

RESEARCH REPORT

Are Leaders Still Presumed White by Default? Racial Bias in Leader Categorization Revisited

Christopher D. Petsko and Ashleigh Shelby Rosette

Fuqua School of Business, Duke University

The Samuel DuBois Cook Center on Social Equity, Duke University

In the United States, leaders of the highest valued companies, best-ranked universities, and most-consumed media outlets are more likely to be White than what would be expected based on White people's representation in the U.S. population. One explanation for this racial gap is that U.S. respondents' prototype of a leader is White by default—which is, in turn, what causes White (vs. non-White) people to be promoted up the organizational ladder more quickly. Although this explanation has empirical support, its central premise was recently challenged by experimental evidence documenting that U.S. respondents no longer associate leaders, more than nonleaders, with being White. To reconcile these contradictory findings, we conducted three preregistered experiments ($N = 1,316$) on the topic of whether leaders, more than nonleaders, continue to be associated with Whiteness (i.e., being categorized as White or being represented with stereotypically White qualities). Results suggest that associations between leaders and Whiteness hold up to scrutiny, but that detecting them may depend on what methods researchers employ. In particular, when researchers use direct methods of detecting racial assumptions (e.g., self-report measures), there appears to be no evidence of an association between leaders and Whiteness (Experiment 1). Yet, when researchers use more indirect methods of detecting racial assumptions (e.g., a Princeton trilogy task), an association between leaders and Whiteness readily emerges (Experiments 2 and 3). In short, although respondents refrain from freely expressing associations they may harbor between leaders and Whiteness, these associations do not appear to have dissipated with time.

Keywords: racial bias, leadership categorization theory, stereotyping, cultural defaults, leadership prototypes

Supplemental materials: <https://doi.org/10.1037/apl0001020.supp>

In 2020, the *New York Times* produced a list of the most powerful leaders in the United States. The list spanned multiple segments of society, including chief executives of the most profitable companies, presidents of the best-ranked universities, lawmakers in the current U.S. Congress, and the like. Of the 922 most powerful people identified, 80% were White (Lu et al., 2020)—a percentage that far exceeds the base rate of actually being White in the U.S. population (61.1%; U.S. Census Bureau, 2020). On the one hand, this figure reflects undeniable progress that has been made with respect to racial equity. For example, a comparison of today's U.S. Congress to the U.S. Congress of 20 years ago reveals that the percentage of non-White lawmakers has nearly doubled over time. On the other hand,

this figure reflects a continued racial gap in terms of who holds, and who does not hold, positions of social influence in the United States.

Racial gaps in the highest echelons of leadership are impossible for any single perspective to fully explain. On the one hand, racial gaps like these can be explained, at least in part, by structural forces: for example, by White individuals having disproportionate access to wealth (e.g., Kraus et al., 2019); by White individuals benefitting from systems of institutional segregation (e.g., Anicich et al., 2021); or by network effects that result in greater career opportunities for White than for non-White individuals (e.g., Pedulla & Pager, 2019). On the other hand, racial gaps like these can be explained by interpersonal forces: for example, by White managers' tendency to exhibit in-group favoritism in their personnel evaluations (e.g., Phillips & Jun, 2021); or by overt racial prejudice against racial outgroup members in the workplace (e.g., Quillian et al., 2017). The present article advocates for one such interpersonal force that may help to explain why White individuals are disproportionately represented in positions of influence—namely, that in the United States, Whiteness itself may be a prototypic attribute of leaders (Rosette et al., 2008).

According to leadership categorization theory, perceivers have a clear prototype in their minds of what a *leader* is (Lord et al., 1984, 2020). This prototype is an abstraction—a fuzzy representation of the traits, behaviors, and beliefs that define a typical leader and that

Christopher D. Petsko  <https://orcid.org/0000-0002-8795-0003>

The authors have no known conflict of interest to disclose.

Preregistration plans, data files, codebooks, survey materials, and R scripts are available on the Open Science Framework (OSF) website: <https://osf.io/p4ny8/>.

Correspondence concerning this article should be addressed to Christopher D. Petsko, Fuqua School of Business, Duke University, 100 Fuqua Drive, Durham, NC 27701, United States. Email: christopher.petsko@duke.edu

differentiate typical leaders from nonleaders. For example, people in the U.S. harbor implicit leadership theories that define leaders as dedicated, charismatic, and intelligent (Offermann & Coats, 2018; Offermann et al., 1994). When someone possesses traits that are consistent with this prototype (e.g., charisma), they are not only more likely to be categorized as leaders, but are also evaluated more positively when they actually hold positions of leadership (Foti et al., 2017; Nye & Forsyth, 1991; Quaquebeke et al., 2011). Thus, understanding the traits that perceivers regard as prototypic of leaders is itself a consequential endeavor. Indeed, understanding who is likely to seem leaderly, even if by default, can help answer the question of who is likely to acquire genuine social influence over time.

The idea that Whiteness may be a prototypic feature of leaders made its debut in the management literature nearly 15 years ago with a series of experiments suggesting that merely describing someone as a leader (vs. not) could increase their likelihood of being presumed White (Rosette et al., 2008). However, a recent set of experiments that sought to directly replicate this phenomenon has called this earlier finding into question (Ubaka et al., 2022). Of relevance to the present article, this latter set of experiments found no evidence that describing someone as a leader (vs. not) could increase their likelihood of being presumed White. The offered explanation for this pattern of findings was that increasing levels of racial diversity in the labor force—and in particular in positions of leadership—may have eliminated any tendency U.S. respondents have to associate leaders with stereotypically White attributes.

Contradictory findings such as these raise questions for management scholars. Do perceivers continue to associate leaders, more than nonleaders, with stereotypically White qualities (Rosette et al., 2008)? Or, as recent data suggest, is it the case that associations between Whiteness and leaders have dissipated with time (Ubaka et al., 2022)? Answering these questions is theoretically consequential. In particular, answering these questions is consequential for the issue of whether organizational scholars should continue to use—or whether they should perhaps jettison—leadership categorization theory as a framework for making sense of racial gaps in positions of leadership. If prototypes of leaders are no longer particularly White by default, it may be worthwhile to dedicate less scientific attention toward leadership categorization theory as an explanatory framework for racial gaps in positions of leadership, and it may be worthwhile to dedicate more scientific attention toward alternative explanatory frameworks for these racial gaps. If in contrast prototypes of leaders continue to be particularly White by default, then leadership categorization theory remains a promising theoretical framework for understanding which psychological processes give rise to these racial gaps.

The purpose of the present article is to test—across three preregistered experiments—whether Whiteness (i.e., being categorized as White or being associated with stereotypically White qualities) is a part of perceivers' prototype of leaders.¹ We propose that people in the U.S. continue to associate leaders, more than nonleaders, with stereotypically White attributes, but that detecting these associations may depend on what methods researchers employ. In particular, we suggest that participants may no longer associate leaders with the concept of Whiteness when race is explicitly mentioned or asked about, but that they may continue to associate leaders with the concept of Whiteness when these racial assumptions are measured indirectly (i.e., without participants' awareness or ability to engage in socially desirable responding). This possibility, if supported,

would help to organize the existing literature on White-leader associations, which has at times observed evidence of these associations (e.g., Gündemir et al., 2014) and which at other times has not (e.g., Ubaka et al., 2022). Most critically, the experiments in this article can weigh in on the timely question of whether White-leader associations have truly dissipated with time, or whether people have instead become unwilling—or perhaps unable—to report on these associations.

The Present Experiments, Transparency, and Openness

We conducted three preregistered experiments ($N = 1,316$) on the topic of whether leaders, more than nonleaders, continue to be associated with Whiteness (i.e., being categorized as White or being represented with stereotypically White qualities). In Experiment 1, participants' racial assumptions were measured directly. In Experiments 2 and 3, participants' racial assumptions were measured indirectly. The experiments themselves received institutional approval (Duke University Institutional Review Board; "Racial Bias in Leader Categorization"; Protocol No. 2021-0294), and they report all conditions, manipulations, and exclusions. All hypotheses, sample sizes, and analysis plans were preregistered in advance of data collection. Raw effect sizes are encompassed by 95% confidence intervals (CIs), and standardized effect sizes are reported as standard betas (β s). Statistical power for detecting main effects and interactions will be presented in the Results sections of each experiment.² Preregistration plans, data files, codebooks, survey materials, and R scripts associated with this article are available on the Open Science Framework (OSF) website: <https://osf.io/p4ny8/>.

Experiment 1

Experiment 1 employed a direct method of assessing participants' racial assumptions. Participants were instructed to read a fictitious news article about a company called Selcom, Inc. The news article itself featured an interview with a company member who was either described as a "leader" or as an "employee." In addition, the article included or did not include, by random assignment, information on the base rate of company members being White at Selcom. The original experiment that used this paradigm (Rosette et al., 2008, Experiment 1) found that participants presumed that a company member was more likely to be White when they were described as a leader (vs. when described as an employee), and that this tendency held regardless of what participants were told about company base rates. A replication experiment that used this same paradigm (Ubaka et al., 2022, Study 1), however, found that participants did *not* presume that a company member was more likely to be White when described as a leader—a finding that likewise held regardless of company base rates. Our goal with Experiment 1 was to perform our

¹ Of note, leadership prototypes can be White either because people presume that leaders are White (which we test in Experiment 1 and Experiment 2) or because people presume that leaders possess attributes that are stereotypically associated with White people (which we test in Experiment 3). Our view is that these two phenomena co-occur. As such, this article uses the term Whiteness to refer both to race as a category as well as the traits that this category implies.

² All power analyses were conducted using the "simr" package in R (Green & MacLeod, 2016), which runs Monte Carlo simulations on one's models to arrive at power estimates (for more on this technique, see Bolger et al., 2012).

own preregistered replication attempt of the original experiment on this topic (that of Rosette et al., 2008; Experiment 1). If White-leader associations are no longer detectable on this task—either because these associations have dissipated with time or because direct measures obscure their measurement—then we, like others (Ubaka et al., 2022), should find no evidence that participants think of leaders as especially White by default.

Method

Participants in Experiment 1 were randomly assigned to one condition in a 2 (target role: leader, employee) \times 3 (race base rate: no information, 50% White, 20% White) between-person experiment.

Participants

Experiment 1 had the goal of obtaining $n = 125$ people per condition. A total of $N = 752$ participants completed this study, of whom 17 (2.26%) were excluded either because: (a) they did not reply “yes” to the question, “Did you take this survey seriously?”; or (b) they did not pass at least three out of four questions that were designed to screen out nonhuman respondents. See Table 1 for demographic information on remaining participants.

Procedure

Participants were recruited online via <https://www.CloudResearch.com> (Litman et al., 2017). All participants were told that they would

Table 1

Final Participant Demographics (After Exclusions) Across All Three Experiments

Demographics	Experiment 1	Experiment 2	Experiment 3
	$N = 735$	$N = 253$	$N = 300$
Age			
<i>M</i>	40.42	39.88	42.07
<i>SD</i>	12.93	12.01	13.71
Range	18–75	21–77	19–83
Gender			
Male	50.6%	68.8%	56.0%
Female	48.2%	30.4%	43.7%
Other	01.1%	00.8%	00.3%
Race			
White	72.5%	81.8%	78.3%
Asian	11.4%	06.3%	08.7%
Black	8.7%	6.7%	5.3%
Latinx	04.8%	2.4%	04.7%
Islander	00.4%	00.4%	01.0%
Native	00.1%	00.4%	00.3%
Other	02.0%	01.6%	01.3%
Ideology			
<i>M</i>	4.13	4.64	3.98
<i>SD</i>	2.86	3.23	2.83
Range	0–10	0–10	0–10
Education			
Bachelor’s or higher	62.2%	66.8%	62.7%

Note. All participants were U.S. citizens. Employment data were not collected on the participants in these experiments (but see Supplemental Material, for projected employment demographics). *Ideology* was measured on a scale from 0 = *extremely liberal* to 10 = *extremely conservative*.

read a newspaper article, and that afterward they would respond to a variety of questions on the basis of what they had read. The newspaper article itself contained the original text from that of Rosette et al., 2008 (Experiment 1), which described either an employee or a leader who was being interviewed about Selcom, Inc.’s “project NOVA,” and which either reported Selcom’s workforce as being 50% White, 20% White, or did not include information on the racial demographics of Selcom (see materials on OSF, for exact wording). After reading the article, participants completed a variety of questions about the article, including manipulation checks.³ The critical dependent variable in this experiment was how participants responded to the question, “What do you think is the race of the person interviewed?” Response options included: Black/African American, Hispanic/Latin American, Asian/Asian American/Pacific Islander, Native American/Alaskan Native, White/European American, and other (ordering of racial groups was randomized for each participant).

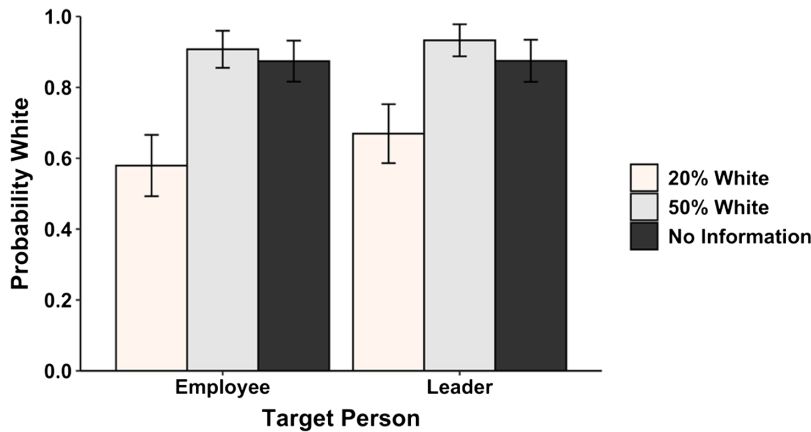
Results

Because Experiment 1 was a preregistered replication of the original study on this topic (that of Rosette et al., 2008, Experiment 1), the key hypotheses under scrutiny were those of the original article: namely (a) that participants would be more likely to presume the target person was White when described as a leader versus as an employee and (b) that this bias would hold equally across base-rate conditions. In order to examine these hypotheses, we regressed judgments of the target person’s race (coded as 1 = *White*, 0 = *non-White*), in a simultaneous regression with robust standard errors, onto contrast codes representing the 2 (target role: leader, employee) \times 3 (race base rate: no information, 50% White, 20% White) factorial design of this experiment. This analytic method is functionally equivalent to a standard 2 \times 3 analysis of variance (ANOVA), but it has the added benefit of being able to accommodate binary outcomes with greater precision than a 2 \times 3 logistic ANOVA can (see Gomila, 2021). According to Monte Carlo simulations, this experiment had more than 80% power to detect an effect of target role (leader vs. employee) as small as $\beta = 0.20$ and to detect a two-way interaction (between target role and base-rate condition) as small $\beta = 0.42$.

Were participants more likely to presume the target person was White when the target person was described as a leader (vs. as an employee)? This did not appear to be the case. There was no main effect of target role, meaning that the target person was equally likely to be presumed White regardless of whether they were described as a leader ($M = .83$, $SE = 0.02$) versus as an employee ($M = .78$, $SE = 0.02$), $M_{diff} = .04$, 95% CI $[-.02, .09]$, $\beta = 0.10$, $F(1, 728) = 1.99$, $p = .158$. Moreover, this null effect was unmoderated by base-rate condition: $F \leq 1.07$, $p \geq .301$. The only significant effect in this analysis was a significant main effect of base-rate condition (see Figure 1, for means; see Table 2, for model estimates), revealing that participants in the 20% White condition were

³ The vast majority of participants appeared to have been paying attention. For example, on average, participants estimated that 26.66% of Selcom was White in the 20% White condition, that 50.37% of Selcom was White in the 50% White condition, and that 60.35% of Selcom was White in the no base-rate condition. Notably, the conclusions reported here do not depend on whether we do or do not exclude participants who performed poorly on manipulation checks (see Supplemental Analyses, for more on this issue).

Figure 1
Probability of Presuming Leaders (vs. Employees) Are White by Default (Experiment 1)



Note. The probability that the target person (in Experiment 1) was categorized as White, broken down by whether the person was described as an employee versus leader, and as well by which of three base-rate conditions participants had been assigned. Probabilities are encompassed by 95% CIs. See the online article for the color version of this figure.

less likely to presume that targets were White ($M = .62$, $SE = 0.03$) than were participants in the 50% White condition ($M = .92$, $SE = 0.02$), $M_{diff} = -.30$, 95% CI $[-.37, -.23]$, $\beta = -0.75$, $F(1, 483) = 70.18$, $p < .001$.⁴ In short, then, this experiment failed to replicate the original finding that leaders—more than followers—would be presumed White by default.⁵ Instead, we found very simply that leaders and followers were presumed White to equal degrees.

Discussion

Consistent with what others have reported (Ubaka et al., 2022), Experiment 1 revealed no evidence that participants think of leaders as particularly White—at least not when they are asked about race explicitly. This observation goes against the original finding from this experimental paradigm (Rosette et al., 2008, Experiment 1), which documented that simply learning that someone is a “leader” (vs. an “employee”) can be sufficient to cause that person to seem

more likely to be White. Yet, as we noted previously, it is unclear why paradigms like that used in Experiment 1 fail to detect associations between leaders and Whiteness. On the one hand, it could be the case that White-leader associations have dissipated with time. On the other hand, it may be the case that—as we have been arguing—people in the U.S. continue to associate leaders with Whiteness, but that they are less willing to report on these associations than they were back in 2008. Given the importance of adjudicating between these possibilities for research on leadership categorization, Experiment 2 aimed to test these two possibilities against one another.

Experiment 2

In recent years, the reverse-correlation procedure has become a reliable method of detecting whether mental representations (of various social targets) are themselves imbued with racial content (Brown-Iannuzzi et al., 2017; Lei & Bodenhausen, 2017; Nelsen & Petsko, 2021). For this reason, Experiment 2 implemented a reverse-correlation procedure (Dotsch & Todorov, 2012). In reverse-correlation procedures, perceivers are instructed to view pairs of

Table 2
Model Estimates From Experiment 1

Fixed effects	<i>b</i>	<i>SE</i>	<i>p</i>
Intercept	.80	0.01	<.001
Target role	.04	0.03	.158
Base-rate contrast 1	.10	0.03	<.001
Base-rate contrast 2	.30	0.04	<.001
Target role × Base-rate contrast 1	-.06	0.06	.301
Target role × Base-rate contrast 2	-.06	0.07	.360

Note. Target role was contrast coded such that *leader* = 1/2 and *employee* = -1/2. Base-rate contrast 1 was contrast coded such that *no-information condition* = 2/3, *50% White condition* = -1/3, and *20% White condition* = -1/3. Base-rate contrast 2 was contrast coded such that *no-information condition* = 0, *50% White condition* = 1/2, and *20% White condition* = -1/2. This analysis was run as simultaneous regression with robust standard errors (Gomila, 2021).

⁴ Participants in the no-information condition did not differ from participants in the 50% White condition ($p = .099$). In addition, we preregistered a test of the ancillary prediction (from Rosette et al., 2008) that participants in the leader conditions, but not employee conditions, would presume Whiteness at rates that exceeded base rates. Instead, we found that both leaders and employees were presumed White at rates that exceeded base rates (all $ps < .001$).

⁵ Notably, our preregistered analytic strategy (based on the recommendations of Gomila, 2021) differs from the analytic strategy of Rosette et al. (2008, Experiment 1). However, our conclusions remain robust to which analytic strategy we employ. Subjecting our data to the same hierarchical logistic regression described in Rosette et al. (2008, Experiment 1) reveals the same pattern of effects reported above—no main effect of target role ($p = .157$), a main effect of base-rate condition ($p < .001$), and no interaction between target role and base-rate condition ($p = .713$).

black-and-white facial images that are overlaid with random visual noise. Their task is to select the face in each pair that looks most similar to a given target group (e.g., leaders). Perceivers do this across hundreds of trials. Afterward, researchers can create aggregated composite images of the faces perceivers choose during the task. These aggregated composite images reflect, in principle, how it is that perceivers mentally represent the target group(s) in question. Once created, composite images can then be rated by naïve participants on any dimension of interest. Here, we used the reverse-correlation task to examine whether leaders (vs. followers) are indeed mentally represented in ways that are seen by naïve raters as more stereotypically White.

An advantage of the reverse-correlation procedure, relative to the procedure used in Experiment 1, is that this procedure can be used to measure participants' racial associations indirectly—in this context, without explicitly mentioning the topic of race to participants themselves. Our preregistered hypothesis was that we would find evidence of an association between leaders and stereotypically White qualities. Such a pattern, if supported, would be consistent with the idea that White-leader associations have not disappeared with time so much as they have become more challenging to detect.

Method

Experiment 2 occurred in two phases. In Phase 1, participants called to mind their mental representations of either leaders or followers, by random assignment, while completing a 300-trial reverse-correlation experiment. In Phase 2, the composite images that were generated in Phase 1—which approximate participants' mental representations of leaders and followers, respectively—were rated on how stereotypically White (vs. Black) they appear.

Participants

Phase 1 had the goal of obtaining $N = 100$ people in total and Phase 2 had the goal of obtaining $N = 150$ people in total.⁶ A total of $N = 259$ people completed Experiment 2, of whom we excluded $n = 6$ (2.32%) for failing to respond “yes” to the question, “Did you take this survey seriously?” See Table 1 for demographic information on remaining participants.

Procedure

Phase 1. Participants in Phase 1, all of whom were recruited via <https://www.CloudResearch.com>, completed a standard reverse-correlation experiment that was designed to approximate how stereotypically White (vs. Black) their mental representations appeared (task adapted from Lei & Bodenhausen, 2017). In particular, participants viewed 300 pairs of blurry faces, and their task was to choose the face in each pair that looked more like a leader, or that looked more like a follower (by random assignment). The face pairs themselves were presented in a randomized order for each participant, and were generated by imbuing a racially ambiguous base image with random visual noise (Dotsch, 2016; Dotsch & Todorov, 2012; see Figure 2). The base image that was used was a morphed average of 100 White men's faces and 100 Black men's faces (Krosch & Amodio, 2014).

After the task was over, the research team computed composite images of all the faces that participants chose in each of the

experimental conditions. To create the composite image of how leaders are mentally represented, we created a morphed average of all the faces that participants chose while thinking of a “leader.” To create the composite image of how followers are mentally represented, we created a morphed average of all the faces that participants chose while thinking of a “follower.” These composite images reflect how it is that participants mentally represent the faces of leaders and followers, respectively, given the racially ambiguous base image that was used as a starting place for these representations (see Figure 3).

Phase 2. In Phase 2, new participants, who knew nothing about the composite images or where they came from, rated them on how White versus Black they appeared in a within-person experiment. In particular, Phase 2 participants (recruited from <https://www.CloudResearch.com>) viewed both images in Figure 3 in a randomized order, and they were asked, “On a scale from 1 = *Very Black (Afrocentric)* to 7 = *Very White (Eurocentric)*, where would you place the following face?”

Results

The preregistered hypothesis in Experiment 2 was that leaders would be mentally represented with Whiter facial features than followers. To examine whether this was the case, ratings of how stereotypically White (vs. Black) the composite images appeared were regressed, in a multilevel model, onto a within-person contrast code representing whether the image being rated was the leader composite versus the follower composite. This model included a random effect of participant intercept, which adjusted for the fact that image ratings were nested within person. According to Monte Carlo simulations, this analytic approach gave us more than 80% power to detect an effect of image type (leader vs. follower) as small as $\beta = 0.30$.⁷


Confirming the preregistered hypothesis, the composite image of a leader was rated as looking stereotypically Whiter ($M = 4.49$, $SE = 0.13$) than the composite image of a follower ($M = 3.55$, $SE = 0.13$), $M_{diff} = 0.94$, 95% CI [0.61, 1.26], $\beta = 0.58$, $F(1, 146) = 31.91$, $p < .001$.⁸ In addition, exploratory analyses revealed that the leader image was rated as looking significantly more stereotypically White than Black, one-sample t test against the scale midpoint of 4: $\beta = 0.30$, $t(286) = 3.82$, $p < .001$, and that the follower image was rated as looking significantly more stereotypically Black than White, one-sample t test against the scale midpoint of 4: $\beta = -0.28$, $t(286) = -3.52$, $p < .001$. Thus, Experiment 2 provided support for

⁶ Previous research indicates that $n = 50$ people per condition is sufficient to arrive at stable composite images from reverse-correlation tasks (Petsko et al., 2021).

⁷ Significance tests for multilevel models were run using the “lmerTest” package in R (Kuznetsova et al., 2017), which estimates degrees of freedom using the Satterthwaite method.

⁸ In a supplemental study (see Study S1), we found that the follower image is also rated as more babyfaced than the leader image ($p < .001$), and that the effect reported here gets larger when babyface ratings are covaried out of the model. In addition, one reviewer raised the interesting question of whether the effect reported here would hold if we were to measure stereotypic Whiteness and Blackness on separate rating scales, rather than as endpoints on the same rating scale. Study S1 supports this possibility, with the leader image being rated as “more White” ($p < .001$) and “less Black” ($p < .001$) than the follower image even when these constructs are measured independently.

Figure 2
Base Image on Which Reverse-Correlation Task Trials Were Based



Note. One of 300 possible reverse-correlation trials is depicted on the right. Each trial is a forced-choice task between two faces: one resulting from adding random visual noise to a base image, the other resulting from subtracting that same visual noise from the same base image. The base image that was used in this task, depicted on the left, is a morphed average of 100 White men's faces and 100 Black men's faces (borrowed with permission from Krosch & Amodio, 2014).

the idea that participants' prototype of a leader may be White by default (see Figure 4, for condition means; see Table 3, for model estimates). Participants not only mentally represented leaders (vs. followers) with stereotypically Whiter facial features, but they also mentally represented leaders as looking more stereotypically White than non-White.

Discussion


A major difference between Experiment 1 and Experiment 2 is that in Experiment 1, participants were explicitly asked about their racial assumptions. In Experiment 2, in contrast, racial assumptions were measured indirectly (i.e., without participants' awareness or ability to engage in socially desirable responding). Experiment 2 revealed that when participants' racial assumptions are measured with a reverse-correlation procedure, they do seem to associate leaders, more than nonleaders, with stereotypically White qualities. This finding stands in contrast to the idea that associations between leaders and Whiteness have faded with time. Instead, this finding accords with the observation that explicitly asking participants about race can obscure the measurement of racial bias (e.g., Greenwald &

Lai, 2020; Kurdi & Banaji, 2021), and that for this reason, detecting associations between leaders and Whiteness may simply require indirect methods of assessment.

Experiment 3

Experiment 3 was designed to serve as a conceptual replication of Experiment 2. However, rather than relying on a reverse-correlation procedure, Experiment 3 relied on a different indirect method of assessing participants' racial assumptions. In particular, Experiment 3 relied on a trait nomination task (sometimes called a Princeton trilogy task) of assessing participants' stereotypes about leaders and followers, respectively (Devine & Elliot, 1995; Hall et al., 2015; Katz & Braly, 1933). On this task, participants are invited to nominate whatever traits come to mind when thinking of a target group (e.g., leaders). New participants, who know nothing about the target group(s) for whom these traits were nominated, then rate these traits on their racial content. The preregistered hypothesis for Experiment 3 was that participants would nominate stereotypically Whiter traits—according to external trait ratings—when thinking of leaders than when thinking of followers.

Figure 3
Composite Images (i.e., Mental Representations) of Followers and Leaders, Respectively

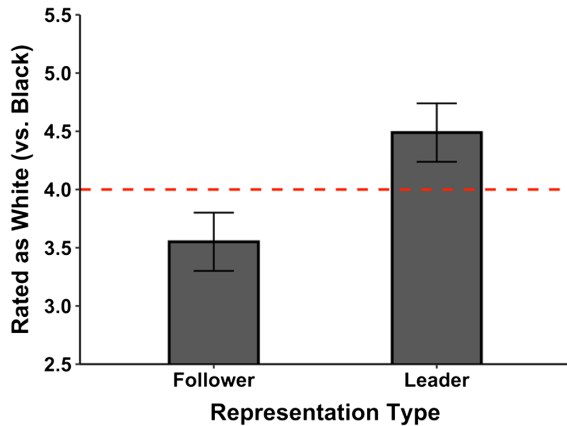


Note. Phase 1 participants' composite images of followers and leaders, respectively. These images were computed by averaging together participants' selections during the reverse-correlation procedure. Note that the base image underpinning these composite images was a morph of 100 White men's faces and 100 Black men's faces (see Krosch & Amodio, 2014).

Method

In Experiment 3, participants nominated traits that came to mind when thinking of either leaders or followers, by random assignment. Participants nominated these traits from a predetermined checklist of 99 attributes (e.g., loud, witty, aggressive; taken from Hall et al., 2015). Unbeknownst to participants, all 99 of these words had already been normed on a variety of stereotype dimensions. In particular, all 99 words had already been normed (by $N = 154$ people, as described in Petsko & Bodenhausen, 2019) on how stereotypically White they seem (from 1 = *not at all* to 7 = *very*), and on how positive versus negative they seem (from 1 = *very negative* to 7 = *very positive*). Thus, we were able to determine whether the traits participants nominated when thinking about leaders (vs. followers), on average, were those that tended to be rated by others as seeming stereotypically Whiter (and as seeming more positive vs. negative, for exploratory purposes).

Figure 4
Rated Whiteness (vs. Blackness) of Follower Versus Leader Representations (Experiment 2)



Note. Rated Whiteness (vs. Blackness) of the follower and leader composite images, respectively. The dotted line corresponds to the race-neutral midpoint of the rating scale (anchored at 1 = *very black* and 7 = *very white*). Means are encompassed by 95% CIs. See the online article for the color version of this figure.

Participants

Experiment 3 had the goal of obtaining $n = 150$ people per experimental condition. A total of $N = 305$ participants completed this study, of whom five (1.64%) were excluded either because: (a) they did not reply “yes” to the question, “Did you take this survey seriously?”; or (b) they did not pass at least three out of four questions that were designed to screen out bots. See Table 1 for demographic information on remaining participants.

Procedure

As in the previous experiments, participants from Experiment 3 were recruited via <https://www.CloudResearch.com>. Participants in this experiment learned that the research team was interested in understanding cultural stereotypes. Following the recommendations of others who have used trait nomination tasks (Devine & Elliot, 1995; Ghavami & Peplau, 2012), participants were told that the research team was not interested in their personal beliefs, but rather in how U.S. respondents *in general* would stereotype either leaders

or followers. Participants were then shown 99 attributes on a checklist (e.g., arrogant, kind, conceited, happy-go-lucky), and they were asked, first, to select all the attributes that were part of the cultural stereotype of either leaders or followers. After having gone over all 99 attributes on the checklist the first time, participants were then shown their trait nominations one more time, and they were asked to narrow their full list of trait nominations down to the 10 most stereotypic attributes of their target group. For illustrative purposes, the top 10 traits that participants selected when thinking of leaders and followers, respectively, are depicted in Table 4.

Results

Our preregistered prediction was that the final attributes participants nominated when thinking about leaders would be stereotypically Whiter—according to independent ratings (taken from Petsko & Bodenhausen, 2019)—than the final attributes participants nominated when thinking about followers. This hypothesis was examined by regressing the average Whiteness score (which could range from 1 = *not at all* to 7 = *very*) of each participant’s final trait nominations onto a contrast code representing whether traits were nominated for leaders versus followers. According to Monte Carlo simulations, this analysis had more than 80% power to detect an effect of target (leader vs. follower) as small as $\beta = 0.26$.

This analysis revealed that leaders were indeed characterized by traits that were rated (by others) as stereotypically Whiter ($M = 4.60$, $SE = 0.03$) than the traits used to characterize followers ($M = 4.00$, $SE = 0.03$), $M_{diff} = 0.60$, 95% CI [0.52, 0.68], $\beta = 1.27$, $F(1, 297) = 202.92$, $p < .001$ (see Figure 5, left-hand panel). Moreover, exploratory analyses suggested that even though participants chose much more positively valenced attributes for leaders than followers ($\beta = 0.98$, $p < .001$), and even though more positively valenced attributes tended to be rated as “Whiter”, $r(97) = .67$, $p < .001$, leaders continued to be characterized by stereotypically Whiter words than followers *even* when controlling for trait valence: $M_{diff} = 0.27$, 95% CI [0.22, 0.33], $\beta = 0.58$, $F(1, 296) = 89.31$, $p < .001$ (see Figure 5, right-hand panel; see Table 5, for all model estimates).⁹ Thus, we found support for the idea that leaders would be stereotyped as possessing stereotypically Whiter qualities than followers.¹⁰ Moreover, we found that this effect was not reducible to the fact that leaders tend to be stereotyped more positively than followers.

Discussion

Experiment 3, like Experiment 2, relied on an indirect method of assessing participants’ racial assumptions. In particular, Experiment

Table 3
Model Estimates From Experiment 2

Fixed effects	<i>B</i>	<i>SE</i>	<i>p</i>	Random effects	Var.	<i>SD</i>
Intercept	4.01	0.10	<.001	Participant intercept	0.37	0.61
Target	0.94	0.16	<.001	Residual	2.02	1.42

Note. Var. = variance estimate. Target was contrast coded such that *leader composite* = 1/2 and *follower composite* = -1/2. This analysis was run as a multilevel regression with one random effect: a random effect of participant intercept. Equality of variance between conditions was observed ($F = 0.16$, $p = .692$) according to a Levene’s test on image ratings.

⁹ Because traits were rated (in Petsko & Bodenhausen, 2019) not just on their stereotypic Whiteness and valence, but also on a variety of other racial dimensions (e.g., stereotypic Blackness, stereotypic Asianness), we were able to perform a variety of other exploratory analyses on the results from Experiment 3. See Supplemental Material for what we examined and found.

¹⁰ Reviewers raised two interesting questions about these findings: First, whether these findings would hold if we were to ask participants about their personal beliefs (rather than about societal stereotypes); and second, whether these findings would be moderated by whether participants were White versus Black. In a supplemental study (see Study S2), we found that (a) all effects hold even when participants report on their personal beliefs ($p < .001$) and (b) the White-leader effect is indeed significantly larger (interaction $p = .019$) among White respondents than among Black respondents—though notably, the effect is present among both groups (all $ps < .001$).

Table 4

Ten Most-Selected Attributes for “Followers” and “Leaders” in Experiment 3, Respectively

Follower stereotypes		Leader stereotypes	
Traits	Percent	Traits	Percent
Naïve	41%	Intelligent	57%
Suggestible	38%	Ambitious	56%
Low in intelligence	37%	Persistent	45%
Uneducated	35%	Efficient	37%
Yielding	31%	Aggressive	36%
Ignorant	30%	Arrogant	31%
Imitative	30%	Industrious	29%
Faithful	28%	Practical	26%
Conventional	25%	Talkative	25%
Lazy	23%	Passionate	23%

Note. Percent = the percentage of people within each condition who nominated a given trait.

3 relied on the checklist (i.e., Princeton trilogy) method of assessing stereotypes, which is an established technique for indexing group stereotypes in a nonreactive way (e.g., Hall et al., 2015). The preregistered hypothesis of Experiment 3 was that the traits participants nominate when thinking about leaders would be rated as more stereotypically White than the traits participants nominate when thinking about followers. Not only did Experiment 3 yield support for this prediction, but it also documented that support for this prediction holds when controlling for the fact that leaders tend to be stereotyped more positively than followers.

General Discussion


The idea that the leader prototype may be a stereotypically White prototype by default is one that gained traction nearly 15 years ago. But recent research has called this finding—whose implications may be critical for understanding racial gaps in the highest echelons of

leadership—into question. Across three preregistered experiments, we found general evidence in favor of the idea that the White-leader effect holds up to scrutiny, but that White-leader associations may be easier to detect when using methods that circumvent participants’ ability to engage in socially desirable responding. Indeed, in our first experiment (Experiment 1), which relied on asking participants directly about whether they thought leaders (vs. nonleaders) were White, we found no evidence of an association in participants’ minds between leaders and Whiteness. However, in our second two experiments (Experiments 2 and 3), we found support for the White-leader effect. Experiment 2 documented that participants’ mental representation of a leader tends to be rated as looking more stereotypically White than their mental representation of a follower. Experiment 3 documented that the attributes that participants think of when imagining a leader tend to be more stereotypically White—according to external ratings—than the attributes participants think of when imagining a follower. Collectively, these findings suggest that associations between leaders and Whiteness have not disappeared from U.S. culture so much as they have become more challenging to detect.

What are the theoretical implications of these findings? As noted previously, there are a variety of theoretical explanations that organizational scholars have used for understanding where racial gaps in positions of leadership come from—some of them socio-structural (e.g., Pedulla & Pager, 2019), and some of them interpersonal (e.g., Phillips & Jun, 2021). A major implication of our analysis is that leadership categorization theory (Lord et al., 2020) remains a promising theoretical framework for explaining the interpersonal processes that give rise to racial gaps in positions of leadership. That is, because leaders continue to be regarded as prototypically White, people who possess stereotypically White attributes may be at an advantage when it comes to being positively evaluated in the context of leadership (Foti et al., 2017; Nye & Forsyth, 1991; Quaquebeke et al., 2011).

Beyond having implications for leadership categorization theory, the present analysis is useful for at least two other reasons. First, this

Figure 5
Average “Whiteness” Rating of Traits Nominated for Followers Versus Leaders



Note. Mean Whiteness, according to external ratings, of the attributes listed for followers versus leaders. Raw effect (without controls) is depicted on the left, the same effect while controlling for trait valence is depicted on the right. Means are encompassed by 95% CIs.

Table 5
Model Estimates From Experiment 3

Fixed effects	Model 1 (no controls)			Model 2 (controlling valence)			
	<i>b</i>	<i>SE</i>	<i>p</i>	Fixed effects	<i>b</i>	<i>SE</i>	<i>p</i>
Intercept	4.30	0.02	<.001	Intercept	4.30	0.01	<.001
Target	0.60	0.04	<.001	Target	0.27	0.03	<.001
				Valence	0.33	0.01	<.001

Note. Target was contrast coded such that *leaders* = 1/2 and *followers* = -1/2. Valence was *z*-standardized. Model 1 was preregistered; Model 2 was exploratory. These models were run as linear regressions. Equality of variance between conditions was not observed ($F = 36.49, p < .001$) according to a Levene's test on trait ratings. However, all reported results remain the same even when robust standard error corrections are used (see Supplemental Material).

analysis helps to integrate divergent research findings in the literature on White-leader effects (Rosette et al., 2008). Second, these findings shed at least a little bit of light on the prototypes people harbor of followers (e.g., Sy, 2010). With respect to former, the data reported here suggest that social desirability concerns may preclude the detection of White-leader effects on direct measures, but not on indirect measures. Thus, the present findings provide an explanation as to why certain replication attempts fail to find evidence of White-leader associations (e.g., those that employ self-report measures; Ubaka et al., 2022) whereas other replication attempts do not (e.g., those that employ implicit association tests; Gündemir et al., 2014). With respect to the latter contribution, Experiments 2 and 3 suggest that prototypes of followers may, like prototypes of leaders, be racialized. Further research should be conducted to identify the contours of these prototypes. It may be the case that racial associations with followers, along with racial associations with leaders, jointly stand to reinforce the racial gaps observed in the highest echelons of organizational leadership.

Limitations and Future Directions

One major limitation of these findings is that they do not answer the question of *why* White-leader associations are detectable when measured indirectly but not directly. In particular, these findings do not answer the question of whether people in the U.S. are genuinely unaware of these associations, or whether instead they are aware of these associations but deliberately conceal them on self-report measures (e.g., Experiment 1). Future research should be dedicated to weighing in on this issue, as awareness of a bias is typically a precondition to regulating the expression of that bias. Nevertheless, we do wish to note that regardless of how (un)aware of this bias participants may be, there is good reason to believe that the biases reported here may be implicated in the processes of leadership selection and evaluation. For example, according to aversive racism theory (Dovidio & Gaertner, 2004), even those who fully reject racial bias at the explicit level—but who nevertheless harbor racial bias at the implicit level—exhibit clear patterns of racial bias in personnel selection (e.g., Dovidio & Gaertner, 2000). It is, therefore, reasonable to presume that White-leader associations—even if fully rejected at the explicit level—have some influence on leadership selection and evaluation processes. Though notably, this hypothesis requires further empirical attention.

A second limitation of the present findings is that they cannot speak to whether Whiteness is more central to leadership prototypes than other trait dimensions. While we did find evidence that White-leader associations hold even when controlling for valence differences between leader and follower prototypes (Experiment 3) and that White-leader associations likewise hold when controlling for differences in how babyfaced leaders versus followers are represented (Livingston & Pearce, 2009; see Study S1), these findings barely scratch the surface of the many ways in which prototypes of leaders and followers differ from one another. A fruitful future direction for research on leader prototypes would be to examine how much variance in leader prototypes Whiteness can explain relative to other trait dimensions (e.g., agency, competence, masculinity). Such an analysis could help to contextualize the importance of Whiteness to leader categorization processes (Lord et al., 2020).

A final limitation of the present article is that these findings cannot be generalized beyond the context of race relations in the United States. Although others have documented an implicit association between Whiteness and leaders outside the U.S. (Gündemir et al., 2014), our findings were based exclusively on U.S. samples. Future research would benefit from examining how general associations between leadership and Whiteness may be. Some perspectives—for example, the social identity theory of leadership (Hogg, 2001)—suggest that the prototype of a leader may only be White to the extent that the prototype of a nation (or of a particular social group) is itself White. Identifying the boundary conditions on White-leader associations would be immensely useful for understanding for whom and under what circumstances organizational scholars should expect racial biases in leader selection to emerge.

Concluding Remarks

The experiments presented here suggest that in the context of the United States, those who are White may be more closely aligned with leadership prototypes than those who are not. Because those who seem more prototypic of leaders are at an advantage in the context of leadership selection and evaluation (Lord et al., 2020), these findings highlight one potential pathway through which the racial gap between White and non-White individuals—as stark as it is (Lu et al., 2020)—emerges in the context of organizational leadership. Beyond that, these findings suggest that White-leader associations have not disappeared with time, as other perspectives suggest. Instead, these findings suggest that White-leader associations are very much alive, even if participants have become unwilling, or perhaps unable, to report on them.

References

- Anicich, E. M., Jachimowicz, J. M., Osborne, M. R., & Phillips, L. T. (2021). Structuring local environments to avoid racial diversity: Anxiety drives Whites' geographical and institutional self-segregation preferences. *Journal of Experimental Social Psychology, 95*(2), Article 104117. <https://doi.org/10.1016/j.jesp.2021.104117>
- Bolger, N., Stadler, G., & Laurenceau, J. P. (2012). Power analysis for intensive longitudinal studies. In M. R. Mehl & T. S. Conner (Eds.), *Handbook of research methods for studying daily life* (pp. 285–301). Guilford Press.

- Brown-Iannuzzi, J. L., Dotsch, R., Cooley, E., & Payne, B. K. (2017). The relationship between mental representations of welfare recipients and attitudes toward welfare. *Psychological Science, 28*(1), 92–103. <https://doi.org/10.1177/0956797616674999>
- Devine, P. G., & Elliot, A. J. (1995). Are racial stereotypes really fading? The Princeton trilogy revisited. *Personality and Social Psychology Bulletin, 21*(11), 1139–1150. <https://doi.org/10.1177/01461672952111002>
- Dotsch, R. (2016). *Rcicr: Reverse-correlation image-classification toolbox* (R package Version 0.3.4.1) [Computer software]. <https://cran.r-project.org/src/contrib/Archive/rcicr/>
- Dotsch, R., & Todorov, A. (2012). Reverse correlating social face perception. *Social Psychological & Personality Science, 3*(5), 562–571. <https://doi.org/10.1177/1948550611430272>
- Dovidio, J. F., & Gaertner, S. L. (2000). Aversive racism and selection decisions: 1989 and 1999. *Psychological Science, 11*(4), 315–319. <https://doi.org/10.1111/1467-9280.00262>
- Dovidio, J. F., & Gaertner, S. L. (2004). Aversive racism. In M. P. Zanna (Ed.), *Advances in experimental social psychology* (Vol. 36, pp. 1–52). Academic Press. [https://doi.org/10.1016/S0065-2601\(04\)36001-6](https://doi.org/10.1016/S0065-2601(04)36001-6)
- Foti, R. J., Hansbrough, T. K., Epitropaki, O., & Coyle, P. T. (2017). Dynamic viewpoints on implicit leadership and followership theories: Approaches, findings, and future directions. *The Leadership Quarterly, 28*(2), 261–267. <https://doi.org/10.1016/j.leaqua.2017.02.004>
- Ghavami, N., & Peplau, L. A. (2012). An intersectional analysis of gender and ethnic stereotypes: Testing three hypotheses. *Psychology of Women Quarterly, 37*(1), 113–127. <https://doi.org/10.1177/0361684312464203>
- Gomila, R. (2021). Logistic or linear? Estimating causal effects of experimental treatments on binary outcomes using regression analysis. *Journal of Experimental Psychology: General, 150*(4), 700–709. <https://doi.org/10.1037/xge0000920>
- Green, P., & MacLeod, C. J. (2016). SIMR: An R package for power analysis of generalized linear mixed models by simulation. *Methods in Ecology and Evolution, 7*(4), 493–498. <https://doi.org/10.1111/2041-210X.12504>
- Greenwald, A. G., & Lai, C. K. (2020). Implicit social cognition. *Annual Review of Psychology, 71*, 419–445. <https://doi.org/10.1146/annurev-psych-010419-050837>
- Gündemir, S., Homan, A. C., de Dreu, C. K. W., & van Vugt, M. (2014). Think leader, think White? Capturing and weakening an implicit pro-White leadership bias. *PLOS ONE, 9*(1), Article e83915. <https://doi.org/10.1371/journal.pone.0083915>
- 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. <https://doi.org/10.1177/0146167215580779>
- Hogg, M. A. (2001). A social identity theory of leadership. *Personality and Social Psychology Review, 5*(3), 184–200. https://doi.org/10.1207/S15327957PSPR0503_1
- Katz, D., & Braly, K. (1933). Racial stereotypes of one hundred college students. *Journal of Abnormal and Social Psychology, 28*(3), 280–290. <https://doi.org/10.1037/h0074049>
- Kraus, M. W., Onyeador, I. N., Daumeyer, N. M., Rucker, J. M., & Richeson, J. A. (2019). The misperception of racial economic inequality. *Perspectives on Psychological Science, 14*(6), 899–921. <https://doi.org/10.1177/1745691619863049>
- Krosch, A. R., & Amodio, D. M. (2014). Economic scarcity alters the perception of race. *Proceedings of the National Academy of Sciences of the United States of America, 111*(25), 9079–9084. <https://doi.org/10.1073/pnas.1404448111>
- Kurdi, B., & Banaji, M. R. (2021). Implicit social cognition: A brief (and gentle) introduction. In A. S. Reber & R. Allen (Eds.), *The cognitive unconscious: The first half-century*. Oxford University Press.
- Kuznetsova, A., Brockhoff, P. B., & Christensen, R. H. B. (2017). ImerTest package: Tests in linear mixed effects models. *Journal of Statistical Software, 82*(13), 1–26. <https://doi.org/10.18637/jss.v082.i13>
- Lei, R. F., & Bodenhausen, G. V. (2017). Racial assumptions color the mental representation of social class. *Frontiers in Psychology, 8*, Article 519. <https://doi.org/10.3389/fpsyg.2017.00519>
- Litman, L., Robinson, J., & Abberbock, T. (2017). TurkPrime.com: A versatile crowdsourcing data acquisition platform for the behavioral sciences. *Behavior Research Methods, 49*(2), 433–442. <https://doi.org/10.3758/s13428-016-0727-z>
- Livingston, R. W., & Pearce, N. A. (2009). The teddy-bear effect: Does having a baby face benefit black chief executive officers? *Psychological Science, 20*(10), 1229–1236. <https://doi.org/10.1111/j.1467-9280.2009.02431.x>
- Lord, R. G., Epitropaki, O., Foti, R. J., & Hansbrough, T. K. (2020). Implicit leadership theories, implicit followership theories, and dynamic processing of leadership information. *Annual Review of Organizational Psychology and Organizational Behavior, 7*, 49–74. <https://doi.org/10.1146/annurev-orgpsych-012119-045434>
- Lord, R. G., Foti, R. J., & De Vader, C. L. (1984). A test of leadership categorization theory: Internal structure, information processing, and leadership perceptions. *Organizational Behavior and Human Performance, 34*(3), 343–378. [https://doi.org/10.1016/0030-5073\(84\)90043-6](https://doi.org/10.1016/0030-5073(84)90043-6)
- Lu, D., Huang, J., Seshagiri, A., Park, H., & Griggs, T. (2020, September 9). Faces of power: 80% are White, even as U.S. becomes more diverse. *New York Times*. <https://www.nytimes.com/interactive/2020/09/09/us/powerful-people-race-us.html>
- Nelsen, M. D., & Petsko, C. D. (2021). Race and white rural consciousness. *Perspectives on Politics, 19*(4), 1205–1218. <https://doi.org/10.1017/S1537592721001948>
- Nye, J. L., & Forsyth, D. R. (1991). The effects of prototype-based biases on leadership appraisals: A test of leadership categorization theory. *Small Group Research, 22*(3), 360–379. <https://doi.org/10.1177/104649649123005>
- Offermann, L. R., & Coats, M. R. (2018). Implicit theories of leadership: Stability and change over two decades. *The Leadership Quarterly, 29*(4), 513–522. <https://doi.org/10.1016/j.leaqua.2017.12.003>
- Offermann, L. R., Kennedy, J. K., Jr., & Wirtz, P. W. (1994). Implicit leadership theories: Content, structure, and generalizability. *The Leadership Quarterly, 5*(1), 43–58. [https://doi.org/10.1016/1048-9843\(94\)9005-1](https://doi.org/10.1016/1048-9843(94)9005-1)
- Pedulla, D. S., & Pager, D. (2019). Race and networks in the job search process. *American Sociological Review, 84*(6), 983–1012. <https://doi.org/10.1177/0003122419883255>
- Petsko, C. D., & Bodenhausen, G. V. (2019). Racial stereotyping of gay men: Can a minority sexual orientation erase race? *Journal of Experimental Social Psychology, 83*, 37–54. <https://doi.org/10.1016/j.jesp.2019.03.002>
- Petsko, C. D., Lei, R. F., Kunst, J. R., Bruneau, E., & Kteily, N. (2021). Blatant dehumanization in the mind's eye: Prevalent even among those who explicitly reject it? *Journal of Experimental Psychology: General, 150*(6), 1115–1131. <https://doi.org/10.1037/xge0000961>
- Phillips, L. T., & Jun, S. (2021). Why benefiting from discrimination is less recognized as discrimination. *Journal of Personality and Social Psychology*. Advance online publication. <https://doi.org/10.1037/pspi0000298>
- Quaquebeke, N. V., Knippenberg, D. V., & Brodbeck, F. C. (2011). More than meets the eye: The role of subordinates' self-perceptions in leader categorization processes. *The Leadership Quarterly, 22*(2), 367–382. <https://doi.org/10.1016/j.leaqua.2011.02.011>
- Quillian, L., Pager, D., Midtbøen, A. H., & Hexel, O. (2017, October 11). Hiring discrimination against Black Americans hasn't declined in 25 years. *Harvard Business Review*. <https://hbr.org/2017/10/hiring-discrimination-against-black-americans-hasnt-declined-in-25-years>
- Rosette, A. S., Leonardelli, G. J., & Phillips, K. W. (2008). The White standard: Racial bias in leader categorization. *Journal of Applied Psychology, 93*(4), 758–777. <https://doi.org/10.1037/0021-9010.93.4.758>

Sy, T. (2010). What do you think of followers? Examining the content, structure, and consequences of implicit followership theories. *Organizational Behavior and Human Decision Processes*, 113(2), 73–84. <https://doi.org/10.1016/j.obhdp.2010.06.001>

Ubaka, A., Lu, X., & Gutierrez, L. (2022). Testing the generalizability of the White leadership standard in the post-Obama era. *The Leadership Quarterly*. Advance online publication. <https://doi.org/10.1016/j.leaqua.2021.101591>

U.S. Census Bureau. (2020). *Quick facts: United States* [Data set]. <https://www.census.gov/quickfacts/fact/table/US/PST045219>

Received June 30, 2021

Revision received March 18, 2022

Accepted March 24, 2022 ■