Identical applicant but different outcomes:

The impact of gender versus race salience in hiring

Aneeta Rattan
London Business School

Jennifer Steele
York University

Nalini Ambady
Stanford University

Please address correspondence to: Aneeta Rattan, Organisational Behaviour, London Business School, Sussex Place, London NW1 4SA, UK, ph: +44 (020) 7000 8907, email: arattan@london.edu, and Jennifer Steele, Psychology Department, York University, 4700 Keele St., Toronto, Ontario, M3J 1P3, Canada, ph: +1 (416) 736 2100, email: steeleje@yorku.ca.

Identical applicant but different outcomes:

The impact of gender versus race salience in hiring

Abstract

People belong to multiple social groups, which may have conflicting stereotypic associations. A manager evaluating an Asian woman for a computer programming job could be influenced by negative gender stereotypes or by positive racial stereotypes. We hypothesized that evaluations of job candidates can depend upon what social group is more salient, even when both are apparent. In three studies, using student (Study 1) and non-student (Studies 2 and 3) samples, we compared ratings of an Asian-American female applicant, after subtly making her race or gender salient in stereotypically male employment contexts. Consistent with our predictions, we found evidence that men rated her as more skilled (Studies 1 and 3), more hirable (Studies 1-3), and offered her more pay (Study 2) in science and technology-related positions when her race, rather than gender, was salient. The theoretical implications for person perception and practical implications in employment contexts are discussed.

Identical applicant but different outcomes:

The impact of gender versus race salience in hiring

Imagine a hiring manager at a technology company who is evaluating applications for a computer programming position. One applicant is an Asian-American woman. How might this applicant's social group memberships affect the hiring manager's evaluation of her and his hiring decision? Controlled experiments have consistently documented gender bias against women in stereotypically male employment contexts, including management as well as science, technology, engineering, and math (STEM) fields, on outcomes ranging from ratings of women's suitability for a job to salary offers (Davidson & Burke, 2000; Moss-Racusin, Dovidio, Brescoll, Graham, & Handelsman, 2012; Olian, Schwab, & Haberfeld, 1988; Terborg & Ilgen, 1975; Uhlmann & Cohen, 2007). Such experimental data is consistent with real world evidence of the gender wage gap and gender representation gaps (World Economic Forum, 2015). For example, estimates of median weekly earnings show that among full-time wage and salary workers in the United States, women make only eighty-one percent of White men's earnings (Table 2, Quarter 3, Bureau of Labor Statistics, 2014). This research and real-world data suggests that our hypothetical hiring manager's decision about the applicant might be biased by negative stereotypes associated with her gender identity in the context of a computer programming position.

However, people belong to multiple social groups (Macrae, Bodenhausen, & Milne, 1995; Markus & Kitayama, 1991; Sinclair & Kunda, 1999), and both theory and research suggests that interpersonal perception can depend on which aspects of someone's identities are salient. Social categorization happens within milliseconds during person perception, often unconsciously and virtually effortlessly (Ito & Urland, 2003), activating networks of associations through a dynamic interactive process that integrates bottom up perceptual cues with top-down social cognitive processes to form coherent impressions (for a review see

Freeman & Ambady, 2011). In the case of the hiring manager and the Asian woman's application for a computer programming position, then, is the manager *only* influenced by her gender? What about her race, and its associated positive stereotypes in STEM (Cheryan & Bodenahusen, 2000; Shih, Pittinsky, & Ambady, 1999; Wong, Nagasawa, & Lin 1998)? In the present research we address this question by examining whether an Asian-American woman, whose gender and racial identities are associated with competing stereotypes in STEM employment contexts, receives different evaluations and outcomes depending on the relative salience of these identities.

Employment Discrimination

The majority of past research on employment discrimination cannot inform this question because it has focused exclusively on a single aspect of identity. For example, research on gender discrimination in stereotypically male employment contexts has typically made use of White women as the gender exemplar (e.g., by using a stereotypically White name such as Joan McKay in study materials, Goldberg, 1968; Swim, Borgida, Maruyama, & Myers 1989). Research on race and employment evaluations has often focused exclusively on stigmatized status, showing, for example, a bias toward positively stereotyped White applicants compared with Black applicants (e.g., Bertrand & Mullainathan, 2004; McConahay, 1983) or compared with applicants from other stigmatized groups (e.g., Arabs, Derous, Ryan, Nguyen, 2012). However, theories of intersectionality argue that gender must always be conceptualized in the context of the other group memberships that women possess (Kulik, Roberson, & Perry, 2007; Reid & Comas-Dias, 1990; Purdie-Vaughns & Eibach, 2008; Sesko & Biernat, 2010; Shields, 2008). In fact, women's percentage of White men's median weekly earnings fluctuates from sixty-two percent among Hispanic or Latino women to eighty-nine percent among Asian women for full-time wage and salary workers in the United States (Table 2, Quarter 3, Bureau of Labor Statistics, 2014). In order to bridge this

gap in the literature, we asked whether evaluators would exhibit systematically different degrees of biased decision making toward the same person – an Asian American woman applying for a STEM job – depending upon the salience of her relevant social categories.

One possibility is that the hypothetical hiring manager's evaluations would be negatively biased by the candidate's gender group membership, regardless of her race. This prediction draws on classic theories of prejudice, which suggested that the stigma associated with negatively stereotyped social categories is wholly devaluing for perceivers' impressions (Goffman, 1963). That is, in line with the typical discourse on gender discrimination in employment settings, simply belonging to the category may be considered a sufficient condition to elicit biased judgments from evaluators, regardless of other relevant social category information.

Another possibility is that the candidate's race would also have an impact on perceivers' impressions. Recent research suggests that there could be a co-activation of categories that the drives social perceptions of individuals who represent double outgroups (Kang & Chasteen, 2009; Migdal, Hewstone, & Mullen, 1998; Remedios, Chasteen, Rule, & Plaks, 2011; Urban & Miller, 1998). Some research in this tradition has documented a link between the Asian racial category and femininity (Galinsky, Hall, & Cuddy, 2013; Hall, Galinsky, & Phillips, 2015; Johnson, Freeman, & Pauker, 2011; McMahon & Kahn, 2016; Shug, Alt, & Klauer, 2015). Taken together, the same type of directional hypothesis arises – a hypothetical hiring manager's evaluation of an Asian woman for a technical position would still be negatively biased, but in this case due to *both* her gender and racial group memberships.

Fit models of workplace discrimination instead suggest that when expectations of the employment context are at odds with job candidates' group memberships, evaluators may devalue job candidates' skills and abilities, and ultimately come to a poorer overall

evaluation (Eagly & Karau, 2002; Heilman, 1984; Heilman & Eagly, 2008). When considering the example of the Asian woman who applied to the computer programming position, if the hiring manager views her through the lens of her negatively stereotyped gender, there should be a perceived misfit with the demands of the job. By contrast, if he considers her through the lens of her positively stereotyped racial group membership, there should be relatively greater perceived fit. There is some research to suggest that ethnic stereotypes can at times outweigh gender stereotypes (Levin, Sinclair, Veniegas, & Taylor, 2002), and that the salience of a single individual's social categories fluctuates based on the context (Turner, Oakes, Haslam, & McGarty, 1994). This raises the possibility that, despite the fact that her gender and race are both apparent to evaluators at all times, the *salience* of her racial (vs. gender) identity might raise (vs. lower) evaluations of her in a stereotypically male employment context.

Multiple Identities and Person Perception

Theory and research from a multiple identities perspective (Ambady, Shih, Kim, & Pittinsky, 2001; Sinclair & Kunda, 1999; Shih et al., 1999; Pittinsky, Shih, & Ambady, 2000) provide support for this possibility. There is evidence that both situational and motivational factors can impact which identities play a primary role in the impression formation process in the face of multiple social identities (Crisp & Hewstone, 2007; Macrae, et al., 1995; Markus & Kitayama, 1991; Sinclair & Kunda, 1999). Macrae et al. (1995) argue that although people can be viewed along a variety of dimensions (sex, race, age, among countless others), we are unable to simultaneously consider every, or even several, possible identities when perceiving others. In support of this claim, in one seminal paper Macrae et al. (1995) demonstrated that people who viewed an Asian woman eating with chopsticks were more likely to spontaneously activate Asian stereotypes, and were less likely to activate female stereotypes, relative to those who viewed an Asian woman applying make-up. Once the Asian woman

was categorized by race or gender, relevant stereotypes were more accessible, whereas competing stereotypes associated with her other identity were inhibited (see also Gilbert & Hixon, 1991).

Subsequent research has confirmed that this is not restricted to perceivers. Increasing the salience of their negatively stereotyped gender identity led Asian-American girls and women to perform worse on math problems, while increasing the salience of their positively stereotyped racial identity led them to perform better (Ambady, et al., 2001; Shih, et al., 1999). Research also suggests that these types of findings extend to spontaneously activated attitudes. Studies suggest that focusing on one of targets' multiple categories can directly impact the implicit attitudes that are activated for perceivers (Barden, Maddux, Petty, & Brewer, 2004; Mitchell, Nosek, & Banaji, 2003; Steele, George, Cease, Fabri, & Schlosser, 2017; Wittenbrink, Judd, & Park, 2001). For example, across several studies, participants had more favourable implicit attitudes toward liked Black athletes relative to disliked White politicians when they were categorized by their occupation as opposed to by their race (Mitchell et al., 2003).

In addition, one study by Pittinsky et al. (2000) tested the differential effects of multiple social group memberships on majority group members' memory for stereotype-relevant information. In this study, participants were presented with the college application of a fictional Asian-American woman, and were later asked to recall the applicant's math SAT scores while being reminded of her Asian identity, her female identity, or neither identity. Consistent with societal stereotypes, increasing the salience of gender led to lowered estimates of an Asian-American woman's math test scores, while increasing the salience of her race led to higher estimates of her math test scores (Pittinsky, et al., 2000), suggesting that memory for people's qualifications can be shaped by the relative salience of their identities. Taken together, these studies suggest that the relative salience of competing

identities can impact stereotype activation, implicit attitudes, memory, and even a person's test performance. However, it is unclear whether such effects extend to hiring decisions and evaluations. In the present research, we extend previous research on multiple identities and person perception by empirically testing whether the differential salience of competing group memberships goes beyond memory for information to also influence evaluators' assessments of women in employment contexts.

Overview of Studies

To summarize, in the present research we investigate whether perceivers will exhibit differentially biased evaluations toward an individual who belongs to multiple relevant social categories, depending upon which category is salient. While previous experimental investigations of biased decision making in employment contexts have generally compared women's outcomes to men's, or the outcomes of racial minorities to the majority group, to address the present research question we instead compare the consequences of highlighting one individual's different relevant social identities. Specifically, we test whether the evaluations of one person, an Asian female, depend on whether her race or gender is made salient. We situate our investigation in the context of the devaluation that women face in STEM employment settings, which has received continued attention in social psychology (Heilman & Parks-Stamm, 2007; Moss-Racusin, et al., 2012) as well as in society (DeNavas-Walt, Proctor, & Smith, 2013; Obama, 2014). Given that men are more frequently the evaluators and decision makers in STEM contexts (Reuben, Sapienza, & Zingales, 2014), we focus our investigation on male perceivers.

Thus, we investigate whether different aspects of *the same person's* identity can affect how much employment and wage discrimination male evaluators exhibit. In three experiments we test whether male decision makers evaluate and pay an Asian-American female applicant in a stereotype-relevant domain differently depending upon the salience of

her gender as opposed to her racial identity. In Study 1, we asked male undergraduates to interview an Asian female applicant for a computer technician position. In Study 2, we asked an online non-student sample of men to evaluate an Asian female applicant for a tutoring job in either a stereotypically male field (computer programming) or a stereotypically female field (English literature). To test the possibility that these effects would emerge only for men, we also collected a non-student sample of women in this study. In Study 3, we recruited a larger online non-student sample of men and asked them to evaluate an Asian woman for a tutoring job in computer programming. Across all studies, we predicted that the applicant's evaluation would be more positive when her racial, rather than gender, identity was salient in a technical field.

Study 1

In Study 1 we asked male undergraduates to interview an Asian-American female candidate for a computer technician position. We decided a priori to include only non-Asian men to ensure participants would not share their racial identity in common with the candidate. We hypothesized that participants' ratings would be higher in the race prime condition compared to the gender prime condition. Given that this was an initial investigation, we also included a control condition that might offer insight into whether any effects were systematically due to the relative salience of one of these two identities. Thus, Study 1 had a three-condition between-participants experimental design.

Method

Participants

We aimed to recruit at least 20 participants per cell but set a strict stopping rule of the end of the school year, and therefore our sample size was determined by participant availability. A total of 53 men (4 African American/Black, 37 European American/White, 3 Latino Americans, 7 biracial/mixed/other, 2 unreported) were recruited from a private

university in the northeastern United States. Participant age was not recorded, but all participants were undergraduate students. Prior to analysis, we excluded 6 participants who expressed suspicion that the research examined stereotyping and/or that the other participant was a confederate and 4 participants for whom there was a procedural error, leaving 43 participants in the analyses. For this and all subsequent studies, all conditions and measures included in the study are described.

Procedure

Participants signed up for an interview study to be conducted with another undergraduate. The other undergraduate was always one of two Asian-American female confederates. Once the participant and confederate arrived at the lab, a White male experimenter seemingly randomly assigned them to the roles of interviewer and interviewee, respectively, and led them to different rooms. After providing consent, the participant read a job description for a "user assistant" computer technician position while the confederate purportedly filled out a job application. An application, which had been previously completed by the experimenter, was then given to the participant to review and this served as our manipulation. The confederate remained blind to experimental condition and did not see the application.

The application included demographic information and provided hand-written ratings in response to four questions assessing technical skills (e.g., experience with software programs, software and hardware installation) and four questions assessing interpersonal expertise (e.g., teaching skills; resolving co-worker conflicts). The ratings communicated ambiguously positive qualifications (average rating of 4, range 3-5) on a 7-point Likert-scale ranging from "no experience" (1) to "very experienced" (7).

All of the information was the same across conditions, except: In the Female Prime condition, the application listed the confederate's name as 'Gloria,' email as

gloria33@hotmail.com, sex was provided in a large, centered box, and race was to the side in smaller print. In the Asian Prime condition, the application listed the confederate's name as 'Chia-Jung Gloria Tsay,' email as tsay33@hotmail.com, Chinese and English as spoken languages, race was provided in a large, centered box, and sex was to the side in smaller print. In the Control Prime condition, the application listed her name as 'Gloria Tsay', email as gt33@hotmail.com, and both sex and race were provided in smaller print to the side. After being randomly assigned to review one of the primes, participants completed a filler questionnaire rating the quality of potential job interview questions.

Participants next completed the in-person interview. They were instructed to ask 3 standard interview questions (e.g., "Why are you interested in this job?") to which the confederate provided scripted responses (e.g., "Well, like any other college student, I'd love to make some extra money. Um, this job is a great opportunity for me to make some money, work with computers, and meet new people. Um, I'd also like to get a job on campus so I don't waste too much travel time. The flexible hours that the position offers also appeals to me, since my schedule varies from day to day"). The confederates learned the identical script and were trained to ensure that their delivery of the script was comparable; confirming this, there were no significant differences between confederates, and no confederate by condition interactions on any of the dependent measures (Fs < 1.5, ps $> .22$).

After the interview, participants completed the following dependent measures: they rated the applicant's technical skills (e.g., ability to solve computer problems; 3 items, $\alpha =$ .92), interpersonal skills (e.g., ability to interact with others; 3 items, $\alpha = .85$), and overall impression of the candidate (1 item) on scales from "not very skilled" (1) to "very skilled" (7). They also rated how willing they would be to hire the candidate on a scale from "not very willing" (1) to "very willing" (9), and indicated the pay rate they would recommend, if she were to be hired, on a scale from $7–$12 USD/hour in $1 increments.

In order to maintain the cover story, we also included a questionnaire asking participants to respond to free-response questions about the applicant's strengths, weaknesses, and other potential careers that participants would recommend for the applicant. Participants also completed a word-stem completion task before and after the interview, which was included as an exploratory measure. No significant differences were found in our coding of participants' responses to the free response questions or word stem completions, and therefore no results are reported for these measures. Finally, participants completed a demographics questionnaire and were debriefed.

Results and Discussion

We conducted one-way ANOVAs with a contrast testing the proposed linear trend to test our hypothesis that participants in the Asian prime condition would evaluate the applicant the more positively than participants in the Female prime condition (see Table 1 for means and standard deviations).

No significant differences by condition emerged for men's ratings of the confederate's technical skills, $t(40)=1.15$, $p=.26$, $n^2_p=.032$, or interpersonal skills, $t(40)=1.10$, $p=.28$, $n^2_p=.030$. This result fits with other research that has found specific evaluation criteria (as opposed to ambiguous or subjective evaluation criteria) do not produce biased judgments in employment contexts (e.g., structured interviews, Bragger, Kutcher, Morgan, & Frith, 2002; clear criteria weights, Norton, Vandello, & Darley, 2004).

However, the anticipated main effect of condition emerged for participants' evaluations of the confederate's overall skill, $t(40)=2.43$, $p=.02$, $n^2_p=.13$, with participants in the Asian prime condition ($M = 5.69$, $SD = .63$) providing the most positive ratings and those in the Female prime condition ($M = 4.67$, $SD = .92$) providing the least positive ratings. Contrasts comparing the Control prime ($M = 5.47$, $SD = 1.06$) condition to the Female prime

condition, $t(40)=1.83$, $p=.08$, $n^2_p=.08$, as well as the Asian prime condition, $t(40)=.66$, $p=.51$, $n^2_p=.01$, were not significant.

A similar pattern also emerged for participants' willingness to hire the confederate, $t(40)=2.50$, $p=.02$, $n^2_p=.14$. Participants reported the greatest willingness to hire following the Asian prime ($M = 7.31$, $SD = 1.18$) and least willingness following the Female prime ($M = 5.87$, $SD = 1.36$). Again, the contrasts comparing the Control prime ($M = 6.47$, $SD = 1.88$) to the Female prime condition, $t(40)=1.08$, $p=.29$, $n^2_p=.03$, as well as the Asian prime condition, $t(40)=1.46$, $p=.15$, $n^2_p=.05$, were not significant.

In addition, the predicted effect emerged for pay, $t(40)=2.28$, $p=.03$, $n^2_p=.12$. Participants primed with her Asian identity ($M = 9.42$, $SD = 1.63$) recommended higher pay than participants primed with her female identity ($M = 8.33$, $SD = 1.13$). Again only the hypothesized contrast was significant; the contrasts comparing the Control prime condition ($M = 8.53$, $SD = .99$) to the Female prime, $t(40)=.44$, $p=.67$, $n^2_p=.005$, and Asian prime, $t(40)=1.86$, $p=.07$, $n^2_p=.08$, conditions were not significant.

Although participants rated an Asian-American female applicant as equally skilled technically and interpersonally for the user assistant computer technician job, they rated her overall skill and hirability most favorably when they had been primed with her race as opposed to her gender. In addition, participants primed with the candidate's positively stereotyped racial identity recommended an average $2,267.20 per year more, for a 40-hour-per-week job, than participants reminded of her female identity. Because neither the Asian nor Female prime conditions were systematically different from the control condition, this initial evidence suggests neither identity alone drives the difference. Rather, as theorized by fit models of discrimination, the relatively lower fit of her gender identity and relatively better fit of her racial identity each contribute to divergent perceptions. However, the

conclusions that we are able to draw from this initial study are limited by the small sample size (observed power for ratings of overall skill = .62, hirability = .62, and pay = .58).

Study 2

In Study 2 we aimed to conceptually replicate these effects with a larger online sample. In addition, we designed this study to test two potential boundary conditions of our initial effects. First, we theorized that differential evaluations of an Asian-American female job candidate would arise because of field-specific stereotypes, not a general positivity towards Asian-Americans or a general negativity towards women. In Study 2, we tested this directly by exposing some participants to a traditionally male domain (computer programming) and others to a traditionally female domain (English literature). We expected the predicted effects to arise in the context of computer programming, but not English literature[1].

In addition, in Study 1 we found evidence that non-Asian men, who share neither gender nor race with the applicant, are influenced by the relative salience of her competing identities. In Study 2, we expected that non-Asian men would again show these effects, even in absence of a direct interaction with the candidate. However, we also explored whether non-Asian women, who share their gender identity, would show a similar pattern of results. Previous research has yielded inconsistent results for whether women are less likely to exhibit employment discrimination against in-group members (Goldberg, 1968; Davidson & Burke, 2000; Moss-Racusin, et al., 2012; Swim, et al., 1989). Therefore, as an exploratory step, we also recruited women into this study in order to separately assess whether they show the same, null, or perhaps different patterns of evaluation depending upon the salience of the target woman's race versus gender.

Method

Participants

We aimed to recruit at least 30 men and 30 women per condition, but given that we planned to exclude participants who failed the manipulation check, we oversampled slightly. Thus, an mTurk hit requested 275 participants and received 269 responses. A total of 7 participants were excluded from the study a priori as they reported being Asian American. An additional 44 participants were excluded prior to any analyses as they failed an instructional manipulation check testing their attention to a simple instruction (i.e., "This page is a test, to confirm that you are reading the instructions carefully. For the three questions which follow this paragraph, please give the following answer to each question: reader. Please just ignore the text of the questions, and type the word reader as your answer for all three items. Thank you for reading carefully."; Oppenheimer, Meyvis, & Davidenko, 2009), leaving a total of 225 participants (106 female, 90 male, 29 unreported; $Mean_{age}$=35.14, SD_{age}=12.31; 12 African American/Black, 174 European American/White, 6 Latino American, 7 biracial/mixed/other, 13 other/unreported).

Procedure

After providing informed consent, participants were randomly assigned to view a job description for a part-time tutoring position in Introduction to either Computer Programming (negative gender, positive race stereotype condition) or English Literature (positive gender stereotype condition). Participants were then randomly assigned to view a LinkedIn profile that served as the identity salience manipulation. The applicant was always the same (fictitious) Asian-American female Stanford University undergraduate. Her profile picture (showing an Asian woman in business attire), work experience (e.g., "peer counselor at Stanford University Career Development Center," "intern at the Hoover Institute Public Affairs Office," "volunteer dance instructor at the Bay Area Boys and Girls Club"), honors (e.g., "presented at the Stanford Undergraduate Conference"), and expertise (e.g., "skilled in Mac and PC, Microsoft Word, Excel, Powerpoint, Publisher") were the same across

conditions. The Female Prime condition, however, listed her name as "Gloria Tsay," mentioned a limited working proficiency in French, and listed study abroad experience as Stanford in Paris. The Asian Prime condition listed her name as "Chia-Jung Gloria Tsay," mentioned a limited working proficiency in Chinese, and listed study abroad experience as Stanford in Beijing.

After reviewing this profile alongside her purported profile picture, participants rated the candidate overall (not very skilled "1"–very skilled "7"), rated their willingness to hire her (not very willing "1"–very willing "7"), and indicated how much she should be paid if hired ($10–$15 USD/hour in $1 increments), as they had in Study 1. Finally, participants completed questions asking about the content and purpose of the study materials they had just read and a demographic questionnaire. Participants were then debriefed.

Results and Discussion

We first conducted the focal 2 (Domain: Computer Programming vs. English Literature) x 2 (Prime: Female vs. Asian) ANOVAs, testing the a priori hypothesis for each dependent variable for men's ratings. As expected, the predicted interaction emerged for each dependent variable, see Table 2. For men's ratings of the candidate's overall skill, there was a main effect of Domain, $F(1,82) = 11.51$, $p = .001$, $n^2_p=.13$, no main effect of prime, $F(1,82) = .46$, $p = .50$, $n^2_p=.006$, but as predicted, a significant Domain X Prime interaction, $F(1,82) = 4.05$, $p = .048$, $n^2_p = .05$. The key pairwise comparison is whether men rated the Asian woman as less skilled when they had been primed with her gender rather than racial identity in the context of computer programming, and indeed exposure to the Asian ($M = 5.26$, $SD = 1.39$), rather than Female ($M = 4.53$, $SD = 1.28$), prime led to marginally higher ratings of the applicant's skill, $F = 3.45$, $p = .07$, $n^2_p = .04$. The salience of her different identities had no effect on men's evaluations when she had applied for the English literature tutoring job, $F(1,79) = .94$, $p = .34$, $n^2_p=.01$ (Female Prime, $M = 6.0$, $SD = 1.14$; Asian Prime, $M = 5.64$,

$SD = 1.09$). Interestingly, when her racial group membership was salient, the applicant's skills were evaluated similarly highly in the context of computer programming and English literature, $F(1,79) = 1.05$, $p = .31$, $n^2_p = .01$, but when her gender was salient, men rated her as significantly less skilled in the context of computer programming, than in the context of English literature, $F(1,82) = 13.4$, $p < .001$, $n^2_p = .15$.

Men's responses to the question of how likely they would be to hire the candidate followed a similar pattern. There was again a main effect of Domain, $F(1,82) = 9.98$, $p = .002$, $n^2_p = .11$, no main effect of prime, $F(1,82) = .70$, $p = .41$, $n^2_p = .009$, and as predicted, a Domain X Prime interaction, $F(1,82) = 5.70$, $p = .02$, $n^2_p = .067$. The key comparison was significant: in the context of computer programming, men primed with the candidate's Asian ($M = 4.83$, $SD = 1.9$), rather than female ($M = 3.53$, $SD = 1.97$), identity rated her to be significantly more hirable, $F = 4.96$, $p = .03$, $n^2_p = .06$. By contrast, the priming manipulation did not alter men's perceptions of her hirability in the context of the English literature job, $F(1,79) = 1.27$, $p = .26$, $n^2_p = .02$ (Female Prime, $M = 5.76$, $SD = 1.45$; Asian Prime, $M = 5.14$, $SD = 1.93$). Again, when primed with her race, men rated her as equally hirable in the context of computer programming and English literature, $F(1,79) = .33$, $p = .57$, $n^2_p = .004$, but when primed with her gender, men rated her as significantly less hirable for the Computer programming job as compared to the English literature job, $F(1,82) = 14.12$, $p < .001$, $n^2_p = .15$.

The same pattern emerged for men's pay recommendations. There was a main effect of domain, $F(1, 82) = 7.24$, $p = .009$, $n^2_p = .08$, no main effect of priming $F(1,82) = 1.86$, $p = .18$, $n^2_p = .02$, and the predicted Domain X Prime interaction was significant, $F(1, 82) = 6.86$, $p = .01$, $n^2_p = .08$. The key comparison again emerged: in the context of the computer programming job, men offered the applicant significantly more pay when they had been primed with her Asian identity ($M = 12.57$, $SD = 1.75$) rather than her female identity ($M = $

11.12, $SD = 1.16$), $F(1, 79) = 7.57$, $p = .007$, $n^2_p = .09$. There were no differences by priming condition in men's pay recommendations in the context of English literature, $F(1,79) = .83$, $p = .37$, $n^2_p = .01$ (Female Prime, $M = 13.05$, $SD = 1.6$; Asian Prime, $M = 12.59$, $SD = 1.87$). When primed with her race, men did not differ in their assigned pay between a computer programming and English literature position, $F(1,79) = .003$, $p = .96$, $n^2_p < .001$, but when primed with her gender, men suggested that she received significantly less pay for the Computer programming job as compared to the English literature job, $F(1, 79) = 12.94$, $p = .001$, $n^2_p = .14$.

Unlike the results obtained for male participants, exploratory 2 (Domain: Computer Programming vs. English Literature) x 2 (Prime: Female vs. Asian) ANOVAs for female participants revealed no significant main effects or interactions for women's ratings of overall skill ($M=5.47$, $SD=1.20$, $Fs<.9$, $ps>.34$), willingness to hire ($M=5.45$, $SD=1.64$, $Fs<.4$, $ps>.5$), or suggested pay ($M=\$12.70$, $SD=1.58$, $Fs<1.4$, $ps>.25$). These results suggest that women are either insensitive to the priming manipulation, given that they share an ingroup identity with the target of evaluation (i.e., their gender), or that women do not exhibit similar biases under these conditions.

Overall, men rated the candidate's abilities lower, were less willing to hire her, and recommended paying her significantly less when she had applied for a tutoring job in Computer Programming rather than in English Literature. These findings replicate the extant literature documenting negatively biased evaluations of women in stereotypically male fields (Heilman & Parks-Stamm, 2007). The predicted interaction emerged for each of the dependent variables and the pattern of results support to our hypotheses[2]. Specifically, when the Asian-American woman candidate had applied to the Computer Programming job, even though she was the same person with the same skills, male evaluators rated her skill and hirability as being lower, and recommended less pay, when they had been primed with her

gender identity rather than her racial identity. These results support our theory that the degree of men's employment discrimination can depend upon which of targets' multiple stereotype-relevant identities are salient and on the domain of judgment being stereotype-relevant, in this case STEM-related. The consequences of this psychological process are perhaps most striking in men's wage discrimination in the context of the computer programming job: the same Asian American woman with the same qualifications was offered, on average, $1.45 less when the evaluators had been primed with her gender rather than her race, which would amount to a difference of over $3,000 per year given a 40-hour work week. At the same time, the conclusions we can draw from these results are limited because, despite the recruitment target we set, the final sample size cannot be considered large under current field conventions (though observed power ranged from .54 for overall skill, .67 for hirability, to .76 for pay).

Study 3

The results of Studies 1 and 2 provide initial evidence that the differential salience of competing group memberships can influence male evaluators' assessment of women in stereotyped employment contexts. Specifically, across these two studies we found evidence that men evaluated an Asian female applicant for a computer technology position more favorably when her positively stereotyped Asian identity was made salient as opposed to her negatively stereotyped female identity. However, these studies also share some limitations. While we created our manipulations to avoid confounds, it is possible that the inclusion of the applicant's last name, which was Asian in origin, unintentionally raised the salience of her race in both the control condition (Study 1) and Female prime condition (Studies 1 and 2). There is also the possibility that the second language and study abroad information included in the manipulation for Study 2 changed participants' impressions of the applicant, possibly by altering their impressions of her generalized linguistic ability. In addition, in Study 2 we included a test of attention, an instructional manipulation check, but not a classic

manipulation check. Given these considerations, we conducted a third experiment with more streamlined manipulations and a more standard manipulation check. In addition, in Study 3 we build on the previous studies by recruiting an even larger sample.

Study 3 had a three-condition between-participants experimental design. As in Study 2, participants again learned about an Asian American woman who had applied for a computer programming tutoring position. Similar to Study 1, we varied whether the description of the applicant raised the salience of her race, her gender, or neither of these identities. We again hypothesized that men's ratings of the applicant would differ depending on which of her two identities were primed, with lower ratings being received when the candidate's gender, as opposed to her race, was salient. Given the lack of consistent evidence from the comparisons of the Female and Asian conditions to the Control condition in Study 1, we were agnostic as to the predictions for these contrasts. Based on the results of Study 1, and given that the control condition represents an aggregate of natural variance in people's attention to the woman's race versus gender identities, we anticipated that evaluations provided in this condition would again fall in-between the other two conditions.

Method

Participants

We conducted an a priori power analysis to estimate a minimum appropriate sample size for this study. The design was closest to that of Study 1, so we averaged the effect sizes for the three key dependent variables (Study 1 average Cohen's $d = .77$) and estimated the sample required to detect an effect with 80% power. From this analysis, the minimum sample size recommended was 69 participants total. However, this estimate is based on a "small sample-big effect" study and therefore vulnerable to underestimating the appropriate sample size. Given this, we reconsidered the target sample size given current field standards, which recommend much higher sample sizes per cell, and ultimately settled on 100 participants per

cell to guide our recruitment. Since we planned a priori to exclude participants who failed the manipulation check, we oversampled slightly.

An mTurk hit requested a single wave of 350 non-Asian U.S. American men. A total of 7 participants were excluded from the study a priori as they identified either as not being male or being Asian American at either the outset survey landing page (N=2) or in the end-of-survey demographics form (N=5). An additional 36 participants were excluded from the analyses as they failed the manipulation check, leaving a total of 307 male participants ($Mean_{age}$=34.19, SD_{age}=9.46; 32 African American/Blacks, 239 European American/Whites, 14 Latino Americans, 5 Native Americans, 6 Middle Eastern Americans, and 11 biracial/mixed/other).

Procedure

After providing informed consent, participants were told to imagine that they were "the hiring manager for a tutoring company that works with high school students." All participants were informed that they would review applicants' materials to hire a tutor for intermediate to advanced computer programming in HTML Java, and C++. Next, participants were randomly assigned to see one of three versions of the applicant's profile, which included some text and a profile picture. In the female prime condition, the applicant was referred to as "a woman" named "Christine". In the Asian prime condition, the applicant was referred to as "an Asian American" named "Chang". In the Control condition, the applicant was referred to simply as "the applicant". All other information, and the standard profile picture (of an Asian woman in business attire), was identical. In the Female, Asian, and Control conditions, respectively, participants read:

> "The next materials you review are from [a woman/an Asian American/an
> undergraduate] who was born and raised in Freemont, CA. [Christine/Chang/the
> applicant] has a 3.8 (out of 4) GPA. [Her/The/The] LinkedIn profile includes [her/a/a]

picture (below) and the following information. [after the image] [Christine's/ Chang's/The applicant's] prior work experience includes 9 months working as a peer counselor at the university career development center and 2 years working as a volunteer dance instructor at the boys and girls club. [Christine/ Chang/The applicant] lists moderate skills in GIS, C, C++, HTML, Java, and Visual basic. You see that [her/Chang's/the applicant's] other skills include familiarity with PC and Mac platforms, word/excel/powerpoint/pages/numbers/keynote, and adobe creative suite (including photoshop and publisher)."

After reviewing this profile alongside her purported profile picture, participants rated the candidate overall (not very skilled "1"–very skilled "7"), rated their willingness to hire the candidate (not very willing "1"–very willing "7"), and indicated how much she should be paid if hired ($10–$15 USD/hour in $1 increments). Participants then completed four questions asking about the content of the scenario they had read. The key item amongst these was the manipulation check question in which participants were asked to select whether the applicant's name was Chang (accurate in the Asian condition), Christine (accurate in the Female condition), or not stated (accurate in the Control condition). As noted above, participants who selected the incorrect answer for their condition were excluded prior to conducting our analyses. Finally, participants completed questions asking about the clarity of the study materials and purpose of the study, as well as a demographic questionnaire. Participants were then debriefed.

Results and Discussion

As in Study 1, we conducted a contrast testing the hypothesis that participants in the Asian prime condition would evaluate the applicant the most positively while participants in the Female prime condition would evaluate her the least positively (see Table 3 for means, standard deviations).

The hypothesized effect emerged for participants' evaluations of the applicant's overall skill, $t(304)=2.67$, $p=.008$, $n^2_p=.02$, with participants in the Asian prime condition ($M = 5.45$, $SD = 1.14$) providing the highest ratings and those in the Female prime condition providing the least positive ratings ($M = 5.24$, $SD = 1.18$). Contrasts comparing the Control prime condition ($M = 5.46$, $SD = 1.14$) to the Female prime condition, $t(304)=1.42$, $p=.16$, $n^2_p=.007$, as well as the Asian prime condition, $t(304)=1.22$, $p=.22$, $n^2_p=.005$, were not significant.

This pattern also emerged for participants' willingness to hire the applicant, $t(304)=3.21$, $p=.001$, $n^2_p=.03$. Participants reported the greatest willingness to hire following the Asian prime ($M = 5.67$, $SD = 1.31$) and least willingness following the Female prime ($M = 5.03$, $SD = 1.49$). Again, the contrasts comparing the Control prime ($M = 5.31$, $SD = 1.51$) to the Female prime condition, $t(304)=1.42$, $p=.16$, $n^2_p=.007$, as well as the Asian prime condition, $t(304)=.1.74$, $p=.08$, $n^2_p=.01$, were not significant.

Contrary to our hypotheses, there was no effect of condition on recommended wages. The contrast comparing the Asian and Female prime was nonsignificant, $t(304)=.23$, $p=.82$, $n^2_p<.001$, as were contrasts comparing the Control and Female, $t(304)=.57$, $p=.57$, $n^2_p=.001$ and the Control and Asian, $t(304)=.34$, $p=.73$, $n^2_p<.001$, prime conditions (Asian Prime, $M = 12.99$, $SD = 1.5$; Neutral Prime, $M = 13.06$, $SD = 1.5$; Female Prime, $M = 12.94$, $SD = 1.37$). Indirect effect can emerge even if the initial predictor variable (condition) does not have a total effect on the outcome variable (see Kenny, Kashy, & Bolger, 1998). Given the pattern of results in the earlier two studies, we examined the correlations between pay ratings and participants' ratings of the candidate overall ($r=.48$, $p<.001$) and their willingness to hire her ($r=.48$, $p<.001$). While beyond our initial hypotheses, these correlations justified an exploratory analysis of indirect effects through judgments of the candidate's overall skill and hirability. We analyzed the indirect effects using Process (Hayes, 2012) Model 6, 5000

Bootstrap samples, with condition as the predictor, pay as the dependent variable, overall ratings as the first proposed mediator, and willingness to hire as the second proposed mediator (see Figure 1). Given that the response scales for pay differed from overall ratings and hirability, we focused on the completely standardized indirect effects. The indirect effect of condition on pay through overall ratings was supported, effect size = .04, SE = .02, CI [.008, .09], as was the indirect effect through both overall ratings and hirability ratings, effect size = .03, SE = .02, CI [.006, .08]. The indirect effect of condition on pay only through hirability ratings was not supported, CI [-.0004, .05], see Figure 1. Together, these results suggest that the salience of an Asian woman's race versus gender can affect the wages offered to her, but only through perceivers' impressions of her overall skill, and in turn hirability.

General Discussion

In the current research we aimed to determine whether an identical applicant could be evaluated differently, and receive different outcomes, depending on the relative salience of her negatively stereotyped gender, and positively stereotyped Asian, identities. Three studies offer evidence in support of this possibility. In each study, men rated an Asian American woman as less skilled and less hirable for a technology position when they had been primed with her gender rather than her race. Men's wage estimates showed inconsistency across studies and must be interpreted with caution. In Studies 1 and 2, men recommended different pay when evaluating the candidate for a technology position, depending on which of her stereotyped identities were primed. In Study 3, only an indirect effect emerged, which may offer insight into how pay disparities could arise: Applicants whose negatively stereotyped identities are salient may be perceived as being less skilled and hirable, which may in turn reduce their offered pay. Given this, further research should unpack precisely the conditions under which wage discrimination based on identity will be most likely to arise.

Research has begun to show that the expression of discrimination can vary between individuals who differ along a single identity dimension, such as phenotypic stereotypicality (Eberhardt, Davies, Purdie-Vaughns, & Johnson, 2006; Kahn & Davies, 2011; Ma & Correll, 2011). However, to our knowledge this is the first research to document these processes in the expression of employment discrimination toward a single individual who possesses multiple identities that are associated with conflicting relevant stereotypes. Thus, this research advances theoretical and empirical approaches to intersectionality that argue that social group memberships must always be considered mutually constitutive (Goff, Thomas, & Jackson, 2008; Kulik et al., 2007; Purdie-Vaughns & Eibach, 2008; Shields, 2008) and adds to emerging research examining how intersections of identity characteristics affect social perception (Kang & Chasteen, 2009; Levin et al., 2002; Remedios et al., 2011; Sesko & Biernat, 2010; Settles, 2006). Thus, the present research takes an important step in acknowledging and empirically testing the consequences of people's multiple social identities for our understanding of discrimination.

Limitations

The current research investigated only one intersection of identities, where one identity (gender) was negatively stereotyped and another (race) was positively stereotyped. Future research should investigate the consequences for decision making when applicants possess two negatively stereotyped group memberships, as would be the case for African American or Latino American women applicants to management and/or STEM jobs. In these cases, the relevant negative racial stereotypes may result in a diminished evaluation of ability more generally (e.g., "unintelligent," or "low competence," Fiske, Cuddy, Glick, Xu, 2002), while the negative gender stereotypes may be more domain specific, relevant only to stereotypically male fields of employment. In the context of these specific social identities, then, it may be that evaluators' decisions are most negatively biased when applicants' racial

(rather than gender) identities are more salient. Alternatively, because the stereotypic associations are both negative in this context and the cognitive structures underlying them may overlap (Freeman & Ambady, 2011), it could be that dual category membership has an additive negative effect on evaluators' decisions ("double jeopardy", Beale, 1970; Reid & Comas-Diaz, 1990; but see Levin et al., 2002).

A second limitation is that our results do not provide a clear indication of whether one identity drives the differential evaluations. Studies 1 and 3 included a control condition, but comparisons with this condition yielded mixed or non-significant results. This suggests either that both identities contributed to the results, or that the current experimental design was not ideal for capturing which identity drives the effects. Early work by Macrae et al. (1995) offers a basis to suggest the former. According to their theorizing, when one identity is salient, associations with other, competing identities can be inhibited (see also Freeman & Ambady, 2011), and this may explain the contrast between the two experimental conditions. More research, with larger samples and professional hiring managers, would offer firmer conclusions about whether one identity drives people's different reactions to the job candidate, as well as the magnitude, mechanism, and generalizability of these effects.

Implications and Future Directions

One ultimate question to arise from these and other related findings is how to reduce biased decision making in employment contexts. The present research introduces a new dimension of complexity, pointing out the importance of considering perceivers' subjective construals of applicants' social identities. Identifying the influence of salient social identities on employment and wage discrimination is, of course, an important first step. Yet, decision makers may be unaware of which of applicants' many social group memberships are salient to them in any given moment, and previous research suggests that many factors, including motivations (Sinclair & Kunda, 1999), the physical environment (e.g., Cheryan, Plaut,

Davies, & Steele, 1999; Murphy, Steele, & Gross, 2007), and the presence of others (Inzlicht & Ben-Zeev, 2000) can raise or lower the relative salience of social identities. Indeed, targets of prejudice may be motivated by findings such as these to try to strategically represent their social identities in order to avoid receiving bias. Thus, the next step for research is to explore practical interventions that organizations can use to undermine the biasing effects of social group memberships on employment and pay evaluations. Mindsets or diversity ideologies that emphasize the intrinsic value of diversity (for a review, see Rattan & Ambady, 2013) could offer less vulnerability to the inappropriate influence of social identities. Our first study suggests that using specific, rather than general, criteria may also be effective. Our results with female participants (Study 2) further suggest that ingroup members may be less susceptible to inappropriate shifts in social perceptions, suggesting diverse hiring teams may be key.

Alternatively, organizations could establish clear and specific policies for employment evaluations and wage offers, which some research has shown can reduce bias in employment contexts (Bragger, et al., 2002; Norton, et al., 2004). But where should organizations set their standards for evaluations and pay? Research on employment discrimination has compared evaluations of a target against evaluations of a White woman or man to establish whether biased decision making has occurred (Davidson & Burke, 2000; Moss-Racusin, et al., 2012; Olian, et al., 1988; Terborg & Ilgen, 1975; Uhlmann & Cohen, 2007). However, in these cases, the standard of comparison may be benefitting from positive associations with their race (and for men, their gender), insofar as this identity is salient to decision makers. In the context of the present research, the question arises of whether these targets represent an appropriate normative standard. Indeed, the actual standard may reside somewhere between the negatively biased evaluations that members of stigmatized groups receive and the positively biased evaluations that members of dominant groups receive.

Determining whether biased employment and pay decision-making is driven by disadvantage, advantage, or both, will inform organizations' attempts at solutions (Lowery, Chow, & Crosby, 2009; Lowery, Chow, Knowles, & Unzueta, 2012).

Conclusions

The three experiments reported here provide evidence that the same individual can receive different evaluations, depending on which of her stereotyped identities is salient to perceivers. These findings could have important implications for organizations, policymakers, and lawmakers who want to ensure an equal access to jobs, pay, and resources, regardless of gender or racial group memberships, or even the relative salience of these identities. We hope that this research encourages further investigations into diminishing employment and wage discrimination on the part of decision makers.

References

Ambady, N., Shih, M., Kim, A., & Pittinsky, T. L. (2001). Stereotype susceptibility in children: Effects of identity activation on quantitative performance. *Psychological Science, 12,* 385-390. doi: 10.1111/1467-9280.00371.

Barden, J., Maddux, W. W., Petty, R. E., & Brewer, M. B. (2004). Contextual moderation of racial bias: The impact of social roles on controlled and automatically activated attitudes. *Journal of Personality and Social Psychology, 87*(1), 5. doi : 10.1037/0022-3514.87.1.5

Bertrand, M. & Mullainathan, S. (2004). Are Emily and Greg more employable than Lakisha and Jamal? A field experiment on labor market decisions. *The American Economic Review, 94,* 991-1013. doi: 10.3386/w9873.

Bragger, J.D., Kutcher, E., Morgan, J., & Firth, P. (2002). The effects of structured interview on reducing bias against pregnant job applicants. *Sex Roles, 46*, 215–226. doi: 10.1023/A:1019967231059.

Beale, F. (1970). Double jeopardy: To be black and female. In T. Cade (Ed.), *The black woman* (pp. 90–100). New York: New American Library.

Bureau of Labor Statistics. (2014). *Usual weekly earnings of wage and salary workers third quarter 2014.* Accessed through: http://www.bls.gov/news.release/pdf/wkyeng.pdf.

Cheryan, S. & Bodenhausen, G.V. (2000). When positive streotypes threaten intellectual performance: The psychological hazards of "model minority" status. *Psychological Science, 11,* 399-402.

Cheryan, S., Plaut, V. C., Davies, P. G., & Steele, C. M. (2009). Ambient belonging: How stereotypical cues impact gender participation in computer science. *Journal of Personality and Social Psychology, 97,* 1045–1060. doi: 10.1037/a0016239.

Crisp, R. J., & Hewstone, M. (2007). Multiple social categorization. In M. P. Zanna (Ed.),

Advances in experimental social psychology (Vol. 39, pp. 163–254). Orlando, FL: Academic Press.

Davidson, H.K. & Burke, M.J. (2000). Sex discrimination in simulated employment contexts: A meta-analytic investigation. *Journal of Vocational Behavior, 56,* 225-248. doi: 10.1006/jvbe.1999.1711.

DeNavas-Walt, C., Proctor, B.D., & Smith, J.C. (2013). U.S. Census Bureau, Current Population Reports, P60–245. *Income, Poverty, and Health Insurance Coverage in the United States: 2012.* Washington, DC: U.S. Government Printing Office.

Derous, E., Ryan, A.M., & Nguyen, H.H.D. (2012). Multiple categorization in resume screening: Examining effects on hiring discrimination against Arab applicants in field and lab settings. *Journal of Organizational Behavior, 33,* 544-570.

Eagly, A. H. & Karau, S. J. (2002). Role congruity theory of prejudice toward female leaders. *Psychological Review, 109*, 573–598. doi: 10.1037/0033-295X.109.3.573.

Eberhardt, J.L., Davies, P.G., Purdie-Vaughns, V., & Johnson, S.L. (2006). Looking Deathworthy: Perceived stereotypicality of Black defendants predicts capital-sentencing outcomes. *Psychological Science, 17,* 383-386. doi: 10.1111/j.1467-9280.2006.01716.x.

Fiske, S.T., Cuddy, A.J.C., Glick, P., Xu, J. (2002). A model of (often mixed) stereotype content: Competence and warmth respectively follow from perceived status and competition. *Journal of Personality and Social Psychology, 82*, 878-902. doi: 10.1037/0022-3514.82.6.878.

Freeman, J.B. & Ambady, N. (2011). A dynamic interactive theory of person construal. *Psychological Review, 118,* 247-279. doi: 10.1037/a0022327.

Freeman, J.B., Nakayama, K., & Ambady, N. (2013). Finger in flight reveals parallel categorization across multiple social dimensions. *Social Cognition, 31,* 792-805.

Galinsky, A.D., Hall, E.V., & Cuddy, A.J.C. (2013). Gendered races: Implications for

interracial marriage, leadership selection, and athletic participation. *Psychological Science, 24,* 498-506.

Gilbert, D. T., & Hixon, J. G. (1991). The trouble of thinking: activation and application of

stereotypic beliefs. *Journal of Personality and Social Psychology, 60*(4), 509. doi: 10.1037/0022-3514.60.4.509

Goff, P.A., Thomas, M.A., & Jackson, M.C. (2008). "Ain't I a Woman?": Towards an

intersectional approach to person perception and group-based harms. *Sex Roles, 59,* 392-403. doi: 10.1007/s11199-008-9505-4.

Goffman, E. (1963). *Stigma: Notes on the management of spoiled identity*. Englewood Cliffs,

NJ: Prentice-Hall.

Goldberg, P. (1968). Are women prejudiced against women? *Transaction, 5,* 28-30. doi:

10.1007/BF03180445.

Hall, E.V., Galinsky, A.D., & Phillips, K.W. (2015). Gender profiling: A gendered race

perspective on person-position fit. *Personality and Social Psychology Bulletin, 41,* 853-868.

Heilman, M.E. & Eagly, A.H. (2008). Gender stereotypes are alive, well, and busy producing

workplace discrimination. *Industrial and Organizational Psychology, 1,* 393-398. doi: 10.1111/j.1754-9434.2008.00072.x.

Heilman, M.E. & Parks-Stamm, E.J. (2007). Gender stereotypes in the workplace: Obstacles

to women's career progress. *Advances in Group Processes, 24,* 47-77. doi: 10.1016/S0882-6145(07)24003-2.

Heilman, M.E. (1984). Information as a deterrent against sex discrimination: The effects of

applicant sex and information type on preliminary employment decisions. *Organizational Behaviour and Human Performance, 33,* 174-186. doi: 10.1016/0030-5073(84)90019-9.

Inzlicht, M. & Ben-Zeev, T., (2000). A threatening intellectual environment: Why females are susceptible to experiencing problem-solving deficits in the presence of males. *Psychological Science, 11*, 365-371. doi: 10.1111/1467-9280.00272.

Ito, T.A. & Urland, G.R. (2003). Race and gender on the brain: Electrocortical measures of attention to the race and gender of multiply categorizable individuals. *Journal of Personality and Social Psychology, 85,* 616-626. doi: 10.1037/0022-3514.85.4.616.

Johnson, K.L., Freeman, J.B., & Pauker, K. (2011). Race is gendered: How covarying phenotypes and stereotypes bias sex categorization. *Journal of Personality and Social Psychology, 102,* 1116-131.

Kahn, K.B. & Davies, P.G. (2011). Differentially dangerous? Phenotypic racial stereotypicality increases implicit bias among ingroup and outgroup members. *Group Processes and Intergroup Relations, 14,* 569-580. doi: 10.1177/1368430210374609.

Kang, S.K. & Chasteen, A.L. (2009). Beyond the double-jeopardy hypothesis: Assessing emotion on the faces of multiply-categorizable targets of prejudice. *Journal of Experimental Social Psychology, 45,* 1281-1285.

Kulik, C.T., Roberson, L., & Perry, E.L. (2007). The multiple-category problem: Category activation and inhibition in the hiring process. *Academy of Management Review, 32,* 529-548.

Levin, S., Sinclair, S., Veniegas, R.C., & Taylor, P.L. (2002). Perceived discrimination in the context of multiple group memberships. *Psychological Science, 13,* 557-560.

Lowery, B.S., Chow, R.M., & Randall-Crosby, J. (2009). Taking from those that have more

and giving to those that have less: How inequity frames affect corrections for

inequity. *Journal of Experimental Social Psychology, 45*, 375–378. doi:

Lowery, B.S., Chow, R.M., Knowles, E.D., & Unzueta, M.M. (2012). Paying for positive

group esteem: How inequity frames affect Whites' responses to redistributive

policies. *Journal of Personality and Social Psychology, 102*, 323-336. doi:

10.1016/j.jesp.2008.09.010.

Ma, D.S. & Correll, J. (2011). Target prototypicality moderates racial bias in the decision to

shoot. *Journal of Experimental Social Psychology, 47*, 391-396. doi:

10.1016/j.jesp.2010.11.002.

Macrae, C. N., Bodenhausen, G. V., & Milne, A. B. (1995). The dissection of selection in

social perception: Inhibitory processes in social stereotyping. *Journal of Personality

and Social Psychology, 69*, 397–407. doi: http://psycnet.apa.org/doi/10.1037/0022-

3514.69.3.397.

Markus, H.R. & Kitayama, S. (1991). Culture and the self: Implications for cognition,

emotion, and motivation. *Psychological Review, 98,* 224-253. doi: 10.1037/0033-

295X.98.2.224.

McConahay, J.B. (1983). Modern racism and modern discrimination: The effects of race,

racial attitudes, and context on simulated hiring decisions. *Personality and Social

Psychology Bulletin, 9*, 551-558. doi:

McMahon, J.M. & Kahn, K.B. (2016). Benevolent racism? The impact of target race on

ambivalent sexism. *Group Processes and Intergroup Relations, 19,* 169-183.

Migdal, M.J., Hewstone, M., & Mullen, B.M. (1998). The effects of crossed categorization

on intergroup evaluations: A meta-analysis. *British Journal of Social Psychology, 37*,

303-324.

Mitchell, J. P., Nosek, B. A., & Banaji, M. R. (2003). Contextual variations in implicit

evaluation. *Journal of Experimental Psychology: General*, *132*(3), 455. doi:

10.1037/0096-3445.132.3.455

Moss-Racusin, C. A., Dovidio, J. F., Brescoll, V. L., Graham, M., & Handelsman, J. (2012).

Science faculty's subtle gender biases favor male students. *Proceedings of the*

National Academy of Sciences, 109, 16474-16479. doi: 10.1177/0146167283094004.

Murphy, M. C., Steele, C. M., & Gross, J. J. (2007). Signaling threat: How situational cues

affect women in math, science, and engineering settings. *Psychological Science, 18*,

879–885. doi: 10.1111/j.1467-9280.2007.01995.x.

Norton, M. I., Vandello, J. A., & Darley, J. M. (2004). Casuistry and social category bias.

Journal of Personality and Social Psychology, 87, 817–831. doi: 10.1037/0022-

3514.87.6.817.

Obama, B.O. (2014). Ensuring equal pay for equal work. accessed through:

http://www.whitehouse.gov/the-press-office/2014/04/12/weekly-address-ensuring-

equal-pay-equal-work

Olian, J.D., Schwab, D.P., & Haberfeld, Y. (1988). The impact of applicant gender compared

to qualifications on hiring recommendations: A meta-analysis of experimental studies.

Organizational Behavior and Human Decision Processes, 41, 180-195. doi:

10.1016/0749-5978(88)90025-8.

Oppenheimer, D.M., Meyvis, T., & Davidenko, N. (2009). Instructional manipulation

checks: Detecting satisficing to increase statistical power. *Journal of Experimental*

Social Psychology, 45, 867-872.

Pittinsky, T. L., Shih, M., & Ambady, N. (2000). Will a category cue affect you? Category

cues, positive stereotypes, and reviewer recall for applicants. *Social Psychology of*

Education, 4, 53-65. doi: 10.1023/A:1009656413789.

Purdie-Vaughns, V. & Eibach, R.P. (2008). Intersectional invisibility: The distinctive

advantages and disadvantages of multiple subordinate-group identities. *Sex Roles, 59,* 377-391. doi: 10.1007/s11199-008-9424-4.

Rattan, A. & Ambady, N. (2013). Diversity ideologies and intergroup relations: An examination of colorblindness and multiculturalism. *European Journal of Social Psychology, 43*, 12-21. doi: 10.1002/ejsp.1892.

Reid, P.T. & Comas-Diaz, L. (1990). Gender and Ethnicity: Perspectives on Dual Status. *Sex Roles, 22,* 397-408.

Remedios, J.D., Chasteen, A.L., Rule, N.O., & Plaks, J.E. (2011). Impressions at the intersection of ambiguous and obvious social categories: Does gay + Black = likeable? *Journal of Experimental Social Psychology, 47,* 1312-1315.

Reuben, E., Sapienza, P., & Zingales, L. (2014). How stereotypes impair women's careers in science. *Proceedings of the National Academy of Sciences*, *111*(12), 4403-4408.

Sesko, A.K. & Biernat, M. (2010). Prototypes of race and gender: The invisibility of Black women. *Journal of Experimental Social Psychology, 46,* 356-360.

Shields, S.A. (2008). Gender: An intersectionality perspective. *Sex Roles, 59*, 301-311. doi: 10.1007/s11199-008-9501-8.

Shih, M., Pittinsky, T. L., & Ambady, N. (1999). Stereotype susceptibility: Identity salience and shifts in quantitative performance. *Psychological Science, 10,* 80-83. doi:

Shug, J., Alt, N.P., & Klauer, K.C. (2015). Gendered race prototypes: Evidence for the non-prototypicality of Asian men and Black women. *Journal of Experimental Social Psychology, 56,* 121-125.

Sinclair, L. & Kunda, Z. (1999). Reactions to a black professional: Motivated inhibition and activation of conflicting stereotypes. *Journal of Personality and Social Psychology, 77*, 885-904. doi: 10.1111/1467-9280.00111.

Steele, J. R., George, M., Cease, M., Fabri, T.L., & Schlosser, J. (2017). Not always Black

and White: The effect of race and emotional expression on implicit attitudes. *Unpublished Manuscript*.

Swim, J., Borgida, E., Maruyama, G., Myers, D.G. (1989). Joan McKay versus John McKay: Do gender stereotypes bias evaluations? *Psychological Bulletin, 105,* 409-429. doi: 10.1037/0033-2909.105.3.409.

Terborg, J.R. & Ilgen, D.R. (1975). A theoretical approach to sex discrimination in traditionally male occupations. *Organizational Behavior and Human Performance, 13,* 352-376. doi: 10.1016/0030-5073(75)90056-2.

Turner, J.C., Oakes, P.J., Haslam, S.A., McGarty, C. (1994). Self and Collective: Cognition and Social Context. *Personality and Social Psychology Bulletin, 20,* 454-463.

Uhlmann, E.L. & Cohen, G.L. (2007). "I think it, therefore it's true": Effects of self-perceived objectivity on hiring discrimination. *Organizational Behavior and Human Decision Processes, 104*, 207-223. doi: 10.1016/j.obhdp.2007.07.001.

Urban, L.M. & Miller N. (1998). A theoretical analysis of crossed categorization effects: A meta-analysis. *Journal of Personality and Social Psychology, 74,* 894-908.

Wittenbrink, B., Judd, C. M., & Park, B. (2001). Spontaneous prejudice in context: variability in automatically activated attitudes. *Journal of Personality and Social Psychology*, *81*(5), 815. doi: 10.1037//0022-3514.81.5.815

Wong, P., Lai, C. F., Nagasawa, R., & Lin, T. (1998). Asian Americans as a model minority: Self-perceptions and perceptions by other racial groups. *Sociological Perspectives, 41*, 95-118.

World Economic Forum (2015). The Global Gender Gap Report. Retrieved from http://www3.weforum.org/docs/GGGR2015/cover.pdf.

Yee, A.H., (1992). Asians as stereotypes and students: Misperceptions that persist. *Educational Psychology Review, 4,* 95-132.

Footnotes

[1] Priming the applicant's Asian identity might raise negatively stereotyped expectations in the context of English literature, but we anticipated that any negative expectations would be offset because the applicant had a high GPA at a highly competitive American university (Stanford University).

[2] To understand whether men who hold a stigmatized racial identity (African American, Latino American, Native American, biracial men) respond differently to the study materials than men who do not (White American men), we re-ran the analyses reported in the main text for Studies 2 and 3 controlling for majority versus minority racial status. In both Studies 2 and 3, all of the results were essentially unchanged when controlling for participant race, which did not significantly predict any of the outcome variables.

Figure 1: The indirect effects of condition on pay through overall skill and willingness to hire (Study 3). Path 1 represents the effect of condition through overall skill on wages. Path 2 represents the effect of condition through overall skill and willingness to hire on wages. Path 3 represents the effect of condition through willingness to hire on wages. Dashed lines represent the unsupported indirect paths. Path 1 CI (.008, .09), Path 2 CI (.006, .08), Path 3 CI (-.0004, .05).
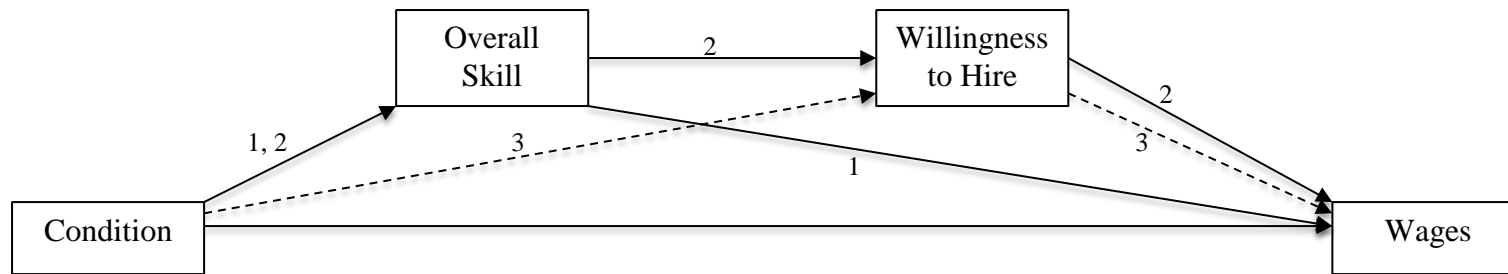
Table 1: Participants' mean ratings by priming condition (Study 1). Standard deviations are shown in parentheses. *p<.05.

	Asian Prime	Neutral Prime	Female Prime	$t_{contrast}$	*p*	$n^2{}_p$
Technical Abilities	4.85 (1.14)	4.91 (1.10)	4.40 (.81)	1.15	.26	.03
Interpersonal Abilities	5.69 (.90)	5.34 (1.10)	5.27 (1.02)	1.10	.28	.03
Overall Skill	5.69 (.63)	5.47 (1.06)	4.67 (.92)	2.43	.02*	.13
Willingness to Hire	7.31 (1.18)	6.47 (1.88)	5.87 (1.36)	2.50	.02*	.14
Wages	$9.42 (1.63)	$8.53 (.99)	$8.33 (1.13)	2.28	.03*	.12

Table 2: Participants' mean ratings by priming condition and domain condition (Study 2). Standard deviations are shown in parentheses. No significant differences emerged for women's ratings. For men's ratings, means within a row that do not share a letter are marginally or significantly different from each other (see main text for details).

Men's Ratings	English Literature		Computer Programming	
	Female Prime	Asian Prime	Female Prime	Asian Prime
Overall Skill	6.00 (1.14) a	5.64 (1.09) a	4.53 (1.28) b	5.26 (1.39) a
Willingness to Hire	5.76 (1.45) a	5.14 (1.93) a	3.53 (1.97) b	4.83 (1.90) a
Wages	$13.05 (1.60) a	$12.59 (1.87) a	$11.12 (1.16) b	$12.57 (1.75) a
Women's Ratings	English Literature		Computer Programming	
	Female Prime	Asian Prime	Female Prime	Asian Prime
Overall Skill	5.57 (1.29)	5.59 (1.33)	5.26 (1.10)	5.44 (1.08)
Willingness to Hire	5.36 (1.55)	5.73 (1.86)	5.35 (1.58)	5.40 (1.68)
Wages	$12.86 (1.43)	$12.91 (1.87)	$12.35 (1.37)	$12.68 (1.68)

Table 3: Participants' mean ratings by priming condition (Study 3). Standard deviations are

shown in parentheses.

	Asian Prime	Neutral Prime	Female Prime	$t_{contrast}$	*p*	$n^2{}_p$
Overall Skill	5.45 (1.14)	5.46 (1.14)	5.24 (1.18)	2.67	.008	.02
Willingness to Hire	5.67 (1.31)	5.31 (1.51)	5.03 (1.49)	3.21	.001	.03
Wages	$12.99 (1.50)	$13.06 (1.50)	$12.94 (1.37)	.23	.82	.00