Implicit Gender Stereotyping in Judgments of Fame

Mahzarin R. Banaji
Yale University

Anthony G. Greenwald
University of Washington

Implicit (unconscious) gender stereotyping in fame judgments was tested with an adaptation of a procedure developed by L. L. Jacoby, C. M. Kelley, J. Brown, and J. Jasechko (1989). In Experiments
1-4, participants pronounced 72 names of famous and nonfamous men and women, and 24 or 48
hr later made fame judgments in response to the 72 familiar and 72 unfamiliar famous and nonfa-mous names. These first experiments, in which signal detection analysis was used to assess implicit stereotypes, demonstrate that the gender bias (greater assignment of fame to male than female names) was located in the use of a lower criterion (P) for judging fame of familiar male than female names. Experiments 3 and 4 also showed that explicit expressions of sexism or stereotypes were uncorrelated-with the observed implicit gender bias in fame judgments.

Some of the earliest achievements of information processing methods in social cognition research demonstrated the involve-ment of memory and other cognitive processes in stereotyping. This research has documented, for example, errors in fre-quency estimation of behaviors performed by members of vary-ing social groups (Hamilton & Gifford, 1976), superior memory for stereotype-confirming than -disconfirming events (Rothb-art, Evans, & Fulero, 1979), and superior memory for informa-tion about members of minority than majority groups (see Tay-lor, 198 1). Empirical effects such as these revealed the cognitive basis of stereotyping, and these discoveries changed existing views of the stereotyping process and of the social perceiver. Rather than portraying stereotyping as an intentionally moti-vated act to discriminate, or as an attribute of a select group of prejudiced individuals, social cognition

182 MAHZARIN R. BANAJI AND ANTHONY G. GREENWALD

Lindsay, & Toth, 1992; Jacoby & Witherspoon, 1982; Roediger,
Weldon, & Challis, 1989; Schacter, 1987). As an illustration of
implicit memory, consider a result that is readily obtained with
the task of identifying words that are presented as brief tachis-
toscopic flashes. Successful identifications of words occur more
often for words to which participants were casually exposed ear-lier in the experiment than for new words that were not pre-viously presented. This effect of prior exposure occurs despite participants' poor performance on tasks of recalling or recog-nizing words from the earlier list. The tachistoscopic identifi-cation task provides an indirect measure of the effect of the prior experience. That is, even though the participant is not in-structed to retrieve the earlier presented material and is pre-sumably not trying to do so (and may well be incapable of such retrieval), the participant's responses indicate a residual effect. (For further reviews of indirect measurement in memory re-search, see Richardson-Klavehn & Bjork, 1988; Roediger, 1990; Roediger & Blaxton, 1987; Tulving & Schacter, 1990.) Extending beyond the memory literature, Greenwald and Ba-naji (1995) identified implicit social cognition as occurring when past experiences affect social judgment or behavior, but the nature of this influence is introspectively unidentified (or inaccurately identified) by the actor. They offered examples of several implicit social cognition phenomena, focusing on the research on attitudes, self-esteem, and stereotyping. In each of these domains, Greenwald & Banaji (1995) (a) showed that so-cial behavior operates in an implicit or unconscious fashion and (b) called for greater use of indirect measures of social cogni-tion. Applying the definition of implicit social cognition to the domain of stereotyping, they defined implicit stereotypes as in-trospectively unidentified (or inaccurately identified) traces of past experience that mediate attributions of qzdalities to mem-bers of a social category. Such a definition has new implications for both theory and research. For theory, it provides explicit recognition of the unconscious manner in which stereotypic ex-pressions may be produced and influence judgment and action, and the possible dissociation between conscious and uncon-scious components of stereotyping. For method, the implica-tions are substantial. Such a definition highlights the previously ignored role of unconscious processes in dominant measures such as adjective checklists and rating scales, and encourages the use of new, indirect measures in investigations of stereotyping.
Obvious attributes that distinguish members of social cate-gories are race and ethnicity, gender, age, class, and correlated attributes of skin color, physical features, dress, accent, physical abilities, and personality attributes. A goal of research on im-plicit stereotyping is to identify the unconscious use of such at-tributes in social judgment and the consequences of such judg-ments for target and perceiver.
Gender is known to be associated with strong and widely shared stereotypes (Basow, 1986; Broverman, Vogel, Brover-man, Clarkson, & Rosenkrantz, 1972; Eagly & Mladinic, 1989; Ruble & Ruble, 1982; Williams & Best, 1990). A well-known gender difference involves achievement, with men more than women being associated with achievement in many fields, such as politics, sports, and science. One source of data revealing the gender-specific association of achievement is from archives in-

dicative of fame: Men constitute 93% of the entries in Who's Who in America (1988), which purports to identify outstanding achievers in such arenas as business, government, arts, science, law, and education. Additionally, as a barometer of popular cul-ture, a board game called Claim to Fame, which tests the ability to recognize famous people, shows the imbalance of having a 3.85: 1 ma1e:female ratio among its 912 famous names.'
The discrepancy in male and female fame in archival and cul-tural indicators is matched by the cognitive availability of male names compared with female names. We asked Yale undergrad-uates to generate the first instance that came to mind in re-sponse to various categories (e.g., a book, a piece of furniture, an article of clothing) including "a famous person" and found that, of the 87 participants, 73 (84%) generated a male name, and only 14 (16%) generated a female name (a 5.21:l ratio). Thus, discrepancies in male-female fame, which are well estab-lished in archival data, are similarly indicated by a straightfor-ward measure of cognitive availability.

A Procedure to Test Implicit Gender Stereotyping

In the present experiments we used an implicit memory pro-cedure developed by Jacoby, Kelley, Brown, and Jasechko (1989) to observe implicit effects of prior exposure to names on subsequent fame judgments. In one experiment, Jacoby et al. (1989) exposed participants to names of famous and nonfa-mous people by asking them to pronounce each name. Then, either immediately or after a 24-hr delay, participants were pre-sented with previously seen (old) and previously unseen (new) famous and nonfamous names and were asked to make the fol-lowing judgment for each name: is this person famous? Jacoby et al. (1989) hypothesized that, in the delayed judgment condi-tion, episodic memory for the familiar names would dissipate, but perceptual fluency for the name would persist. This dissoci-ation of (explicit) episodic memory from (implicit) perceptual fluency for the names was predicted to lead to a higher proba-bility of incorrect judgments of familiar nonfamous names as famous, compared with otherwise equivalent unfamiliar nonfa-mous names. In other words, the familiarity evoked by a name due to seeing it recently was confused with the familiarity asso-ciated with a famous name. Consistent with this expectation, Jacoby et al. (1989) obtained a reliably higher rate of false alarms (i.e., false fame) for familiar nonfamous than unfamiliar nonfamous names in their delayed-judgment condition.
In the present research we adapted the false-fame procedure to examine a hypothesized gender bias in fame judgments. Names communicate social information, such as race and

GENDER BIAS IN FAME JUDGMENTS 183

obvious that an effect of gender was being investigated. Partici-pants were exposed to

184 MAHZARIN R. BANAJI AND ANTHONY G. GREENWALD

the availability of knowledge about the association between gen-
der and fame in the judging of each name. In the present varia-
tion of the false-fame experiment, male versus female gender of
names is thus included as an additional independent variable,
meaning that d' and B can be determined separately for male
and female name stimuli. Use of signal detection methods, therefore, permits description of the stereotypic association of gender with achievement in terms of differences in sensitivity, bias, or both, between male and female names. The simplest finding would take the form of greater readiness to assign the judgment of fame to male than to female names. This finding would indicate differential criteria for judging fame for male and female names, with the prediction of a stricter criterion applied to female than male names.
The signal detection analysis also permits the conclusion that the perceptual difference consists exclusively of a sensitivity difference. Such a finding would indicate a difference in dis-criminability of famous from nonfamous names for male versus female names. If such a sensitivity difference is obtained, the direction of difference might reveal lower sensitivity for male names-that is, greater fame accorded to male names may stem from reduced ability to separate famous male from nonfamous male names. Of course, these experiments may also show that both sensitivity and bias effects underlie gender stereotyping. The power of signal detection analyses comes from the oppor-tunity to separate bias from sensitivity effects and to test the contribution of both to producing the gender bias in false fame. In addition, these analyses will inform about the nature of the false-fame bias itself. Jacoby et al. (1989) provided only false alarm data, and hence it is not known whether the false-fame bias is rooted in sensitivity or bias.
Added to these possibilities is the use of both male and female participants in the present research, which allows a measure of differences in implicit stereotypes held by men and women. The research on outgroup homogeneity should predict that men should show lower sensitivity for female names, and women should show lower sensitivity for male names. However, if these outgroup homogeneity effects are restricted to explicit stereo-typing, such 68 Tc 1.-0 Td.6 iFm e loof f Flse

s(0429)snumbr if

GENDER BIAS IN FAME JUDGMENTS 185

a between-subjects variable. The independent variable of name fame (famous or nonfamous) is not specified as a design variable, rather being incorporated into the signal detection dependent variables, d' and 0.

Procedure

In Session I, each participant received a list of 72 names preceded by 2 buffer names (1 famous, 1 nonfamous). Participants were informed that the purpose of the study was to obtain ratings of the ease of pro-
nounceability of common names. Twenty-six names appeared on the
first page, and 24 names appeared on each oftwo additional pages. Each
name was accompanied by a 5-point scale (I = very easy to pronounce
through 5 = very hard to pronounce). Participants were asked to read
each name silently and to rate the ease of pronouncing the name "for a
person who speaks little English." The only0016 Tm ((I)Tj /T1_5 the nounce48(judgmeTd ((co9(t01 103(study)]TJ -07Tj 6 Tc 1.0178 0 Td (to)T465 Tc 1.74384 0 Td (to)Tj -

MAHZARIN R. BANAJI AND ANTHONY G. GREENWALD

PERCENT

70.4 1 Experiment 1 (N = 28) hits false

alarms
54.6

100 Experiment 2 (N = 49)
73.2

50

I Experiment 3 (N = 38) 1

Familiar Familiar Unfamiliar Unfamiliar male female male female

WITHIN-SUBJECT CONDITION

Figure I. Mean hit rates and false alarm rates as a function of within-subject conditions of name famil-iarization and gender. The unfilled extensions of bars show the upper half of the 95% confidence interval for each mean.

decompositions considering patterns of statistical significance in two-treatment contrasts among the four means is not given here. Rather, this more detailed analysis is deferred to the Gen-eral Discussion section, where we present a combined analysis of the four experiments. This strategy, which we also used for Experiments 2, 3, and 4, serves to avoid populating each exper-iment's Results section with multiple significance tests that may be rendered obsolete by the more powerful summary analysis.

Discussion

The results of Experiment 1 indicated that gender stereotypes operated with greatest force for nonfamous names that were

given a boost in familiarity by their presentation in the first ses-sion. When encountered in the second session, the familiarity of these nonfamous names was more likely to be interpreted as arising from fame when the name was male than when it was female. In Experiment 1, using a paper-and-pencil version of Jacoby etal.3 (1989) procedure, we replicated their finding that familiar nonfamous names are indeed more likely to be incor-rectly judged as famous than comparable, new nonfamous names. The signal detection analysis further indicated that par-ticipants did not differ in their sensitivity (d? to fame for male and female names but did discriminate between male and fe-

male names by having a lower threshold or criterion (j3) for as-signing the judgment of fame to male names.

GENDER BIAS IN FAME JUDGMENTS

dprime
3.00 4 1 Erpetlment 1 (N = 28) 1

(Experiment 3 (N = 38) /
3.00

2.00

1 .oo

0.00

2.00

1 .oo

0.00
Famlllar Familiar
male female

Unfamiliar
male

Unfamlllar
female

WITHIN-SUBJECT CONDITION

Figure 2. Mean d' scores as a function of within-subject conditions of name familiarization and gender, showing upper halves of the 95% confidence interval for each mean.

This finding of a gender bias in the threshold for fame judg-ments qualifies as an implicit stereotyping effect. Although the experiment provided no direct evidence that the observed gen-der bias in fame judgments occurred without part 10u9nts'

MAHZARIN R. BANAJI AND ANTHONY G. GREENWALD

log beta

I Experiment 2 (N = 49) /
n n

I Experiment 3 (N = 38) 1

Fami iar Familiar Unfamiliar Unfamiliar male female male female
WITHIN-SUBJECT CONDITION

Figure 3. Mean log B scores as a function of within-subject conditions of name familiarization and gender, showing upper halves of the 95% confidence interval for each mean.

The main result apparent in this analysis was that more false alarms occurred on names derived from male names (64%) than on ones derived from female names (36%). This observa-tion led to a revised procedure for obtaining nonfamous names in Experiment 2, the main purpose of which was to provide a replication that could bolster the finding of implicit gender stereotyping from Experiment 1. The famous names used in Experiment 2 came from three of the six categories used in Ex-periment 1: writers, actors, and musicians. The other three cat-egories (scholars, athletes, and politicians) were dropped be-cause of relative difficulty in identifying women who matched men in actual fame. The only other substantial change was the use of a 48-hr (rather than 24-hr) delay between sessions, neces-

sitated by the meeting schedule of the class that provided par-ticipants. Because of previous findings that perceptual fluency effects survive extended periods between exposure and test (Kolers, 1976), this increased delay of Session 2 was not ex-pected to alter findings.

Method

Participants and Design

Participants were tested in their classroom during both sessions. Fifty-seven students in an introductory psychology class participated in Session 1. Of these participants, data were used for only 49 (23 men and 26 women) 6fthe 50 who were also present during Session 2,48 hr later.

GENDER BIAS IN FAME JUDGMENTS 189

Data for 1 male subject were discarded because of an unusually high rate of false alarms on nonfamous names (23% approximately equal for male and female names, compared with a mean of 5.5% for other participants). The design was a 2 (name gender: male or female) X 2 (name exposure: familiar or unfamiliar) X 2 (participant sex: male or female) factorial, with only participant sex as a between-subjects variable.

Materials and Procedure

From the famous names used in Experiment 1, 12 famous names in each of these three categories (writers, musicians, and actors) were used. Twelve additional names were added to each category to create a total of 24 famous names (1 2 female and 12 male) in each. The resulting set of 72 famous names was divided into two lists of 36, each containing 12 names (6 male and 6 female) in each of the three fame categories. Nonfamous last names were selected from the New Haven, Connect-icut phone book, each matching one of the 72 famous names in length and ethnicitv. To each nonfamous last each

190 MAHZARIN R. BANAJI AND ANTHONY G. GREENWALD

hr apart. For both sessions, participants were tested in their classroom.
The design was a 2 (name gender: male or female) X 2 (name exposure:
familiar or unfamiliar) X 2 (participant sex: male or female) factorial,
with all but participant sex as within-subjects variables.

Materials and Procedure

The materials (name lists) and procedure were identical to those of Experiment 2, with the addition of the explicit measure of sexism fol-lowing Session 2's fame judgment task. In measuring explicitly held be-liefs about gender equality, two scales have seen frequent use in recent research: the Feminism Scale (EM; Smith, Ferree, &Miller, 1975) and the Attitudes Toward Women Scale (ATW; Spence & Helmreich, 1972). Unfortunately, both the FEM and ATW scales contain items that ap-pear to be outdated for use with samples of contemporary college stu-dents (for example, a FEM item reads, "The 'clinging vine' wife is jus-tified provided she clings sweetly enough to please her husband"). Pilot data indicated that some participants simply refused to answer several items, and most students were clustered at the feminist or belief-in-equality ends of the scales. Instead of using these existing measures, in Experiment 3 we used a set of face-valid items that we had selected on the basis of both their meaningfulness to the present participant popu-lation and the variability of responses they elicited.
Eleven participants generated several statements representing a range of opinions about gender equality. From the resulting pool of more than 30 statements. we selected 8 that we considered to reflect gender issues of interest to students. Of these, a "strongly agree" response to 4 ofthem (Numbers 1,3,5, and 7) indicated beliefs in gender equality (feminism), and a "strongly agree" response to the remaining 4 (Numbers 2, 4, 6, and 8) indicated endorsement of gender inequality (se~ism).~ Partici-pants made their judgment for each item on a 5-point Likert scale with options labeled Strongly Agree, Agree, Neutral, Disagree, and Strongly Disagree. Items were scored in the direction of sexism, such that the greater the sexist endorsement, the higher the score. The overall sexism score was computed as the mean score across these 8 items. An addi-tional, separate dichotomous score was obtained from the ninth item, which asked whether or not the participant accepted a self-description offeminist.

Results

As in Experiments 1 and 2, we classified data from judgments of fame into four within-subjects categories (familiar male, fa-miliar female, unfamiliar male, unfamiliar female) and reduced them to hit and false alarm rates. As Figure 1 shows, hit rates were higher for male names than for female names, F(1, 37) = 16.17, p < .0005; and for familiar names than for unfamiliar names, F(1, 37) = 6.80, p = .013; with the interaction effect being nonsignificant, F(1, 37) = 0.02. The false alarm rate was higher for familiar names than for unfamiliar names, F(1, 37) = 39.32, p < .0005; and for male names than for female names, F(1, 37) = 5.15, p = .029; with the interaction effect of name familiarity and gender being nonsignificant, F(l,37) = 1.04. Signal detection analyses again indicated an absence of effects on either d' or log fi for sex of participant, leading to the elimi-nation of this design variable from the analyses to be reported. Mean values of d' and log @ are presented in Figures 2 and 3. The results ford' indicated that participants were more sensitive to the fame variation for unfamiliar than for familiar names, F(1, 37) = 14.78, p < .0005, with no other effects significant. The results for log p showed, as before, that participants more readily assigned fame to familiar names compared with unfa-

miliar names, F(l,37) = 33.89, p < .0005. Ofgreatest relevance to the assessment of implicit stereotyping, analysis of log ,8 also showed, as before, that participants attributed fame more readily to male than to female names, F(1, 37) = 5.93, p = .02. The interaction effect on log /3 was nonsignificant, F(1, 37) = 1.29. However, the pattern of data replicated previous experi-ments, with the comparison of familiarized male and female names replicating the difference favoring male names, t(37) = 2.62, p = .O 1, and the unfamiliarized names showing no gender bias, t(37) = .84, p

192 MAHZARIN R. BANAJI AND ANTHONY G. GREENWALD

can easily make judgments about which group has more or less of
some quality.

We are interested in your judgment of how typical some traits are
of men and wc 1J 0.059 3.B3-umTd (53955 0 Td (We)Tj 0.0298 Tc 1.60978 0 Td (are)Tj 0.01489 Tc 1.678 0 Tdnotnd)Tj 0 Tc 7480739 0 Td [(interested 2

GENDER BIAS IN FAME JUDGMENTS 193

tion thing was a red herring and this study is somehow gender based." No participants correctly produced hypotheses con-

MAHZARIN R. BANAJI AND ANTHONY G. GREENWALD

PERCENT RESPONSES
Hits Hits Hits Hits

196 MAHZARIN R. BANAJI AND ANTHONY G.

GENDER BIAS IN FAME JUDGMENTS 197

miliarity) about the cause of familiarity under which judgments about people are required and provided. These results suggest that under such conditions, perceivers use social category infor-mation (elicited by even seemingly superficial markers such as a name) by using criteria of differing strictness in judging indi-vidual members. The contribution of the present findings is in locating such a difference in the criterion used in judgment as the basis of stereotyping.

Conclusions and Implications

The main empirical discovery of the present experiments is that a lower criterion is used in fame judgments of familiar men than of women. This benefit to men is assumed to be uncon-scious and produced by an interaction of two factors: (a) an inability to distinguish two sources of familiarity (actual fame vs. fluency from previous exposure) and (b) the use of social category knowledge in individual judgment. The main method-ological contributions of this article are (a) the introduction of signal detection analyses to isolate the criterion component of implicit stereotyping and (b) introduction of the name-judg-ment procedure as one with potential broad applicability for examining implicit cognition in stereotyping. The discovery of such implicit stereotyping effects bolsters the recommendation for a departure in thinking about individual responsibility in the everyday perpetration of stereotyping. In particular, such discoveries seriously question the wisdom of assuming individ-ual responsibility and extracting individual punishment for un-conscious acts that may occur outside the social perceiver's awareness. Such findings render moot the possibility of sponta-neous, individually devised compensatory strategies. Rather, efforts must focus on alternative strategies to (a) generate pro-cedures by which individuals can routinely recognize the effects of implicit stereotyping and (b) modify environments such that the operation of implicit stereotypes is inhibited.
The pernicious consequences of persistent implicit stereotyp-ing for the social target, however, raise questions about the avail-ability of adequate measures to safeguard against the occur-rence of such actions. In particular, they raise the issue of target protection against judgments of which the target is unaware and which therefore render moot standard legal procedures such as due process. Such thinking has led us to discuss in another arti-cle the relative merits of blinding, consciousness raising, and affirmative action procedures (Greenwald & Banaji, 1995). Al-though adequate solutions to such issues lie in the future, dis-covering the cognitive conditions that produce implicit stereo-typing can inform about the operations of social knowledge and contribute to proposals for social engineering.

References

Allport, G. W. (1954). The nafure ofprejudice. Cambridge, MA: Addi-son- Wesley.
Aronson, E. (1988). The social animal. New York Freeman.
Banaji, M. R., & Greenwald, A. G. (1991, June). Measuring uncon-
scious attitudes. Paper presented in a symposium on the psychologi-
cal unconscious at the third annual meetings of the American Psy-
chological Society, Washington, DC.
Banaji, M. R., & Greenwald, A. G. (1994). Implicit stereotyping and
unconscious prejudice. In M. P. Zanna & J. M. Olson b1 Tf ss.),j /T1_2 1 Tf 8.10154 Tc (a.10 Tc /.625 94.8394 82.3198 Tm (The)Tj45 262 725 Tc460