Original Manuscript

Social Psychological and Personality Science
1–12

DOI: 10.1177/19485506221127493
journals.sagepub.com/home/spp
SAGE

Asian Men and Black Women Hold Weaker Race–Gender Associations: Evidence From the United States and China

Jordan R. Axt[1,2], S. Atwood[3], Thomas Talhelm[4], and Eric Hehman[1]

Abstract

Prior work finds a consistent association between race and gender: People associate Asian with female and Black with male. We used mouse-tracking to examine whether different U.S. racial/ethnic groups hold this same association (Study 1) and compared Asian-American participants to ethnically Chinese participants in China (Study 2). In Study 1, White and Hispanic participants showed the expected "race is gendered" effect, and the strength of the effect did not differ between men and women. However, participants with a counter-stereotypical racial-gender identity (Black women and Asian men) showed weaker race–gender associations. The same pattern emerged for East Asian participants in Study 2, both among people living in the United States and China. These data provide the first evidence of moderation in Asian-female, Black-male associations and further reveal the importance of considering intersectional identities in social cognition and social perception.

Keywords

stereotypes, race, gender, intersectionality, culture, mouse-tracking

People hold multiple social identities, such as those tied to age, religion, and sexual orientation. Social cognition research has focused on the consequences of these identities but has largely examined them in isolation. More recently, psychology has begun to appreciate the intersectional relationships between identities (Purdie-Vaughns & Eibach, 2008). One prominent example is the association between race and gender. Several studies have found that people associate Black people with the concept of male and Asian people with the concept of female (Goff et al., 2008; Schug et al., 2015). Researchers have found evidence of this race–gender association in explicit, self-reported stereotypes (Lei et al., 2020) and also in indirect or "implicit" outcomes that are under less conscious control. For example, when asked to categorize faces by gender, participants responded slower to Black female than Asian female faces, as well as to Asian male than Black male faces (Lei et al., 2020).

A popular method to test for implicit race–gender associations is mouse-tracking (Freeman & Ambady, 2009), which uses computer mouse trajectories to infer the strength of competing response tendencies. One mouse-tracking task where participants categorized Asian and Black faces by gender found that while categorizations were largely correct, mouse paths deviated more toward "Men" for Black female faces than for Asian female faces and more toward "Women" for Asian male faces than for Black male faces (Johnson et al., 2012).

Although this race–gender association is well-supported, the forces involved in the effect are less clear. Here, we extend prior work by investigating the strength of these mouse-tracking effects across racial/ethnic groups in the United States (Study 1) and between Asian-American versus ethnic Chinese participants in China (Study 2). No prior studies have investigated cultural variation in race–gender associations, and little work has explored the role of participant race. Exceptions come from recent studies in developmental psychology that used racially diverse samples (Lei et al., 2020, 2022). However, these prior studies on race–gender associations either did not test for or found no evidence of moderation by participant race, such as when participants categorized Black and Asian targets by

[1]McGill University, Montreal, Quebec, Canada
[2]Project Implicit, Seattle, WA, USA
[3]Princeton University, NJ, USA
[4]The University of Chicago Booth School of Business, IL, USA

Corresponding Author:
Jordan R. Axt, Department of Psychology, McGill University, 2001 McGill College Avenue, Montreal, Quebec, Canada H3A 1G1.
Email: jordan.axt@mcgill.ca

gender (Lei et al., 2020) or when selecting among prototypes of different races to represent the categories of *men* or *women* (Lei et al., 2022).

Exploring race–gender associations across racial/ethnic groups and between cultures has several benefits. For one, increasing sample diversity creates more generalizable knowledge about social categorization and reduces overreliance on WEIRD (Western, educated, industrialized, rich and democratic) samples (Cheon et al., 2020). In addition, this type of comparative study can shed light on the relative importance of certain factors in the formation of race–gender associations.

Ingroup and Cultural Influences on Implicit Associations

Prior research has highlighted several moderators of the strength of people's implicit associations. In particular, these studies focused on implicit *attitudes* (the degree to which a group is associated with positive versus negative), which are distinct from implicit *stereotypes* (the degree to which a group is associated with a certain trait). However, studies suggest that implicit attitudes and implicit stereotypes are related (Kurdi et al., 2019; Phills et al., 2020), and it is therefore possible that moderators of implicit attitudes also emerge for implicit stereotypes.

Two prominent moderators of implicit attitudes are ingroup identity and group status (i.e., being a majority versus stigmatized group member). Research has shown that people have more positive implicit evaluations toward members of their own group (Nosek et al., 2007). For example, in a study on implicit attitudes toward Black versus Asian people, White Americans showed relative neutrality in implicit attitudes, whereas Asian American participants ($d = -.50$) and Black American participants ($d = .46$) showed evidence of ingroup favoritism (Axt et al., 2014). In terms of group status, prior work has also found that members of stigmatized groups have weaker levels of implicit ingroup favoritism than members of majority groups (Essien et al., 2021). Even participants sharing an ingroup identity (e.g., Judaism) express more implicit ingroup positivity when living in a Jewish-majority than Jewish-minority culture (Axt et al., 2018). Researchers have attributed such majority/minority effects on implicit attitudes to mechanisms like intergroup contact (MacInnis et al., 2017) or the experience of social power (Axt et al., 2018).

The present work explores whether ingroup identity and group status influence race–gender associations. In Study 1, if ingroup identity or being a member of a stigmatized group influences race–gender associations, we should find weaker race–gender associations among East Asian and Black participants than White participants. Or, when keeping racial identity constant in Study 2, comparing Asian-American to participants in China allows us to contrast the role of majority versus minority group status. If being a minority group influences race–gender associations, then race–gender associations should be stronger among East Asian participants living in the United States.

Although such results cannot conclusively support or refute certain mechanisms behind the formation of race–gender associations, they would be broadly consistent or inconsistent with various accounts. For instance, if race–gender associations are influenced by intergroup contact, and Asian and Black Americans have higher levels of contact with Asian and Black people on average than White people (Logan et al., 2004), race–gender associations should be weaker for Asian and Black than White participants. Similarly, if race–gender associations are influenced by exposure to counter-stereotypical exemplars (Dasgupta & Greenwald, 2001), Chinese participants should show smaller race–gender associations than Asian-American participants because Chinese participants live in a culture where more Asian men occupy positions associated with masculinity (Lu et al., 2020). Failure to find such differences—either between White versus Asian or Black participants in the United States or between Asian-American versus ethnically Chinese participants—would cast doubt on mechanisms like intergroup contact or exposure to counter-stereotypes as influences in the formation of race–gender associations.

However, a competing perspective focuses on participants' racial *and* gender identities. When considering race and gender, Asian men and Black women can be considered "counter-stereotypical" because their race and gender identities are not aligned with the stereotyped association (at least in the United States). Here, we adopt language from prior research (e.g., Johnson & Ghavami, 2011; Pietri et al., 2021). The formation of race–gender associations may be more related to whether such associations align with one's self-concept, meaning only counter-stereotypical participants would show weaker associations because the stereotype is inconsistent with their self-identity and therefore not encoded as strongly (Greenwald et al., 2002). That is, Asian men are less likely to perceive themselves as feminine, and Black women less likely to perceive themselves as masculine, creating weaker race–gender associations.

Support for this view comes from a math-gender IAT (Implicit Association Test) modified for mouse-tracking (Smeding et al., 2016) that produced stronger male–math, female–language associations among stereotypical group members (female humanities students and male engineers) than counter-stereotypical group members (male humanities students and female engineers). These results align with data from gender-science IATs (Smyth & Nosek, 2015), which show that women with degrees in "hard" sciences (biology, engineering) had weaker male-science associations than men holding the same degrees.

We tested predictions from this perspective by seeing whether mouse-tracking performance was moderated by participants' racial and gender identities. Such results would indicate that race–gender associations are not shaped by factors shared among members of a racial group, but instead that race–gender associations are formed through more nuanced and intersectional forces; specifically, that the strength of such associations depends on the degree to which they align with one's identity.

Deviations From Preregistered Analyses

We preregistered methods, target sample sizes, and analyses. Original analyses focused only on moderation by participants' racial identity, meaning analyses incorporating participants' racial and gender identities are exploratory.

Study 1

Method

Participants. Participants completed all studies online. In Study 1, we recruited 1,134 participants from Project Implicit (PI; M_{age} = 34.3, SD = 15.3). Participants completed demographics when first registering for PI. The study was available to U.S. citizens/residents whose racial/ethnic identity was "White," "Black," "East Asian," or "Hispanic." We excluded mouse-tracking trials with long responses (>2,000 ms; + 3 SD from the average response time; Freeman et al., 2016), which removed 5.66% of responses, including all data from 63 participants. See https://osf.io/jqtye/ for Study 1's preregistration and https://osf.io/47mw9/ for data, materials, and analysis syntax. We included Hispanic participants to examine whether race–gender associations were moderated by a broader minority status, based on past work showing that members of racial minority groups sometimes have a shared sense of identity (Cortland et al., 2017).

We targeted a minimum sample of 200 for each racial/ethnic subsample. Our sample of 1,071 eligible participants was slightly larger: N_{White} = 329 (89 men, 234 women, 6 other or not reported), N_{Asian} = 202 (69 men, 130 women, 3 other or not reported), N_{Black} = 279 (62 men, 214 women, 3 other or not reported), $N_{Hispanic}$ = 261 (60 men, 196 women, 5 other). This provided greater than 80% power for detecting a medium-level effect (γ = .30) in a cross-level interaction between target race, target gender, and participant gender within each racial/ethnic group, based on prior simulations (Mathieu et al., 2012).

Procedure and Measures. In a fixed order, participants completed a mouse-tracking task, a race–gender stereotypes questionnaire, and a priming task. These latter measures were included to examine whether mouse-tracking performance was moderated by explicit and another implicit measure of race–gender stereotypes.

Mouse-Tracking Task. The mouse-tracking task involved two 60-trial blocks, where participants categorized an equal number of Black and East Asian male and female faces using labels "Male" and "Female." Stimuli came from the Chicago Face Database (Ma et al., 2015), and were selected such that male faces were as similar as possible in masculinity ratings (Black M = 4.95, SD = 0.27; East Asian M = 4.22, SD = 0.35) and female faces as similar as possible in femininity ratings (Black M = 4.93, SD = 1.74; East Asian M = 4.85, SD = 1.63). Like in prior work, participants received a warning to respond more quickly if the mouse did not move for 600 ms following trial onset.

Self-Reported Race–Gender Associations. Participants completed a 5-item measure assessing the degree to which Black and East Asian people were associated with masculinity or femininity. The first question was a relative assessment of Black versus East Asian people (1 = "I strongly associate Black people with masculinity and East Asian people with femininity," 7 = "I strongly associate East Asian people with masculinity and Black people with femininity'). Next, were two thermometer items measuring the extent to which East Asian or Black people are likely to behave in a masculine and feminine manner (1 = "East Asian people, much more," 7 = "Black people, much more"). Finally, two slider items measured the degree to which East Asian and Black people are associated with masculinity and femininity (−100 = "strongly with masculinity," 100 = "strongly with femininity"). Difference scores were made from the slider and thermometer items. Difference scores and the relative preference item were standardized and averaged to make a composite (McDonald's ω = .77).

Implicit Race–Gender Associations. Participants completed a three-block, 96-trial priming task. In each trial, a Black or East Asian face (the same as the mouse-tracking task) appeared for 200 ms and was replaced by a male-associated (e.g., son) or female-associated word (e.g., daughter), which participants categorized as "male" or "female" as quickly as possible. Performance was scored using a D algorithm (Bar-Anan & Nosek, 2014). Participants were excluded from analyses using the priming measure if >40% of responses were incorrect. Online supplement Appendix A reports more information on the task.

Results

Mouse-Tracking Performance. Analyses focused on four measures of decisional conflict (Hehman et al., 2015; March & Gaertner, 2021): maximum deviation, area under the curve (AUC), response latency, and spatial disorder. We calculated maximum deviation and AUC independently, but both are considered measures of response conflict. We analyzed both for converging evidence, but as expected based on prior research (Freeman et al., 2008), they were highly correlated (e.g., r = .89 in Study 1). Response latency has also been used to measure response conflict, under the logic that categorization takes longer when there is greater

Table 1. Target Race × Target Gender Interaction Terms by Mouse-Tracking Outcome and Participant Race.

Outcome	Estimate	*SE*	*df*	*t*	*p*
Black participants					
Maximum deviation	−.05	.02	89.44	−2.54	.013
Latency	−18.86	11.24	82.14	−1.68	.097
AUC	−.03	.009	108.27	−2.76	.007
Spatial disorder	−.01	.003	88.50	−2.37	.020
East Asian participants					
Maximum deviation	−.09	.02	97.82	−4.18	<.0001
Latency	−35.07	11.55	75.97	−3.04	.003
AUC	−.04	.01	99.90	−4.10	<.0001
Spatial disorder	−.01	.003	96.43	−4.31	<.0001
Hispanic participants					
Maximum deviation	−.08	.02	88.95	−4.28	<.0001
Latency	−23.52	11.64	84.65	−2.02	.046
AUC	−.05	.01	108.31	5.38	<.0001
Spatial disorder	−.01	.003	103.17	−3.80	.0003
White participants					
Maximum deviation	−.06	.02	89.60	−2.41	.018
Latency	−24.03	11.35	93.17	−2.12	.037
AUC	−.02	.02	77.06	−1.42	.160
Spatial disorder	−.01	.004	74.70	−2.51	.014

Note. All values rounded to 2 decimal places unless additional places are needed to accurately express numerical value. AUC = area under the curve.

competition between response options (Johnson et al., 2012). Finally, spatial disorder (calculated as entropy) is a metric of response complexity: trajectories that alternate back and forth at different speeds are high in complexity (March & Gaertner, 2021). We included each outcome so readers can assess the robustness of results (Appendix B in the online supplement presents a correlation matrix).

We group-mean centered all trial-level outcomes by participant and ran multi-level models using the lme4 package (Kuznetsova et al., 2017) predicting each trial-level outcome by target race and target gender. Initial models focused on whether race–gender associations existed among White, East Asian, Black, and Hispanic participants, which included fixed effects for target race, target gender, and their interaction along with random intercepts for subject and stimulus.[1] Analyses found reliable target by target gender interactions, which were significant in 14 of 16 tests (the exceptions being latency for Black participants and AUC for White participants; see Table 1 for interaction terms and Online Appendix C for full models). Exploratory follow-up analyses looking at differences in decisional outcomes by target gender found that Black male faces produced less decisional conflict than Black female faces in 11 of 16 tests, and similar analyses by target gender for Asian faces found that female faces produced less decisional conflict than male faces in 11 of 16 outcomes (see Online Appendix D). Broadly, the pattern of results was consistent with the "race is gendered" effect in previous work (Johnson et al., 2012). We also investigated whether the magnitude of any target race by target gender interactions differed for White vs. Black, White vs. Asian, and White vs. Hispanic participants. Only two significant contrasts at $p < .05$ emerged across 12 analyses (Table 2). Results were also consistent when comparing White and non-White participants (see Online Appendix E).

Finally, we incorporated both participants' racial and gender identities (limited to the 98.4% reporting gender as "man" or "woman"). Among Hispanic and White participants, social groups not included as targets in the task, participant gender did not moderate mouse-tracking outcomes. However, for Black participants, gender moderated all four outcomes (see Table 3 and Figure 1) and for Asian participants, gender moderated three of four outcomes. Specifically, Black women (vs. Black men) and Asian men (vs. Asian women) showed weaker race–gender associations (see Online Appendix F). Exploratory follow-up analyses revealed that effects were largely driven by own-race faces. For Black participants, the interaction between target gender and participant gender was significant at $p < .05$ for none of the mouse-tracking outcomes in Asian targets, but all four outcomes in Black targets. That is, relative to Black men, Black women showed less gender-biased processing of Black targets than for East Asian targets. Similarly, for East Asian participants, participant gender moderated performance for three of four outcomes in East Asian targets but only one of four outcomes in Black targets (see Online Appendix G).

Explicit and Implicit Race–Gender Associations. Self-report was also consistent with the race is gendered effect, indicating stronger associations between Black people with masculinity and Asian people with femininity (Relative preference: $M = 4.48$, $SD = 0.82$, Thermometer difference: $M =$

Table 2. Target Race × Target Gender × Minority Interaction Terms by Mouse-Tracking Outcome and Participant Race.

Outcome	Estimate	*SE*	*df*	*t*	*p*
White vs. Asian participants					
Maximum deviation	−.03	.02	59,433.16	−1.81	.070
Latency	−11.35	6.37	59,429.70	−1.78	.075
AUC	−.02	.01	59,443.18	−1.66	.097
Spatial disorder	−.01	.003	58,898.02	−2.17	.030
White vs. Black participants					
Maximum deviation	.01	.01	66,379.47	.65	.514
Latency	4.75	6.12	66,375.19	.78	.438
AUC	−.005	.01	66,386.99	−.43	.670
Spatial disorder	.001	.002	65,772.93	.48	.629
White vs. Hispanic participants					
Maximum deviation	−.03	.01	65,280.37	−1.71	.087
Latency	.31	6.23	65,275.55	.05	.961
AUC	−.03	.01	65,288.73	−2.67	.008
Spatial disorder	−.002	.002	64,686.68	−.64	.522

Note. All values rounded to 2 decimal places unless additional places are needed to accurately express numerical value. AUC = area under the curve.

Table 3. Target Race × Target Gender × Participant Gender Interaction Terms by Mouse-Tracking Outcome and Participant Race.

Outcome	Estimate	*SE*	*df*	*t*	*p*
Black participants					
Maximum deviation	−.09	.03	29,498.45	−3.20	.001
Latency	−23.84	11.38	29,493.50	−2.10	.036
AUC	−.06	.02	29,461.11	−2.85	.004
Spatial disorder	−.01	.004	29,498.99	−3.31	.001
East Asian participants					
Maximum deviation	.07	.03	22,436.49	2.42	.015
Latency	39.91	10.60	22,439.32	3.77	.0002
AUC	.04	.02	22,378.42	2.13	.033
Spatial disorder	.01	.01	22,436.92	1.00	.317
Hispanic participants					
Maximum deviation	.01	.03	28,045.63	.30	.767
Latency	−4.73	11.62	28,042.68	−.41	.684
AUC	.01	.02	27,960.87	.64	.525
Spatial disorder	−.001	.004	28,041.48	−.24	.809
White participants					
Maximum deviation	−.02	.02	35,855.95	.80	.424
Latency	−.79	8.92	35,854.03	−.09	.930
AUC	−.02	.02	35,865.11	−1.29	.200
Spatial disorder	.002	.004	35,864.36	.41	.684

Note. All values rounded to 2 decimal places unless additional places are needed to accurately express numerical value. AUC = area under the curve.

1.13, SD = 1.55, Slider difference: M = 18.30, SD = 36.40; Online Appendix H lists descriptive statistics of individual items). Using the composite measure, there were no reliable differences between White versus Hispanic (d = .09), White versus Black (d = .06) and White versus Asian (d = .14) participants (see Online Appendix I). Within each racial/ethnic group, male and female participants did not reliably differ in self-reported stereotypes, though effect sizes were moderate in some cases (White d = .16, Black d = .07, Asian d = .23, Hispanic d = .28; see Online Appendix J).

The priming task was scored such that greater values indicated faster identification of male words following Black (vs. Asian) faces and of female words following Asian (vs. Black) faces. D scores did not reliably differ from 0 (M = .002, SD = .40, d < .0001), likely because participants relied on prime gender over race. Priming task performance did not reliably differ between White versus Black (d = .07) or White versus Hispanic (d = .12) participants, though White participants showed a slightly stronger effect than Asian participants (d = .19; see Online Appendix K). Participant gender did not moderate priming

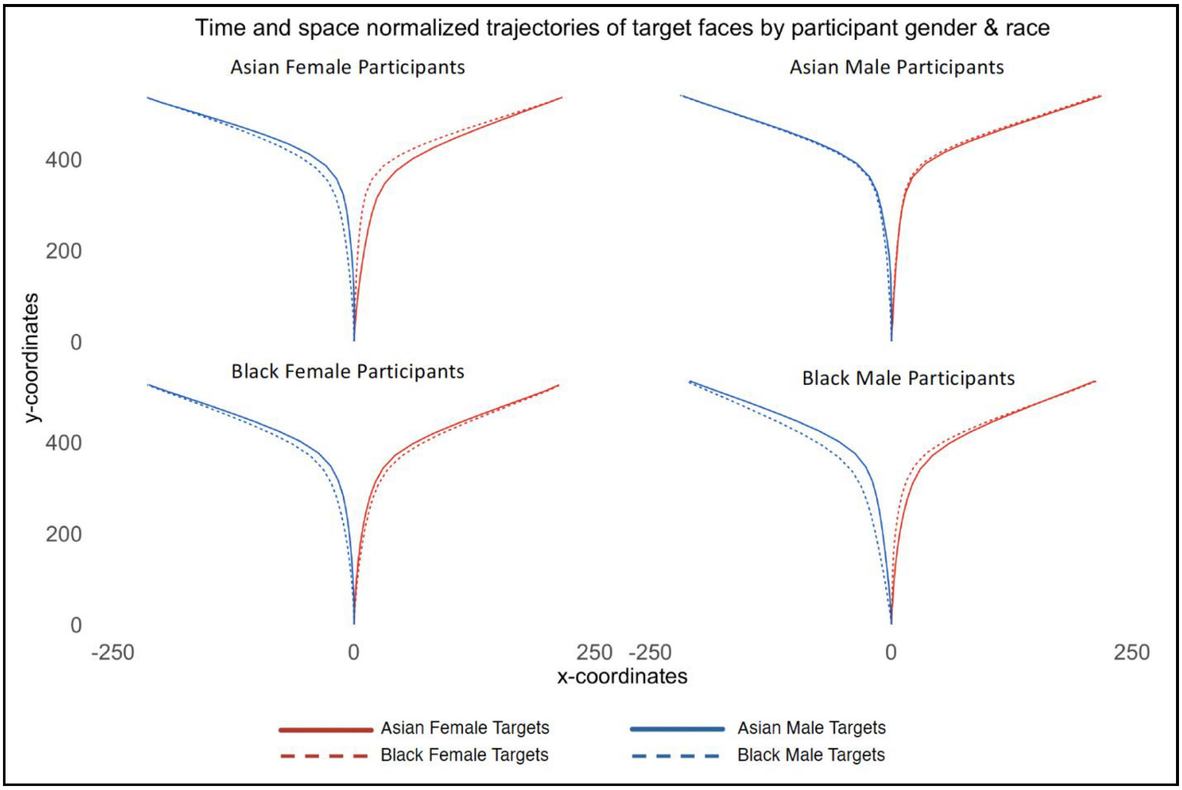

Figure 1. Time and Space Normalized Trajectories for Each Combination of Participant Gender (Columns) and Participant Race (Rows). More Direct Paths Represent Less Decisional Conflict.

task performance (White d = .14, Black d = .11, Asian d = .09, Hispanic d = .18; see Online Appendix L).

Given that only the self-report measure showed evidence of race–gender associations, exploratory analyses investigated whether explicit race–gender stereotypes moderated mouse-tracking performance. Self-reported stereotypes did not moderate the race by gender interaction for any mouse-tracking outcome (see Online Appendix M).

Discussion

Ignoring gender, White, Black, Asian, and Hispanic participants demonstrated race–gender associations in mouse-tracking, and non-White participants did not differ from White participants in the strength of these effects. These results depart from past studies in implicit intergroup attitudes, which found robust differences among these groups (Axt et al., 2014, 2018) and reveal how this same moderator was not present in a measure of implicit stereotypes.

Yet analyses that incorporated participant gender revealed a more complex pattern. For White and Hispanic participants, gender did not affect mouse-tracking. However, gender consistently moderated performance among Black and Asian participants: Counter-stereotypical group members (Asian men and Black women) showed weaker effects than stereotypical group members (Black men and Asian women). This effect was largely due to the processing of same-race faces, not other-race faces.

Study 2 replicates and extends the results by exploring cultural context. We compared mouse-tracking performance between Asian-American participants and ethnically Chinese participants in China. Study 2 enables another test of whether participants' racial and gender identities moderate race–gender associations, and most critically, investigates this question in contexts where Asian people are a minority (United States) or majority group (China).

Study 2

Method

Participants. Participants came from three sources. American participants were undergraduates at a midwestern university (N = 87, M_{age} = 22.0, SD = 2.90; 75.9% female) and volunteers from Project Implicit (N = 81, M_{age} = 20.0, SD = 3.06; 60.5% female, restricted to under 27 years old to better match the undergraduate sample). Prescreening could only identify "Asian" participants (instead of "East Asian" in Study 1), so we added a question differentiating between East and South Asian. American participants were eligible if they (a) reported race solely as "East Asian," (b) identified as Asian-American, and (c) lived in the United States for at least 5 years. These

Table 4. Target Race × Target Gender Interaction Terms by Mouse-Tracking Outcome and Participant Sample.

Outcome	Estimate	*SE*	*df*	*t*	*p*
US participants					
Maximum deviation	−.02	.05	11,483.00	−.44	.660
Latency	−38.00	10.32	87.42	−3.68	.0004
AUC	−.03	.01	118.88	−2.08	.040
Spatial disorder	−.02	.004	88.92	−3.98	.0001
Chinese participants					
Maximum deviation	−.17	.03	83.29	−6.52	<.0001
Latency	−61.56	12.18	64.03	−5.05	<.0001
AUC	−.07	.01	53.09	−6.30	<.0001
Spatial disorder	−.03	.005	64.89	−5.88	<.0001

Note. All values rounded to 2 decimal places unless additional places are needed to accurately express numerical value. AUC = area under the curve.

Table 5. Target Race × Target Gender × Location Interaction Terms by Mouse-Tracking Outcome.

Outcome	Estimate	*SE*	*df*	*t*	*p*
Chinese vs. U.S. participants					
Maximum deviation	.14	.04	31,631.75	3.37	< .001
Latency	23.33	7.85	31,719.29	2.97	.003
AUC	.04	.02	31,625.61	2.54	.011
Spatial disorder	.01	.003	31,723.63	3.39	< .001

Note. All values rounded to 2 decimal places unless additional places are needed to accurately express numerical value. AUC = area under the curve.

exclusions dropped 56.7% of the PI sample and 25.9% of the undergraduate sample, primarily for identifying as South Asian. We excluded mouse-tracking trials using the criteria of Study 1, leading to an additional exclusion of 0.73% of data in the PI sample and 1.56% of data in the undergraduate sample.

We recruited participants in China through a research pool in Wuhan, Hubei (N = 201, M_{age} = 21.6, SD = 3.01; 69.2% women). Participants were eligible if their reported ethnicity was "Han" (92.46% of data), the majority ethnic group in China. Trial-level exclusions dropped an additional 2.03% of data. See https://osf.io/pvru2/ for preregistration.

We targeted a minimum sample of N = 200 but were able to recruit 99 American (71 women, 27 men, 1 other) and 184 Chinese participants (128 women, 56 men). This sample provided a minimum of 70% power for detecting a medium-level interaction effect within the American and Chinese samples, using the Study 1 criteria.

Procedure. Participants completed the same mouse-tracking task and measure of self-reported race–gender associations (ω = .78) as Study 1. All participants completed a 14-item demographics questionnaire. American participants reported years spent living in the United States, whether they were born in the United States, whether their parents were born in the United States, whether they identify as Chinese-American and/or Asian-American, and the extent to which they identify with American culture as well as their heritage culture. Project Implicit participants completed Study 1's priming measure (data available in online datasets).

Results

Mouse-Tracking Analyses. Ignoring participant gender, we investigated race–gender associations separately among U.S. and Chinese samples. Among Asian-American participants, three of four outcomes showed a reliable race-by-gender interaction. Among participants in China, all outcomes showed reliable race-by-gender interactions (Table 4 reports interaction terms; Online Appendix N reports full results). Follow-up exploratory analyses found that Black male faces produced less decisional conflict than Black female faces in five of eight tests, and Asian female faces produced less decisional conflict than Asian male faces in six of eight tests (see Online Appendix O).

Next, we investigated whether the race-by-gender interaction was moderated by location. Analyses found a race-by-gender-by-location interaction for each outcome. Chinese participants showed *stronger* associations between Asian with female and Black with male (see Table 5 and Online Appendix P). This pattern is the reverse of what would be expected if the "race is gendered" effect was

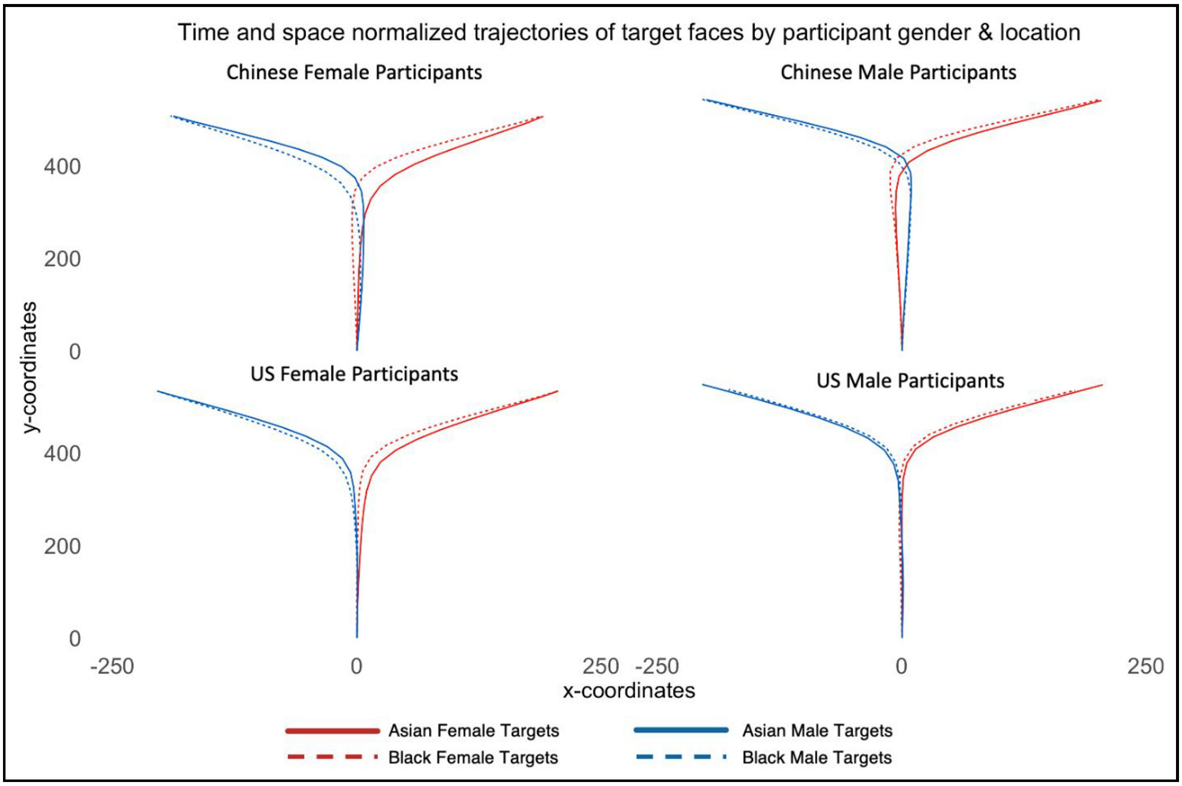

Figure 2. Time and Space Normalized Trajectories for Each Combination of Participant Gender (Columns) and Participant Location (Rows) in Study 2. More Direct Paths Represent Lower Levels of Decisional Conflict.

Table 6. Target Race × Target Gender × Participant Gender Interaction Terms by Mouse-Tracking Outcome and Participant Location.

Outcome	Estimate	*SE*	*df*	*t*	*p*
U.S. participants					
Maximum deviation	−.07	.11	11,362.00	−0.62	.540
Latency	−16.44	13.62	11,263.22	−1.21	.228
AUC	−.05	.03	11,264.20	−1.89	.059
Spatial disorder	−.02	.01	11,263.17	−3.87	.0001
Chinese participants					
Maximum deviation	−.10	.03	20,232.73	−3.32	.001
Latency	−23.28	10.35	20,233.58	−2.25	.025
AUC	−.05	.02	20,002.52	−2.34	.019
Spatial disorder	−.01	.01	20,237.55	−2.19	.030

Note. All values rounded to 2 decimal places unless additional places are needed to accurately express numerical value. AUC = area under the curve.

weakened by participants living in cultures where their racial group is the majority.

Finally, we incorporated participant gender. For participants in China, gender moderated each target-race-by-target-gender interaction. For American participants, gender moderated the race-by-gender interaction for spatial disorder only, though the result was marginal for AUC and directionally consistent for all outcomes. The four-way interaction between target gender, target race, participant gender, and location was not reliable for any outcomes. Men had weaker race–gender associations than women (see Figure 2, Table 6 and Online Appendix Q).

Exploratory follow-up analyses again suggested that these effects were from the processing of same-race faces, though evidence was stronger in the Chinese sample. Among Chinese participants, the target gender by participant gender interaction was significant for all four outcomes in East Asian targets, and zero outcomes in Black targets. Among American participants, the target gender by participant gender interaction was significant for two of four outcomes in East Asian targets, and one outcome in Black targets (see Online Appendix R).

Explicit Race–Gender Associations. Replicating Study 1, participants self-reported a stronger association of Black people with masculine and Asian people with feminine (Relative preference: M = 4.58, SD = 1.03, Thermometer difference: M = 1.45, SD = 1.70, Slider difference: M = 30.71, SD = 43.21). The composite measure did not reveal differences between U.S. and Chinese participants (d = .01), or between male and female participants in China (d = .02) or the United States (d = .37), though effect sizes were moderate in the latter case (see Online Appendix S).

Exploratory Moderator Analyses. We ran exploratory analyses among Asian-American participants examining whether strength of American or heritage identity moderated the race-by-gender interaction, and neither variable moderated race–gender associations (see Online Appendices T and U). Finally, exploratory analyses investigating whether explicit race–gender stereotypes moderated mouse-tracking performance found no evidence of moderation (see Online Appendix V).

Meta-Analysis of Asian American Participants. Although Study 1 found that Asian-American men showed weaker race–gender associations than Asian-American women for three (of four) outcomes, the same was true for only one outcome in Study 2. To better understand these results, we combined mouse-tracking data for Asian-American men and women in Studies 1 and 2 and examined whether participant gender moderated target race-by-gender interactions. This individual data meta-analysis (Tierney et al., 2019) found weaker race–gender associations for Asian-American men in all mouse-tracking outcomes except maximum deviation (all ps $< .01$; see Online Appendix W).

Discussion

Chinese participants showed a stronger race–gender association in mouse-tracking than Asian-American participants. Participant gender moderated the strength of these associations, such that male participants showed weaker associations than female participants, though moderation was more consistent in the Chinese sample. Weaker race–gender associations were driven more by less biased processing of own-race than other-race faces, though again support for this finding was stronger in the Chinese sample.

General Discussion

Two studies explored whether race–gender associations in mouse-tracking differed across racial/ethnic groups and between cultures. Prior work had demonstrated the "race is gendered" pattern primarily among White Americans, and the underlying causes of the effect are unclear. In an American sample (Study 1), racial/ethnic minorities did not differ from White participants in the interaction between target race and target gender on measures of decisional conflict. In Study 2, both Asian-American and Chinese participants showed the effect and results were actually stronger in China, where Asian ethnicity is the majority. These results are inconsistent with common moderators in the implicit attitudes literature because ingroup identity (here, racial identity) and majority vs. minority status did not moderate implicit race–gender stereotypes. Results then fail to support the idea that mechanisms like intergroup contact or exposure to counter-stereotypical exemplars are integral in the formation of race–gender associations.

Subsequent analyses found that effects were weaker based on participants' gender—specifically for participants whose racial and gender identities made them counter-stereotypical. Black women showed weaker effects than Black men, and Asian men showed weaker effects than Asian women. This is the first evidence of any moderation in the implicit association between male with Black and female with Asian, and highlights the importance of examining variation among subcategories (e.g., Black men and women) of larger groups.

Results are consistent with work showing that weaker implicit associations were limited to group members that were themselves counter-stereotypical (Smeding et al., 2016). The present data are an important extension of this perspective by providing stronger—though not conclusive—evidence for the causal role of having a counter-stereotypical identity. One important difference from prior research is that earlier studies used identities that were more self-selected, like university major (Smeding et al., 2016; Smyth & Nosek, 2015), allowing for results to be explained by selection effects (i.e., women with initially weaker male–science associations pursue science degrees).

This account is considerably less plausible with the present data. These studies then provide the strongest evidence to date that the experience of being counter-stereotypical is directly related to weaker implicit stereotypes about one's own identity. In further support of this view, weaker race–gender associations among counter-stereotypical participants was largely driven by the processing of faces that matched participants' racial and gender identities. For example, among Black participants in Study 1, Black women (versus Black men) showed weaker gender effects on decisional conflict for Black targets, but not East Asian targets.

Notably, self-reported race–gender associations were strong in all samples. Participants willingly reported a connection between Asian with female and Black with male. That the combination of participant racial and gender identity did not moderate self-reported associations in the same manner as mouse-tracking behavior suggests that there are differing mechanisms behind the creation of self-reported versus more automatic expressions of race–gender associations.

We did not preregister analyses incorporating participant gender. Therefore, while the gender findings are generally consistent, they are ultimately exploratory. Confirmatory analyses will provide stronger evidence.

Cross-Cultural Emergence of Race–Gender Associations

The Study 2 finding of stronger race–gender effects among Chinese participants can illuminate factors that give rise to associations between Asian with female and Black with male. For one, these data indicate that race–gender associations exist beyond contexts where Asian people are a minority or Black people are a significant portion of the population. In addition, results may speak to the pervasiveness of the association between Asian with female. For instance, one analysis of Chinese movie posters found an increase in the feminization of male characters over time (Hu, 2018), a trend that may have spurred the Chinese government's recent ban of "effeminate men" on television (Associated Press, 2021). In all, Study 2 results suggest that associations between race and gender are not only a function of minority group membership but are instead sensitive to broader cultural factors.

Future Directions

There are several intriguing directions for this work. One improvement is for future studies to randomize all measures. Studies here used a fixed order, which leaves open the possibility that participants may have become aware of the study's purpose after the mouse-tracking task, potentially contaminating responses on subsequent measures. It would also be informative to examine associations with masculinity and femininity (rather than male or female), as such labels may reveal stronger effects.

Another direction is to use samples from majority-Black countries. We anticipate that comparing Black participants in such countries to those in the United States would show a similar pattern as our Study 2 analysis of Asian-American versus Chinese participants, such that Black women would show weaker race–gender associations relative to Black men in both contexts. However, data confirming (or refuting) this prediction would increase generalizability of our findings. In addition, including White faces would allow for tests of whether Black female or Asian male faces create greater response conflict than White counterparts. Finally, future work could explore whether these associations influence important outcomes, like evaluation of job candidates or romantic partners (e.g., Galinsky et al., 2013).

Conclusion

Associations between race and gender do not exist equally across social groups or cultural contexts. In particular, the experience of a counter-stereotypical identity may weaken the formation of such associations. The findings add nuance to prior studies concerning how stigmatized identities express implicit associations. In particular, these insights came from adopting an intersectional approach to the study of stereotypes and social categorization, and an intersectional perspective should continue to influence future research in this area.

Declaration of Conflicting Interests

The author(s) declared the following potential conflicts of interest with respect to the research, authorship, and/or publication of this article: This research was partly supported by Project Implicit. J. R. Axt is Head of Data and Methodology for Project Implicit, Inc., a nonprofit organization with the mission to "develop and deliver methods for investigating and applying phenomena of implicit social cognition, including especially phenomena of implicit bias based on age, race, gender, or other factors." There are no other potential conflicts of interest with respect to authorship or the publication of this article.

Funding

The author(s) received no financial support for the research, authorship, and/or publication of this article.

ORCID iDs

Jordan R. Axt https://orcid.org/0000-0003-3944-8331
Thomas Talhelm https://orcid.org/0000-0002-0954-5758
Eric Hehman https://orcid.org/0000-0003-2227-1517

Data Availability

All data and study materials are available at the project page on the Open Science Framework (https://osf.io/47mw9/).

Supplemental Material

Supplemental Material for this article is available online.

Notes

1. All models with two random intercepts yielded boundary fit warnings, so we removed the random intercept of subject and found the same pattern of results. In-text analyses report results from the regressions with both intercepts, but the full model of both versions of each regression is included in the online supplement.

References

Associated Press. (2021, September 2). China bans effeminate men from TV. *National Public Radio*. https://www.npr.org/2021/09/02/1033687586/china-ban-effeminate-men-tv-official-morality

Axt, J. R., Ebersole, C. R., & Nosek, B. A. (2014). The rules of implicit evaluation by race, religion, and age. *Psychological Science*, *25*, 1804–1815.

Axt, J. R., Moran, T., & Bar-Anan, Y. (2018). Simultaneous ingroup and outgroup favoritism in implicit social cognition. *Journal of Experimental Social Psychology*, *79*, 275–289.

Bar-Anan, Y., & Nosek, B. A. (2014). A comparative investigation of seven indirect attitude measures. *Behavior Research Methods*, *46*, 668–688.

Cheon, B. K., Melani, I., & Hong, Y. Y. (2020). How USA-centric is psychology? An archival study of implicit assumptions of generalizability of findings to human nature based on origins of study samples. *Social Psychological and Personality Science*, *11*, 928–937.

Cortland, C. I., Craig, M. A., Shapiro, J. R., Richeson, J. A., Neel, R., & Goldstein, N. J. (2017). Solidarity through shared disadvantage: Highlighting shared experiences of discrimination improves relations between stigmatized groups. *Journal of Personality and Social Psychology*, *113*, 547–567.

Dasgupta, N., & Greenwald, A. G. (2001). On the malleability of automatic attitudes: Combating automatic prejudice with images of admired and disliked individuals. *Journal of Personality and Social Psychology*, *81*, 800–814.

Essien, I., Calanchini, J., & Degner, J. (2021). Moderators of intergroup evaluation in disadvantaged groups: A comprehensive test of predictions from system justification theory. *Journal of Personality and Social Psychology*, *120*, 1204–1230.

Freeman, J. B., & Ambady, N. (2009). Motions of the hand expose the partial and parallel activation of stereotypes. *Psychological Science*, *20*(10), 1183–1188.

Freeman, J. B., Ambady, N., Rule, N. O., & Johnson, K. L. (2008). Will a category cue attract you? Motor output reveals dynamic competition across person construal. *Journal of Experimental Psychology*, *137*(4), 673–690.

Freeman, J. B., Pauker, K., & Sanchez, D. T. (2016). A perceptual pathway to bias: Interracial exposure reduces abrupt shifts in real-time race perception that predict mixed-race bias. *Psychological Science*, *27*, 502–517.

Galinsky, A. D., Hall, E. V., & Cuddy, A. J. (2013). Gendered races: Implications for interracial marriage, leadership selection, and athletic participation. *Psychological Science*, *24*, 498–506.

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.

Greenwald, A. G., Banaji, M. R., Rudman, L. A., Farnham, S. D., Nosek, B. A., & Mellott, D. S. (2002). A unified theory of implicit attitudes, stereotypes, self-esteem, and self-concept. *Psychological Review*, *109*, 3–25.

Hehman, E., Stolier, R. M., & Freeman, J. B. (2015). Advanced mouse-tracking analytic techniques for enhancing psychological science. *Group Processes & Intergroup Relations*, *18*, 384401.

Hu, L. (2018). Is masculinity "deteriorating" in China? Changes of masculinity representation in Chinese film posters from 1951 to 2016. *Journal of Gender Studies*, *27*, 335–346.

Johnson, K. L., Freeman, J. B., & Pauker, K. (2012). Race is gendered: How covarying phenotypes and stereotypes bias sex categorization. *Journal of Personality and Social Psychology*, *102*, 116–131.

Johnson, K. L., & Ghavami, N. (2011). At the crossroads of conspicuous and concealable: What race categories communicate about sexual orientation. *PLOS ONE*, *6*, Article e18025.

Kurdi, B., Mann, T. C., Charlesworth, T. E., & Banaji, M. R. (2019). The relationship between implicit intergroup attitudes and beliefs. *Proceedings of the National Academy of Sciences*, *116*, 5862–5871.

Kuznetsova, A., Brockhoff, P. B., & Christensen, R. H. (2017). lmerTest package: Tests in linear mixed effects models. *Journal of Statistical Software*, *82*, 1–26.

Lei, R. F., Leshin, R. A., Moty, K., Foster-Hanson, E., & Rhodes, M. (2022). How race and gender shape the development of social prototypes in the United States. *Journal of Experimental Psychology*, *151*(8), 1956–1971.

Lei, R. F., Leshin, R. A., & Rhodes, M. (2020). The development of intersectional social prototypes. *Psychological Science*, *31*, 911–926.

Logan, J. R., Stults, B. J., & Farley, R. (2004). Segregation of minorities in the metropolis: Two decades of change. *Demography*, *41*, 1–22.

Lu, D., Huang, J., Seshagiri, A., Park, H., & Griggs, T. (2020, September 9). Faces of power: 80% are White even as US becomes more diverse. *The New York Times*. https://www.nytimes.com/interactive/2020/09/09/us/powerful-people-race-us.html

Ma, D. S., Correll, J., & Wittenbrink, B. (2015). The Chicago face database: A free stimulus set of faces and norming data. *Behavior Research Methods*, *47*, 1122–1135.

MacInnis, C. C., Page-Gould, E., & Hodson, G. (2017). Multilevel intergroup contact and antigay prejudice (explicit and implicit) evidence of contextual contact benefits in a less visible group domain. *Social Psychological and Personality Science*, *8*, 243–251.

March, D. S., & Gaertner, L. (2021). A method for estimating the time of initiating correct categorization in mouse-tracking. *Behavior Research Methods*, *53*, 2439–2449.

Mathieu, J. E., Aguinis, H., Culpepper, S. A., & Chen, G. (2012). Understanding and estimating the power to detect cross-level interaction effects in multilevel modeling. *Journal of Applied Psychology*, *97*, 951–966.

Nosek, B. A., Smyth, F. L., Hansen, J. J., Devos, T., Lindner, N. M., Ranganath, K. A., & . . . Banaji, M. R. (2007). Pervasiveness and correlates of implicit attitudes and stereotypes. *European Review of Social Psychology*, *18*, 36–88.

Phills, C. E., Hahn, A., & Gawronski, B. (2020). The bidirectional causal relation between implicit stereotypes and implicit prejudice. *Personality and Social Psychology Bulletin*, *46*, 1318–1330.

Pietri, E. S., Johnson, I. R., Majid, S., & Chu, C. (2021). Seeing what's possible: Videos are more effective than written portrayals for enhancing the relatability of scientists and promoting black female students' interest in stem. *Sex Roles*, *84*, 14–33.

Purdie-Vaughns, V., & Eibach, R. P. (2008). Intersectional invisibility: The distinctive advantages and disadvantages of multiple subordinate-group identities. *Sex Roles*, *59*, 377–391.

Schug, 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.

Smeding, A., Quinton, J. C., Lauer, K., Barca, L., & Pezzulo, G. (2016). Tracking and simulating dynamics of implicit stereotypes: A situated social cognition perspective. *Journal of Personality and Social Psychology*, *111*, 817–834.

Smyth, F. L., & Nosek, B. A. (2015). On the gender–science stereotypes held by scientists: Explicit accord with gender-ratios, implicit accord with scientific identity. *Frontiers in Psychology*, *6*, Article 415.

Tierney, J. F., Stewart, L. A., & Clarke, M. (2019). Individual participant data meta-analysis. In J. P. T. Higgins, J. Thomas, J. Chandler, M. Cumpston, T. Li, M. J. Page, & V. A. Welch (Eds.), *Cochrane handbook for systematic reviews of interventions* (pp. 643–658). Cochrane.

Author Biographies

Jordan R. Axt is an Assistant Professor at McGill University. His research explores how people form and express intergroup bias in attitudes and behavior.

S. Atwood is a graduate student in psychology and social policy at Princeton University. Their research explores the development, perceptions, and implications of gender diversity in adults and children.

Thomas Talhelm is an Associate Professor of Behavioral Science at the University of Chicago Booth School of Business. Thomas researches culture and how rice farming gave southern China a very different culture from northern China. He has lived in China for seven years as a freelance journalist and a Fulbright scholar.

Eric Hehman is an Associate Professor at McGill University. His research examines how individuals perceive and evaluate one another across group boundaries (e.g., race, gender, sexual orientation, occupation, etc.), and the downstream consequences of such perceptions.

Handling Editor: Margo Monteith