

Psychological Outcomes of Prototypicality in Marginalized Group Members

Rebecca I. Mohr

Submitted in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy
in the Graduate School of Arts and Sciences

Columbia University
2018

©2018
Rebecca I. Mohr
All Rights Reserved

ABSTRACT

Psychological Outcomes of Prototypicality in Marginalized Group Members

Rebecca I. Mohr

Social psychologists have long been interested in judgments of stereotyping, prejudice, and discrimination towards people with marginalized identities. However, the majority of past social psychological studies have focused on understanding how perceivers view one dimension of marginalized identity in isolation from other marginalized identities. Specifically, past studies typically focus on the group members who are believed to be the most prototypical of marginalized groups in order to examine processes associated with discrimination (e.g., using Black men as targets when studying prejudice towards Black people). Because previous work largely examines the perceptions and experiences of prototypical marginalized group members, our understanding of stereotyping, prejudice, and discrimination is incomplete. I report five studies that use the intersectional invisibility framework as a scaffold to explore how perceivers' judgments of prototypical versus non-prototypical group members differ. In Study 1, I explore how non-prototypical marginalized group members are represented in the media relative to their prototypical counterparts. Study 2 measures how explicit perceived stereotypes of prototypical and non-prototypical marginalized group members differ. Study 3 investigates how perceivers make attributions about prototypical and non-prototypical marginalized groups. Studies 4 and 5 examine how perceivers detect discrimination towards prototypical and non-prototypical marginalized group members. These studies empirically demonstrate that non-prototypical

marginalized group members are perceived differently than their prototypical counterparts. These differences are associated with downstream consequences including reduced representation in popular culture and enhanced perceiver attention towards non-prototypical group members when compared to their prototypical counterparts.

TABLE OF CONTENTS

	Page
List of Figures	ii
List of Tables	iv
Acknowledgements	v
Dedication	vii
Introduction	1
Study 1: Archival Study of Group Representation	8
Study 2: Explicit Perceived Differences Between Prototypical and Non-Prototypical Group Members	13
Study 3: Prototypical and Non-Prototypical Marginalized Group Attributions	17
Study 4: Race-Based Discrimination Detection of Prototypical and Non-Prototypical Black Targets	28
Study 5: Gender-Based Discrimination Detection of Prototypical and Non-Prototypical Black Targets	55
Discussion	73
References	80
Appendices	90

LIST OF FIGURES

	Page
Figure 1: Representation of minority men, minority women, White men, and White women on the covers of Time Magazine.	10
Figures 2 & 3: Representation of white women compared to minority women and minority women compared to minority men on both the covers of Time Magazine and relative to U.S. demographic proportions.	11
Figure 4: Reported perceived masculinity/femininity of different prototypical and non-prototypical groups.	15
Figures 5: Number of ideas in participants' written statements that focused on a non-prototypical group (lesbian women) relative to a prototypical group (straight women) differing by sexual orientation.	22
Figure 6: Scores indicating which group (lesbian women or straight women) participants believed would be more likely to change in the future.	22
Figure 7: Number of ideas in participants' written statements that focused on a non-prototypical group (lesbian women) relative to prototypical LGBTQ individuals (gay men) differing by gender.	23
Figure 8: Scores indicating which group (lesbian women or gay men) participants believed would be more likely to change in the future.	23
Figure 9: Number of ideas in participants' written statements that focused on a non-prototypical group (Black women) relative to a prototypical group (White women) differing by race.	24
Figure 10: Scores indicating which group (Black women or White women)	25

participants believed would be more likely to change in the future.

Figure 11: Number of ideas in participants' written statements that focused on a non-prototypical group (Black women) relative to a prototypical group (Black men) differing by gender. 26

Figure 12: Scores indicating which group (Black women or Black men) participants believed would be more likely to change in the future. 26

Figure 13: Percentage of written responses mentioning discrimination, by discriminatory comparison (Study 4). 42

Figure 14: Perceived fairness by hiring comparison (Study 4). 45

Figure 15: Scores of perceived deliberate action by hiring manager (Study 4). 48

Figure 16: Self-reported concern over hiring decision (Study 4). 50

Figure 17: Percentage of correct recognition responses by discriminatory comparison (Study 4). 51

Figure 18: Percentage of written responses mentioning discrimination by discriminatory comparison (Study 5). 63

Figure 19: Perceived fairness by hiring comparison (Study 5). 65

Figure 20: Scores of perceived deliberate action by hiring manager (Study 5). 68

Figure 21: Self-reported concern over hiring decision (Study 5). 70

Figure 22: Percentage of correct recognition responses by discriminatory comparison (Study 5). 71

Figure A1: Sample quantitative summaries used in Studies 4 & 5. This example is from a non-discriminatory decision. 90

LIST OF TABLES

	Page
Table 1: Scores of Perceived Masculinity-Femininity by Social Group	15
Table 2: Correlations Between Individual Difference Measures (Study 4)	38
Table 3: Percentage of Participants Who Wrote about Discrimination by Condition (Study 4)	41
Table 4: Log Odds that a Participant Wrote About Discrimination by Condition (Study 4)	41
Table 5: Scores of Perceived Fairness by Hiring Comparison (Study 4)	43
Table 6: Scores of Perceived Deliberate Action by Hiring Manager (Study 4)	46
Table 7: Scores of Self-Reported Concern About Each Hiring Decision (Study 4)	48
Table 8: Correlations Between Individual Difference Measures (Study 5)	60
Table 9: Percentage of Participants Who Wrote about Discrimination by Condition (Study 5)	62
Table 10: Log Odds that a Participant Wrote About Discrimination by Condition (Study 5)	62
Table 11: Scores of Perceived Fairness by Hiring Comparison (Study 5)	63
Table 12: Scores of Perceived Deliberate Action by Hiring Manager (Study 5)	66
Table 13: Scores of Self-Reported Concern About Each Hiring Decision (Study 5)	68
Table B1: Qualitative Coding Guidelines for Discrimination Type Coding (Studies 4 & 5).	91

Acknowledgments

As I wrap up my graduate student career, there are countless people I need to thank for helping me reach this point.

First, I must thank my parents, Lynn and Andrew, who instilled in me a love for learning, encouraged me to ask questions about the world, and were always my biggest cheerleaders. Thank you for all of your help in more ways than I could possibly count—from reading over my application materials when I applied to grad school, to listening to my practice talks as I prepared to reach this point.

Of course I must also thank my academic family and mentors who helped me develop at various stages of my career. Sam, Keith, and Evan introduced me to the field of social psychology as an undergrad, got me started in research, and provided me with invaluable advice as I applied for graduate school. Michele and the postdocs and graduate students from the University of Maryland introduced me to a new subfield of psychology and taught me new ways to ask psychological questions. And my mentors at Columbia both official and unofficial helped me hone in and focus my research (Valerie), remember the importance of studying real world questions (Geraldine), and approach research using new methods (Niall). Finally, I'd like to thank my committee for their thoughtful comments and guidance in this process.

I also want to thank all of the graduate students that I have had the pleasure of meeting and working with throughout the years. I especially want to thank members of CASPR and my lab mates for always being a sounding board for research and non-research challenges alike. In particular, I want to thank Kaytee

Turetsky who has shared this experience with me over the last five years. I could not have hoped to have a more supportive and brilliant lab sibling.

I have been assisted with my research over the years by a team of research assistants and lab managers who have helped me collect data, qualitatively code thousands of participant responses, and design online surveys. Although there are too many to name, I'd especially like to thank Adriana Germano, Yeji Park, and Julia Bottesini for their assistance over the years. I literally could not have done this research without you.

Last but certainly not least, I want to thank my husband Travis Riddle. He really deserves the (Red Sox) MVP (most valuable partner) award for his support throughout grad school. Thank you for cheering me up throughout challenges, providing unlimited statistics consulting, and being the best dog co-parent to Ronnie and Chance I could ever hope for.

For Lynn, Andy & Travis

Introduction

Social psychologists have long been interested in judgments of stereotyping, prejudice, and discrimination towards people with marginalized identities (i.e., Allport, 1954; Blanchard, Adelman, & Cook, 1975; Bobo, 1999; Brewer, 1999; Dovidio, 2001; Duckitt, 1992; Fiske, 1998). However, the majority of these studies have focused on understanding how perceivers view one kind of marginalized identity in isolation from other marginalized identities (Sesko & Biernat, 2010). For instance, psychologists have examined questions about how Black people experience racism and how women experience sexism for decades (Glick, & Fiske, 1996; McConahay, 1986). Moreover, social psychologists typically focus on group members who are perceived as the most representative, or *prototypical* of marginalized groups in order to examine processes associated with discrimination (Bowleg 2008; Crenshaw, 1989; Goff & Kahn, 2013). For instance, Goff, Eberhardt, Williams, & Jackson (2008) presented participants with pictures of Black men to understand how U.S. citizens dehumanize Black people. Researchers who are interested in sexism, such as Kaiser and Miller (2004), focus on the most prototypical women, White women. Taken together, much psychological literature exploring stereotypes and discrimination against marginalized group members has focused on gay men when studying perceptions of the LGBTQ community, White women as the target when studying perceptions of women, and Black men when studying perceptions of Black people (Bowleg 2008, Cole, 2009; Goff & Kahn, 2013; Kang & Bodenhausen, 2015).

This research largely overlooks the perceptions and experiences of individuals with multiple marginalized identities (e.g., Black women, lesbian women), and the fact that the combination of these identities alters the stereotypes and behaviors perceivers employ (Crenshaw, 1989; Goff & Kahn, 2013; Shields, 2008). Because conducting empirical work on non-prototypical marginalized group members is methodologically challenging (see Babbitt, 2013), only in the last decade have social psychologists begun to consider how the combination of multiple marginalized social identities affect judgments of these individuals (Purdie-Vaughns & Eibach, 2008). However, this recent attention has been primarily theoretical rather than empirical.

Prototypical vs. Non-Prototypical Group Members

Purdie-Vaughns & Eibach (2008) proposed a theoretical model that explains how certain group members are understood as prototypical versus non-prototypical of particular social identities. According to their model, the American system-maintaining ideologies of androcentrism (the tendency to view masculinity as normative), heterocentrism (the tendency to view heterosexuality as normative), and ethnocentrism (the tendency to view whiteness as normative) define the prototypical person in the United States as a straight white man (Bem, 1994; Bonilla-Silva, 2000; Devos & Banaji, 2005; Eagly & Kite, 1987; Hegarty, Pratto, & Lemieux, 2004; Miller et al., 1991; Smith & Zárate, 1992). In turn, when considering members of marginalized groups (e.g., women, Black people, LGBTQ individuals), people continue to define prototypical members of these groups through these ideologies. These ideologies thus establish the prototypical woman as a straight

White woman because, although she deviates from the norm of maleness defined by androcentrism, she is still defined by heterocentrism and ethnocentrism, which designate heterosexuality and whiteness as prototypical. Similarly, straight Black men are perceived as prototypical of Black people due to androcentrism and heterocentrism, and gay White men are perceived as prototypical of LGBTQ individuals because of androcentrism and ethnocentrism.

In contrast, according to this model, individuals who deviate from the prototypical identity defined by androcentrism, heterocentrism, and ethnocentrism on more than one dimension (e.g., both gender and race, or both race and sexual orientation) are viewed as non-prototypical of their constituent identity groups. For example, a White lesbian woman deviates from the norms of maleness and straightness defined by androcentrism and heterocentrism, respectively, and thus is perceived as a non-prototypical LGBTQ individual (prototypical group: White gay men) and a nonprototypical woman (prototypical group: straight White women). Similarly, a straight Black woman deviates from the norms defined by the ideologies of androcentrism and ethnocentrism because she is both female and nonwhite, making her a non-prototypical woman (prototypical group: straight White women) and a non-prototypical Black person (prototypical group: straight Black men). For the White lesbian woman and the straight Black woman, the combination of their identities renders them non-prototypical across multiple dimensions.

Work in person perception buttresses this idea, providing empirical evidence that individuals whose identities deviate in multiple ways from the ideologies of androcentrism, heterocentrism, and ethnocentrism are perceived as non-

prototypical. Specifically, perceivers are more likely to make mistakes in categorizing the gender of Black female faces relative to Black male and White male and female faces (Goff et al., 2008). Additionally, participants take much longer to categorize the race and gender of faces of Black women relative to those of Black men, White men, and women (Zárate & Smith, 1990). These differences in categorization accuracy and latencies indicate that participants must override their automatic defaults for gender and Blackness when identifying Black female faces, suggesting that Black women are non-prototypical of either category.

Intersectional Invisibility Theory: Invisibility and Hypervisibility

Purdie-Vaughns & Eibach's (2008) theory of intersectional invisibility predicts that perceivers view individuals who are non-prototypical across multiple dimensions (sometimes termed intersectional individuals) differently from individuals who do not have multiple non-prototypical identities. Specifically, the theory makes two distinct predictions regarding how perceivers view non-prototypical individuals compared to others.

Invisibility outcome. On the one hand, the theory predicts that perceivers could pay less attention to prototypical individuals (e.g. racial minority women, lesbian women) because perceivers have a more difficult time recognizing these non-prototypical individuals as members of their constituent identity groups. As a result, perceivers should overlook and misremember non-prototypical group members. This result is referred to as the *invisibility outcome*.

The invisibility outcome can confer disadvantages. Sesko & Biernat (2010) demonstrated a disadvantage of the invisibility outcome where Black women were

recalled less accurately, and their statements were more likely to be misattributed, relative to Black men or White women in memory tasks. Being omitted and misremembered can signify that one does not belong and can be psychologically painful to the overlooked person (Fryberg & Eason, 2017; Fryberg & Townsend, 2008; Mohr & Purdie-Vaughns, 2015).

On the other hand, non-prototypical marginalized individuals may glean some advantages from this invisibility. Because non-prototypical individuals are invisible and overlooked as members of a marginalized identity, they may escape some of the disadvantages that prototypical marginalized group members incur as a result of the system maintaining ideologies of androcentrism, ethnocentrism, and heterocentrism (Purdie-Vaughns & Eibach, 2008). For example, Livingston, Rosette, and Washington (2012), found that Black female bosses (non-prototypical women and Black people) were evaluated less harshly for acting like a leader than were both White women (prototypical women) and Black men (prototypical Black people). Additionally, Pedulla (2014) found that participants evaluated the resumes of White gay men (prototypical of gay people) more negatively than the resumes of Black gay men (non-prototypical gay people). Thus, invisibility due to having multiple non-prototypical identities can sometimes be advantageous by providing cover from damaging stereotypes, prejudice, and discrimination.

Hypervisibility outcome. Alternatively, perceivers could attend to non-prototypical group members (e.g. racial minority women, lesbian women) more than their prototypical counterparts (e.g. White women, racial minority men)

because non-prototypical group members differ from the norm and are thus more conspicuous. This outcome is known as *hypervisibility*.

Miller, Taylor, & Buck (1991) demonstrated the hypervisibility outcome when they asked participants to explain gender differences in the voting patterns of Americans. Participants were asked to explain why they believed this gender difference occurred and indicate which group (men who are prototypical for gender or women who are non-prototypical for gender) they believed would change should this gendered voting gap disappear in the future. Participants overwhelmingly attributed the voting difference to women and believed that women would be more likely to change their behavior in the future, demonstrating the hypervisibility outcome for the non-prototypical group. While Miller, Taylor, & Buck's paradigm found this outcome when comparing non-prototypical group members with only a single marginalized identity (women) to their prototypical counterparts (men), the researchers proposed that the perceived difference of a group from the norm is what leads to hypervisibility. Therefore, this outcome should extend even when non-prototypical group members have multiple marginalized identities.

Invisibility versus hypervisibility. While Purdie-Vaughns and Eibach's (2008) theory makes predictions about the different outcomes of perceiving non-prototypical group members, it does not make any predictions regarding which contexts should lead to invisibility or hypervisibility. What, then, determines whether non-prototypical individuals are invisible or hypervisible relative to prototypical group members? Previous research has found that invisibility versus hypervisibility outcomes are contingent on perceivers' motivation to identify and

combat the system maintaining ideologies of androcentrism, ethnocentrism, and heterocentrism that determine prototypicality and non-prototypicality. For instance, individual differences in perceivers' endorsement of egalitarian ideals and support for minority groups has been shown to moderate whether perceivers view non-prototypical group members as invisible or hypervisible, such that more egalitarian people attend to non-prototypical groups more (Federico, Sidanius, & Rabinowitz, 2002; Jost, 2006; Lei, Mohr, Rucker & Richeson). Additionally, features of a task, such as asking perceivers to directly compare prototypical and non-prototypical groups, may increase or decrease this motivation by highlighting disparities between groups (Marti, Bobier, & Baron, 2000; Miller, Taylor & Buck, 1991), leading non-prototypical group members to be hypervisible, although this has not yet been empirically tested.

Dissertation Aim

The majority of work on non-prototypical individuals has focused on understanding how perceivers understand and behave towards non-prototypical targets with a single marginalized identity (Miller, Taylor, & Buck, 1991; Sesko & Biernat, 2010). The work within psychology that focuses on targets who are non-prototypical of marginalized identities is largely theoretical and proscriptive (Babbitt, 2013; Cole, 2009; McCall, 2005; Purdie-Vaughns & Eibach, 2008), or focuses on bottom up processes associated with non-prototypical marginalized individuals, such as differences in person perception (Carpinella, Chen, Hamilton, & Johnson, 2015; Freeman, Johnson, Ambady, & Rule, 2010; Lick, & Johnson, 2014;

Lick, Johnson, & Riskind, 2015), without exploring the downstream consequences of perceiving non-prototypical marginalized group members.

Because non-prototypical group members of marginalized groups have largely been excluded in social psychological research, this dissertation aims to answer the overarching question of how perceivers' judgments of prototypical versus non-prototypical group members differ. Moreover, this dissertation seeks to elucidate downstream consequences of these judgments, examining implications for representation, stereotypes, and discrimination detection.

Study 1

Study 1 measured invisibility by examining representation of different groups in popular culture. Because prototypes are more accessible (Rosch & Mervis, 1975), people tend to be biased in favor of prototypical group members when deciding who should be represented. As such, Study 1 empirically tested whether racial minority women's non-prototypicality would manifest as invisibility, as predicted by Purdie-Vaughns and Eibach (2008).

I coded covers of *Time Magazine* to examine the relative representation of non-prototypical individuals (e.g., racial minority women) compared to their prototypical counterparts (e.g., White women or racial minority men). Specifically, I examined covers of *Time Magazine* because it is one of the longest running magazines and has been printed weekly since 1969. Additionally, because it is a general interest magazine, the covers should appeal to a wide swath of Americans. If the invisibility hypothesis were correct, I expected to see that *Time Magazine* would feature prototypical individuals on its cover more frequently than non-prototypical

individuals.

Sample and Procedure

I randomly selected 809 covers of *Time Magazine*. Because I was interested in the representation of different marginalized groups within the United States, I excluded any covers that did not contain people, or featured an international issue or person. This resulted in a final sample of 475 covers, which I coded for the number of people on the cover and the gender and racial/ethnic identity of the people on the cover. Interrater reliability was satisfactory, $\kappa=.855$.

Results

Descriptive statistics. Overall, there were 708 people total coded on the covers and the average number of people on a cover was 1.53 ($SD = 1.17$). Of these 708 people, 80.2% were White, 9.6% were Black, 3.3% were Latino American, 1.7% were Asian American, .4% were Arab American, and .1% were other/bi-racial; the ethnicity of the remaining 4.7% of people could not be determined. Additionally, 72% of people on the cover were male, 24% were female, and 4% could not be determined.

Following the hypothesis that White men are the most prototypical group and therefore would be the most represented, the majority of the people on the covers were White men (63%). Additionally, White women (21%) and racial minority men (12%) were shown on the covers more than racial minority women (4%), lending support for the hypothesis that White women (prototypical women) and racial minority men (prototypical people of color) were more visible than racial minority women (non-prototypical women and people of color).

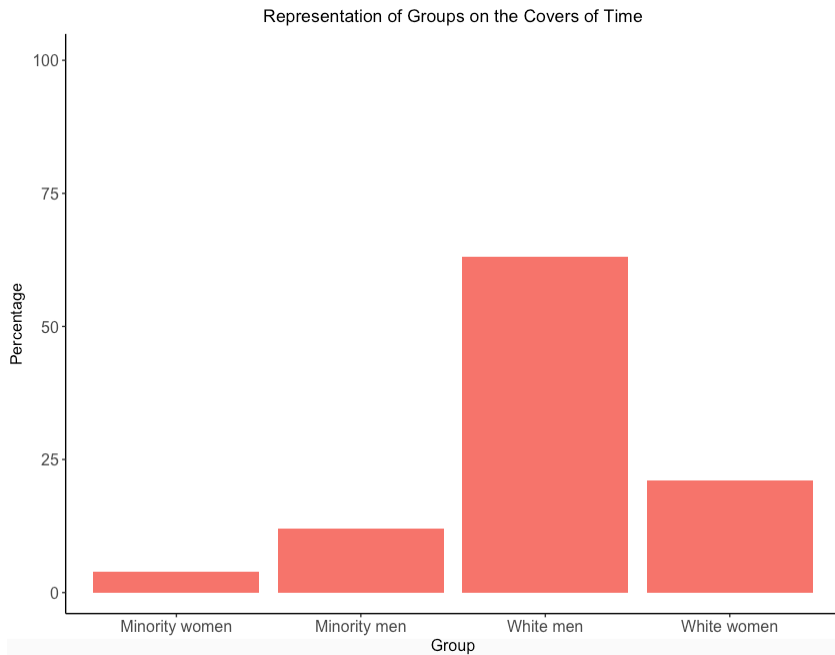



Figure 1. Representation of minority men, minority women, White men, and White women on the covers of *Time Magazine*.

Main analyses. Because White women are prototypical women and racial minority women are non-prototypical women, I predicted that, when covers depicted a woman, racial minority women should be less likely to be depicted than White women. Of the people on the covers coded as female, 141 (85%) were White and 25 (15%) were racial minority women (Black, Asian, Latina, Arab, bi-racial, and Native Americans). I performed a chi-squared test and found that White women (141 total) were shown on the covers of *Time* significantly more often than racial minority women (25 total), $\chi^2 (1) = 82.39, p < .01$. I also predicted that when covers depicted racial minorities, minority women would be less likely to be depicted on the cover than minority men. Of the 15% of people coded on covers as racial minority members, racial minority men comprised 76% (80 people total) of the people of color shown and minority women only comprised 24% (25 people total). I

again performed a chi-squared test and found that racial minority men (80) were shown on the covers of *Time* significantly more often than racial minority women (25 total), $\chi^2(1) = 27.77, p < .01$.

One possibility is that the differences in representation outlined above could simply reflect the base rates of White women and racial minority men relative to racial minority women within the United States. In order to test this alternative explanation, I obtained the population proportions for White women, minority men, and minority women in the United States (U.S. Census Bureau, 2014). A Chi-squared test examining representation of White and minority women on *Time Magazine* covers versus the population revealed a significant difference $\chi^2(1) = 21.01, p < .01$, such that minority women were underrepresented and White women were overrepresented on the covers relative to their actual proportion in the United States. Similarly, minority women were underrepresented and minority men overrepresented relative to their population proportions in the United States, $\chi^2(1) = 85.80, p < .01$. See supplementary material for analyses of how representations of these different groups change over time.



Figures 2 (left) and 3 (right). Representation of White women compared to minority women (left) and minority women compared to minority men (right) on both the covers of *Time Magazine* and relative to U.S. demographic proportions (left).

Discussion

Study 1 provides empirical support for the invisibility outcome of Purdie-Vaughns' and Eibach's (2008) intersectional invisibility theory. Specifically, this study supports the idea that individuals who are not prototypical of their various constituent identities, such as women of color on *Time Magazine* covers, are rendered invisible or are not fully recognized within society. In the case of Study 1, this invisibility specifically manifests as underrepresentation of racial minority women on magazine covers relative to both prototypical marginalized group members (White women, racial minority men) and relative to racial minority women's proportion within the population of the United States. Both White women and racial minority men were overrepresented on the covers relative to their actual proportions in the population of the United States, whereas racial minority women were underrepresented on the covers relative to their proportion within the population.

While Study 1 is an important first step in demonstrating a context in which the invisibility outcome occurs, this study does not provide insight into how or why prototypical and non-prototypical groups differ from each other. Study 2 addresses this by exploring the underlying beliefs that affect perceptions of prototypical and non-prototypical group members.

Study 2

Having empirically demonstrated an instance where non-prototypical individuals were invisible in Study 1 (as well as in other work replicating the same pattern of findings using covers of the *Advocate*, an American LGBT-interest magazine; Lei, Mohr, Rucker & Richeson, *under review*), I conducted Study 2 to examine how explicit beliefs about the basic categorizations of different group members would influence perceptions and consequences of prototypicality.

Most previous psychological research explores differences in perceptions between groups (e.g., perceptions of Black versus White individuals) without considering how stereotypes may differ within groups (e.g., perceptions of Black men versus Black women). In order to better understand perceptions of non-prototypical individuals, it is important to explore how beliefs about prototypical and non-prototypical individuals from the same marginalized group differ. Specifically, understanding these differences allows for a more complete understanding of stereotyping, prejudice, and discrimination towards individuals from that group.

Perceivers quickly and automatically make judgments of others on the basis of race, age, and gender, and these judgments are central to how others are perceived (Allport 1954; Bem, 1981; Brewer, 1988; Divine, 1989; Fiske & Neuberg; Stangor, Lynch, Duan, & Glass, 1992; Zarate & Smith, 1990). In particular, judgments about the gender of targets can dramatically alter beliefs about individuals (Carpinella, Chen, Hamilton, & Johnson, 2015; Freeman, Johnson, Ambady, & Rule, 2010; Lick, & Johnson, 2014; Lick, Johnson, & Riskind, 2015). Given how central

judgments of gender are to forming beliefs about others, Study 2 examines perceivers' explicit gender beliefs about prototypical and non-prototypical group members to understand the underlying beliefs that affect prototypicality.

Sample and Procedure

I recruited 262 participants from Mechanical Turk. Of those, 226 participants met inclusion criteria (were born in and socialized within the United States). Participants indicated how masculine or feminine (1= *feminine* to 7 = *masculine*) they believed each of the following six groups to be: straight Black men, straight Black women, gay White men, lesbian White women, straight White men, straight White women. Participants indicated masculinity/femininity scores for all groups, and social groups were presented in random order to avoid potential order effects.

Results

Overall, straight Black men were perceived to be the most masculine ($M=6.12, SD=0.97$), followed by straight White men ($M=5.86, SD=1.04$), White lesbian women ($M=4.84, SD=1.29$), straight Black women ($M=3.72, SD=1.32$), White gay men ($M=2.85, SD=1.24$), and straight White women ($M=2.25, SD=1.25$). I ran a multilevel model that predicted both the participant level variance (random effects) and the effects of each social group on participants' scores of masculinity and femininity (fixed effects) for each social group. All social groups (straight Black men, straight Black women, gay White men, lesbian White women, straight White men, straight White women) were statistically different from all other social groups on perceived masculinity/femininity. The results of Study 2 are displayed in Figure 4.

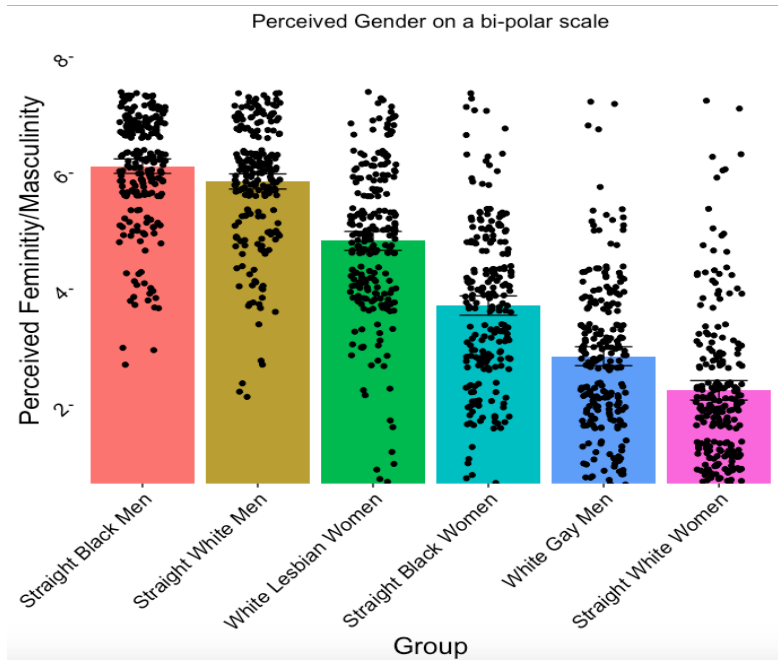


Figure 4. Reported perceived masculinity/femininity of different prototypical and non-prototypical groups.

Table 1

Scores of Perceived Masculinity-Femininity by Social Group

Straight Black men <i>Est.</i> = 6.121 <i>SE</i> =.08			
Straight White men <i>Est.</i> = 5.858 <i>SE</i> =.08	$p=.02^*$		
White lesbian women <i>Est.</i> = 4.836 <i>SE</i> =.08	$p<.01^{**}$	$p<.01^{**}$	
Straight Black women <i>Est.</i> = 3.716 <i>SE</i> =.08	$p<.01^{**}$	$p<.01^{**}$	$p<.01^{**}$

White gay men <i>Est.</i> = 2.845 <i>SE</i> =.08	<i>p</i> <.01**	<i>p</i> <.01**	<i>p</i> <.01**	<i>p</i> <.01**	
Straight White women <i>Est.</i> = 2.252 <i>SE</i> =.08	<i>p</i> <.01**	<i>p</i> <.01**	<i>p</i> <.01**	<i>p</i> <.01**	<i>p</i> <.01**
	Straight Black men	Straight White men	White Lesbian women	Straight Black women	White gay men

Discussion

Study 2 demonstrated clear differences between prototypical and non-prototypical group members on a basic dimension of person perception, gender. Specifically, the prototypical women, straight White women, were perceived as statistically more feminine than women who are non-prototypical because of their sexual orientation, White lesbian women, and women who are non-prototypical because of their race, straight Black women.

Straight White men, who are prototypical of the category men, were also perceived as statistically different in masculinity relative to their non-prototypical counterparts, White gay men and straight Black men. However, straight Black men were considered more masculine than the prototypical male group, straight White men. While on the surface it may seem puzzling that a group of non-prototypical men would be perceived as more masculine than the prototypical group, this difference aligns with cognitive work on prototypes showing that prototypes are determined by the most common concept (White men's level of masculinity as the norm) rather than the most extreme case of an attribute (Black men's higher level of

masculinity) (Mervis & Rosch, 1981; Rosch, 1975; Rosch & Mervis, 1975).

Additionally, this finding that Black men are perceived as hyper-masculine relative to the norm is consistent with other work on implicit associations between Blackness and masculinity (Galinsky, Hall, & Cuddy, 2013; Hall, Galinsky, & Phillips, 2013; Johnson, Freeman, & Pauker, 2012; Schug, Alt, & Klauer, 2015; Schug, Alt, Lu, Gosin, & Fay, 2017).

Study 2 establishes that stereotypes differ within marginalized groups (e.g., perceptions of Black men versus Black women) but does not examine how these different stereotypes affect judgments about these groups. Study 3 addresses this by exploring how perceivers differentially make attributions about prototypical and non-prototypical marginalized group members.

Study 3

Study 1 demonstrated an instance in which non-prototypical marginalized individuals are invisible relative to their prototypical marginalized counterparts. However, theoretical work by Purdie-Vaughns & Eibach (2008) suggests that there are instances in which non-prototypical marginalized individuals should be more visible than their prototypical counterparts. Past research shows that when individuals encounter a direct comparison between groups, they are more likely to attribute the difference between the groups to the non-normative group members (Miller, Taylor, & Buck, 1991) because comparing a non-prototypical group to a prototypical group highlights that the non-prototypical group is different. Additionally, in comparing groups, perceivers should be motivated to notice and address the system maintaining ideologies of androcentrism, heterocentrism, and

ethnocentrism that define these groups as non-prototypical to begin with. However, this previous work (Miller, Taylor, & Buck, 1991) has only compared group members who are prototypical and non-prototypical on a single marginalized identity (for example, comparing men to women).

Study 3 thus presented participants with a scenario that asked them to explain a difference between two groups—one that was prototypical of a marginalized group and one that was non-prototypical of a marginalized group (had multiple marginalized identities)—to explore whether perceivers are more likely to attribute between group differences to non-prototypical race, gender, and sexual orientation groups. If, as predicted, comparing prototypical and non-prototypical groups leads to hypervisibility of the non-prototypical group, I would expect to find that participants would focus their explanations on the non-prototypical comparison group (e.g., Black women, lesbian women), rather than the prototypical comparison group (e.g., Black men or straight White women).

Sample and Procedure

I recruited 265 participants from Amazon's Mechanical Turk. Of those, I excluded 29 participants because they failed at least one of two attention checks, reported not being a native English speaker, or did not reside in the United States. This left 236 participants total. Participants were on average 35 years old ($SD = 11.79$) and the majority of participants were White (White American = 85%, Asian/Asian American=3%, African American= 7% Latino American = 5%). Participants reported sexual orientations in line with nationally representative estimates for the United States (91% heterosexual, 8% non-heterosexual), and

skewed more liberal (14% reported being “Very Liberal,” 31% “Moderately Liberal,” 25% “Moderate,” 13% “Moderately Conservative,” 5% “Very Conservative,” and 12% did not respond). These demographics were generally in line with those of other MTurk samples (Buhrmester, Kwang, & Gosling, 2011; Paolacci, Chandler, & Ipeirotis, 2010; Mason & Suri, 2012).

In a paradigm adapted from Miller, Taylor, and Buck (1991), participants were randomly assigned to one of four conditions and were presented with ostensible differences in voting patterns between prototypical and non-prototypical group members. Specifically, this was a 4 (1. Voting difference between Black men and Black women, 2. Voting difference between White women and Black women, 3. Voting difference between straight women and lesbian women, or 4. Voting difference between gay men and lesbian women) x 2 (1. Prototypical group was presented first, 2. Prototypical group was presented second) between-subjects design.

Because two of the voting patterns compared prototypical and non-prototypical groups that assessed differences based on sexual orientation (1. lesbian women compared to straight women and 2. lesbian women compared to gay men) and two compared prototypical and non-prototypical groups that assessed differences based on race (1. Black women compared to Black men and 2. Black women compared to White women), these conditions are categorized into sexual orientation and race based judgments, respectively.

Measures

I measured attribution using two measures: 1. Qualitative explanations of group differences and 2. Likert scales measuring how much participants believed each group would change their behavior in the future.

Qualitative explanations. Participants were asked to write about why they thought the presented voting difference occurred. Two research assistants coded all responses ($\kappa = .98$). Responses were divided into “distinct ideas,” adopting a modified coding scheme based on prior work (Cousins, 1989; Kanagawa et al., 2001; Ross et al., 2005). I defined such distinct ideas as phrases that conveyed a complete idea, thought, or feeling, and were unique from other ideas within the same explanation. For example, “maybe White women vote because of their husbands or maybe Black women don’t care” would be divided into 2 thought units: 1. “maybe White women vote because of their husbands” and 2. “maybe Black women don’t care.”

Those ideas were then coded for which group participants attributed the voting difference to (if any). For example, the first statement, “maybe White women vote because of their husbands,” would be coded as attributing the voting difference to White women, whereas the second statement, “maybe Black women don’t care,” would be coded as attributing the voting difference to Black women. Social groups were only coded if the race and gender or gender and sexual orientation of the group was identifiable in the explanation. As a result, only the following social groups were coded: lesbian women, Black women, straight women, White woman, Black men, gay men, White men, and straight men. Any distinct thoughts that did not attribute the voting difference to any of these social groups were not coded.

Attributions about prototypical and non-prototypical groups were therefore measured by the number of participants' distinct ideas in their written responses that focused on one group relative to the other in their explanation of the voting difference.

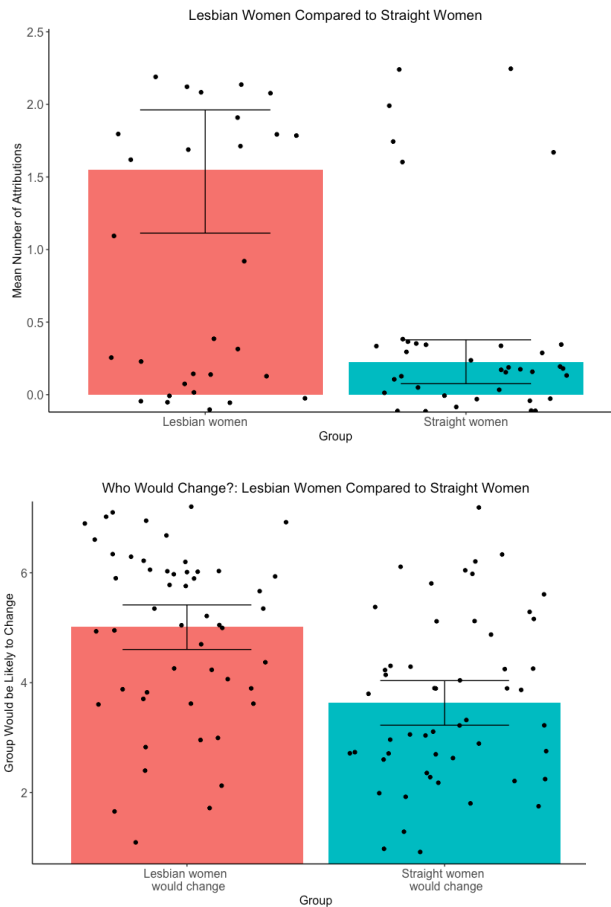
Likert scales measuring future group change. Participants also indicated the extent to which they believed each of the two groups that were compared would be more likely to change to resemble the voting patterns of the other group on a 1 (*strongly disagree*) to 7 (*strongly agree*) Likert scale.

Results

This analysis excluded participants who did not focus any of their written explanation on any social group, excluding 45 participants total in this dataset. The four different comparison types were dummy coded and entered into a linear regression model along with order to predict how often participants attributed voting differences to the non-prototypical marginalized group members. There was a significant effect of condition, $F(3, 230) = 2.653, p = .05$, meaning that participants attributed voting differences to the various social groups differently depending on comparison condition. There was no significant effect of order, $F(1, 230) = 0.129, p = .720$, indicating that participants were not more likely to attribute voting differences to the group presented first.

Sexual orientation comparisons. When participants were asked to compare prototypical women (straight women) with non-prototypical women (lesbian women), participants' written responses focused significantly more on the non-prototypical group (lesbian women, $M = 1.547, SD = 1.612$) relative to the

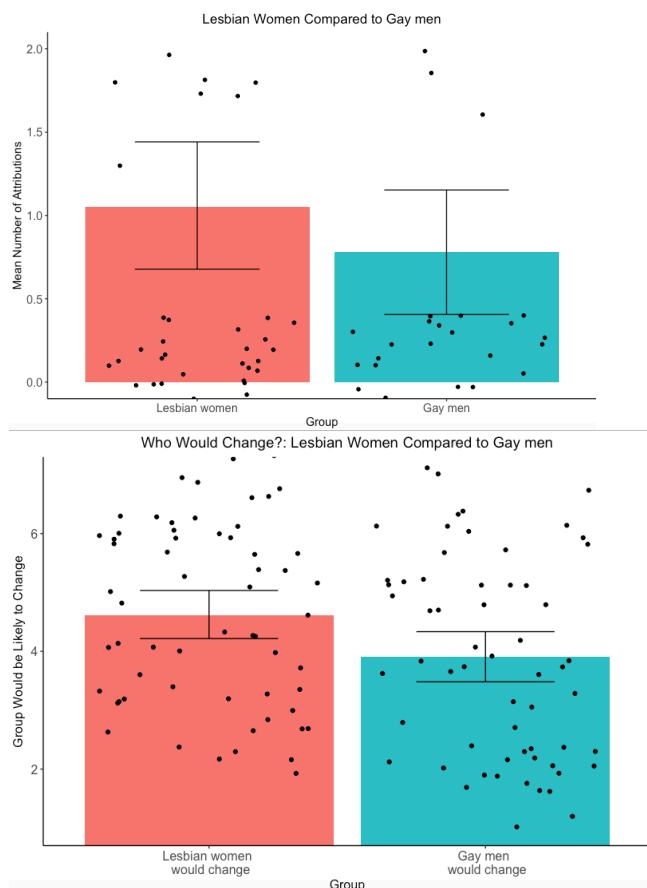
prototypical group (straight women, $M=0.226$, $SD=0.640$), $t(52) = 6.000$, $p < .01$ (figure 5). Additionally, participants' self-reported change scores indicated that the non-prototypical group (lesbian women, $M= 5.019$, $SD= 1.562$) would be more likely to change relative to the prototypical group (straight women, $M= 3.642$, $SD= 1.482$), $t(52)= 4.702$, $p < .01$. In short, when considering the context of women, lesbian women were hypervisible relative to straight women, as measured by both focus of explanation and change scores (figures 5 and 6).



Figures 5 & 6. Number of ideas in participants' written statements that focused on a non-prototypical group (lesbian women) relative to a prototypical group (straight women) differing by sexual orientation (top). Scores indicating which group (lesbian


women or straight women) participants believed would be more likely to change in the future (bottom).

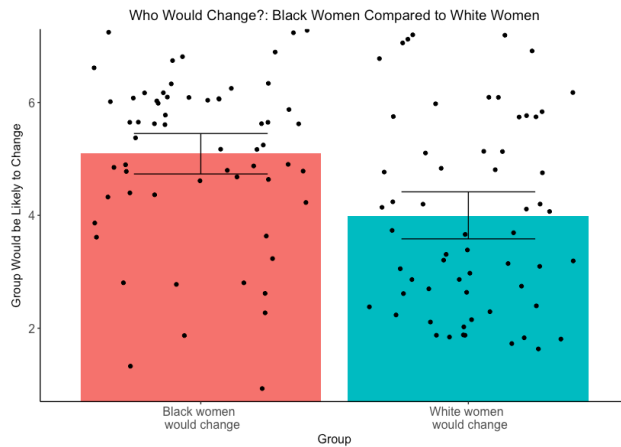
In contrast, in the context of LGBTQ individuals, when comparing prototypical LGBTQ individuals (gay men, $M=0.780$, $SD=1.52$), with non-prototypical LGBTQ individuals (lesbian women, $M=1.051$, $SD=1.634$), participants were not more likely to focus their responses on the non-prototypical group, attributing voting differences to lesbian women 58% of the time, $t(58) = .938$, $p = .35$. However, participants were marginally more likely to indicate that the non-prototypical group (lesbian women, $M=4.617$, $SD=1.574$) would be more likely to change than the prototypical group (gay men, $M=3.9$, $SD=1.714$), $t(59) = 1.862$, $p = .06$ (figures 7 and 8).



Figures 7 & 8. Number of ideas in participants' written statements that focused on a non-prototypical group (lesbian women) relative to a prototypical LGBTQ individuals (gay men) differing by gender (top). Scores indicating which group (lesbian women or gay men) participants believed would be more likely to change in the future (bottom).

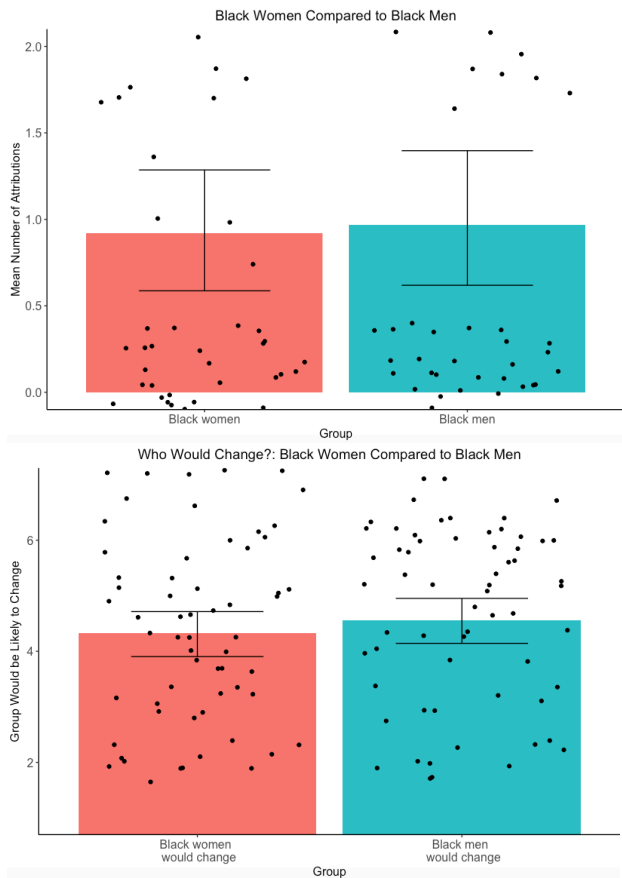
Race comparisons. When looking at the race comparisons, hypervisibility emerged when comparing prototypical women (White women) with non-prototypical women (Black women), figures 9 and 10. Participants' written responses focused statistically significantly more on the non-prototypical group (Black women, $M=1.33$, $SD=1.79$) compared to the prototypical group (White women, $M=0.583$, $SD=1.279$), $t(59) = 2.873$, $p < .01$. Additionally, participants believed that the non-prototypical group (Black women, $M=5.1$, $SD=1.469$) would be more likely than the prototypical group (White women, $M= 3.983$, $SD= 1.702$) to change, $t(59)=3.07$, $p < .01$. In short, when considering the context of women, Black women were hypervisible compared to White women.





Figures 9 & 10. Number of ideas in participants' written statements that focused on a non-prototypical group (Black women) relative to a prototypical group (White women) differing by race (top). Scores indicating which group (Black women or White women) participants believed would be more likely to change in the future (bottom).

However, when considering the context of Black people, participants' written responses did not focus more on the non-prototypical group (Black women, $M=0.921$, $SD=1.484$), relative to the prototypical group (Black men, $M=0.968$, $SD=1.502$), $t(62) = 0.197$, $p = .84$. There was also no statistically significant difference between whether participants thought the non-prototypical group (Black women, $M= 4.333$, $SD= 1.732$) or the prototypical group (Black men, $M= 4.555$, $SD= 1.643$) would be more likely to change, $t(62) = .576$, $p = .57$, figures 11 and 12.



Figures 11 & 12. Number of ideas in participants' written statements that focused on a non-prototypical group (Black women) relative to a prototypical group (Black men) differing by gender (top). Scores indicating which group (Black women or Black men) participants believed would be more likely to change in the future (bottom).

Discussion of Study 3

Generally, the results from the Study 3's attribution task support the hypervisibility hypothesis whereby when comparing prototypical and non-prototypical groups, participants are more likely to focus their qualitative explanation on the non-prototypical group, and they more highly endorse the idea that the non-prototypical group will change their behavior to look like the

comparison group in the future. However, participants attributed the difference to the non-prototypical group with some comparisons but not others. Specifically, I found hypervisibility of non-prototypical individuals when comparing lesbian women to straight women and Black women to White women. However, participants were just as likely to attribute differences to lesbian women as they were to gay men, and similarly, equally likely to attribute differences to Black women and Black men.

This lack of difference in attributions for two of these comparisons is worth noting. One reason for this pattern of results could be that these groups were not viewed as particularly different from each other in the context of voting. There could be several reasons for this lack of difference. For example, this could be due to the ethnic-prominence hypothesis (Levin et al., 2002), which proposes that race is a more influential social category than gender, and therefore women of color would be viewed similarly to men of color. As a result, participants may have viewed Black men and Black women similarly because they shared the same racial identity and therefore did not overwhelmingly attribute group differences to the non-prototypical marginalized group. Similarly, sexual orientation may be viewed as a more important social identity than gender in the comparison where participants were asked to compare lesbian women and gay men.

Additionally, the two comparisons in which participants did not differentially attribute group differences to the non-prototypical group were both comparisons in which the groups differed by gender. Participants may have been more motivated to address social differences based on race and sexual orientation (as was the case in

the comparisons in which participants attributed differences to the non-prototypical group) rather than gender (i.e. the two comparisons where participants attributed behavior to the prototypical and non-prototypical groups equally).

Although the mechanisms underlying why the hypervisibility outcome only occurred for certain comparisons between prototypical and non-prototypical marginalized groups is left for future work, this study established that the ideologies shaping non-prototypical and prototypical marginalized group membership matters for perceivers' judgments and attributions of groups.

Study 4

Study 3 explored how perceivers make attributions about prototypical and non-prototypical race, gender, and sexual orientation groups, and established when perceivers differentially attribute behaviors to prototypical and non-prototypical marginalized group members. Studies 4 and 5 sought to understand the implications of judging non-prototypical marginalized group members as different from their prototypical peers on real world outcomes. Specifically, one instance in which being judged as a non-prototypical member of a marginalized group could have serious implications is in how people think about and detect discrimination. Thus, Studies 4 and 5 explore the real-world behavioral consequences of non-prototypicality by using a hiring paradigm that mimics real-world hiring decisions to understand how detection of discrimination against individuals with marginalized identities may differ depending on whether the target is prototypical or non-prototypical of the marginalized group.

Previous research on discrimination detection has established that perceivers tend to appraise discrimination based on the race and gender of the target more than other social identities (Weiner, Perry, & Magnusson, 1988; Marti, Bobier & Baron, 2000). When participants are asked to imagine incidents of these salient forms of discrimination, they overwhelmingly imagine a Black man when thinking of instances of racism (Goff & Kahn, 2014). Moreover, when people are asked to think about instances of sexism, they overwhelmingly picture a White woman as the target (Goff & Kahn, 2013). This work suggests that individuals who are prototypical of a marginalized identity group, such as Black men and White women, are also the most accessible when thinking of targets of discrimination.

In addition to thinking of prototypical individuals as the likely targets of discrimination, people also tend to think about discrimination based on singular marginalized identities rather than considering discrimination on the basis of compound identities (Crenshaw, 1989). For example, people tend to envision discrimination as targeting a specific single identity, such as racism *or* sexism, as opposed to targeting multiple identities, such as discrimination based on both race *and* gender. Indeed, within the United States legal system, plaintiffs with multiple marginalized identities (i.e., non-prototypical marginalized group members like Black women or lesbian women) who sue for discrimination must chose a single identity on which to build a case (Crenshaw, 1989; Purdie-Vaughns & Eibach, 2008). For example, in 1976, five Black women sued their employer on the basis of discrimination against Black women. Instead of being allowed to argue that they experienced compound discrimination as Black women (the non-prototypical

marginalized identity) their case was instead decided using discrimination against White women (the prototypical marginalized group) as the “historical base” for determining discrimination (Crenshaw, 1989). Additionally, previous research tends to use these prototypical members of marginalized groups to understand how perceivers detect sexism and racism, respectively, excluding non-prototypical marginalized group members (Goff & Kahn, 2013; Sesko & Biernat, 2010).

Study 4 therefore tested the hypothesis that perceivers would detect racism against prototypical marginalized targets (i.e., Black men) differently than non-prototypical marginalized targets (i.e., Black women). As described previously, perceivers’ detection of discrimination of non-prototypical marginalized targets could differ from detection of prototypical marginalized targets in one of two ways. The *invisibility outcome* would occur if perceivers do not recognize an instance of discrimination when it happens to a person who is a non-prototypical member of a group (Purdie-Vaughns & Eibach, 2008). Specifically, if perceivers are actively searching for instances of racism, they may miss instances when the target of racism is not a prototypical group member. On the other hand, non-prototypical marginalized targets might be *hypervisible*—that is, they may attract attention because they are viewed as different and non-normative, relative to prototypical group members (Kahneman & Miller, 1986). Thus, when perceivers are actively searching for instances of racism, they may particularly attend to instances when the target of racism is not a prototypical group member.

Method

I studied participants' ability to detect discrimination within a hiring context. Participants viewed a series of hiring decisions that they believed were made by a technology ("tech") company over several years. In recent years, the tech industry has come under increasing scrutiny for a homogenized work culture (Chang, 2018) in which women and African Americans are underrepresented ("Diversity in high tech", 2016). The tech industry offers a timely context to study judgments of discrimination where the targets vary on both race and gender dimensions.

Participants evaluated the results of a series of 10 ostensible hiring decisions made by a hiring manager for different coding positions within a tech company. For each position, they were presented with information on the top two finalists. One finalist was always objectively superior to the other finalist, as indicated by their background experience and scores on different relevant skills. In four of these ten decisions, the hiring manager selected the less qualified White candidate over the more qualified Black candidate (either a Black woman or a Black man). Accordingly, the hiring manager engaged in racial discrimination during the 10 hiring decisions over which he ostensibly presided.

Participants

Fifty-nine participants were recruited from the participant pool of a private university in the Northeast and participated in exchange for course credit. Participant enrollment numbers were determined based on previous discrimination detection work (Marti, Bobier, & Baron, 2000) and on the availability of participants through the university's participant pool. Only participants who grew up in the United States qualified for the study, leaving 51 participants (30 women, 19 men, 2

other, $M_{age}=21.16$, $SD=4.27$). Selection criteria of participants were consistent with previous work on discrimination detection, using only participants socialized within the U.S. and who are therefore likely to be familiar with stereotypes in the U.S. (Cuddy et al., 2011).

Design

Study 4 was a four-condition within-subjects design, where each participant was exposed to all conditions. The discrimination conditions were: 1. Black woman not hired compared to a White woman hired, 2. Black woman not hired compared to a White man hired, 3. Black man not hired compared to a White man hired, 4. Black man not hired compared to a White woman hired. Thus, I manipulated the gender of the Black and White candidates across conditions, but in all four conditions, the hiring manager selected the inferior White candidate over the superior Black candidate.

These four discriminatory hiring decisions were among 10 total decisions that the participants viewed. Participants also examined six additional hiring decisions that the hiring manager made. These decisions did not include instances of discrimination. The candidates were White men and Asian American men and both finalists were always equally qualified for the position in all six decisions. Thus, participants evaluated a total of 10 hiring decisions in which four were our conditions of interest in which discrimination occurred against a Black (male or female) finalist and six filler conditions in which discrimination did not occur and Black finalists were not included. The order in which participants viewed the hiring decisions were randomized.

Materials

Hiring evaluation system. The hiring evaluation system was used to embed our experimental manipulations. The task required participants to review the numeric ratings for each pair of finalists for each position (i.e., 10 positions) and review the hiring manager's final decision (hire or not). Participants received a binder with all 10 hiring comparisons and the manager's hiring decisions. Each finalist was presented on a single page. Participants always reviewed the pair of finalists together so they could compare the ratings of one candidate to the other.

Photographs. Photographs were used to depict the ethnicity and gender of the finalist. There were 2 Black women, 2 Black men, 2 White women, 4 White men, and 8 Asian men. Photographs were taken from the Chicago Face Database (Ma, Correll & Wittenbrink, 2015), which is normed, includes a diverse range of photos, and is widely used in social science research. In order to rule out the effects of attractiveness on discrimination detection, all pictures of finalists were matched with other finalists who were equivalently attractive, and perceived as unequivocally White or Black, and male or female (Ma, Correll & Wittenbrink, 2015), depending on the comparison.

Quantitative summary. Each finalist's quantitative summary included the name of the hiring manager whose name was pre-tested to be perceived as a White male, a photograph of the applicant (allowing us to manipulate the race and gender of each applicant) ostensibly taken by security prior to their interview, and a list of the applicant's qualifications and work experience. The quantitative evaluation summary included numeric scores indicating how qualified the finalist was on each

skill where scores of 7 indicated that a job applicant was extremely qualified in that skill and a score of 1 indicated that a candidate was very weak in that skill. Finally, the summary included the hiring manager's decision of which of the two applicants he hired and notes from the hiring manager about the applicant (see Appendix A for sample hiring task materials).

Participants were instructed that there was no time limit to reviewing the decisions, but were instructed that once they turned the page to assess the quality of the hiring manager's decision for the pairing, they should not return and look back at the prior pairings.

Operationalization of discriminatory hiring decisions. Job applicant finalists who were targets of discrimination had higher scores on their quantitative evaluation summary than the comparison finalist who was hired. Pilot testing by a separate group of participants confirmed that each individual who was the target of discrimination was rated as more qualified for the position than the finalist who was hired¹. Additionally, to bolster the believability of the decision, the hiring manager's notes included comments suggesting that his rationale for his hiring choice was ambiguous (e.g., "not sure she fits our image").

Measures

Survey questions. I adapted a measure from Marti, Bobier, & Baron's work on non-prototypical forms of prejudice (2000) in order to assess how accurately participants were able to detect discrimination. This measure consisted of three

¹ Pilot tests indicated that discriminated against targets were viewed an average of 3.5 points higher ($SD=1.2$) on a 1-7 scale where 1 indicates *no difference* and 7 indicates *much more qualified*.

constructs: perceived fairness of the decision, perceptions of deliberate action on the part of the hiring manager, and negative affect experienced by the perceiver. All items except for the question measuring negative affect were rated on a Likert scale ranging from 1 (*Completely Disagree*) to 6 (*Completely Agree*). Negative affect was rated from 1 (*Small Concern*) to 6 (*Serious Concern*). Perceived fairness included two items “I believe that the hiring manager made the best hiring decision given the specific details of the job information” and “the hiring decision seemed fair.” Perceptions of the deliberate action on the part of the hiring manager and negative affect experienced by the perceiver were single items: “The hiring manager stuck to the hiring guidelines provided by the company, such as not discriminating based on social group membership such as race/ethnicity” and “If you listed any concerns, how serious are these concerns?” respectively. Four filler items were interspersed between discrimination detection items (e.g. “The hired candidate has met the required educational requirement for the job”). All filler items were collapsed and used to assess baseline attitudes about the hiring decision (see supplementary materials for scale construction). After viewing the summaries of both finalists for a position, participants completed the discrimination detection measure to assess their attitudes towards the manager’s hiring decision. Accordingly, they completed this measures 10 times.

Qualitative measures of discrimination. Participants were asked to write about any concerns they had about the hiring decision. Coding about concerns about discrimination took place in two stages. First, two research assistants independently coded each written response for whether or not participants mentioned any form of

discrimination ($\kappa=.89$). In the second stage of coding, in order to understand how participants differentially viewed the discrimination in the different conditions, two additional independent coders coded the content of each written response for the type of discrimination mentioned (e.g. racism, sexism) and negative affect expressed by the participant about the decision. See Appendix B for discrimination type coding information and the supplementary material for qualitative coding results.

Surprise recall task. Because attention (or lack thereof) is associated with both invisibility and hypervisibility (Sesko & Biernat, 2010) and people attend to negative social information such as discrimination (Pratto, & John, 1991), we measured which finalist participants attended to in the hiring evaluation task. After completing the hiring evaluation task, participants were presented with 16 photographs in random order that depicted individuals who were either previously seen in the hiring evaluation task or were novel. Participants were presented with each face one at a time on a computer screen and pressed a button to identify whether they previously saw each face. Participants completed the task at their own pace. Participants viewed: 4 Black women (2 previously seen in the hiring summaries, 2 novel Black women), 4 Black men (2 previously seen in the hiring summaries, 2 novel), 4 White men (2 previously seen in the hiring summaries, 2 novel), and 4 White women (2 previously seen in the hiring summaries, 2 novel).² The novel pictures were matched to the previously seen pictures in attractiveness (Ma, Correll & Wittenbrink, 2015).

² Additionally, participants also viewed 4 previously seen Asian male faces taken from the filler hiring cycles, however, there were no matched novel Asian faces. As a result, correct recognition of these faces was not explored.

Social dominance orientation (SDO). Social Dominance Orientation is associated with favoritism for high status groups, and therefore may serve as an individual level difference for discrimination detection (Federico, 1999; Levin, Federico, Sidanius & Rabinowitz, 2002). Social Dominance Orientation was assessed with a shortened, 4-item version (Pratto et al., 2013) of the original scale (Pratto et al., 1994; e.g., “Inferior groups should stay in their place”) $\alpha = .8$; 1 = (*very negative*) to 7 (*very positive*), $M = 1.45$, $SD = .63$.

Political orientation. Political orientation is associated with ideologies that predict discrimination detection including holding an egalitarian worldview and support for minority groups (Jost, 2006; Jost, Nosek & Gosling, 2008). Participants’ political orientation was measured using a 5 point scale where 1 indicated *strongly liberal* and 5 indicated *strongly conservative*, $M = 2.14$, $SD = .825$.

Gender centrality. Previous work in discrimination detection has found that the strength of participants’ gender identity moderates perceptions of discrimination (Inman & Baron, 1996). Participants indicated how central their gender identity was to their sense of self by answering “How important is your gender identity to you?” from 1 (*not at all important*) to 7 (*extremely important*), $M = 4.92$, $SD = 1.89$.

Race centrality. Because previous work has found that the strength of perceivers’ racial identity moderates perceptions of discrimination (Inman & Baron, 1996), we asked participants to indicate how central their race/ethnic identity was to their sense of self by answering “How important is your race/ethnic identity to you?” from 1 (*not at all important*) to 7 (*extremely important*), $M = 4.78$, $SD = 1.57$.

Table 2

Correlations Between Individual Difference Measures

	Gender Importance	Race Importance	Socio-Economic Status	Political Orientation	Social Dominance Orientation
Gender Importance	1.000	0.117	-0.281	-0.104	-0.170
Race Importance	0.117	1.000	-0.416	0.065	-0.178
Socio-Economic Status	-0.281	-0.416	1.000	0.180	0.206
Political Orientation	-0.104	0.065	0.180	1.000	0.292
Social Dominance Orientation	-0.170	-0.178	0.206	0.292	1.000

Procedure

Participants arrived individually to the lab, indicated consent, and were told that the study investigated how managers in the tech industry make hiring decisions based on a new method of assessing job candidates. The purpose of this “novel” hiring system was to allow managers to quickly and efficiently highlight the strengths and weaknesses of each candidate, thus increasing the efficiency and speed with which they make hiring decision. To bolster the cover story, participants received information about the ostensible tech company including marketing information, a job description, and details about the credibility of the rating system (See Appendix C for materials). In order to better understand how this new system affects hiring decisions, participants were led to believe that they would review the results of actual hiring materials from a tech company.

Participants were told that prior to being interviewed, the human resources department rated each candidate's credentials (e.g., coding skills, leadership skills, level of education) on a Likert-type scale from 1-7. The hiring manager assessed the qualities of the candidates and made a decision to hire based on this numeric system and an interview with the candidate.

Next, participants completed the hiring evaluation task. They received a binder with all 10 hiring comparisons and the manager's hiring decisions. The hiring evaluation task required participants to review the numeric rating hiring summary for each pair of finalists for each position and review the hiring manager's final decision. Each finalist was presented on a single page. Participants always reviewed the pair of finalists so they could compare the ratings of one candidate to the other.

After reviewing the results of the hiring managers' decision, participants completed the assessment surveys at their own pace. After the hiring evaluation task, participants completed a measure of Social Dominance Orientation (Pratto, 1994), a surprise recall task, and demographic information. Participants were then debriefed and told about the fabricated nature of the hiring materials as well as the true purpose of the study.

Results

Given that our data was within-subjects and consisted of repeated measures for each individual, we used multilevel modeling to estimate effects. Multilevel modelling predicts participant level self-report measures (random effects) and the overall effects of the average hiring comparison differences on self-report measures (fixed effects). Data were explored using a variety of models, but the slope and

intercept were allowed to vary randomly by participant in all models. Models were fit using 'lme4' and 'lmerTest' packages in R (Bates, Maechler, Bolker, & Walker, 2014; Kuznetsova, A., Brockhoff, P.B., & Christensen R.H.B., 2017).

Because of the within-subjects nature of the study design, participants evaluated every hiring comparison, resulting in completely equal cells. As a result, for each analysis, standard errors for conditions are the same because of the balanced nature of the data. Additionally, all graphs show 95% confidence intervals and individual data points.

Discrimination conditions and non-discriminatory decisions. Confirming the validity of the discrimination manipulation, all four discriminatory hiring comparisons were always viewed as statistically more problematic (less fair, higher in deliberately discriminatory action, and higher negative affect) than the non-discriminatory comparisons. Because this pattern was consistent for all self-report questions and participants' evaluation of the non-discriminatory decisions does not add to our understanding of how perceivers detect discrimination, the relationship of the non-discriminatory decisions relative to the discriminatory hiring decisions is not discussed in detail further.

Qualitative coding for presence or absence of discrimination detection.

In the Black woman discriminated against/White woman hired decision, a statistically greater proportion of participants wrote about discrimination than not, $t(47)= 2.42, p=.02$. Participants were also marginally less likely to mention discrimination in the Black man discriminated against/White woman hired decision than to mention it, $t(49)=1.73, p=.08$. There was no statistically significant

difference in mentions of discrimination for the other two comparisons (Black woman/White man, $t(49)=.56, p=.58$; Black man/White man, $t(48) = 1.29, p=.20$).

See table 3 for the percentage of participants who mentioned discrimination by condition.

Table 3

Percentage of Participants Who Wrote about Discrimination by Condition

	Black woman/ White woman	Black woman/ White man	Black man/ White woman	Black man/ White man
Percentage of Participants Mentioning Discrimination	66.66%	54%	38%	59.18%

Table 4

Log Odds that a Participant Wrote About Discrimination by Condition

Black woman/ White woman <i>Log odds Est.=1.257</i> <i>SE= 0.587</i>			
Black woman/ White man <i>Log odds Est.= 0.267</i> <i>SE= 0.545</i>	$p=.099.$		
Black man/ White woman <i>Log odds Est. =-1.020</i> <i>SE= 0.565</i>	$p<.001^{**}$	$p=.030^*$	
Black man/ White man <i>Log odds Est.= 0.641</i> <i>SE= 0.556</i>	$p= 0.304$	$p=0.512$	$p<.001^{**}$
	Black woman/ White woman	Black woman/ White man	Black man/ White woman

Additionally, using a logistic multi-level model to compare the different discriminatory decisions to each other revealed that the Black man discriminated against/White woman hired decision had significantly fewer mentions of discrimination compared to the other discriminatory conditions: Black woman/White woman, $p < .001$; Black woman/White man, $p = .030$; Black man/White man, $p < .001$.

Participants were marginally more likely to mention discrimination in the Black woman/White woman hiring decision than in the Black woman/White man decision ($p = .099$). However, the number of mentions of discrimination in the Black woman/White woman decision was not statistically different than in the Black man/White man comparison. See table 4 for all log odds estimates and comparisons between conditions.

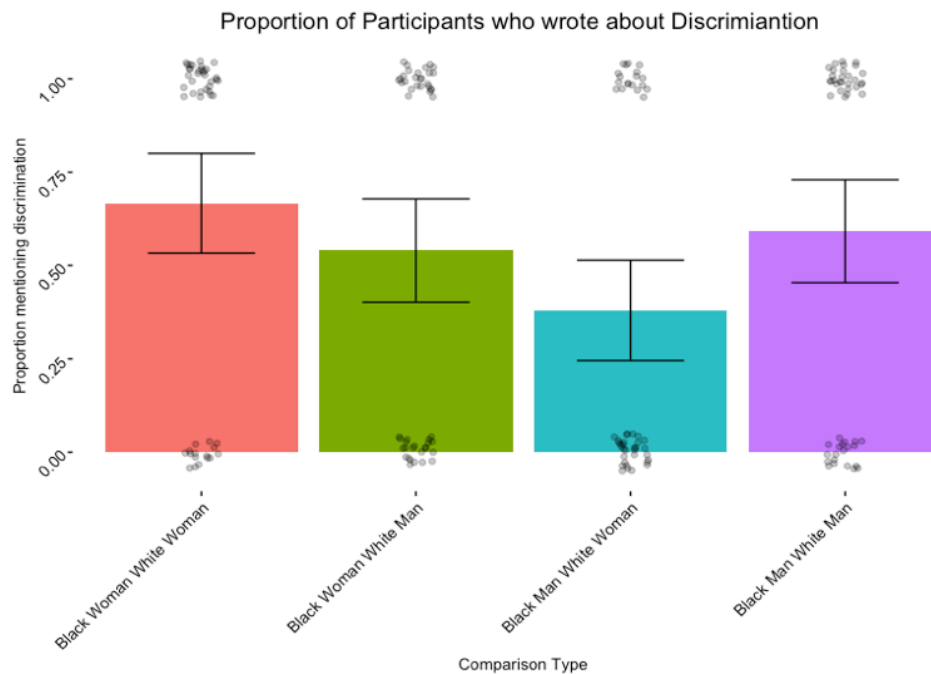


Figure 13. Percentage of written responses mentioning discrimination, by discriminatory comparison.

Self-report measures of discrimination detection.

Participants evaluated 8 Likert scale questions after viewing each hiring decision. These questions measured three precedents for discrimination detection: 1. Perceived fairness of decisions 2. Perceptions of deliberate action on the part of the hiring manager, and 3. The negative affect experienced by the perceiver. See supplementary materials for analysis of all Likert scale questions and factor analyses.

Because this study had a within-subjects design and we had multiple data points for each participant, we used multilevel modelling to test our hypotheses that self-reports of discrimination would differ by discrimination comparison (the race and gender of the target of discrimination compared to the race and gender of the less qualified candidate who was instead hired).

Perceived fairness of decision.

Table 5

Scores of Perceived Fairness by Hiring Comparison

Black woman/ White woman <i>Est.</i> = 1.902 <i>SE</i> = .1504		
Black woman/ White man <i>Est.</i> = 2.177 <i>SE</i> = .1504	<i>p</i> = .069 .	
Black man/ White woman <i>Est.</i> = 2.519	<i>p</i> < .01**	<i>p</i> = .024*

<i>SE</i> =.1504				
Black man/ White man <i>Est.</i> = 2.284 <i>SE</i> =.1504	<i>p</i> =.012*	<i>p</i> =.474	<i>p</i> =.119	
Non-Discriminatory Decisions <i>Est.</i> = 3.857 <i>SE</i> =.1504	<i>p</i> <.01**	<i>p</i> <.01**	<i>p</i> <.01**	<i>p</i> <.01**
	Black woman/ White woman	Black woman/ White man	Black man/ White woman	Black man/ White man

To test how participant's perceptions of fairness differed by hiring decision, we fit a multilevel regression with hiring decision as a predictor of perceived fairness, see table 5. Because non-discriminatory comparisons were collapsed, we had 5 ratings of fairness for each participant (the four discriminatory hiring comparisons and a score for the 6 collapsed non-discriminatory decisions), leaving 255 judgments of fairness total.

Overall, the Black woman discriminated against/White woman hired decision was viewed as the least fair (*Est.*=1.902, *SE*=.150), followed by the Black woman discriminated against/White man hired decision (*Est.*=2.177, *SE*=.1504), and the Black man discriminated against/White man hired decision (*Est.*=2.284, *SE*=.1504). The Black man discriminated against/White woman hired decision was viewed as the most fair of all of the hiring decisions (*Est.*=2.519, *SE*=.1504). Additionally, the non-discriminatory decisions were viewed as significantly more fair than all of the discriminatory decisions (*Est.*=3.857, *SE*=.1504), indicating that

participants viewed all discriminatory hiring decisions to be less fair than the non-discriminatory decisions.

Participants perceived the overall fairness of the Black woman/White woman hiring decision to be statistically significantly lower than both of the decisions in which Black men were discriminated against (Black man/White woman, $p < .01$; Black man/White man, $p = .012$).

The Black woman/White man decision was perceived as marginally more fair than the Black woman/White woman decision ($p = .069$) and was also perceived as significantly less fair than the Black man/White woman comparison ($p = .024$).

The two decisions where Black men were discriminated against (Black man/White woman, Black man/White man) were not statistically significantly different than each other in perceived fairness ($p = .119$).

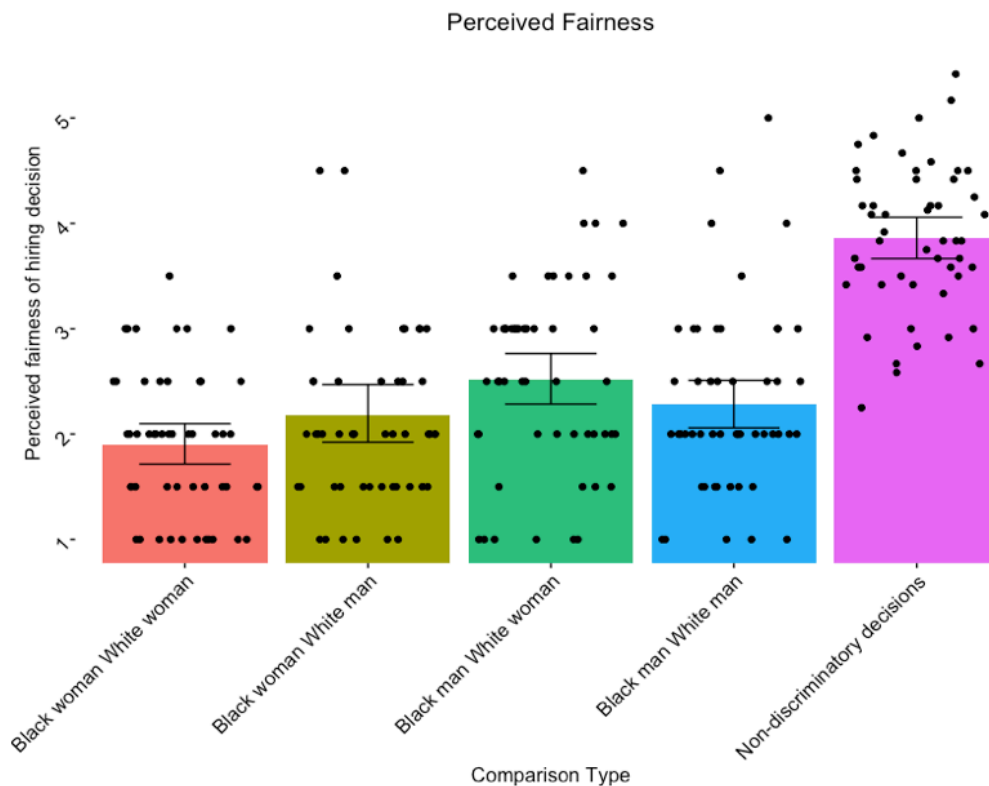


Figure 14. Perceived fairness by hiring comparison.

Perception of deliberate discriminatory action by hiring manager.

Table 6

Scores of Perceived Deliberate Action by Hiring Manager

Black woman/ White woman Est.= 5.255 SE=0.156				
Black woman/ White man Est.= 4.961 SE= 0.156	$p = .061.$			
Black man/ White woman Est.= 4.598 SE= 0.156	$p < .01^{**}$	$p = .021^*$		
Black man/ White man Est.= 4.853 SE= 0.156	$p = .011^*$	$p = .490$	$p = .104$	
Non-Discriminatory Decisions Est.= 2.905 SE= 0.156	$p < .01^{**}$	$p < .01^{**}$	$p < .01^{**}$	$p < .01^{**}$
	Black woman/ White woman	Black woman/ White man	Black man/ White woman	Black man/ White man

To understand participant's perceptions that the hiring manager acted in a deliberately discriminatory way, we fit a multilevel regression with hiring decision condition as a predictor of perceived deliberateness of discriminatory action, see table 6. Because non-discriminatory hiring decisions were collapsed, we again had 5 ratings for each participant total leaving 254 judgments of deliberate discriminatory action total (one participant left one response blank).

Following the same pattern as perceived fairness responses, participants overall found the Black woman discriminated against/White woman hired decision to be the most discriminatory ($Est.=5.255, SE=0.156$) followed by the Black woman discriminated against/White man hired decision ($Est.=4.961, SE=0.156$) and the Black man discriminated against/White man hired decision ($Est.=4.853, SE=0.156$). The Black man discriminated against/White woman hired decision was viewed as the least deliberately discriminatory of all of the hiring decisions ($Est.= 4.598, SE=0.156$). Additionally, the non-discriminatory decisions were viewed as significantly lower than all of the discriminatory decisions ($Est.=2.905, SE=0.180$), indicating that participants viewed all discriminatory hiring decisions to be more deliberately discriminatory than the non-discriminatory decisions.

Participants perceived the Black woman/White woman hiring decision to be significantly more deliberately discriminatory than both of the decisions in which Black men were discriminated against (Black man/White woman, $p<.01$; Black man/White man, $p=.011$).

The Black woman/White man decision was again perceived as marginally less deliberate than the Black woman/White woman decision ($p=.061$) and was also perceived as significantly more deliberately discriminatory than the Black man/White woman comparison ($p=.021$).

The two decisions where Black men were discriminated against (Black man/White woman, Black man/White man) were not significantly different than each other in perceived deliberate discrimination ($p=.104$).

Perceived Deliberate Action by Hiring Manager




Figure 15. Scores of perceived deliberate action by hiring manager.

Reported concern by hiring decision.

Table 7

Scores of Self-Reported Concern About Each Hiring Decision

Black woman/ White woman Est.= 5.423 SE= .180			
Black woman/ White man Est.= 5.034 SE= .180	$p = .031^*$		
Black man/ White woman Est.= 4.719 SE= .180	$p < .001^{**}$	$p = .080.$	
Black man/ White man Est.= 4.980	$p = .013^*$	$p = .742$	$p = .152$

<i>SE</i> = .180				
Non-Discriminatory Decisions <i>Est.</i> = 3.171 <i>SE</i> = .180	<i>p</i> <.001**	<i>p</i> <.001**	<i>p</i> <.001**	<i>p</i> <.001**
	Black woman/ White woman	Black woman/ White man	Black man/ White woman	Black man/ White man

Following the same pattern found in the other two self-report measures of discrimination detection, participants reported that the Black woman discriminated against/White woman hired decision was the most concerning (*Est.*=5.423, *SE*=.180) followed by the Black woman discriminated against/White man hired decision (*Est.*= 5.034, *SE*=.180) and the Black man discriminated against/White man hired decision (*Est.*=4.980, *SE*=0.180). The Black man discriminated against/White woman hired decision was viewed as the least concerning of all of the hiring decisions (*Est.*= 4.719, *SE*=.180). Additionally, the non-discriminatory decisions were viewed as significantly lower than all of the discriminatory decisions (*Est.*=3.171, *SE*=.180), indicating that participants viewed all discriminatory hiring decisions to be more concerning than the non-discriminatory decisions. See table 7 for all estimates and comparisons by condition.

Participants perceived the Black woman/White woman hiring decision to be significantly more concerning than both of the decisions in which Black men were discriminated against (Black man/White woman, *p*<.01; Black man/White man, *p*=.013). Additionally, participants also viewed the Black woman/White woman

comparison to be significantly more concerning than the Black woman/White man decision ($p=.031$).

The Black woman/White man decision was also viewed as marginally more concerning than the Black man/White woman comparison ($p=.080$).

The two decisions where Black men were discriminated against (Black man/White woman and Black man/White man) were not significantly different than each other in perceived concern ($p=.152$).

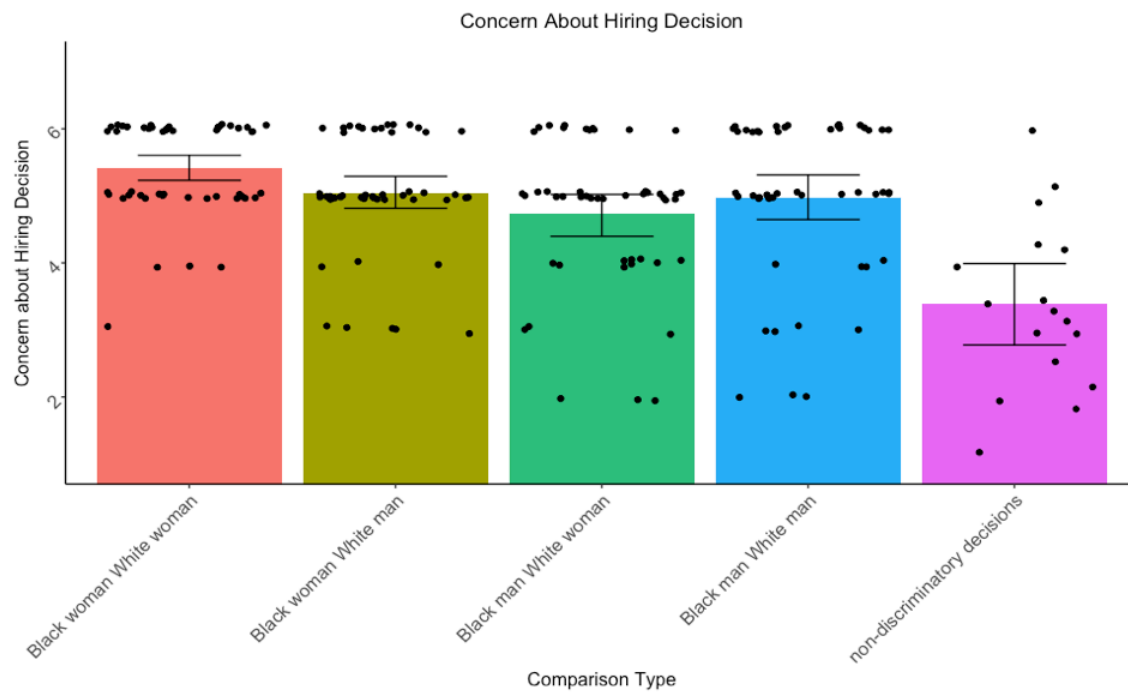
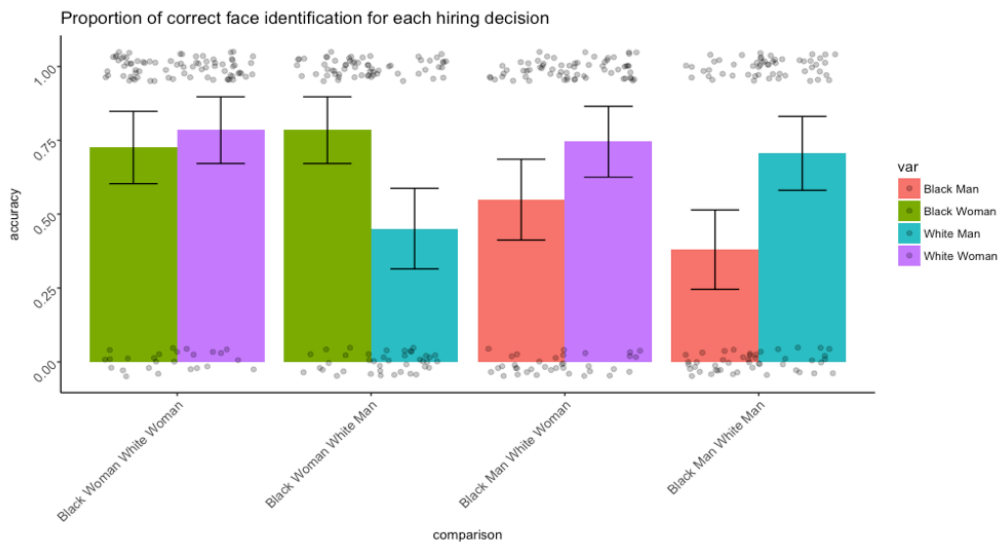


Figure 16. Self-reported concern over hiring decision.

Surprise cued recall. Participants were presented with a series of faces and asked whether or not they previously saw each face in one of the job applications. Participants' accuracy for the Black woman ($M=.726$, $SD= 0.451$)/white woman comparison ($M= 0.784$, $SD= 0.415$) was equivalent for both faces, $t(50)= 0.685$, $p=.497$, indicating that participants' attention for both applicants in this comparison

were similar. This was not the case with the other three discriminatory comparisons.³ For the two comparisons in which Black man were discriminated against (Black man $M= 0.549$, $SD= 0.503$ /White woman $M= 0.745$, $SD= 0.440$, $t(50)=-2.21$, $p=.03$; Black man $M= 0.38$, $SD= 0.490$, White man $M= 0.706$, $SD= 0.460$), $t(49)= -3.653$, $p<.01$), participants' memory for the Black male applicant was significantly lower than participants' memory for the other applicant, indicating that participants did not pay as much attention to the Black male applicant in these hiring comparisons. In the Black woman ($M= 0.784$, $SD= .415$)/White man ($M= 0.451$, $SD= 0.502$) condition, participants were significantly more accurate in recognizing the Black woman, $t(50)= 3.828$, $p<.01$, indicating that they paid greater attention to the discriminated against Black woman than the White man.



³ Importantly, these results suggest that these differences are not just due to the cross race effect, which would predict that White participants would be generally less accurate when recognizing the faces of outgroup members (in this case both Black men and women). Instead, participants were always good at recognizing Black women and White men were less likely to be recognized.

Figure 17. Percentage of correct recognition responses by discriminatory comparison.

Discussion of Study 4

Consistent with the hypervisibility outcome, the two comparisons in which Black women (non-prototypical Black people) were discriminated against were viewed as the most problematic of the discriminatory hiring decisions. For example, in the condition where a Black woman was discriminated against and a White woman (prototypical woman) was hired, participants were more likely to mention discrimination in their written assessments, more likely to accurately recall the faces of the job candidates from that hiring decision, and reported that this condition was the most problematic (least fair, most deliberately discriminatory, and most concerning). In contrast, the hiring decisions where a Black man was discriminated against (prototypical Black person) were viewed as less worrisome. These results suggest that the comparisons that involved discrimination against Black women stood out to participants, causing greater recognition of discrimination than in comparisons where the target was a prototypical marginalized group member.

Interestingly, participants' perception of discrimination not only varied as a function of who was discriminated against (Black woman or Black man), but whom the target was compared to in the hiring decision (White man or White woman). The Black woman/White man comparison was viewed as marginally less problematic than the Black woman/White woman comparison for all self-report measures, despite the fact that in both hiring decisions the target had the same non-

prototypical identity (i.e., Black woman). Similarly, while there was no statistical difference between the self report scores for the two hiring comparisons where a Black man was the target (Black man/White woman and Black man/White man), the Black man/White woman comparison was always viewed as the least problematic of the discriminatory decisions and was consistently viewed as less problematic on the self-report measures than the Black woman/White man hiring comparison (while this pattern was not the case for the other Black male target comparison).

These results mirror those of Study 3, which found that non-prototypical group members were hypervisible only in certain comparisons between prototypical and non-prototypical marginalized groups. For example, Study 3 found that Black women (non-prototypical) were hypervisible relative to White women (prototypical) when those two groups were compared. Similarly, participants in Study 4 found the discriminatory hiring decision where these two social groups were compared to be the most problematic and hypervisible (the Black woman/White woman hiring decision).

As with Study 3, there are several underlying psychological processes that could explain why the Black woman/White woman discrimination comparison is hypervisible relative to the Black woman/White man comparison and the discriminatory hiring comparisons where Black men were the targets. Participants may detect discrimination by making judgments about the perceived difference between the identity groups to which the job candidates belong, and the degree of

the perceived difference between groups could drive judgments about discrimination.

Alternatively, as with Study 3 where participants did not differentially attribute group differences to the non-prototypical group when the social groups differed only by gender, perhaps comparisons between groups that differ by gender are less visible than groups that differ by one of the other system maintaining ideologies, such as heterocentrism or ethnocentrism.

Finally, even though targets in Study 4 were discriminated against on the basis of race, Sanchez et al. (2017) found that perceivers often believe that people who hold racist beliefs tend to also hold sexist beliefs, known as stigma transfer. As such, participants may have attended to and reacted more negatively towards the Black woman/White woman comparison relative to the Black woman/White man comparison because they expected the hiring manager would discriminate in ways that would benefit both men and White people. Therefore, in the Black woman/White woman hiring decision when the hiring manager acts in a way that benefits a woman, participants may attend to this decision more (i.e., this decision is more hypervisible) relative to the Black woman/White man comparison. Similarly, this idea of stigma transfer also could explain participants' judgments of discrimination against prototypical targets (Black men). Specifically, participants who expect the hiring manager to be racist and sexist may view the Black man/White woman as less problematic because the decision benefits a marginalized group (women). Additionally, the Black man/White man comparison would be viewed as more problematic because it benefits men and White people.

Study 5

Study 4 explored how perceivers differentially detect discrimination against Black targets on the basis of race. Using the same non-prototypical marginalized group as in Study 4 (Black women), Study 5 explores how perceivers understand discrimination on the basis of gender when the target is a prototypical woman (White woman) compared to when the target is a non-prototypical woman (Black woman).

As in Study 4, I was interested in how having a non-prototypical marginalized identity would affect perceivers' ability to detect discrimination. The previous study revealed evidence of hypervisibility. In Study 5, I assessed whether hypervisibility persists when perceivers are making judgments of gender-based discrimination.

Given that people tend to think about racism and sexism in disparate and distinct ways (Crenshaw, 1989), it may be the case that perceivers would use different standards when detecting race-based rather than gender-based discrimination. Study 5 therefore expands our understanding of discrimination detection against non-prototypical marginalized group members by changing the type of discrimination to which participants are exposed.

Method

Using the same hiring paradigm that was used in Study 4, we measured participants' ability to detect gender-based discrimination within a hiring context. Because both women and Black people are underrepresented ("Diversity in high tech", 2016) in the tech industry, this remains an appropriate context within which

to study how judgments of discrimination vary when the targets of discrimination differ by gender, in addition to race.

As with Study 4, participants evaluated 10 different hiring decisions for different coding positions within a tech company. For each hiring decision, participants viewed information on the top two finalists and one finalist was always more qualified and thus superior to the other. In four of these ten decisions, the hiring manager selected the less qualified male candidate (either a Black man or a White man) over the more qualified female candidate (either a Black woman or a White woman). Thus, the hiring manager engaged in discrimination in these four hiring decisions.

Participants

Fifty-eight participants were recruited from the participant pool of a private university in the Northeast and participated in exchange for course credit. Similar to Study 4, participant enrollment numbers were decided based on previous work on discrimination detection (Inman, & Baron, 1996; Marti, Bobier, & Baron, 2000) and on the availability of participants through the university's participant pool. Because stereotypes and beliefs about discrimination are culture specific (Cuddy et al., 2011), only participants who grew up in the United States qualified for the study, leaving 52 participants (32 women, 20 men, $M_{age}=20.538$, $SD=3.744$).

Design

Study 5 employed the same methods as Study 4 with one main difference: All discriminatory hiring decisions in Study 5 were gender-based (rather than the race-based decisions of Study 4). As such, Study 5 was a four-condition within-subjects

design, where each participant was exposed to all conditions. The four discrimination conditions were: 1. Black woman not hired compared to a White man hired, 2. Black woman not hired compared to a Black man hired, 3. White woman not hired compared to a Black man hired 4. White woman not hired compared to a White man hired. Thus, I manipulated the race of the male and female candidates across conditions, but in all four conditions, the hiring manager engaged in discrimination by selecting the inferior male candidate over the superior female candidate. These four discriminatory hiring cycles were the only hiring cycles to include Black and female candidates.

As in Study 4, participants also examined six additional hiring decisions that the hiring manager made. These decisions did not include instances of discrimination. The finalists were White men and Asian American men and both finalists were equally qualified for the position. Thus, participants evaluated a total 10 hiring decisions in which four were our conditions of interest in which discrimination occurred against a female (Black or White) finalist and six filler conditions in which discrimination did not occur. Again, the order in which participants viewed the hiring decisions was randomized.

Materials

Hiring evaluation system. The hiring evaluation system used to embed our experimental manipulations was identical to the one used in Study 4. However, the photos of the candidates and names and gender pronouns of candidates were changed in the candidate review documents to reflect the sexism discrimination conditions.

Again, each finalist's quantitative summary included the name of the hiring manager, a photograph of the applicant, and a list of the applicant's qualifications and work experience. Each finalist also had a quantitative summary of numeric scores that indicated how qualified the finalist was on each skill. Finally, the summary included the hiring manager's decision of which of the two applicants he hired and notes from the hiring manager about the applicant.

Participants were instructed that there was no time limit to reviewing the decisions but that they should not reference the candidate summaries when evaluating the hiring decision.

Photographs. Again, photographs were used to depict the ethnicity and gender of the finalist. There were 2 Black women, 2 Black men, 2 White women, 4 White men, and 8 Asian men associated with the candidate documents. Again, all pictures of finalists were matched with other finalists who were equivalently attractive, and perceived as unequivocally White or Black, and male or female (Ma, Correll & Wittenbrink, 2015), depending on the comparison.

Operationalization of discriminatory hiring decisions. As in Study 4, finalists who were targets of discrimination (women) had higher scores on their quantitative evaluation summary than the comparison finalist who was hired. Additionally, the hiring manager's notes included comments suggesting that his rationale for his hiring choice was ambiguous (e.g., "not sure she fits our image").

Measures

Survey questions. Survey questions for Study 5 were identical to those of Study 4 with one exception. Rather than asking participants about perceived

deliberate action by the hiring manager on the basis of race, Study 5 asked participants about perceived deliberate action based on gender (“The hiring manager stuck to the hiring guidelines provided by the company, such as not discriminating based on social group membership such as sex/gender”).

Qualitative measures of discrimination. I used the same coding procedure from Study 4 to understand participants’ concerns about each hiring decision. Coding about concerns about discrimination took place in two stages. First, two research assistants independently coded each written response for whether or not participants mentioned any form of discrimination ($\kappa = .99$).

In the second round of coding, two additional independent coders coded the content of each written response for the type of discrimination mentioned (e.g., racism, sexism) and negative affect expressed by the participant about the decision in order to understand how participants differentially viewed the discrimination in the different conditions. See supplementary material for discrimination type coding information and qualitative coding results.

Surprise recall task. As in Study 4, participants were presented with faces that were either novel or previously seen in the hiring comparison task and were asked to identify whether or not they previously saw each face. Participants viewed: 4 black women (2 previously seen in the hiring comparisons, 2 novel black women), 4 black men (2 previously seen in the hiring comparisons, 2 novel), 4 white men (2 previously seen in the hiring comparisons, 2 novel) and 4 white women (2 previously seen in the hiring comparisons, 2 novel).

Social Dominance Orientation (SDO). Social Dominance Orientation was assessed with a shortened, 4-item version (Pratto et al., 2013) of the original scale (Pratto et al., 1994) (e.g., “Inferior groups should stay in their place” $\alpha = .8$; 1 (*very negative*) to 7 (*very positive*), $M = 1.572$, $SD = .905$).

Political Orientation. Participants’ political orientation was measured using a 5 point scale where 1 indicated *strongly liberal* and 5 indicated *strongly conservative*, $M = 2.235$, $SD = .971$.

Gender Centrality. Participants indicated how central their gender identity was to their sense of self by answering “How important is your gender identity to you?” from 1 (*not at all important*) to 7 (*extremely important*), $M = 5.116$, $SD = 1.464$.

Race Centrality. Participants indicated how central their racial identity was to their sense of self by answering “How important is your race/ethnic identity to you?” from 1 (*not at all important*) to 7 (*extremely important*), $M = 4.788$, $SD = 1.673$.

Table 8

Correlations Between Individual Difference Measures

	Gender Importance	Race Importance	Socio-Economic Status	Political Orientation	Social Dominance Orientation
Gender Importance	1.000	0.475	0.231	0.008	-0.018
Race Importance	0.475	1.000	0.284	0.078	-0.122
Socio-Economic Status	0.231	0.284	1.000	0.094	-0.052
Political Orientation	0.008	0.078	0.094	1.000	0.496
Social Dominance Orientation	-0.018	-0.122	-0.052	0.496	1.000

Results

Because the data was within-subjects and consisted of repeated measures, I again used a multi-level modeling approach to estimate all effects discussed below. For all models, both the slope and intercept were allowed to vary randomly by participant.

Discrimination conditions and non-discriminatory decisions. All four discriminatory hiring comparisons were always viewed as statistically more problematic (less fair, higher in deliberately discriminatory action, and higher negative affect) than the non-discriminatory comparisons. Because including the non-discriminatory decisions does not add to our understanding of how perceivers detect discrimination, the relationship of the non-discriminatory decisions relative to the discriminatory hiring decisions is not discussed in detail further.

Qualitative coding of written hiring evaluation question. The Black woman discriminated against/White man hired decision had the greatest number of mentions of discrimination, and the proportion of participants who wrote about discrimination was statistically greater than the proportion of those who did not $t(51) = 3.33, p < .01$. For the other three discriminatory hiring comparisons, the proportion of written mentions of discrimination was not statistically different than the number of responses that did not mention discrimination: Black woman discriminated against/Black man hired ($t(50) = 1.268, p = 0.21$), White woman/Black man ($t(51) = .830, p = .41$); White woman/White man ($t(50) = 0.139,$

$p = 0.89$). See table 9 for the percentage of participants who mentioned discrimination by condition.

Table 9

Percentage of Participants Who Wrote about Discrimination by Condition

	Black woman/ White man	Black woman/ Black man	White woman/ Black man	White woman/ White man
Percentage of Participants Mentioning Discrimination	71.15%	58.82%	55.77%	50.98%

Table 10

Log Odds that a Participant Wrote About Discrimination by Condition

Black woman/ White man <i>Log odds Est.</i> = 1.397 <i>SE</i> = 0.479				
Black woman/ Black man <i>Log odds Est.</i> = 0.533 <i>SE</i> = 0.443	$p = .10$			
White woman/ Black man <i>Log odds Est.</i> = 0.345 <i>SE</i> = 0.435	$p = .05^*$	$p = 0.71$		
White woman/ White man <i>Log odds Est.</i> = 0.029 <i>SE</i> = 0.436	$p = .01^*$	$p = 0.32$	$p = 0.53$	
	Black woman/ White man	Black woman/ Black man	White woman/ Black man	

A logistic multi-level model to compare the different discriminatory decisions to each other revealed that participants were more likely to mention

discrimination in the Black woman discriminated against/White man hired condition compared to the discriminatory conditions where White women were discriminated against: White woman/Black man, $p=.05$; White woman/White man, $p=.01$. See table 10 for all log odds estimates and comparisons between conditions.

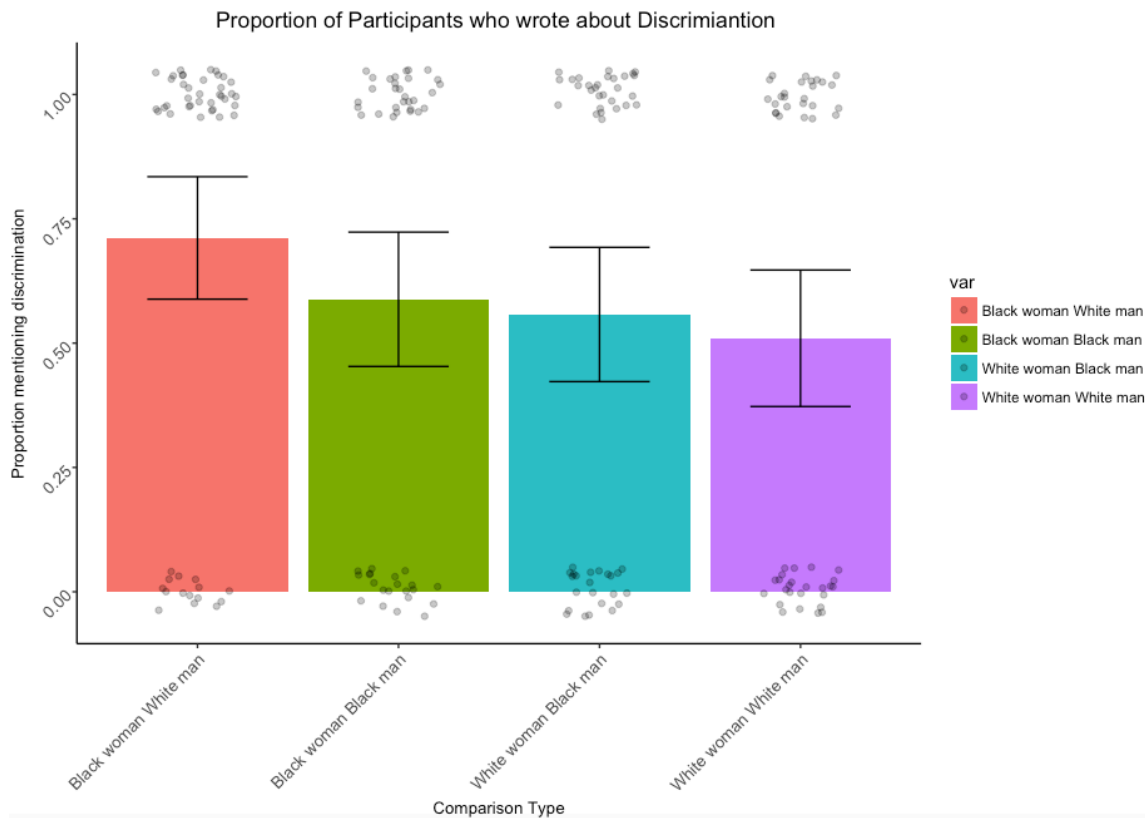


Figure 18. Percentage of written responses mentioning discrimination by discriminatory comparison.

Perceived fairness of decision.

Table 11

Scores of Perceived Fairness by Hiring Comparison

Black woman/ White man Est.= 1.942 SE= 0.124	
Black woman/ Black man	$p<.001^{**}$

Black man <i>Est.</i> = 2.663 <i>SE</i> = 0.124				
White woman/ Black man <i>Est.</i> = 2.587 <i>SE</i> = 0.124	$p < .001^{**}$	$p = 0.575$		
White woman/ White man <i>Est.</i> = 2.519 <i>SE</i> = 0.124	$p < .001^{**}$	$p = 0.294$	$p = 0.624$	
Non-Discriminatory Decisions <i>Est.</i> = 3.853 <i>SE</i> = 0.124	$p < .001^{**}$	$p < .001^{**}$	$p < .001^{**}$	$p < .001^{**}$
	Black woman/ White man	Black woman/ Black man	White woman/ Black man	White woman/ White man

To test how participant's perceptions of fairness differed by hiring decision, I fit a multilevel regression with hiring decision as a predictor of perceived fairness, see table 11. Because non-discriminatory comparisons were collapsed, I had five ratings of fairness for each participant total (the four discriminatory hiring comparisons and an averaged score for the six collapsed non-discriminatory decisions), leaving 260 judgments of fairness total.

Overall, the Black woman discriminated against/White man hired decision was viewed as the least fair (*Est.*= 1.942, *SE*=0.124) followed by the White woman discriminated against/White man hired decision (*Est.*= 2.519, *SE*=0.124) and the White woman discriminated against/Black man hired decision (*Est.*= 2.587, *SE*=0.124). The Black woman discriminated against/Black man hired decision was viewed as the most fair of all of the hiring decisions (*Est.*= 2.663, *SE*=0.124).

Additionally, the non-discriminatory decisions were viewed as significantly more fair than all of the discriminatory decisions ($Est.= 3.853, SE=0.124$).

Participants perceived the fairness of the Black woman/White man hiring decision to be significantly less fair than all of the other hiring discrimination decisions (Black woman/Black man, $p<.01$; White woman/Black man, $p<.01$; White woman/White man, $p<.01$).

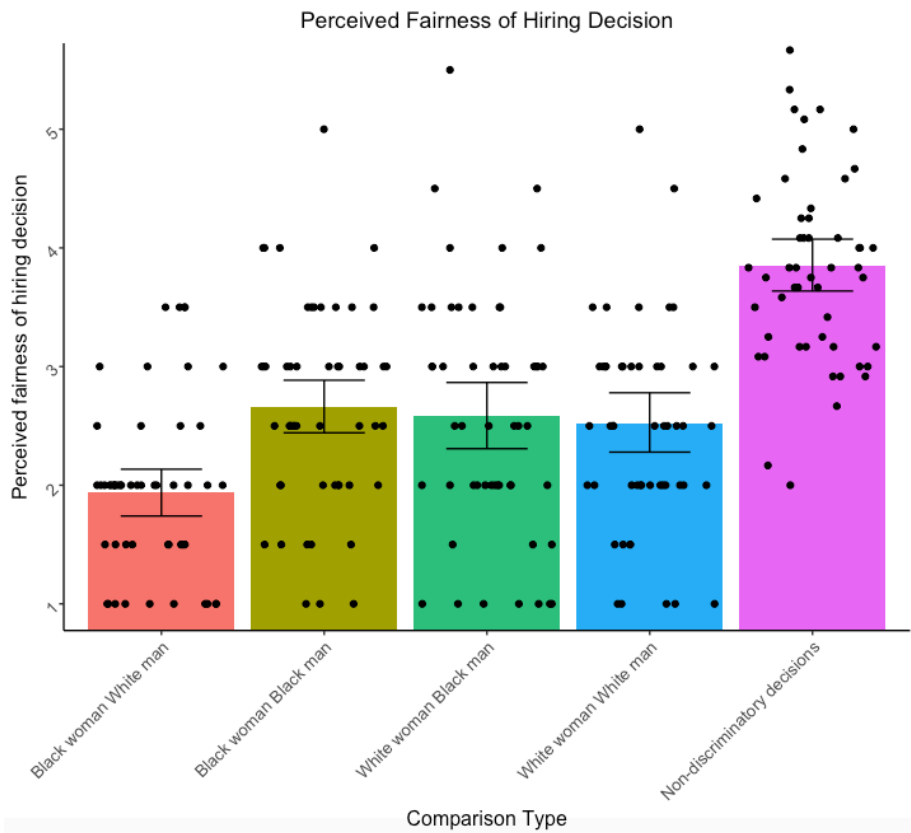


Figure 19. Perceived fairness by hiring comparison.

Perception of deliberate discriminatory action by hiring manager.

Table 12

Scores of Perceived Deliberate Action by Hiring Manager

Black woman/ White man <i>Est.</i> = 5.173 <i>SE</i> = 0.150				
Black woman/ Black man <i>Est.</i> = 4.822 <i>SE</i> = 0.150	$p = 0.053.$			
White woman/ Black man <i>Est.</i> = 4.346 <i>SE</i> = 0.150	$p < .01^{**}$	$p = .009^*$		
White woman/ White man <i>Est.</i> = 4.635 <i>SE</i> = 0.150	$p < .01^{**}$	$p = .301$	$p = .110$	
Non-Discriminatory Decisions <i>Est.</i> = 2.484 <i>SE</i> = 0.150	$p < .01^{**}$	$p < .01^{**}$	$p < .01^{**}$	$p < .01^{**}$
	Black woman/ White man	Black woman/ Black man	White woman/ Black man	White woman/ White man

I fit a multilevel regression with hiring decision as a predictor of ratings of perceived deliberately discriminatory action to understand participant's perceptions that the hiring manager acted in a deliberately discriminatory way (see table 12). Because non-discriminatory hiring decisions were collapsed, I again had five ratings of fairness for each participant total leaving 260 total judgments of deliberate action by the hiring manager.

Following the same pattern of perceived fairness responses, participants overall viewed the Black woman discriminated against/White man hired decision as the most discriminatory ($Est.= 5.173, SE=0.150$) followed by the Black woman discriminated against/Black man hired decision ($Est.= 4.822, SE=0.150$) and the White woman discriminated against/White man hired decision ($Est.= 4.635, SE=0.150$). The White woman discriminated against/Black man hired decision was viewed as the least deliberately discriminatory of all of the hiring decisions ($Est.= 4.346, SE=0.150$). Additionally, the non-discriminatory decisions were viewed as significantly lower than all of the discriminatory decisions ($Est.=2.484, SE=0.150$), indicating that participants viewed all discriminatory hiring decisions to be more deliberately discriminatory than the non-discriminatory decisions.

Participants perceived the Black woman/White man hiring decision to be significantly more deliberately discriminatory than the hiring decisions where White women were discriminated against (White woman/Black man, $p<.01$; White woman/White man, $p<.01$). Additionally, the Black woman/White man decision was also perceived as marginally more deliberately discriminatory than the Black woman/Black man decision ($p=0.053$).

The second condition with a non-prototypical target, the Black woman/Black man condition, was also viewed as statistically more deliberate than the White woman/Black man hiring decision ($p=.009$). The two decisions where White women (prototypical marginalized group members) were discriminated against (White woman/Black man, White woman/White man) were not significantly different than each other in perceived deliberate discriminatory action ($p= 0.11$).

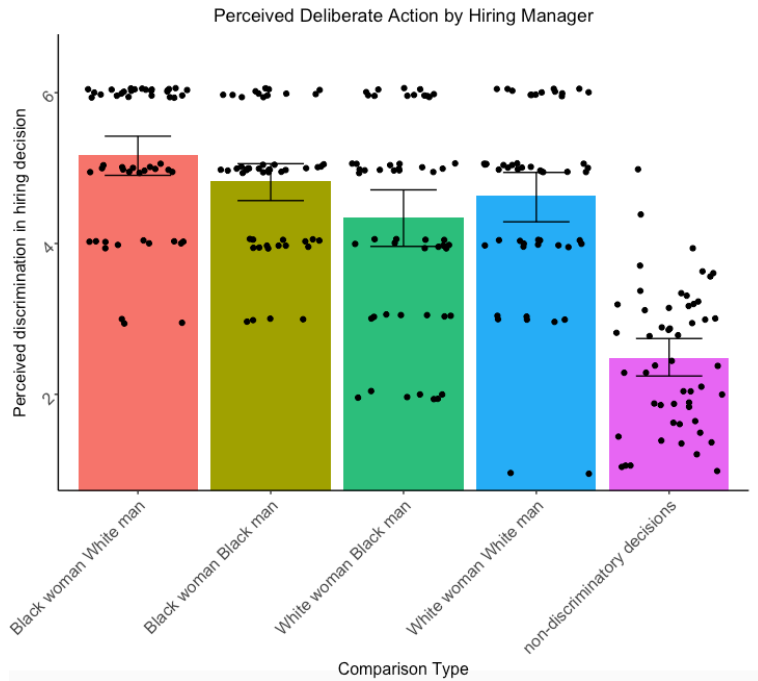


Figure 20. Scores of perceived deliberate action by hiring manager.

Reported concern by hiring decision.

Table 13

Scores of Self-Reported Concern About Each Hiring Decision

Black woman/ White man Est.= 5.260 SE= 0.163				
Black woman/ Black man Est.= 4.804 SE=0.163	$p = .01^*$			
White woman/ Black man Est.= 4.730 SE= 0.164	$p < .01^*$	$p = .678$		
White woman/ White man Est.= 4.631 SE= 0.165	$p < .01^{**}$	$p = 0.33$	$p = 0.58$	
Non-Discriminatory Decisions	$p < .01^{**}$	$p < .01^{**}$	$p < .01^{**}$	$p < .01^{**}$

<i>Est.</i> = 3.318 <i>SE</i> = 0.209				
	Black woman/ White man	Black woman/ Black man	White woman/ Black man	White woman/ White man

Following the same pattern found in the other two self-report measures of discrimination detection, participants found the Black woman discriminated against/White man hired decision to be the most concerning (*Est.*= 5.260, *SE*=0.163) followed by the Black woman discriminated against/Black man hired decision (*Est.*= 4.804, *SE*=0.163) and the White woman discriminated against/Black man hired decision (*Est.*= 4.730, *SE*=0.164). The White woman discriminated against/White man hired decision was viewed as the least concerning of all of the hiring decisions (*Est.*= 4.631, *SE*=.165). Additionally, participants viewed all discriminatory hiring decisions to be more concerning than the non-discriminatory decisions (*Est.*= 3.318, *SE*=0.209). See table 13 for all estimates and comparisons of conditions.

Participants perceived the Black woman/White woman hiring decision to be significantly more concerning than all of the other hiring decisions (Black woman/Black man, $p=.01$; White woman/Black man, $p<.01$; White woman/White man, $p<.01$).

The other three discriminatory decisions were not significantly different than each other in level of concern.

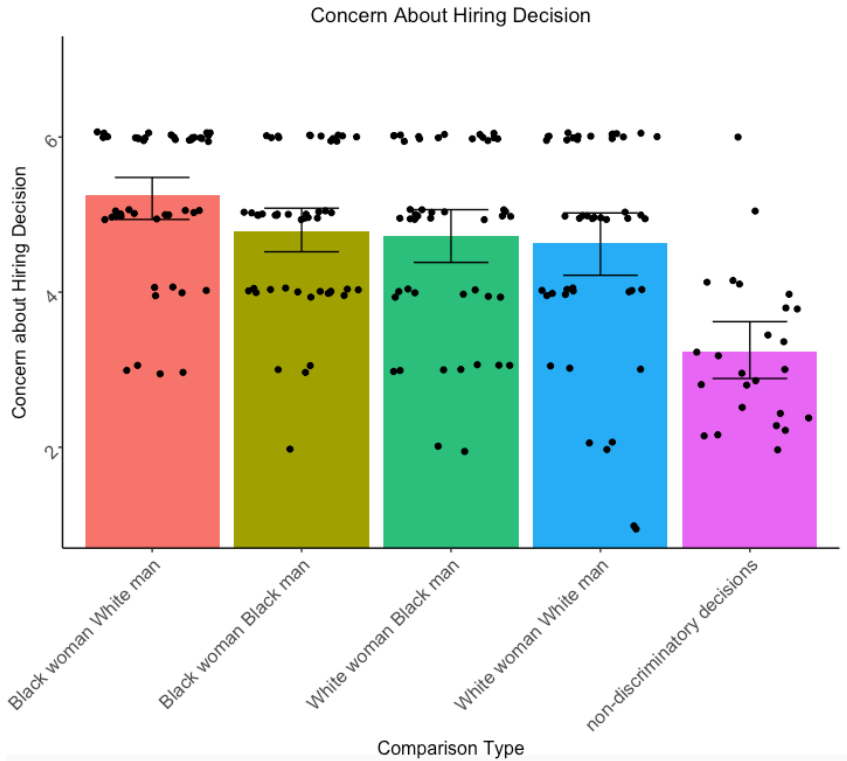


Figure 21. Self-reported concern over hiring decision.

Surprise cued recall. Participants saw a series of faces and were asked whether they previously saw each face in a hiring decision application. Participants' accuracy for the Black woman ($M= 0.712, SD= 0.46$)/White man comparison ($M= 0.615, SD= .491$) was equivalent for both faces, $t(51)=1, p=.322$, indicating that participants' attention for both applicants in this comparison were similar. This was not the case with the other three discriminatory comparisons. Instead, participants were significantly more accurate in recognizing the discriminated against women compared to the non-discriminated against male faces, indicating that participants paid greater attention to the discriminated against women than the non-discriminated against men in these three comparisons: 1. Black woman ($M=.788, SD=.412$)/Black man ($M=.615, SD =.491$), $t(51)= 2.134, p=.04$; 2. White woman ($M=$

0.788, $SD=0.412$)/Black man ($M= 0.461$, $SD= 0.503$), $t(51)=3.818$, $p<.01$; 3. White woman ($M= 0.788$, $SD= 0.412$)/White man ($M= 0.596$, $SD= .495$), $t(51)= 2.210$, $p= 0.032$.

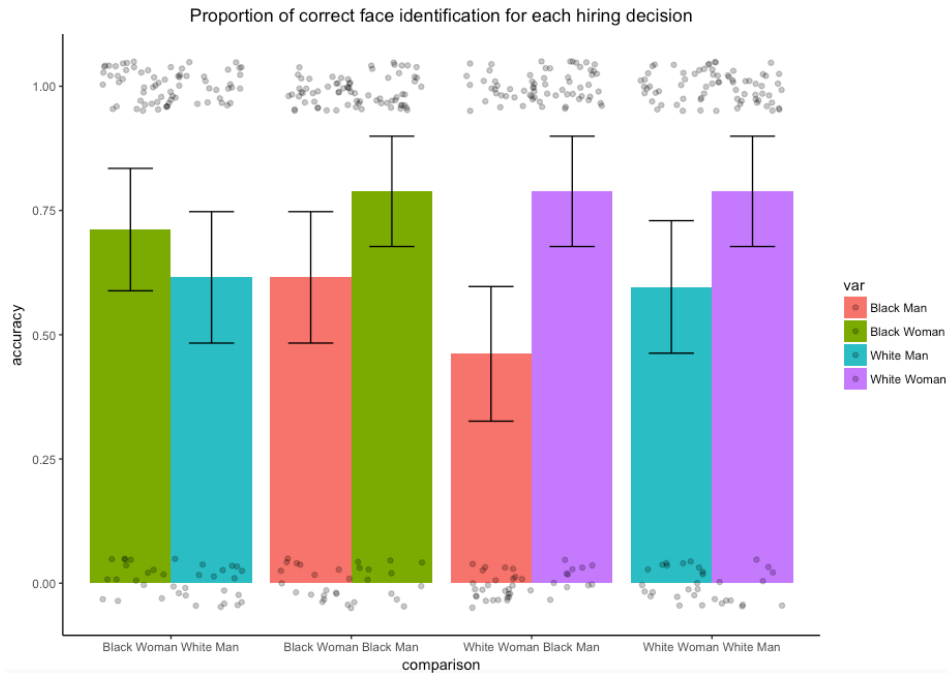


Figure 22. Percentage of correct recognition responses by discriminatory comparison.

Study 5 Discussion

Supporting the hypervisibility outcome of non-prototypicality, the Black woman discriminated against/White man hired condition was viewed as the most problematic of the discriminatory hiring decisions. Specifically, participants were more likely to mention discrimination in their written assessments, more likely to accurately recall the faces of the job candidates, and reported that this condition was the most problematic (least fair, most deliberate discriminatory action, and most concerning) relative to the other discriminatory conditions.

However, with the exception of perceived deliberate action, participants did not view the other comparison with a Black woman (the Black woman discriminated against/Black man hired condition) as any more problematic than the two conditions where prototypical women (White women) were the targets of discrimination (White woman/Black man and White woman/White man). Because discrimination towards the non-prototypical marginalized targets were not perceived similarly, these results mirror those of Studies 3 and 4 and reinforce the idea that participants' perception of discrimination varies both as a function of the target's group (Black woman or White woman) and the comparison candidate's group (White man or Black man).

There are several psychological processes that could explain why the Black woman/White man discrimination comparison was hypervisible relative to the other discriminatory conditions. Participants may detect discrimination by comparing the difference in expected disadvantages between group members. Thus, participants may have been more concerned about discrimination against the Black woman in the Black woman/White man comparison because participants perceived the greatest difference between these two groups.

Sanchez et al.'s (2017) work on perceived prejudice transfer may also explain these results. Specifically, if, as Sanchez suggests, perceivers believe that someone who perpetuates gender based discrimination (such as the hiring manager in this study) also has racist beliefs, it then follows that participants would be most concerned by the discriminatory hiring decision that benefits both men and White people (the Black woman discriminated against/White man hired condition).

General Discussion

Five studies empirically demonstrated that non-prototypical marginalized group members are perceived differently than their prototypical counterparts and that these differences are associated with downstream cognitive consequences. Specifically, Study 1 used archival data to empirically test Purdie-Vaughns & Eibach's (2008) theoretical model and demonstrated an instance of the invisibility outcome. Non-prototypical individuals were shown on the covers of Time magazine less frequently than their prototypical counterparts, and were represented less relative to their proportion in the population. Study 2 established that beliefs about prototypical and non-prototypical individuals from the same marginalized group differ in fundamental ways.

Studies 3, 4, and 5 demonstrated the hypervisibility outcome. All of these studies presented participants with non-prototypical marginalized group members and their prototypical counterparts in direct one-on-one comparisons. Study 3 found that when comparing differences between groups, participants were most likely to attribute these differences to the non-prototypical marginalized group—but only when comparing groups that differed only by sexual orientation and race. Studies 4 and 5 explored how participants differentially detected instances of racism and sexism when the target was a prototypical or non-prototypical marginalized group member and found that in both studies, participants most easily recognized discrimination when the target was non-prototypical. Importantly however, this hypervisibility outcome was not consistently observed when comparing non-prototypical marginalized group members to their prototypical

counterparts. For Studies 3-5, the hypervisibility outcome was dependent on both the identity of the target and the identity of the comparison partner, a finding that has not previously been demonstrated in previous work on prototypical and non-prototypical identities.

While the current set of studies did not actively test and explore why certain comparisons between non-prototypical and prototypical groups are hypervisible while other comparisons are not, the findings from this dissertation inform when and why this may occur. Previous research (Miller, Taylor, & Buck, 1991) has established that hypervisibility occurs in direct comparisons because the non-normative group member is perceived as aberrant or different (and therefore worth paying attention to) relative to the normative, prototypical group. Therefore, in instances in which direct comparisons do not lead to the hypervisibility outcome, the two groups being compared are not viewed as sufficiently different from each other.

Study 2 established that in the absence of a specific context (simply asking participants their views about different groups), stereotypes of prototypical and non-prototypical marginalized groups differ significantly. Study 3 however, demonstrated that the context in which groups are considered may minimize these perceived differences, leading in turn to a failure to focus on the non-prototypical group (hypervisibility of the non-prototypical marginalized group members). For example, in Study 3, issues like gay marriage and civil rights are often frequently discussed in relation to major presidential elections (Miller, Taylor & Buck, 1991). As such, participants asked to explain differences between groups in the context of

presidential voting patterns (Study 3) may have minimized the differences between gay men and lesbian women and Black men and Black women much more so than they would in a different context, leading to a lack of hypervisibility in these comparisons.

Further evidence for this is found by comparing Studies 4 & 5. The context of the tech industry, where most people anticipate that both women and Black people are marginalized (“Diversity in high tech”, 2016), is another example of how the context may affect perceivers’ judgments of difference between prototypical and non-prototypical groups. In addition to making judgments in a context that is perceived to be unfair for multiple groups, the language of some of the self report questions (“race/ethnicity” vs. “sex/gender”) as well as the pattern of discrimination participants viewed (all Black people vs. all women shown were treated unfairly) potentially further affected impressions that prototypical and non-prototypical groups were more or less similar to each other in this context. For example, in the context of race based discrimination, the Black woman/White man comparison is not as hypervisible as the Black woman/White woman comparison. However, participants paid the most attention to this same comparison (Black woman/White man) in the context of gender based discrimination, suggesting that elements of the context (such as asking about discrimination based on “gender/sex”) affected perceptions of difference between groups.

Had this study explored perceptions of discrimination in a different hiring context, the work in this dissertation suggests that perceivers’ context-dependent judgments about group-based differences should lead to specific (and different)

patterns of hypervisibility. Future studies should further explore the role of perceived difference between groups in perceiving non-prototypical marginalized groups as hypervisible relative to their prototypical counterparts.

Limitations

Because the studies included in this dissertation explore how perceivers view non-prototypical individuals relative to their prototypical counterparts, it is important to examine who the participants in these studies were. While all studies required that participants were socialized within the U.S., there were no other participation restraints (such as participants' age, race, or gender). Regardless, some demographic groups were overrepresented in the data by virtue of the demographic makeup of the participant pools used to recruit participants. Specifically, the majority of perceivers from these studies were White (80% in Study 2, 85% in study 3, 51% in Study 4, & 58% in Study 5, all studies combined—76%). Additionally, all the editors of Time Magazine for the time period during which data for Study 1 was collected are/were White ("Time Magazine," 2018). Because we know that one's own ingroup identification can influence judgments about outgroup members (Hewstone et al, 2002), it is possible that the results presented in the study represent the views of White perceivers more so than generalizing to all perceivers broadly (regardless of race). While exploratory statistical tests indicated that there were no significant differences in the response patterns for White and non-White perceivers, the statistical power was also not great enough to definitively conclude that differences did not exist. Thus, while the studies included in this dissertation seek to understand how perceivers in general stereotype, perceive, and understand

discrimination against non-prototypical marginalized individuals relative to their prototypical counterparts, the views of White perceivers are heavily overrepresented in these studies.

Future Directions

While this dissertation makes important contributions in exploring the downstream consequences of perceiving non-prototypicality, there are many questions that still remain. For example, the current research focuses on non-prototypical identities on the basis of race, gender and sexual orientation. However, there are countless marginalized identities that could be explored such as socio-economic status, religion, or weight. Additionally, the research in this dissertation explored judgments of prototypical and non-prototypical group members with immediately visible and public marginalized identities (Carpinella, Chen, Hamilton, & Johnson, 2015; Freeman, Johnson, Ambady, & Rule, 2010; Goff et al. 2008; Lick, & Johnson, 2014; Lick, Johnson, & Riskind, 2015). Future research should explore how these judgments may change when exploring more concealable marginalized identities such as religion or disability (Quinn, 2006).

Furthermore, the studies in this dissertation were mainly limited to exploring the consequences of specific non-prototypical identities. For example, when trying to understand perceivers' judgments of non-prototypical women, these studies used Black women as targets. However, according to Purdie-Vaughns & Eibach's (2008) model, Latina, Asian American, and Native American women should also be viewed as non-prototypical women. Previous research suggests however, that stereotypes about women differ by race (Carpinella, Chen, Hamilton, & Johnson,

2015; Goff et al. 2008; Mohr & Purdie-Vaughns, 2015; Williams, 2014). Therefore, future studies should include targets with different non-prototypical identities (such as targets from different racial groups) to explore the extent to which the outcomes of invisibility and hypervisibility are generalizable to all non-prototypical marginalized groups.

Lastly, the studies in this dissertation test the outcomes of invisibility and hypervisibility in contexts that likely highlight the multiple marginalized identities of non-prototypical targets to perceivers. For example, the voting patterns of different demographic groups, including both gender and race groups, are often widely discussed before and after national elections (Miller, Taylor, & Buck, 1991) and the tech industry has historically employed very few women or Black people (Diversity in High Tech, 2016). This may have made the race, sexual orientation, and gender identities of non-prototypical marginalized targets especially salient in Studies 3, 4 and 5, respectively. Future studies should explore the extent to which these patterns of invisibility and hypervisibility replicate in contexts in which only one marginalized identity is salient (where gender but not race is normative, for example, as is the case for elementary school teachers) or in contexts that do not have any strong associations with a particular identity.

Overall Contributions

This dissertation makes several key contributions to the field. First, this work highlights the importance of studying individuals who are non-prototypical members of marginalized identity groups, in addition to those who are prototypical. This work clearly demonstrates that patterns in stereotyping and discrimination

detection differ depending on the identity of the target. Examining non-prototypical individuals therefore is important so that the field does not fall prey to intersectional invisibility and only include and understand prototypical marginalized group members as targets. Doing so would lead to an incomplete and inaccurate understanding of how the processes of stereotyping, prejudice, and discrimination affect a significant proportion of marginalized individuals. Additionally, including non-prototypical marginalized group members as targets allows researchers to gain a more complete understanding of when and how perceivers detect discrimination towards all targets because studying discrimination against non-prototypical individuals sheds light into the mechanisms of discrimination generally.

Finally, this work also suggests that discrimination researchers should pay careful attention to whom targets of discrimination are compared (i.e., the person or group that benefits from the discrimination or the benchmark to which an individual or group is judged; Crenshaw, 1989). Most discrimination researchers carefully select the targets of discrimination, but this work demonstrates that the identity of the comparison group is critical to understanding and predicting how much people will detect and care about discrimination. Importantly, work from this dissertation shows that failing to take comparison groups into account should alter perceptions of whether or not discrimination occurred.

References

- Allport, G. W. (1954). *The Nature of Prejudice*. Cambridge, MA: Addison-Wesley.
- Bates, D., Maechler, M., Bolker, B., & Walker, S. (2014). lme4: Linear mixed-effects models using eigen and S4: R package version 1. 1-7. <http://cran.r-project.org/package=lme4>
- Babbitt, L. G. (2013). An intersectional approach to Black/White interracial interactions: The roles of gender and sexual orientation. *Sex Roles, 68*, 791-802.
- Bem, S. L. (1981). Gender schema theory: A cognitive account of sex typing. *Psychological Review, 88*, 354-364.
- Bem, S. L. (1994). *The lenses of gender: Transforming the debate on sexual inequality*. New Haven, CT: Yale University Press.
- Blanchard, F. A., Adelman, L., & Cook, S. W. (1975). Effect of group success upon interpersonal attraction in cooperating interracial groups. *Journal of Personality and Social Psychology, 31*, 1021-1030.
- Bobo, L. D. (1999). Prejudice as group position: Microfoundations of a sociological approach to racism and race relations. *Journal of Social Issues, 55*, 445-472.
- Bonilla-Silva, E. (2000). "This is a White Country": The racial ideology of the Western nations of the world-system. *Sociological Quarterly, 70*, 188-214.
- Bowleg, L. (2008). When Black+ lesbian+ woman ≠ Black lesbian woman: The methodological challenges of qualitative and quantitative intersectionality research. *Sex roles, 59*(5-6), 312-325.

- Brewer, M. B. (1988). A dual process model of impression formation. In T. K. Scrull & R. S. Wyer (Eds.) *Advances in social cognition* (pp. 1–36). Hillsdale, NJ: Erlbaum.
- Brewer, M. B. (1999). The psychology of prejudice: Ingroup love or outgroup hate? *Journal of Social Issues, 55*, 429–444.
- Buhrmester, M., Kwang, T., and Gosling, S.D. (2011). Amazon’s Mechanical Turk: A New Source of Inexpensive, Yet High-Quality, Data? *Perspectives on Psychological Science, 6*(1), 3-5.
- Carpinella, C. M., Chen, J., Hamilton, D., & Johnson, K. L. (2015). Gendered facial cues influence race categorizations. *Personality and Social Psychology Bulletin, 41*, 405 - 419.
- Chang, E. (2018). *Brotopia: Breaking up the boys’ club of Silicon Valley*. New York, NY: Penguin Random House.
- Cole, E. R. (2009). Intersectionality and research in psychology. *American Psychologist, 64*(3), 170-180.
- Crenshaw, K. (1989). Demarginalizing the intersection of race and sex: A black feminist critique of antidiscrimination doctrine, feminist theory and antiracist politics. *University of Chicago Legal Forum, 139*.
- Cuddy, A. J. C., Fiske, S. T., Kwan, V. S. Y. Glick, P., Demoulin, S. Leyens, Bond, M. H., Croizet, J. C., Ellemers, N., Sleebos, E., Htun, T. T., Kim, H. J., Maio, G. Perry, J., Petkova, K., Todorov, V., Rodríguez-Bailón, R., Morales, E., Moya, M., Palacios, M., Smith, V., Perez, R., Vala, J., & Ziegler, R. (2011). Stereotype content

model across cultures: Towards universal similarities and some differences.
British Journal of Social Psychology, 48(1), 1-33.

Cuddy, A. J. C., Fiske, S. T., Kwan, V. S. Y., Glick, P., Demoulin, S., Leyens, J.-P., Ziegler, R. (2009). Stereotype content model across cultures: Towards universal similarities and some differences. *British Journal of Social Psychology*, 48, 1–33.

Devos, T., & Banaji, M. R. (2005). American = white? *Journal of Personality and Social Psychology*, 88, 447–466.

Devine, P. (1989). Stereotypes and prejudice: Automatic and controlled components. *Journal of Personality and Social Psychology*, 56(1), 5-18.

Diversity in high tech. (2016, May) Retrieved from

<https://www.eeoc.gov/eeoc/statistics/reports/hightech/upload/diversity-in-high-tech-report.pdf>

Dovidio, J. F. (2001). On the nature of contemporary prejudice: The third wave. *Journal of Social Issues*, 57, 829–849.

Duckitt, J. (1992). Psychology and prejudice: A historical analysis and integrative framework. *American Psychologist*, 47, 1182–1193.

Eagly, A. H., & Kite, M. E. (1987). Are stereotypes of nationalities applied to both women and men. *Journal of Personality and Social Psychology*, 53, 451–462.

Fiske, S. T. (1998). Stereotyping, prejudice, and discrimination. In D. T. Gilbert, S. T. Fiske, & G. Lindzey (Eds.), *The Handbook of Social Psychology* (4th ed., Vol. 2, pp. 357–411). New York: McGraw-Hill.

- Fiske, S. T., & Neuberg, S. L. (1990). A continuum of impression-formation, from category-based to individuating processes: Influences of information and motivation on attention and interpretation. In M. Zanna (Ed.). *Advances in Experimental Social Psychology* (pp. 1–74). San Diego, CA: Academic Press.
- Federico, C. (1999). The interactive effects of social dominance orientation, group status, and perceived stability on favoritism for high-status groups. *Group Processes and Intergroup Relations*, 2, 119-144.
- Freeman, J. B., Johnson, K. L., Ambady, N., & Rule, N. (2010). Sexual orientation perception involves gendered facial cues. *Personality and Social Psychology Bulletin*, 36, 1318 - 1331.
- Fryberg, S. A. & Eason, A. E. (2017). Making the invisible visible: Acts of commission and omission. *Current Directions in Psychological Science*, 26(6), 554-559.
- Fryberg, S. A. & Townsend, S. M. (2008). The psychology of invisibility. In G. Adams, M. Biernat, N. Branscombe, C. Crandall, & L. Wrightsman (Eds.). *Commemorating Brown: The Social Psychology of Racism and Discrimination*. Washington DC: American Psychological Association.
- Galinsky, A., Hall, E.V., & Cuddy, A. (2013). Gendered races: Implications for interracial marriage, leadership selection, and athletic participation. *Psychological Science*, 24(4), 498-506.
- Glick, P. & Fiske, S. T. (1996). The Ambivalent Sexism Inventory: Differentiating hostile and benevolent sexism. *Journal of Personality and Social Psychology*, 70(3), 491– 512.

- Goff, P. A., Eberhardt, J. L., Williams, M. J., & Jackson, M. C. (2008). Not yet human: implicit knowledge, historical dehumanization, and contemporary consequences. *Journal of personality and social psychology, 94*(2), 292.
- Goff, P. A., & Kahn, K. B. (2013). How psychological science impedes intersectional thinking. *Du Bois Review: Social Science Research on Race, 10*(2), 365- 384.
- Goff, P. A., Thomas, M. A., & Jackson, C. T. (2008). Ain't I a woman? Towards an intersectional approach to person perception and group-based harms. *Sex Roles, 59*, 392–403.
- 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*(6), 853-868.
- Hegarty, P., Pratto, F., & Lemieux, A. (2004). Heterocentric norms and heterosexist ambivalences: Drinking in intergroup discomfort. *Group Processes and Intergroup Relations, 7*, 119–130.
- Hewstone, M., Rubin, M., & Willis, H. (2002). Intergroup bias. *Annual review of psychology, 53*(1), 575-604.
- Inman, M. L., & Baron, R. S. (1996). Influence of prototypes on perceptions of prejudice. *Journal of Personality and Social Psychology, 70*(4), 727-739.
- 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.
- Jost, J.T. (2006). The end of the end of ideology. *American Psychologist, 61*, 651-670.

- Jost, J. T., Nosek, B. A. & Gosling, S. A. (2008). Ideology: Its resurgence in social, personality, and political psychology. *Perspectives on Psychological Science*, 3(2), 126-136.
- Kahneman, D., & Miller, D. T. (1986). Norm theory: Comparing reality to its alternatives. *Psychological review*, 93(2), 136.
- Kaiser, C. R. & Miller, C. T. (2004). A stress and coping perspective on confronting sexism. *Psychology of Women Quarterly*, 28(2), 168-178.
- Kang, S. K., & Bodenhausen, G. V. (2015). Multiple identities in social perception and interaction: Challenges and opportunities. *Annual review of psychology*, 66, 547- 574.
- Kuznetsova, A., Brockhoff, P.B., & Christensen R.H.B. (2017). lmerTest Package: Tests in Linear Mixed Effects Models. *Journal of Statistical Software*, 82(13), pp. 1–26. doi: 10.18637/jss.v082.i13.
- Levin S., Federico C.M., Sidanius J. & Rabinowitz J.L. (2002). Social dominance orientation and intergroup bias: The legitimation of favoritism for high-status groups. *Personality and Social Psychology Bulletin*, 28(2) , pp. 144-157.
- 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.
- Lick, D. J., & Johnson, K. L. (2014). Perceptual roots of anti-gay prejudice: Negative evaluations of targets perceived to be lesbian/gay arise early in person perception on the basis of gender atypical visual cues. *Personality and Social Psychology Bulletin*, 40, 1178 - 1192.

- Lick, D. J., Johnson, K. L., & Riskind, R. G. (2015). Haven't I seen you before? Gender insecure men are vigilant to gender-atypical faces. *Group Processes and Intergroup Relations, 18*, 131 - 152.
- Livingston, R. W., Rosette, A. S., & Washington, E. F. (2012). Can an agentic Black woman get ahead? The impact of race and interpersonal dominance on perceptions of female leaders. *Psychological Science, 23*(4), 354-358.
- Ma, Correll, & Wittenbrink (2015). The Chicago Face Database: A Free Stimulus Set of Faces and Norming Data. *Behavior Research Methods, 47*, 1122-1135.
- Marti, M. W., Bobier, D. M. & Baron, R. S. (2000). Right before our eyes: The failure to recognize non-prototypical forms of prejudice. *Group Processes and Intergroup Relations, 3*(4), 403-418.
- Mason, M. and Suri, S. (2012). Conducting Behavioral Research on Amazon's Mechanical Turk. *Behavior Research Methods, 44*(1), 1-23.
- McCall, L. (2005). The complexity of intersectionality. *Signs: Journal of Women in Culture and Society, 30*, 1771-1800.
- Mervis, C. B., & Rosch, E. (1981). Categorization of natural objects. *Annual Review of Psychology, 32*, 89-115.
- McConahay, J. B. (1986). Modern racism, ambivalence, and the modern racism scale. In J. F. Dovidio, & S. L. Gaertner (Eds), *Prejudice, Discrimination, and Racism* (pp. 91-125). Orlando, FL: Academic Press.
- McGill, A. L. (1989). Context effects in judgments of causation. *Journal of Personality and Social Psychology, 57*(2), 189.
- Miller, D. T., Taylor, B. & Buck, M. L. (1991). Gender gaps: Who needs to be

- explained? *Journal of Personality and Social Psychology*, 61(1), 5-12.
- Mohr, R. I., & Purdie-Vaughns, V. (2015). Diversity within women of color: Why experiences change felt stigma. *Sex Roles*, 73, 391–398.
- Paolacci, G., Chandler, J. & Ipeirotis, P. G. (2010). Running experiments on Amazon Mechanical Turk. *Judgment and Decision Making*, 5(5), 411-419.
- Pedulla, D. S. (2014). The positive consequences of negative stereotypes: Race, sexual orientation, and the job application process. *Social Psychology Quarterly*, 77(1), 7-94.
- Pratto, F., Çidam, A., Stewart, A. L., Bou Zeineddine, F., Aranda, M., Aiello, A., & Henkel, K. E. (2013). Social dominance in context and in individuals: Contextual moderation of robust effects of social dominance orientation in 15 languages and 20 countries. *Social Psychology and Personality Science*, 4, 587–599. doi:10.1177/1948550612473663
- Pratto, F., & John, O. P. (1991). Automatic vigilance: The attention-grabbing power of negative social information. *Journal of Personality and Social Psychology*, 61(3), 380-391.
- Pratto F. & Pitipitan, E. V. (2008). Ethnocentrism and sexism: How stereotypes legitimize six types of power. *Social and Personality Compass*, 2, 2159-2176.
- Pratto, F. Sidanius, J., Stallworth, L. M., Malle, B. F. (1994). Social Dominance Orientation: A personality variable predicting social and political attitudes. *Journal of Personality and Social Psychology*, 67 (4), 741-763.

- Purdie-Vaughns, V. J. & P. Eibach, R. P. (2008). Intersectional Invisibility: The Distinctive Advantages and Disadvantages of Multiple Subordinate-Group Identities. *Sex Roles, 59*(5-6), 377-391. 10.1007/s11199-008-9424-4.
- Quinn, D. M. (2006). Concealable versus conspicuous stigmatized identities. In Shana Levin & Colette Van Laar (Eds.), *Stigma and group inequality: Social psychological perspectives* (pp. 83-103). Mahway, NJ: Lawrence Erlbaum Associates, Inc.
- Rosch, E. (1975). Cognitive representations of semantic categories. *Journal of Experimental Psychology: General, 104*, 192–232.
- Rosch, E., & Mervis, C. B. (1975). Family resemblances: Studies in the internal structure of categories. *Cognitive Psychology, 7*, 573–6.
- Sanchez, D. T., Chaney, K. E., Manuel, S. K. Wilton, L. S. & Remedios, J. D. (2017). Stigma by prejudice transfer: Racism threatens white women and sexism threatens men of color. *Psychological Science, 28*(4), 445-461.
- 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.
- Schug, J., Alt, N. P., Lu, P. S., Gosin, M. & Fay J. L. (2017). Gendered race in mass media: Invisibility of Asian men and Black women in popular magazines. *Psychology of Popular Media Culture, 6*(3), 222-236.
- 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*(5- 6), 301-311.
- Smith, E. R., & Zarate, M. A. (1992). Exemplar-based model of social judgment. *Psychological review, 99*(1), 3.
- Stangor, C., Lynch, L., Duan, C., & Glas, B. (1992). Categorization of individuals on the basis of multiple social features. *Journal of Personality and Social Psychology, 62*(2), 207.
- Stroessner, S.J. (1996). Social categorization by race or sex: Effects of perceived non-normalcy on response times. *Social Cognition, 14*(3), 247-276.
- Time Magazine (2018). Retrieved from [https://en.wikipedia.org/wiki/Time_\(magazine\)](https://en.wikipedia.org/wiki/Time_(magazine))
- U.S. Census Bureau (2016). Quick Facts: United States. Retrieved from <https://www.census.gov/quickfacts/fact/table/US/PST045216>
- Weiner, B., Perry, R. P., & Magnusson, J. (1988). An attributional analysis of reactions to stigmas. *Journal of Personality and Social Psychology, 55*(5), 738-748.
- Williams, J. C. (2014). Double jeopardy? An empirical study with implications for the debates over implicit bias and intersectionality. *Harvard Journal of Law & Gender, 37*, 185-242. doi:10.1016/ b978-012679130-3/50035-1.
- Zarate, M. A., & Smith, E. R. (1990). Person categorization and stereotyping. *Social Cognition, 8*(2), 161-185.

Appendix A

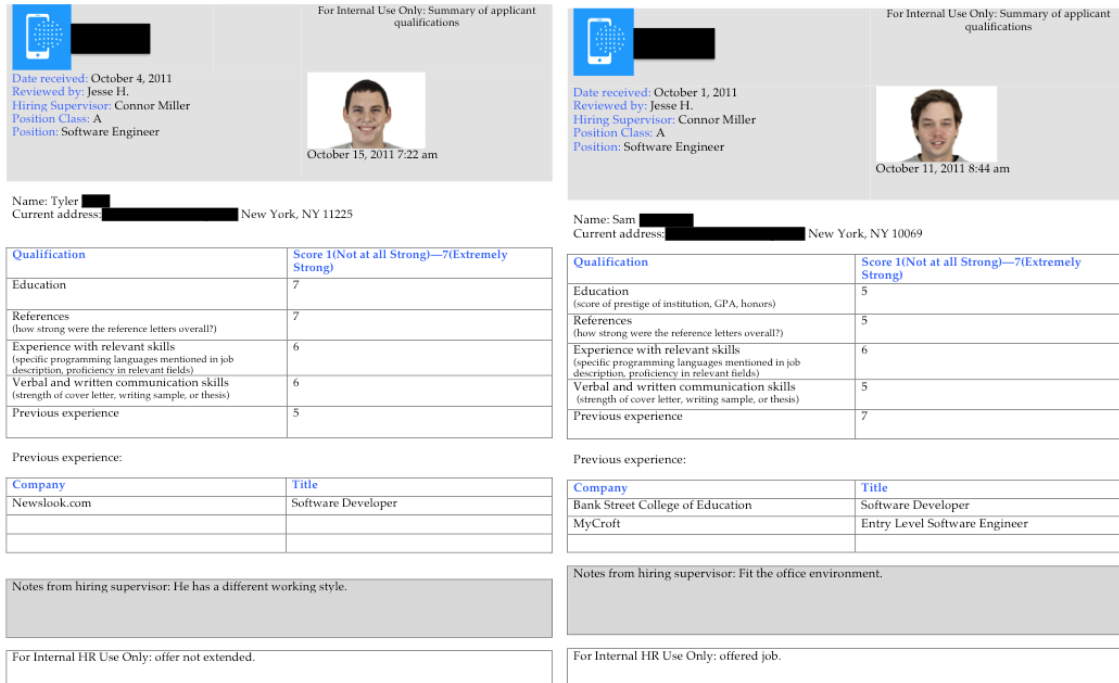


Figure A1. Sample quantitative summaries used in Studies 4 & 5. This example is from a non-discriminatory decision.

Appendix B

Table B1

Qualitative Coding Guidelines for Discrimination Type Coding (Studies 4 & 5)

Code Type	Coding Guidelines	Example
Sexism	<p>0 (No discrimination): No mention of sexism. This includes instances in which participants are just stating facts of the hiring information. Responses were coded as 0 even if there is some exasperation/confusion related to the summary (e.g. “Yet again, the not hired candidate had higher scores.”) as long as that exasperation/confusion does not mention discrimination and/or specific groups.</p> <p>1 (sexism—weak assertion): Participant alludes to sexism/unfairness based on gender, but hedges (or expresses uncertainty) about their judgment of sexism. Words to look for—“maybe” “not sure” “seems” “unclear”</p> <p>2 (sexism—strong assertion): Participant explicitly and unequivocally expresses that they believe that sexism occurred.</p>	<p><i>Weak Assertion:</i> “I do believe, at least at first glance, that gender might have played a role.”</p> <p><i>Strong Assertion:</i> “HR interviewer comments on #2 very troubling and suggest gender and / or bias played a huge role here. “</p>
Racism	<p>0 (No discrimination): No mention of racism. Same guidelines as for sexism above.</p> <p>1 (racism—weak assertion): Participant alludes to racism/unfairness based on race, but hedges (or expresses uncertainty) about their judgment of racism. Words to look for—“maybe” “not sure” “seems” “unclear”</p> <p>2 (racism—strong assertion): Participant explicitly and unequivocally expresses that they believe that racism occurred.</p>	<p><i>Weak Assertion:</i> “I hesitate to say the hiring decision was racially charged, but the notes on the second applicant were problematic.”</p> <p><i>Strong Assertion:</i> “Concerns -- this company doesn’t like</p>

		black people.”
General Discrimination	<p>General discrimination means that participants did not mention racism or sexism specifically, but rather talk about discrimination or unfairness generally.</p> <p>0 (No discrimination): No mention of discrimination. Same guidelines as for racism and sexism above.</p> <p>1 (general discrimination—weak assertion): Participant alludes to discrimination/unfairness, but hedges (or expresses uncertainty) about their judgment. Words to look for—“maybe” “not sure” “seems” “unclear”</p> <p>2 (sexism—strong assertion): Participant explicitly and unequivocally expresses that they believe that discrimination occurred.</p>	<p><i>Weak Assertion:</i> "Unclear fit" seems like a really vague reason not to hire this person; it leads to concerns of alternate reasons."</p> <p><i>Strong Assertion:</i> "This clearly references her appearance and displays discrimination."</p>
Intersectional Discrimination	<p>0 (No discrimination): No mention of discrimination based on race <u>and</u> gender. Same guidelines as above.</p> <p>1-(intersectional discrimination—weak assertion): Participant alludes to discrimination/unfairness based on race and gender, but hedges (or expresses uncertainty) about their judgment. Words to look for—“maybe” “not sure” “seems” “unclear”</p> <p>1-(intersectional discrimination—strong assertion): Participant explicitly and unequivocally expresses that they believe that discrimination based on race & gender occurred.</p>	<p><i>Weak Assertion:</i> “I do believe, at least at first glance, that race and gender might have played a role. “</p> <p><i>Strong Assertion:</i> “Concerns -- discrimination against her for being black and female.”</p>
Negative Affect	<p>0 (No negative affect present): Participant does not use sarcasm, anger, sadness, expresses general confusion in their response.</p> <p>1 (Negative Affect): Participant uses sarcasm, anger, sadness, expresses general confusion in their response</p>	<p>“The hiring decision is not fair.” “A girl!”</p>

Appendix C

Materials Given to Participants about the Tech Company (Studies 4 & 5)

Thank you for participating in this study today!

We're interested in the hiring process and what additional information is gained from face to face interviews. Many technology companies have started using more computational methods for making hiring decisions because they believe that this computational way saves them time and may lead to better hires. The goal of this method is to show the strengths and weaknesses of each candidate in one document so that hiring managers or HR personnel can make rational decisions regarding whom to hire.

The goal of this study today is to evaluate how this new way of representing job candidates compares to more traditional hiring practices. **You have been assigned to the applicant summary materials only condition.** Other participants will only view videotapes of job interviews. Your help in evaluating this method is extremely important as several other industries have started to adopt similar hiring approaches.

We've collaborated with a technology company that uses this method, and they have provided us with their hiring materials from several recent searches they've conducted over the last three years. For the sake of privacy for both the job applicants and the company, we have redacted both the full names of the applicants as well as the real name of the company.

After reading these materials, you will see the hiring summary information for the top two job applicants from different job searches. You'll have a chance to view both summaries at the same time and you'll be provided with which of the two applicants ended up being hired. After viewing both applicants' hiring summaries, your job is to evaluate how well the hiring manager, Connor Miller, performed his job by using the criteria presented in the summaries. You'll do this a total of 10 times for 10 different hiring decisions.

How were these candidates ranked?

Each job applicant receives a score that indicates how qualified they are for every job requirement. So for example, the software developer job that you will soon see requires a bachelor's degree. Someone who does not have this requirement would receive a score of a 1. Someone who is extremely qualified for this particular requirement, for example someone who graduated from MIT with honors, would receive a score of a 7. These ratings are determined for every requirement, and are compiled into a computational hiring summary for each job applicant. In other words, an applicant's rating for a particular requirement can be thought of as an indication of how qualified they are for that particular skill, and their hiring summary compiles all of that qualification information in one place.



About the Company: At [REDACTED] we're developing a game-changing tech platform enabling brands to create, test and optimize micro video content (video less than 15 secs) from a unique community of digital creators.

Brands can tap into our creative network to quickly produce micro video for mobile platforms and digital channels. [REDACTED] goes beyond traditional video production to provide an innovative content creation process powered by creators and driven by technology. We're working with global brands and agencies, while also providing unique opportunities for amazing new creative talent.

[REDACTED] is committed to creating a diverse environment and is proud to be an equal opportunity employer. All qualified applicants will receive consideration for employment without regard to their sex, race, color, marital status, national origin, religious affiliation, disability, sexual orientation, gender identity, or expression of age.

Job Position: Software Engineer

About the specific job: We are looking for a Software Engineer who wants to be an integral part of the [REDACTED] team and ready to take on the technical challenges that come with disrupting the video advertising and creative industries.

This is an opportunity to be part of dynamic, continually growing company where qualified applicants' contributions make a real difference and new ideas are always welcome. Qualified applicants will have strong input into the way code, culture and product are built at [REDACTED].

We have a working system but wish to re-write this from the ground up to accommodate our swiftly growing business. Therefore, qualified applicants will be able to determine most of the technical parameters, including the coding languages, frameworks, cloud provider and other technology choices.

Responsibilities:

- Write well documented and tested python code
- Perform load testing to ensure that new designs can scale
- Work with the team to make sure that software is running optimally in production
- Address software quality issues as they arise
- Contribute ideas for new features and identify areas for improvement proactively

Candidates must have:

- 1) B.S. or higher in Computer Science
- 2) Strong experience with Python or Ruby based technology stack building
- 3) Solid foundation in data structures, algorithms and complexity analysis
- 4) Comfortable working in a small, high-growth, and fast paced development environment
- 5) Very good verbal and written English communication skills