = -.050$, $p = .686$; explicit gender stereotypes; $r[66] = .037$, $p = .767$). This result is consistent with other research demonstrating a lack of correspondence between explicit and implicit measures of stereotyping (Banaji & Greenwald, 1995).

In sum, Experiment 1 provided evidence for automatic gender stereotyping using a broad range of primes and using time and task parameters that reflect automatic information use. The effect occurred regardless of subjects' awareness of the prime–target relation, and independently of explicit beliefs about gender stereotypes. The effect was also obtained for both primes related to gender by definition and primes related to gender by normative base rate, although not surprisingly the effect was larger for primes related to gender by definition.

EXPERIMENT 2

Participants in Experiment 1 judged whether each target was male or female, thereby focusing attention on the gender of the target. This form of the judgment task is quite conventional. For example, when theoretical interest has focused on the semantic link between prime and target, the commonly used judgment task is a lexical decision (word/nonword; Neely, 1991). Likewise, when the interest is in the evaluative component of the prime and target, the task is typically a good/bad judgment (Bargh et al., 1992; Fazio et al., 1986; Greenwald, Klinger, & Liu, 1989; Perdue & Gurtman, 1990). However, stronger evidence for automaticity may be obtained if the effect is observed when the judgment task is unrelated to the dimension of the prime–target relationship. For example, Bargh, Chaiken, Raymond, and Hymes (in press) showed that the automatic evaluative effect is obtained even when the judgment involves mere pronunciation, a task unrelated to evaluation. Hence, in Experiment 2, the judgment task was a pronoun/not pronoun decision, unrelated to gender.

Method

Subjects

Sixty subjects (29 female, 31 male) from Yale University participated in exchange for $5 or in partial fulfillment of a course requirement.

Materials, design, and procedure

For this experiment, 120 of the primes used in Experiment 1, representing male (40 primes), female (40 primes), and neutral (40 primes) categories, were selected. Of the four target pronouns used, *she* and *he* allowed the comparisons of primary interest. The pronoun *it* was included because it is the most frequently occurring gender-neutral pronoun, and *me* was included for exploratory purposes to examine a possible relationship between prime gender and subject gender (cf. Markus, 1977). The four nonpronouns (*is, do, as, all*) were chosen to match the critical targets in length, number of syllables, and frequency (Kuçera & Francis, 1967).

In all, each subject made 720 experimental judgments divided into five blocks of trials, counterbalanced across subjects. For 480 of these judgments, the correct response to the question "Is this a pronoun?" was "yes," and for 240, the correct answer was "no." Each prime was paired with (a) both critical "yes"-response targets (i.e., *she, he*), (b) both noncritical "yes"-response targets (i.e., *it, me*), and (c) two of the four "no"-response targets (i.e., *do, all, is, as*). For each subject within each block, prime and target items were randomly associated. After completing the priming task, subjects were probed for awareness regarding the hypotheses and debriefed.

Results and Discussion

As before, results are based on a log transformation of the raw RT latencies for correct judgments, excluding outliers (RT > 1,300 ms or > 3 SD above the mean; 1.4% of the total). Also as in Experiment 1, the error rate was low (370 of 28,800 "yes" judgments, or 1.3%; 928 of 43,200 total judgments, or 2.1%); 97.7% (28,134) of the "yes" judgments were retained in the reported analyses. Seven of the 60 subjects revealed some knowledge of a possible gender relation between the prime and target words. Again, however, the pattern of results was identical for both aware and unaware subjects, but no statistical significance tests were conducted because of the small number of aware subjects.

As Figure 3 shows, the predicted gender priming effect was obtained, indicating that judgment was faster when target gender matched than mismatched prime gender. The omnibus 3 (prime, gender: male, female, neutral) × 4 (target gender: *she, he, it, me*) × 2 (subject gender: male, female) analysis of variance yielded the predicted Prime Gender × Target Gender interaction, $F(6, 336) = 3.66$, $p < .01$. In addition, a Subject Gender × Target Gender interaction indicated that subjects were faster to respond to targets that matched rather than mismatched their own gender, $F(3, 168) = 3.58$, $p < .02$.[6] No difference in responding to the male and female targets was observed for dubiously neutral primes such as *layman* and *man-*

5. However, we found (Hardin & Banaji, in press) no bias favoring males in a similar experiment using first names as targets.

6. A similar finding was reported by Zarate and Smith (1990). In addition, a main effect of prime gender indicated that subjects' responses were fastest following male primes and slowest following female primes, $F(2, 112) = 3.89$, $p < .03$. A main effect of target gender indicated that subjects were slower to respond to the target *it* than to *she, he,* and *me*, $F(3, 168) = 154.33$, $p < .0001$.

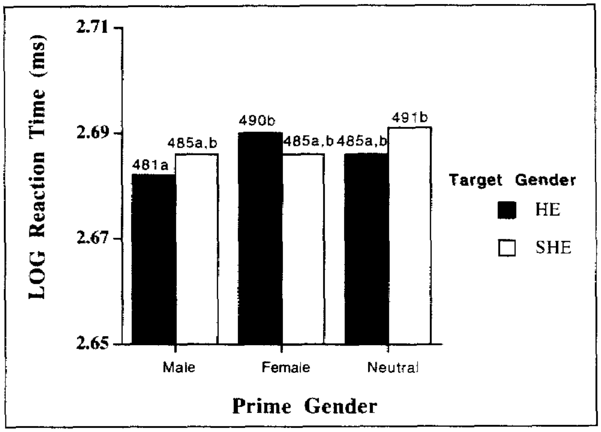

Fig. 3. Mean reaction time to judge words as pronouns or not pronouns as a function of prime gender and target gender (Experiment 2, $n = 58$). Bars with shared subscripts are not significantly different from each other ($p > .05$).

kind. There was also no main effect of subject's sex, $F(1, 56) = 1.49$.

The more specific interaction of prime gender by target gender (excluding neutral primes) was also significant, indicating that subjects were faster to judge targets in gender-congruent prime-target pairs than in gender-incongruent pairs, $F(1, 56) = 4.63$, $p < .04$. Again, the Subject Gender × Target Gender interaction was reliable, indicating that subjects were faster to respond to the target pronouns that were consistent than inconsistent with their own social category, $F(1, 56) = 17.95$, $p < .0001$.

However, these 2 two-way interactions were qualified by a three-way Subject Gender × Prime Gender × Target Gender interaction, $F(1, 56) = 4.15$, $p < .05$. For purposes of clarity, we describe results separately for primes related to gender by definition and primes related to gender by normative base rates. Analyses of primes related to gender by definition (e.g., *mother, father, waitress, waiter*) yielded the gender priming effect unmoderated by subject gender (Fig. 4, left panel). RT was smaller when prime gender and target gender were congruent than incongruent, as indicated by a reliable two-way interaction, $F(1, 56) = 8.70$, $p < .005$. Subjects were faster to identify *he* when primes were male than female, $t(59) = 2.44$, $p < .02$, but faster to identify *she* than *he* when the primes were female, $t(59) = 2.53$, $p < .02$. In addition, subjects were faster to identify targets that matched their own gender, as indicated by the reliable interaction between subject gender and target gender, $F(1, 56) = 7.43$, $p < .01$.

Analyses of primes related to gender by normative base rates (e.g., *secretary, mechanic, doctor, nurse*) suggest limitations to the generality of automatic gender priming under conditions in which the task does not require subjects to focus on the dimension of gender (see Fig. 4, right panel). Although reliable effects were obtained for the Subject Gender × Target Gender interaction, $F(1, 56) = 14.02$, $p < .0001$, and there was a main effect of prime gender, $F(1, 56) = 6.49$, $p = .01$, both were qualified by a marginal three-way Subject Gender × Prime Gender × dTarget Gender interaction, $F(1, 56) = 3.49$, $p < .07$. Male subjects were faster to identify *he* than *she* regardless of prime gender, $F(1, 29) = 7.44$, $p = .01$, and faster to identify targets after male than female primes, $F(1, 29) = 9.21$, $p < .01$. In contrast, female subjects were faster to identify *she* than *he* after female primes, $t(28) = 2.81$, $p < .01$, and faster to identify *he* after male than female primes, $t(28) = 2.18$, $p < .05$. Females were also faster, in general, to identify *she* than *he*, $F(1, 27) = 6.70$, $p < .05$.

GENERAL DISCUSSION

These two experiments provide the first strict tests of a form of automatic stereotyping. Using a large number and wide range of stimuli, we demonstrated that judgments of targets that follow gender-congruent primes are made faster than judgments of targets that follow gender-incongruent primes. This effect was obtained despite subjects' deliberate attempt to ignore the prime, regardless of whether subjects were aware or unaware of the gender relation of prime-target pairings, independently of subjects' explicit beliefs about gender, regardless of whether the judgment was gender relevant or irrelevant, and on both words that are gender related by definition and words that are gender related by normative base rates.

The results, however, also show two moderators of the gender priming effect. First, the effect was stronger when the judgment was gender relevant (e.g., male or female pronoun?) than gender irrelevant (e.g., pronoun or not pronoun?). Further research will investigate whether this difference also obtains on other forms of gender-irrelevant tasks, such as pronunciation. Second, the gender priming effect was stronger for primes related to gender by definition (e.g., *mother, father*) than by normative base rate (e.g., *doctor, nurse*). In Experiment 1, for example, the effect size for definition primes was large (Cohen's $d = .78$), whereas for normative-base-rate primes, the effect size was moderate ($d = .47$). This difference reflects the differential strength of the two types of primes in evoking gender. Words that are exclusively reserved to denote gender will produce stronger priming than words that connote gender (for a replication with names as targets, see Hardin & Banaji, in press).

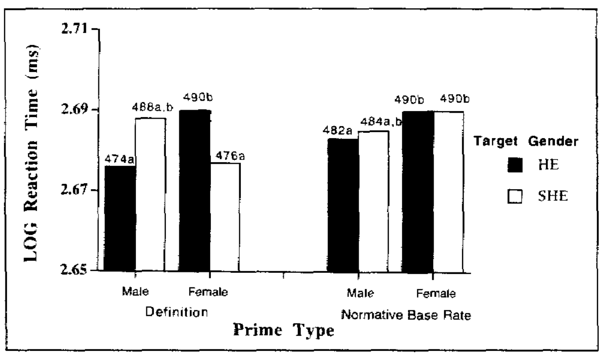

Fig. 4. Mean reaction time to judge words as pronouns or not pronouns as a function of prime type, prime gender, and target gender (Experiment 2, $n = 58$). Bars with shared subscripts are not significantly different from each other ($p > .05$).

Sapir (1963) commented that one of the important functions of language is to repeatedly declare to society the psychological status of its members. These experiments show the automatic effects of such repeated linguistic declarations, in particular, those that convey the social psychological positions that are occupied through gender. A noteworthy aspect of the gender priming effect observed in Experiment 1 is that the effect can obtain not only when the primes denote gender (*man, woman*), but also when they more tacitly connote gender (*mechanic, nurse*). That gender-signifying information permeates thought sufficiently to influence judgment points to the fundamental nature of gender as a category in verbally communicated thought. This article is not the place to catalogue the various ways in which gender is coded in most languages, but we note that English stands out as one language that has received a quite extensive analysis of what might be called the "genitalia of language" (Baron, 1986): gender-signifying words, gender-specific pronouns, and the covert presence of gender in grammatical structure. We expect that such automatic gender priming effects are best observed in languages that provide extensive and deep coding of gender in grammar and semantics (Hardin & Banaji, 1993). Further evidence for the generality of automatic gender-stereotyping effects might be obtained by demonstrating such effects independently of language (i.e., through the use of nonverbal, pictoral stimuli that denote and connote gender). Such effects would be especially important in revealing the degree to which the present effect is a function of gendered language per se or gender stereotypes more generally.

Although research on beliefs and attitudes has usually depended on direct, verbal measures of stereotypes (see Greenwald & Banaji, 1995), response latencies may provide a more indirect measure of stereotype strength. A case for RT as a measure of attitude or evaluation has already been effectively made (see Bargh et al., 1992; Fazio et al., 1986; Perdue & Gurtman, 1990), and other investigators have used RT as an indicator of stereotypes (Dovidio et al., 1986). However, these experiments, in conjunction with others (Blair & Banaji, 1995; Hardin & Banaji, in press), demonstrate the operation of beliefs under conditions that meet currently accepted standards for measuring automatic processes. Such measures are likely to increasingly complement the more traditional measures of evaluation and belief, especially as their validity and feasibility are further established.

Acknowledgments—This research was supported in part by Grants DBC 9120987 and SBR 9422291 from the National Science Foundation. We are grateful to Lisa Driscoll and John Beauvais for assistance with data collection; Kimberly Hinds and Irene Blair for assistance with programming; and R. Bhaskar, Irene Blair, Richard Hackman, and Eliot Smith for comments on a previous draft.

REFERENCES

Ashmore, D.D., & Del Boca, F.K. (1981). Conceptual approaches to stereotypes and stereotyping. In D.L. Hamilton (Ed.), *Cognitive processes in stereotyping and intergroup behavior* (pp. 1–36). Hillsdale, NJ: Erlbaum.

Banaji, M.R., & Greenwald, A.G. (1994). Implicit stereotypes and prejudice. In M.P. Zanna & J.M. Olson (Eds.), *The psychology of prejudice: The Ontario Symposium* (Vol. 7, pp. 55–76). Hillsdale, NJ: Erlbaum.

Banaji, M.R., & Greenwald, A.G. (1995). Implicit stereotyping in false fame judgments. *Journal of Personality and Social Psychology, 68,* 181–198.

Banaji, M.R., Hardin, C., & Rothman, A.J. (1993). Implicit stereotyping in person judgment. *Journal of Personality and Social Psychology, 65,* 272–281.

Bargh, J.A. (1994). The four horsemen of automaticity: Awareness, intention, efficiency, and control in social cognition. In R.J. Wyer & T. Srull (Eds.), *Handbook of social cognition* (2nd ed., pp. 1–40). Hillsdale, NJ: Erlbaum.

Bargh, J.A., Chaiken, S., Govender, R., & Pratto, F. (1992). The generality of the automatic activation effect. *Journal of Personality and Social Psychology, 62,* 893–912.

Bargh, J.A., Chaiken, S., Raymond, P., & Hymes, C. (in press). The automatic evaluation effect: Unconditional automatic attitude activation with a pronunciation task. *Journal of Experimental Social Psychology.*

Baron, D. (1986). *Grammar and gender.* New Haven, CT: Yale University Press.

Blair, I.V., & Banaji, M.R. (1995). *Automatic and controlled processes in gender stereotyping.* Unpublished manuscript, Yale University, New Haven, CT.

Devine, P. (1989). Stereotypes and prejudice: Their automatic and controlled components. *Journal of Personality and Social Psychology, 56,* 5–18.

Dovidio, J.F., Evans, N., & Tyler, R.B. (1986). Racial stereotypes: The contents of their cognitive representations. *Journal of Experimental Social Psychology, 22,* 22–37.

Fazio, R.H., Sanbonmatsu, D.M., Powell, M.C., & Kardes, F.R. (1986). On the automatic activation of attitudes. *Journal of Personality and Social Psychology, 50,* 229–238.

Gaertner, S.L., & McLaughlin, J.P. (1983). Racial stereotypes: Associations and ascriptions of positive and negative characteristics. *Social Psychology Quarterly, 46,* 23–30.

Greenwald, A.G., & Banaji, M.R. (1995). Implicit social cognition: Attitudes, self-esteem, and stereotypes. *Psychological Review, 102,* 4–27.

Greenwald, A.G., Klinger, M.R., & Liu, T.J. (1989). Unconscious processing of dichoptically masked words. *Memory and Cognition, 17,* 35–47.

Hardin, C., & Banaji, M.R. (1993). The influence of language on thought. *Social Cognition, 11,* 277–308.

Hardin, C., & Banaji, M.R. (in press). Gender in language and thought. *Social Cognition.*

Kuçera, N., & Francis, W.N. (1967). *Computational analysis of present day American English.* Providence, RI: Brown University Press.

Markus, H. (1977). Self-schemata and processing information about the self. *Journal of Personality and Social Psychology, 35,* 63–78.

Meyer, D., & Schevaneveldt, R. (1971). Facilitation in recognizing pairs of words: Evidence of a dependence between retrieval operations. *Journal of Experimental Psychology, 90,* 227–234.

Neely, J.H. (1977). Semantic priming and retrieval from lexical memory: Roles of inhibitionless spreading activation and limited-capacity attention. *Journal of Experimental Psychology: General, 106,* 226–254.

Neely, J.H. (1991). Semantic priming effects in visual word recognition: A selective review of current findings and theories. In D. Besner & G. Humphreys (Eds.), *Basic processes in reading: Visual word recognition* (pp. 264–336). Hillsdale, NJ: Erlbaum.

Perdue, C., & Gurtman, M. (1990). Evidence for the automaticity of ageism. *Journal of Experimental Social Psychology, 26,* 199–216.

Posner, M.I., & Snyder, C.R. (1975). Attention and cognitive control. In R.L. Solso (Ed.), *Information processing in cognition: The Loyola Symposium* (pp. 55–85). Hillsdale, NJ: Erlbaum.

Sapir, E. (1963). *Selected writings of Edward Sapir in language, culture, and personality* (D. Mandelbaum, Ed.). Berkeley: University of California Press.

Schneider, W. (1990). *MEL users guide: Computer techniques for real time psychological experimentation.* Pittsburgh: Psychology Software Tools.

Zarate, M.A., & Smith, E.R. (1990). Person categorization and stereotyping. *Social Cognition, 8,* 161–185.

(RECEIVED 12/8/94; ACCEPTED 3/26/95)