



The nature of racial superhumanization bias

Prachi Solanki & Joseph Cesario

To cite this article: Prachi Solanki & Joseph Cesario (2024) The nature of racial superhumanization bias, The Journal of Social Psychology, 164:6, 947-963, DOI: [10.1080/00224545.2023.2218995](https://doi.org/10.1080/00224545.2023.2218995)

To link to this article: <https://doi.org/10.1080/00224545.2023.2218995>

 [View supplementary material](#) 

 Published online: 30 May 2023.

 [Submit your article to this journal](#) 

 Article views: 466

 [View related articles](#) 

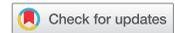
 [View Crossmark data](#) 

 Citing articles: 1 [View citing articles](#) 

 This article has been awarded the Centre for Open Science 'Open Data' badge.

 This article has been awarded the Centre for Open Science 'Open Materials' badge.

 This article has been awarded the Centre for Open Science 'Preregistered' badge.



The nature of racial superhumanization bias

Prachi Solanki  and Joseph Cesario 

Michigan State University

ABSTRACT

A superhumanization bias involves attribution of qualities that are beyond human to a certain group. Waytz and colleagues reported evidence supporting this bias among White Americans wherein Black targets were perceived as more capable of possessing superhuman qualities than White targets. We sought to better understand the nature of this effect by using different response scales (forced choice vs. Likert) and instruction sets (supporting vs. not supporting existence of superhuman abilities). Results across three studies replicate the superhumanization effect and demonstrate the necessity of several key methodological features; however, under the most realistic survey conditions (i.e. allowing unbiased decisions, being truthful about the existence of such abilities), no significant superhumanization bias emerged. Additionally, in conditions with significant bias, the size of the effect was relatively small, suggesting that this bias may not be as widespread as previously believed; indeed, only a minority of participants showed superhumanization in the predicted direction. Findings support the importance of exploring how arbitrary methodological decisions change inferences about psychological phenomena in the population.

ARTICLE HISTORY

Received 10 October 2022
Accepted 8 May 2023

KEYWORDS

racial bias; research methodology; superhumanization bias

Recent work has proposed a novel form of dehumanization of outgroups: Superhumanization, wherein a group is ascribed mental or physical abilities that are “supernatural, extrasensory, and magical” (Waytz et al., 2015). Past work on dehumanization of outgroups has provided evidence that groups may be perceived as lacking humanness by representing them as animalistic or mechanistic (see Haslam & Loughnan, 2014). Extending this work, Waytz et al. (2015) proposed that superhumanization involves the ascription of nonhuman qualities that are *beyond* human capabilities. These authors proposed two types of superhumanization: magical and physical. Magical superhumanization entails magical or mystical abilities, whereas physical superhumanization entails impossible physical strength, speed, and resilience to injury and pain. They define superhumanization as (p. 352):

The representation of others as possessing mental and physical qualities that are supernatural (transcending the laws of nature), extrasensory (transcending the bounds of normal human perception), and magical (influencing or manipulating the natural world through symbolic or ritualistic means).

Waytz et al. (2015) reported striking evidence in support of a “superhumanization bias” wherein White Americans perceived Black targets as more capable of possessing magical and/or physical superhuman qualities than White targets. Although prior work has shown that people perceive Blacks as possessing stronger physical abilities (e.g., Cottrell & Neuberg, 2005; Harris Laceywell, 2001; Payne, 2001), the findings presented by Waytz et al. are particularly striking because they claim to show that Whites perceive Blacks as able to perform *obviously impossible* actions. For instance, Waytz et al. (Study 4) reported that White participants in their studies believed that Black people were capable of *running at the speed of light* and *lifting up a building*. Assuming that participants are responding

CONTACT Prachi Solanki  solanki7@msu.edu  Michigan State University, 16789 Chandler Road, Apt 323C, East Lansing, MI 48823, USA

This article has been corrected with minor changes. These changes do not impact the academic content of the article.

 Supplemental data for this article can be accessed online at <https://doi.org/10.1080/00224545.2023.2218995>.

© 2023 Taylor & Francis Group, LLC

honestly and taking such questions literally--as intended by the researchers--the obvious response would be to reject the possibility of *any* human possessing such abilities. It is impossible to run at the speed of light or lift up a building, and any participant who comprehends the questions correctly would be compelled to answer such questions in the negative. This suggests that it might be fruitful to explore the nature of the reported superhumanization effect in more detail.¹

In this manuscript, we explore three underappreciated elements of the original superhumanization effect reported by Waytz et al. (2015). First, the study design may have introduced experimental demand characteristics. Second, the response format used to obtain superhumanization ratings may have inflated the size of the effect. Third, the statistical characteristics of the sample (e.g., a minority of strong outliers) may have led to significant results. Any of these influences, alone or in conjunction, could produce superhumanization effects even when a majority of “White participants” do not believe that such differences between Blacks and Whites exist. We test whether these factors, independently or in combination, may have played a role in producing a statistically significant superhumanization effect. To better understand each of these factors, we first present a closer look at Waytz et al.’s Study 4.

Superhumanization experiment

To explicitly measure superhumanization bias, Waytz et al. (2015) presented participants with a pair of images (Black male vs. White male; see Figure 1) and asked three questions: 1) Which of these people has supernatural strength that makes them capable of lifting up a building, 2) Which of these people is more capable of using their supernatural powers to suppress bodily needs (food, water, etc.), 3) Which of these people has supernatural quickness that makes them capable of running at the speed of light? Overall, the results indicated that participants were more likely to choose the Black target than the White target ($N = 190$, $p < .001$, $d = 0.52$).

Although the study as just described seems straightforward, several aspects warrant more attention. First, participants were given a unique set of instructions when the task was introduced, specifically:



Figure 1. Target images used in the Greyscale (top row) vs. Color (bottom row) photo conditions.

Recent research has pointed to the possible existence of superhuman or supernatural abilities in some people. That is, evidence suggests that some people may be capable of behaviors that a normal person cannot perform, including displays of superhuman physical and mental skills. Such capabilities can make these people supernatural in some sense—more like spirits, ghosts, or Gods than like human beings.

This instruction falsely informs participants that there is empirical evidence supporting the existence of superhuman abilities, introducing a demand characteristic to report the existence of *some* superhuman abilities even when participants may have believed otherwise. This would explain why participants would have been willing to endorse the existence of *any* superhuman capabilities; the racial difference may be understood by another feature of the study, discussed next.

Second, Waytz et al. (2015) used a forced-choice response format in assessing participants' superhumanization beliefs. Each participant was given three forced-choice items in which they were *required* to select either the Black or White target as certainly possessing superhuman abilities for each item. First, there is no necessary reason why participants should be forced to ascribe superhuman characteristics to the Black target *or* the White target. The clear and obvious response to such questions—“no one can do this”—is simply not an option as the researchers restricted participants' responses in a way that forced them to endorse the existence of superhuman traits toward one of the targets. Second, the odd number of items *requires* that each participant show racial bias in one direction or the other. Although sample-level bias can average out to “no bias” if equal numbers of participants show White versus Black bias, an odd number of items will necessarily produce bias in participant-level analysis. Taken together, these two features mean that even if every participant did not believe in superhumanization *and* believed that there was no racial difference, each participant was *required* to both endorse superhuman abilities *and* some racial bias in such abilities. Thus, the response format prevents accurate assessment of the supposed underlying construct, necessarily skewing that which the measure is designed to assess.

Of course, at the sample level it is possible for the group to show no racial bias in the average superhumanization value, if each participant chooses whether to show Black or White bias at random. But this counter-argument obscures two important facts. First, the measurement of superhumanization bias is an individual-level psychological construct and the individual-level measurement is misleading. When researchers prevent participants from responding in racially neutral or honest ways, broad inferences about a population or claims about psychological processes outside of those constraints (e.g., “White participants show superhumanization bias”) become questionable (see also Hughes, 1969; Ray, 1990). Second, the sample-level analysis may have revealed a Black superhumanization bias not because participants believed in superhumanization of Blacks *per se* (which is the precise argument being made by Waytz et al.) but instead because superhumanization beliefs were inferred from other beliefs, e.g., beliefs about greater physical strength of Blacks. In other words, the nature of the response options means that it is possible (1) no participant believed in any superhumanization of Blacks *or* Whites and (2) given that they were forced to endorse such beliefs, a greater number of participants used preexisting beliefs about strength differences to answer the required question. While this might be interesting on its own, it is important to note that this is *not* the argument advanced by Waytz et al., who are claiming that participants hold in their minds greater superhumanization beliefs about Blacks independent of those forced by the response items.

Third, the sample distributions could contain a minority of outliers that produced the observed statistically significant effects (see Renkewitz & Keiner, 2019). Given the aforementioned concerns, it is possible that a majority of the sample did not endorse any existence of a superhumanization bias even while Waytz et al (Study 4; Waytz et al., 2015) report an overall bias against Black targets.

In sum, Waytz et al. (2015) asked each individual participant to report the existence of impossible characteristics in a manner that prevented each participant from replying in a non-biased manner. These decisions may have caused an overestimation of superhumanization effects or caused such effects to be realized in an experimental setting that fails to reflect most participants' beliefs outside of these constrained, artificial conditions.

Current studies

Across three experiments we tested whether the above factors may have contributed to a superhumanization effect reported by Waytz et al. (2015). In Experiment 1, we replicated Waytz et al.'s Study 4 by using their instruction set and materials but extended their method with one change: In addition to the forced-choice format, we included a condition with an interval scale. We suspected that the forced-choice response format led to an inflation of the superhumanization effect and tested this by adding an interval scale for comparison.

In Experiment 2, we replicate and extend Experiment 1 by testing whether the experimental instructions influenced participants' responses. In Waytz et al. (2015) study, participants were first given false information supporting the existence of superhuman abilities. In Experiment 2, we also included a condition with (true) instructions stating that *no* empirical evidence supports the existence of superhuman abilities.

In Experiment 3, we made two further changes. First, concerned that our (true) instructions in Experiment 2 may have produced an opposing demand characteristic for participants to *not* report superhumanization, we instructed participants in a completely agnostic manner. Second, we provided the most straightforward response options possible by asking participants to directly report (using a yes/no question) if a person can perform tasks that require superhuman abilities.

Across all studies, we also included non-superhumanization measures from the original work to show that more reasonable effects could be replicated in our studies; in this way, any failures to replicate superhumanization effects could not be attributed to experimenter incompetence or sample differences. For instance, Waytz et al. (2015) found race differences in pain tolerance ratings and no race differences in "everyday capabilities" ratings (e.g., walking a dog). Compared to superhumanization items--which are entirely impossible--these effects are more plausible. We replicate the original results on these measures; but as the focus of the current work was on superhumanization bias, we report these additional analyses in the Supplementary Online Material.

Experiment 1

This study had two goals: First, to directly replicate Study 4 of Waytz et al. (2015), which found that participants gave higher ratings of superhumanization toward Black targets relative to White targets. Second, to explore the degree to which results were influenced by the response format by using a between-subjects response format manipulation (Forced-choice vs. Interval scale). The interval scale would allow participants to indicate that *neither* Black nor White targets could have superhuman abilities.

Participants & preregistration

Participants from Amazon's Mechanical Turk were given a monetary reward (\$0.50) for participation ($N = 665$; $M_{age} = 30$ years; $Male = 62.23\%$, $Female = 37.16\%$, $Other = 0.60\%$; $White = 64.04\%$, $Black = 12.15\%$, $Asian = 12.67\%$, $Hispanic/Latinx = 3.08\%$, $American Indian/Alaska Native = 4.96\%$, $Native Hawaiian/Pacific Islander = 0.68\%$, $Multiracial = 2.39\%$). Waytz et al. (2015) tested the main effect of race with 190 White participants in a single-cell design, collapsing across photo type (Greyscale vs. Color photo). The present study had approximately twice as many White participants per cell as the original Waytz et al. study ($N_{Forced\ choice \times Greyscale\ photo} = 98$, $N_{Forced\ choice \times Color\ photo} = 93$, $N_{Likert \times Greyscale\ photo} = 87$, $N_{Likert \times Color\ photo} = 96$). That is, we determined our sample size based on Waytz et al.'s study. This study was not preregistered, but all data and data exclusions (if any), all analysis scripts, all manipulations, and all measures are available as a component on the following OSF page: <https://osf.io/zpr8t>.

Procedure

This study had a 2 (Photo type: Grayscale vs. Color images) \times 2 (Response Format: Forced-choice vs. Likert scale) \times 2 (Target race: Black male vs. White male) mixed design, with the first two factors between-subjects and the last factor within-subjects. All participants provided ratings of Superhumanization as the dependent measure. Although the original Waytz et al. (2015) study found no significant differences between Grayscale vs. Color photos conditions, we included this manipulation to directly replicate the original work. Participants were randomly assigned to one of four between-subjects conditions: forced-choice format with grayscale photos, forced-choice format with color photos, Likert scale format with grayscale photos, or Likert scale format with color photos. As in the original study, upon being presented with each target image, participants were first asked to provide a description of each target and then provide ratings of superhumanization for each target. This experiment included only two pairs of target images, the same as in the original study (see Figure 1): one pair in the Grayscale photo condition and the other in the Color photo condition.

Dependent measure

Prior to providing superhumanization ratings, participants were given the (false) instructions described above that evidence exists of superhuman abilities in humans (taken directly from Waytz et al., 2015).

In the forced-choice condition, participants had to indicate which of two targets was more capable of possessing superhuman qualities. This was rated using the same three questions as in Waytz et al. (mentioned above). Reproducing the original analysis technique, choices were coded such that 0 = White and 1 = Black. Responses were then summed to create a superhumanization score for each participant, ranging from 0 to 3 (midpoint = 1.5). Individual participant scores were used to compute a group average, which was compared to the midpoint (i.e., the value assumed to reflect no racial bias in superhumanization judgments). The average score is interpreted such that scores above the midpoint (1.5) reflect greater superhumanization of Black targets compared to White targets.

In the Likert scale condition, participants rated the same three superhumanization items on an interval scale ranging from 1 to 5 (1 = *not at all able to do this*, 5 = *absolutely able to do this*) where higher ratings indicate greater judgments of superhuman abilities. Participants provided separate ratings for Black vs. White male targets. The responses to all the items were averaged separately for each target, thus allowing participants to indicate that neither Blacks nor Whites were capable of such abilities.

Hypotheses

In the forced-choice condition, we expected to replicate the Waytz et al. (2015) results for judgments of superhumanization. Specifically, we expected that Black targets would be rated significantly higher on superhuman abilities than White targets. In contrast, in the Likert scale condition, we expected that superhumanization ratings would be attenuated. This pattern would provide evidence that the superhumanization effect was inflated by the nature of forced-choice response options.

Results

To provide the most direct replication of the original work, we recreated Waytz et al. (2015) analyses with White participants only, collapsing across photo type (Grayscale vs. Color). Analyses including the photo type manipulation and including the entire sample are reported in the Supplementary Online Materials.

To test whether we replicated the effect of race on judgments of superhumanization in the forced-choice condition, we conducted the same one-sample *t*-test as Waytz et al. (2015). Specifically, we

compared the sample-average of superhumanization scores to the mid-point of the scale. Results showed that participants attributed greater superhuman abilities to Black targets compared to White targets, $M = 1.89$, $SD = 0.93$, 95% $CI [1.75, 2.02]$, $t(185) = 5.67$, $p < .001$, $d = 0.42$, 95% $CI [0.21, 0.63]$. This replicates the original Waytz et al. finding, $M = 1.94$, $SD = 0.84$, $t(189) = 7.30$, $p < .001$, $d = 0.52$, 95% $CI [0.22, 0.82]$.

To determine whether participants displayed similar race effects when assessed on an interval scale, we conducted a paired samples t-test to measure the effect of Target race (Black vs. White) on superhumanization ratings. Supporting the superhumanization account, there was a significant effect of target race, $t(182) = 3.73$, $p < .001$, $d = 0.27$, 95% $[0.08, 0.26]$ such that participants judged Black targets to have greater superhuman abilities ($M = 2.16$, $SD = 1.35$) than White targets ($M = 1.98$, $SD = 1.26$). This was the same directional effect as found in the forced-choice format, though the effect size was approximately half the size in the forced-choice format, suggesting that response format may meaningfully contribute to the size of this effect.

The superhumanization ratings in the Likert scale conditions was highly skewed with a large range (see Figure 2). About half of the sample (53.55%, $n = 98$) averaged under 1.5 on the Likert scale in response to a Black target and 57.37% of participants ($n = 105$) averaged below this point in response to a White target. This suggests that despite the sample-wide average differences in ratings, most participants showed no racial bias and indicated that both Blacks and Whites were incapable of superhuman abilities. Indeed, computing the difference in participants' Likert ratings toward Black and White targets reveals that the claim "White participants" show a superhumanization bias might be an overestimation: 62.84% ($n = 115$) of participants showed no bias in one direction or the other. In other words, the majority of the sample gave exactly the same superhumanization rating toward the Black and White targets. Indeed, on this difference score, the most extreme White superhumanization bias was 1.33, and only nine participants (4.92%) showed a Black superhumanization bias that was greater than 1.33 (see Figure 3). The cutoff of 1.33 was chosen as it helped quantify mean differences in superhumanization responses toward a Black target *over and above* a White target. The most extreme value for White targets was 1.33, thus we used this cutoff value to determine the number of participants whose responses were above 1.33 for Black targets. That is, any values greater than -1.33 were thought to indicate a higher superhumanization bias toward Black targets over and above White targets.

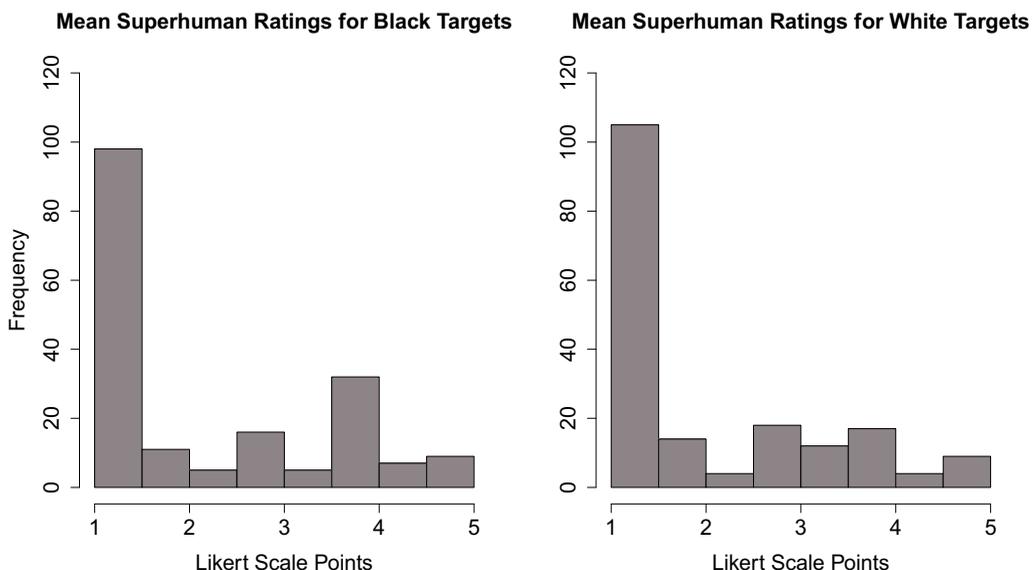


Figure 2. Histograms depicting superhumanization judgements for Black vs. White targets collapsing across the Photo Type condition on a 5-point Likert scale, (1 = not at all able to do this; 5 = very able to do this).

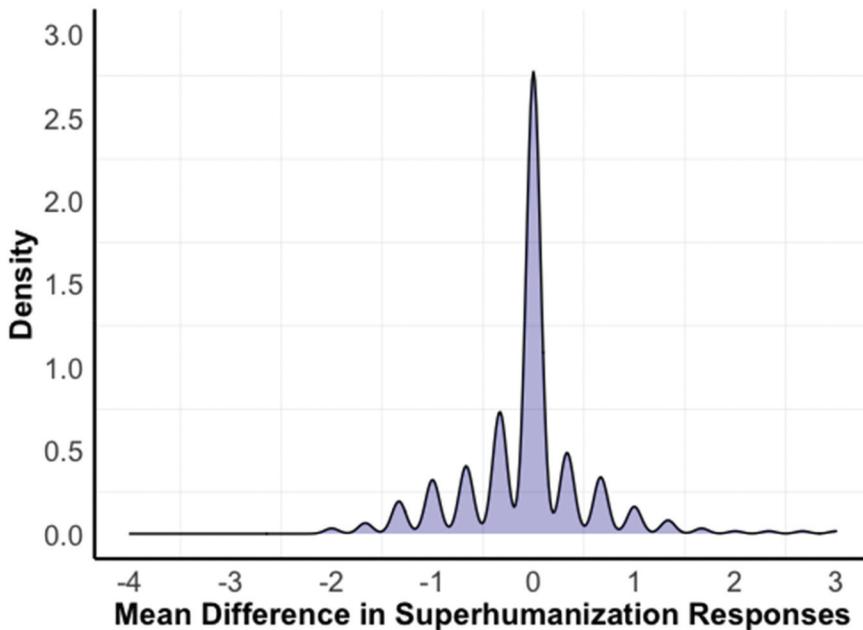


Figure 3. Density plot depicting mean difference in superhumanization judgements for Black vs. White targets in the Likert scale condition in Experiment 1. On the X axis, scores above zero indicate superhumanization bias toward Whites and scores below zero indicate superhumanization toward Blacks. Here, the most extreme value for White targets was 1.33. Hence, values above -1.33 (nine participants) indicate a higher superhumanization bias toward Black targets compared to White targets.

Discussion

We replicated and extended Waytz et al. (2015) and found superhumanization effects in both the forced-choice and Likert scale conditions. However, the size of this effect was drastically reduced in the Likert scale condition, suggesting that the forced-choice response format may meaningfully contribute to finding a superhumanization bias.

Superhumanization effects in the current study were greatly attenuated when using a Likert scale. However, there was still evidence of a superhumanization bias and it is possible that other demand characteristics were at play and inflated the obtained effect. Recall that the instructions explicitly stated that empirical evidence supports the existence of superhumanization effects. Experiment 2 tested whether these instructions influenced the superhumanization bias by explicitly manipulating instructional set.

Experiment 2

In this study, we aimed to replicate and extend findings from Experiment 1. Similar to the prior experiment, we had two response conditions (forced-choice vs. Likert scale) with participants' ratings of superhumanization as the dependent variable. We made four changes in the current experiment: First, we randomly presented one of two instruction sets: one which replicated instructions from the original Waytz et al. (2015) and our Experiment 1 (i.e., participants were told there is empirical evidence supporting superhumanization effects) or one which suggested *no* empirical evidence in support of superhumanization effects. To the extent that the original instructions introduced an experimental demand that produced a superhumanization bias, we thought the latter instructions should reduce or eliminate the obtained superhumanization bias in Experiment 1. Second, we included a "neither" option in the forced-choice condition, which would allow for unbiased responding even with that response format. Third, to ensure that the effects observed were not simply an

artifact of using a single pair of target images, we randomized the image pairs such that a new pair of target images (Black male vs. White male) was randomly presented in each question. Twenty-five image pairs matched on masculinity ($Md_{Black} = 4.48$, $Md_{White} = 4.12$) and attractiveness ($Md_{Black} = 2.90$, $Md_{White} = 2.40$) were chosen from the Chicago Face Database and presented using randomization without replacement. This was different from Experiment 1 and from Waytz et al., both of which used a single pair of target images across all conditions. Fourth, we only included color photos as there was no significant difference observed in the Greyscale vs. Color photo conditions in Experiment 1. Hence, this task had a 2 (Instruction set: Superhuman vs. No superhuman abilities) \times 2 (Response type: Forced-choice vs. Likert scale) \times 2 (Target race: Black male vs. White male) mixed design, with the first two factors between-subjects.

Participants & preregistration

Participants from Michigan State University's Psychology Department Human Subjects system received class credits for participation. Waytz et al. (2015) reported Cohen's $d = 0.52$ in their study, which suggests that 70 participants per group would yield 90% power to detect a true effect. Conducting null hypothesis significant testing to detect differences between independent groups for a small effect (i.e., $d = 0.20$) with 90% power, we would require a sample size of 858 participants. In the current experiment we collected 1161 participants in total ($Md_{age} = 19.53$ years; $Male = 22.06\%$, $Female = 77.33\%$, $Other = 0.26\%$, $Prefer\ not\ to\ say = 0.34\%$; $White = 66.98\%$, $Black = 7.92\%$, $Asian = 16.89\%$, $Hispanic/Latinx = 6.36\%$, $American\ Indian/Alaska\ Native = 0.17\%$, $Native\ Hawaiian/Pacific\ Islander = 0.17\%$, $Multiracial = 6.09\%$, $Other = 2.87\%$) to have approximately twice as many participants for each between-subjects component in our design (i.e., ~ 150 participants per cell). After removing participants who failed to complete the survey or reported that they did not pay attention during the task ($n = 73$), our final sample had 1088 participants ($N_{Forced\ choice \times Superhuman\ effect\ instructions} = 185$, $N_{Forced\ choice \times No\ superhuman\ effect\ instructions} = 178$, $N_{Likert \times Superhuman\ effect\ instructions} = 171$, $N_{Likert \times No\ superhuman\ effect\ instructions} = 198$). The original preregistration documents can be found on Open Science Framework (<https://osf.io/rnq35/>).²

Deviation from pre-registration

As per the pre-registration (linked above), we collected data from ~ 500 participants and upon conducting our analyses found large uncertainties around the estimates. Hence, we performed another round of data collection ($N = \sim 700$) to get more precise estimates and to attain the total sample size of ~ 1200 participants. Although we did not pre-register this extended data collection, the procedure and predictions remained the same as the original pre-registration, and we jointly analyzed all the data as per the pre-registered analysis plan. Data was collected from Michigan State University's Human Participation in Research/SONA research participation system. This link <https://osf.io/rnq35/>³ contains all pre-registration documents, all deviations from the pre-registered data collection plan, all sample size considerations, all data and data exclusions (if any), all analysis scripts, all manipulations, and all measures in the study.

Procedure

Participants were randomly assigned to one of four between-subjects conditions: forced-choice format with instructions supporting superhumanization, forced-choice format with instructions of no evidence for superhumanization, Likert scale format with instructions supporting superhumanization, or Likert scale format with instructions of no evidence for superhumanization. All participants provided ratings of superhumanization for Black vs. White targets. The questions assessing superhumanization as a dependent measure were like those in Experiment 1, with changes elaborated below.

Prior to providing ratings of superhumanization, this experiment randomly presented participants with one of two instructions: those suggesting the existence of a superhumanization bias (from Experiment 1) or those suggesting no such effect, as follows:

Recent research suggests that there is NO evidence supporting the existence of superhuman or supernatural abilities in some people. That is, all empirical evidence suggests that NO ONE is capable of superhuman physical and mental skills. All human beings have the normal range of abilities, and NO ONE has capabilities that make them supernatural—like spirits, ghosts, or Gods.

Dependent measure

The questions were the same as in Experiment 1 but because we added a “neither” option in the forced-choice questions, choices were coded such that $-1 = \text{White}$, $0 = \text{Neither}$, and $1 = \text{Black}$. These responses were then summed to create a superhumanization score for each participant, ranging from -3 to 3 (midpoint = 0). The average score was interpreted such that scores above the midpoint reflected greater superhumanization of Black targets compared to White targets. Like Experiment 1, in the Likert scale condition, participants rated the same three superhumanization items on an interval scale ranging from 1 to 5 ($1 = \text{not at all able to do this}$, $5 = \text{absolutely able to do this}$) where higher ratings indicated greater judgments of superhuman abilities.

Hypotheses

We had separate predictions based on the instruction set (Superhuman vs. No Superhuman abilities) and response type (Forced-choice vs. Likert scale).

Instructions supporting superhumanization effects

In the forced-choice format we predicted that participants would select “neither” as a response more often than the other options. To the extent that instructions supporting superhumanization had an effect, though, it is possible that a small racial bias in superhumanization ratings would still be observed.

In the Likert scale format, we expected participants’ responses to display a bias toward Black targets on the superhumanization items thus replicating findings from Study 1.

Instructions not supporting superhumanization effects

In the forced-choice format, we predicted that there would be no significant effect of target race on participants’ choices. In the Likert scale condition, we expected no significant effect of target race on participants’ responses for the superhumanization items.

Results

To directly replicate the original work and Experiment 1, we ran our analyses with White participants only. Analyses including the entire sample are reported in the Supplementary Online Materials.

Two separate questions can be asked in this study using the forced choice data: 1) was there a superhumanization bias (i.e., a non-zero effect) in both instruction-set conditions, and 2) did changing the instructions reduce the superhumanization bias?⁴ To answer the first question, the same one-sample t -tests as in Waytz et al. (2015) were separately conducted in each instruction set condition, comparing the average superhumanization score in each instruction condition to the midpoint of the scale. Results were significant for both the condition in which instructions supported the existence of superhuman abilities ($M = 0.31$, $SD = 1.02$, $95\% CI [0.16, 0.46]$, $t(175) = 4.05$, $p < .001$, $d = 0.30$, $95\% CI = [0.004, 0.60]$) and the condition where instructions did not support any such effect ($M = 0.19$, $SD = 1.07$, $95\% CI [0.02, 0.35]$, $t(160) = 2.21$, $p = .03$, $d = 0.18$, $95\% CI = [-0.13, 0.48]$). As can be

seen, the size of the superhumanization effect was descriptively smaller in the latter condition. To test whether this was a significant reduction, we conducted exploratory analysis using an independent samples t -test ($t(329.41) = 1.10$, 95% $CI [-0.10, 0.35]$, $p = .27$, $d = 0.12$, 95% $CI = [-0.09, 0.33]$) which showed that the strength of this effect did not significantly reduce in the forced-choice format when instructions supported a superhuman effect vs. when instructions did not support any such effect.

As with the forced choice data, two separate questions can be asked of the interval scale data: 1) was there a superhumanization bias (i.e., a non-zero effect) in both instruction set conditions, and 2) did changing the instructions reduce the superhumanization bias?⁵

We conducted paired samples t -tests to investigate whether there was a significant racial superhumanization bias in either of the instruction sets. Results suggested that the mean responses for Black targets ($M = 2.52$, $SD = 1.12$) and the mean responses for White targets ($M = 2.42$, $SD = 1.05$) in the condition where instructions supported the existence of superhuman abilities were not, technically, different in terms of statistical significance, $M_{diff} = 0.10$, 95% $CI [-0.01, 0.20]$, $t(163) = 1.88$, $p = .06$, $d = 0.15$, 95% $CI = [-0.16, 0.45]$. Similarly, and more clearly, the difference between Black ($M = 2.11$, $SD = 1.01$) and White ($M = 2.06$, $SD = 0.97$) targets did not differ from each other in the condition where instructions did *not* support the existence of such abilities, $M_{diff} = 0.04$, 95% $CI [-0.04, 0.13]$, $t(186) = 1.01$, $p = .31$, $d = 0.07$, 95% $CI = [-0.21, 0.36]$.

To see whether there was an effect of instruction set on the size of the superhumanization bias, we further conducted exploratory analysis using a two-way ANOVA with Target race (Black vs. White) and Instruction set (Superhuman effect vs. No Superhuman effect). First, the interaction between target race and instruction set was not significant ($F(1, 349) = 0.72$, $p = .40$, $d = 0.05$, 95% $CI = [-0.26, 0.16]$) indicating that there was no significant difference in ratings between Black and White targets across the instruction types. There was a main effect of target race, $F(1, 349) = 4.29$, $p = .04$, $d = 0.22$, 95% $CI = [0.01, 0.43]$ such that participants judged Black targets to have greater superhuman abilities ($M = 2.30$, $SD = 1.08$) than White targets ($M = 2.23$, $SD = 1.02$). There was also a main effect of instruction set, $F(1, 349) = 13.60$, $p < .001$, $d = 0.40$, 95% $CI = [0.18, 0.61]$ such that responses in the condition with instructions supporting superhuman abilities were higher ($M = 2.47$, $SD = 1.09$) than responses in the condition where instructions did not support the existence of superhuman abilities ($M = 2.09$, $SD = 0.99$).

In summary, we did replicate the superhuman effect but there was either an attenuation of the superhumanization effect or no significant superhumanization effect under the most realistic conditions with relatively fewer demand characteristics—for example, adding “neither” as an option in the forced-choice condition, including instructions which did not support the existence of the bias, and using an interval scale as opposed to a forced-choice format. As in Experiment 1, the superhumanization ratings in the Likert scale condition were skewed with a large range (see Figure 4) with about 32% participants on average ($n = 112$) selecting under 1.5 on the Likert scale in response to a Black target and about 31% participants on average ($n = 109$) choosing this point in response to a White target. So, despite the sample-wide average differences in ratings, a subset of participants (although not a majority) appeared to show no racial bias when judging superhuman abilities in Black versus White targets. Indeed, exploratory analysis comparing participants’ Likert ratings toward Black and White targets (by subtracting mean Black ratings from mean White ratings) indicated that 48.72% ($n = 171$) showed no bias in one direction or the other. That is, nearly half of the sample gave the same superhumanization ratings toward Black and White targets but more than half gave superhumanization ratings toward one race or another. On this difference score, the most extreme Black superhumanization bias was 2.00 and *three* participants (0.85%) showed a *White* superhumanization bias that was greater than 2.00 (see Figure 5). The cutoff of 2.00 was chosen as it helped quantify mean differences in superhumanization responses toward a Black target *over and above* a White target. However, here we found that relatively more participants displayed a bias toward White target than Black targets. The most extreme value for Black targets was 2.00, thus we used this value to determine the number of participants whose responses were above 2.00 for White targets. That is, any values greater than -2.00 were thought to indicate a higher superhumanization bias toward *White* targets *over and above* Black targets. Recall that in Experiment 1, nine participants had a higher *Black* superhuman bias which was more extreme than the highest White superhumanization

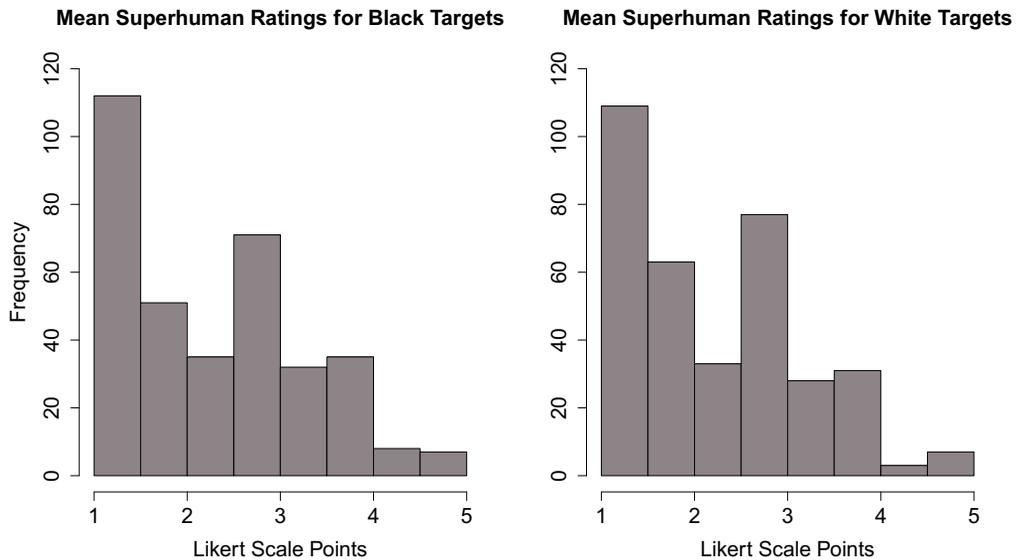


Figure 4. Histograms depicting mean superhumanization judgements for Black vs. White targets on a 5-point Likert scale (1 = not at all able to do this; 5 = very able to do this).

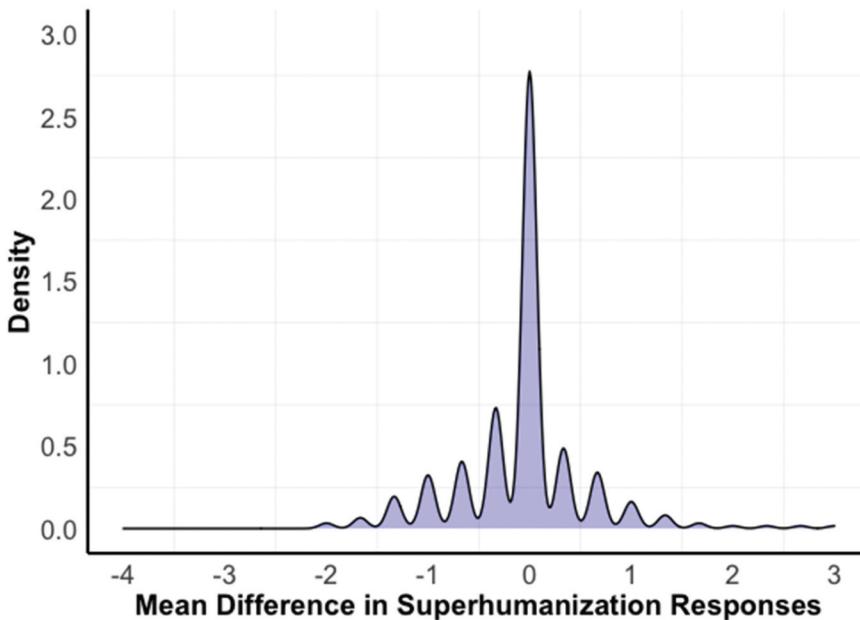


Figure 5. Density plot depicting mean difference in superhumanization judgements for Black vs. White targets in the Likert scale condition in Experiment 2. On the X axis, scores above zero indicate superhumanization bias toward Whites and scores below zero indicate superhumanization toward Blacks. Here, the most extreme value for Black targets was -2.00 . Hence, values above 2.00 (three participants) indicate a higher superhumanization bias towards White targets compared to Black targets.

bias. In comparison, participants in Experiment 2 gave more extreme ratings to White targets than Black targets. Thus, we replicated findings from Experiment 1 and provided evidence which suggests that the superhumanization bias effects might exist within a small subset of people but also that the effect size reported in Waytz et al. (2015) may be somewhat inflated.

Discussion

We replicated and extended Waytz et al. (2015) and Experiment 1 and found superhumanization effects in both the forced-choice and Likert scale conditions. In the forced choice condition, the superhuman bias was attenuated in the condition where participants were given instructions that did not support the existence of such a bias ($d = 0.18$) compared to the condition where instructions supported the existence of such a bias ($d = 0.30$). When using an interval scale, the size of the effect was smaller compared to Waytz et al ($d = 0.52$; Waytz et al., 2015) when looking at the main effect of target race ($d = 0.22$) and instruction set ($d = 0.40$). These findings suggest that response format and instruction set may meaningfully contribute to finding a superhumanization bias.

However, we still found that many participants (51.28%) responded in a way that might be biased toward one race or another. Further, instructing participants that there is no scientific evidence supporting the existence of a superhumanization bias might have introduced a different demand characteristic than Waytz et al.'s experiment in that it nudged participants to respond in a racially unbiased way. To test these possibilities, we conducted another experiment where we asked participants to respond to the superhuman questions in a relatively more straightforward manner.

Experiment 3

In this study, we aimed to replicate and extend findings from Experiments 1 and 2. Like in Experiment 2, participants answered questions about superhuman abilities and saw Black vs. White target matched on masculinity and attractiveness. However, we made two key changes. First, we provided completely agnostic instructions prior to the rating task so as not to lead participants to answer in the affirmative or negative. Recall that in Experiment 2, we instructed participants that no humans could perform superhuman tasks. This may have produced the opposing demand characteristic to that produced by Waytz et al.'s instructions. Second, we changed the response option to be a simple yes/no response format as a means of getting the most straightforward measure of participants' beliefs about this bias. Hence, this study had a between-subjects' design with target race as the between-subjects factor (Black vs. White) and participants' ratings of superhumanization as the dependent variable.

Participants & preregistration

Participants from Michigan State University's Psychology Department Human Subjects system received class credits for participation. To detect a true mean difference between two independent groups at 95% power and Cohen's $d = 0.52$ (as reported by Waytz et al., 2015), we would need to recruit 162 participants total. However, to account for a ~ 15% attrition rate, we collected data from about 200 participants total. After removing participants who failed to complete the survey or reported that they did not pay attention during the task ($n = 6$), our final sample had 192 participants ($Md_{age} = 19.35$ years; $Male = 19.27\%$, $Female = 80.21\%$, $Prefer\ not\ to\ say = 0.52\%$; $White = 67.18\%$, $Black = 8.33\%$, $Asian = 15.62\%$, $Hispanic/Latinx = 3.13\%$, $Middle\ Eastern = 3.64\%$, $Other = 1.56\%$, $Prefer\ not\ to\ say = 0.52\%$). This study was preregistered on Open Science Framework where we report how we determined our sample size, procedure and predictions, all data and data exclusions (if any). In addition, we report all analysis scripts, all manipulations, and all measures in the study (<https://osf.io/a3zxy/>).⁶

Procedure

Participants were randomly assigned to one of two between-subjects conditions: Black vs. White targets. All participants provided ratings of superhumanization for either Black or White targets. The questions assessing superhumanization as a dependent measure were like those in Experiment 1 and 2, with changes elaborated below.

Dependent measure

Prior to providing ratings of superhumanization, this experiment presented participants with instructions as follows:

This is a very simple and straightforward task. We are going to show you a picture of a person. We are then going to ask you whether that person can do some behavior. Some of these behaviors might be things that humans can do. Some of these behaviors might be things that humans can't do. You will simply indicate whether the person pictured CAN or CANNOT do the behavior listed. Just answer honestly. There are no tricks, hidden manipulations, or anything else. The questions are all straightforward and direct, and we simply want your honest and direct answers.

The questions were the same as in Experiment 1 and 2 but we changed the options to yes/no choices, coded such that 1 = Yes and 0 = No. Responses were then summed to create a superhumanization score for each participant, ranging from 0 to 3, where higher ratings indicate greater judgments of superhuman abilities.

Hypothesis

We hypothesized that there would be no significant race effect in participants' judgments of a target's abilities—in either the superhumanization items or the everyday capabilities items.

Results

We tested whether participants reported belief in the existence of a superhumanization bias when directly asked if Black vs. White targets could perform superhuman tasks. To test whether the difference in responses given target race was statistically significant, we ran an independent-samples t-test and found that the difference in responses was indeed significant, $t(176.48) = 2.26$, $p = .02$, $d = 0.34$, 95% CI [0.05, 0.62] such that “yes” responses toward Black targets ($M = 0.36$, $SD = 0.69$) were higher than responses toward White targets ($M = 0.17$, $SD = 0.50$; see Figure 6).

Thus, we replicated findings from the prior two experiments and Waytz et al. (2015) but provided evidence which suggests that the superhumanization bias effects might only exist in a small number of people. Indeed, in a sample of almost 200 participants, there were only *seven* more participants who endorsed superhumanization bias (versus participants who did not endorse a bias in either direction) of Blacks versus Whites.

General discussion

Superhumanization involves the denial of human characteristics and assumes possession of magical qualities or qualities beyond human capabilities. These perceptions support age-old beliefs about extraordinary strength, physicality, and aggression associated with Black people and suggest that people hold relatively above-average perceptions of strength with African Americans (Cottrell & Neuberg, 2005; Harris Lacewell, 2001; Payne, 2001). The aim of these studies was to better understand the nature of the reported superhumanization bias by White participants toward Black targets. Under the most exact conditions resembling those originally used by Waytz et al. (2015), we replicated the superhumanization bias, though the effect size in Experiments 1, 2, and 3 ($d = 0.42$, $d = 0.30$, $d = 0.34$ respectively) were smaller than the original ($d = 0.52$).

Importantly, we obtained evidence that under the most realistic conditions this superhumanization bias may be inflated by two methodological choices. That is, minimal evidence of a superhumanization bias emerged when more appropriate response options were used and when participants were given truthful information about the nature of superhuman abilities. Each of these two factors appeared to contribute some effect to the original demonstration, as the manipulations reduced the effect sizes

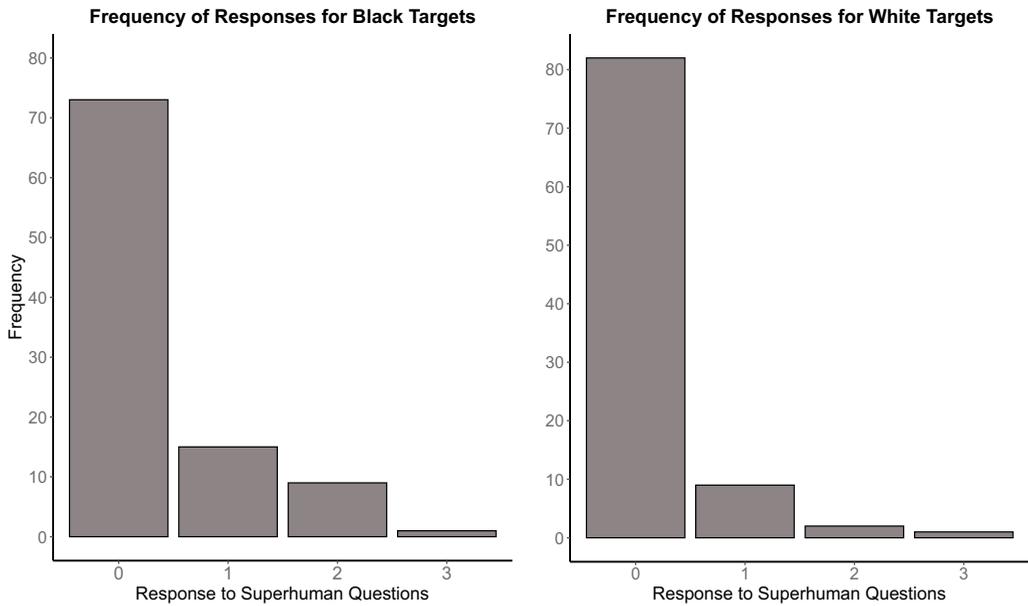


Figure 6. Bar graphs depicting frequency of superhumanization judgements for Black vs. White targets on a yes/no scale (1 = Yes, 0 = No). Scores above zero indicate superhumanization bias.

across the three studies based on response format and instructional set (see Table 1 below). We did not further explore the interactions between response scale and instruction set in the present work, our point remains that they may have contributed substantially to the superhumanization effect discussed in the original work.

Once these artifacts are accounted for, looking at the distributions of responses – rather than just sample averages – we see that the bias replicates but within only a subset of participants, questioning the original effect size in Waytz et al. (2015).

Some critiques of this work may include that we replicated only a single study (when other studies in the original work indicate support for a superhuman bias) and introduced a demand characteristic (in Experiment 2). Indeed, Waytz et al. (2015) presented several studies supporting the existence of a superhumanization bias but here we focused on Study 4 because it included an explicit measure of superhumanization bias along with other items (like everyday capabilities, pain tolerance). By including these additional items, we wanted to ensure there were no issues with our overall replication procedures if the original superhuman effects failed to replicate. Although other methods are included in the original work (methods which could be interpreted as support for superhumanization effects), we recommend caution in forming strong conclusions from such

Table 1. Effect sizes across conditions in Experiments 1, 2 and 3.

Experiment	Response type	Condition	Cohen's <i>d</i>
Waytz et al. (Study 4)	Forced-choice	NA	0.52*
#1	Forced-choice	NA	0.42*
	Likert scale	NA	0.27*
#2	Forced-choice	Instructions: Superhumans exist	0.30*
	Likert scale	Instructions: No superhumans exist	0.18*
		Instructions: Superhumans exist	0.15, <i>ns</i>
		Instructions: No superhumans exist	0.07, <i>ns</i>
#3	Yes/No	NA	0.34*

Note. * indicates statistical significance, *ns* indicates lack of statistical significance.

studies. For instance, in Study 1 participants are given the Implicit Association Test (IAT) and results suggested that superhuman qualities were moderately associated with Black people. However, the meaning of associations assessed with implicit measures is contended and unclear. For instance, believing that *others* hold beliefs about Black Americans having superhuman qualities (while disavowing such beliefs) could equally produce such associations. Thus, the mere presence of “other studies” on the topic does not necessarily provide compelling support for the claim of a superhumanization bias.

Another concern could be that by instructing participants “no superhuman bias exists” may have introduced a different demand characteristic which encouraged giving neutral responses rather than racially biased responses (in Experiment 2) thus suppressing any superhuman bias effects. As there is no evidence of humans possessing superhumanization abilities in the real world, we argue that holding a belief against such biases is the realistic, default view rather than a demand characteristic. In any case, we addressed this issue in Experiment 3 by removing instructions that may have created a demand characteristic.

Prior work has found that White participants provide different judgments of pain tolerance toward Black targets versus White targets (e.g., Hoffman et al., 2016; Waytz et al., 2015), and indeed we replicate these effects in the present studies. This work may have important implications for understanding perceptions of minorities and highlighting the impact of racial bias in judgments across various domains, especially medical decision-making. It is important to note, however, that most existing studies have failed to examine Black targets in isolation (i.e., superhumanization in Black targets has been examined only in comparison to White targets), which leaves the extent (and existence) of this bias unclear. Thus, replication of past and current findings with more diverse samples and within different experimental setups is required to determine the impact of superhumanization bias (if any). Because we did not collect demographic data (e.g., personality, political attitudes), we cannot draw firm conclusions about what caused biased responding in the small subsample that displayed superhumanization tendencies toward Black targets. This further highlights the importance of future replications in this area.

The current set of studies highlight the importance of reporting empirical and statistical findings in a more nuanced manner so as not to suggest the existence of an effect in the general population when, in fact, only a small number of people within the sample display a certain effect under highly arbitrary and laboratory-specific conditions. Empirical claims supporting general existence of a superhumanization bias (or any other bias) toward people belonging to a certain race can lead to false perceptions about the number of people who endorse such biases and should be made with caution (e.g., Hoffman et al., 2016). That is, researchers must focus on utilizing sensible guidelines when designing experiments, analyzing data, and reporting findings when making decisions to conduct any empirical work (see Harder, 2020; Steegen et al., 2016). Here, we focused on Waytz et al.’s Study 4 which used a forced-choice response format. Although this response format is common in the field, it might have been an unjustified decision that led to inflated superhumanization estimates. Indeed, note that in the original study, Waytz and colleagues compare the sample average to the midpoint of 1.5, which is explicitly defined as unbiased responding. However, a value of 1.5 was *literally impossible* for any participant to register, meaning that the ideal, unbiased response simply could not be achieved by any participant. This is not true for any theoretical or principled reasons but only because the methodology chosen artificially requires biased responding. This implies that there are several reasonable ways to design experiments and conduct statistical analyses. The current replication highlights the importance of making justifiable choices and increasing transparency in current research practices.

Notes

1. Waytz et al. (2015) present several studies supporting the existence of a superhumanization bias but here we focus on Study 4. We address other studies in this research area in the General Discussion.

2. Note that time and date stamps appearing on this OSF webpage do not reflect the actual times and dates of pre-registration because the pre-registration documents had to be anonymized and reuploaded for peer-review (which inadvertently changed the stamps). Original pre-registration time and date stamps can be viewed upon opening each document on this OSF page.
3. Note that this webpage contains the documents included in the original pre-registration as well as documents about any deviation from the original pre-registered data collection plan. The original pre-registration contained only 3 documents titled “Overview.docx”, “Data Collection Plan.docx”, and “Procedure and Predictions.docx”. All additional documents regarding data collection were not pre-registered but were subject to the same procedure, predictions, and analysis plan as the original pre-registration.
4. However, we want to note that this second question was an exploratory analysis, and we did not include it in our pre-registration.
5. Again, this was an exploratory question which was not pre-registered.
6. Note that time and date stamps appearing on this OSF webpage do not reflect the actual times and dates of pre-registration because the pre-registration documents had to be anonymized and reuploaded for peer-review (which inadvertently changed the stamps). Original pre-registration time and date stamps can be viewed upon opening each document on this OSF page or by viewing documents using the pre-registration link in the prior sentence.

Disclosure statement

No potential conflict of interest was reported by the author(s).

Funding

The author(s) reported there is no funding associated with the work featured in this article.

Notes on contributors

Prachi Solanki is a rising fifth-year graduate student at Michigan State University. Broadly, she is interested in judgement, decision-making, and social cognition research. Her other research interests include promoting open science practices and replications in psychological research.

Joseph Cesario is professor of psychology at Michigan State University. His social cognition research includes stereotyping, decision-making, and automaticity.

ORCID

Prachi Solanki  <http://orcid.org/0000-0001-9139-2836>

Joseph Cesario  <http://orcid.org/0000-0002-1892-4485>

Data Availability Statement

The data that support the findings of this study are available from the corresponding author, upon reasonable request.

Open scholarship



This article has earned the Center for Open Science badges for Open Data, Open Materials and Preregistered. The data and materials are openly accessible at <https://osf.io/rnq35/> and <https://osf.io/nc3f5>

References

- Cottrell, C. A., & Neuberg, S. L. (2005). Different emotional reactions to different groups: A sociofunctional threat-based approach to "Prejudice". *Journal of Personality & Social Psychology*, 88(5), 770–789. <https://doi.org/10.1037/0022-3514.88.5.770>
- Harder, J. A. (2020). The multiverse of methods: Extending the multiverse analysis to address data-collection decisions. *Perspectives on Psychological Science*, 15(5), 1158–1177. <https://doi.org/10.1177/1745691620917678>
- Harris Lacewell, M. (2001). African American political attitudes and the myth of black women's strength. *Women & Politics*, 23(3), 1–33. https://doi.org/10.1300/J014v23n03_01
- Haslam, N., & Loughnan, S. (2014). Dehumanization and inhumanization. *Annual Review of Psychology*, 65, 399–423. <https://doi.org/10.1146/annurev-psych-010213-115045>
- Hoffman, K. M., Trawalter, S., Axt, J. R., & Oliver, M. N. (2016). Racial bias in pain assessment and treatment recommendations, and false beliefs about biological differences between blacks and whites. *Proceedings of the National Academy of Sciences*, 113(16), 4296–4301. <https://doi.org/10.1073/pnas.1516047113>
- Hughes, G. D. (1969). Some confounding effects of forced-choice scales. *Journal of Marketing Research*, 6(2), 223–226. <https://doi.org/10.1177/002224376900600214>
- Payne, B. K. (2001). Prejudice and perception: The role of automatic and controlled processes in misperceiving a weapon. *Journal of Personality & Social Psychology*, 81, 181–192. <https://doi.org/10.1037/0022-3514.81.2.181>
- Ray, J. J. (1990). Acquiescence and problems with forced-choice scales. *The Journal of Social Psychology*, 130(3), 397–399. <https://doi.org/10.1080/00224545.1990.9924595>
- Renkewitz, F., & Keiner, M. (2019). How to detect publication bias in psychological research: A comparative evaluation of six statistical methods. *Zeitschrift für Psychologie*, 227(4), 261–279. <https://doi.org/10.1027/2151-2604/a000386>
- Steege, S., Tuerlinckx, F., Gelman, A., & Vanpaemel, W. (2016). Increasing transparency through a multiverse analysis. *Perspectives on Psychological Science*, 11(5), 702–712. <https://doi.org/10.1177/1745691616658637>
- Waytz, A., Hoffman, K. M., & Trawalter, S. (2015). A superhumanization bias in whites' perceptions of blacks. *Social Psychological & Personality Science*, 6(3), 352–359. <https://doi.org/10.1177/1948550614553642>