Spring 2017

Does Racial Bias in the Identification of Threatening Stimuli Generalize to Older Black Men?

Gustav Lundberg

Comments

This thesis was revised and published as:
Lundberg GJW, Neel R, Lassetter B, Todd AR (2018) Racial bias in implicit danger associations generalizes to older male targets. *PLoS ONE* 13(6): e0197398. https://doi.org/10.1371/journal.pone.0197398

DOES RACIAL BIAS IN THE IDENTIFICATION OF THREATENING STIMULI GENERALIZE TO OLDER BLACK MEN?

by

Gustav Lundberg

A thesis submitted in partial fulfillment of the requirements for graduation with Honors in the Psychology

Rebecca Neel
Thesis Mentor

Spring 2017

All requirements for graduation with Honors in the Psychology have been completed.

J. Toby Mordkoff
Psychology Honors Advisor

Does Racial Bias in the Identification of Threatening Stimuli Generalize to Older Black Men?

Gustav John Wilhelm Lundberg

University of Iowa

Abstract

We examined whether implicit stereotypes of younger Black men as physically threatening extend to older Black men. In Experiment 1, participants categorized objects as weapons or tools, following briefly presented prime images of men who varied in age (younger, older) and race (Black, White). In Experiment 2, we used new prime images of younger and older Black and White men, and participants categorized words as threatening or safe. Results revealed robust racial biases in object and word identification, replicating prior research: Threatening stimuli were more quickly and accurately identified after Black primes, whereas non-threatening stimuli were more quickly and accurately identified after White primes. Process-dissociation analyses further indicated that these effects were entirely driven by racial biases in automatic processing. Prime age did not moderate any of these effects, suggesting that implicit threat-based racial biases commonly evoked by younger Black men appear to extend to older Black men.

Keywords: age; process dissociation; stereotyping; threat; weapon identification task

The central focus of the current work was to examine whether implicit stereotypes of Black men as threatening generalize to older Black men. The generalization of stereotypes across categories has important implications for our understanding of threat identification and for developing more effective interventions to reduce implicit cognitive biases related to race.

Stereotypes of older adults suggest that older Black men will be seen as less threatening than younger Black men. Older adults are typically stereotyped as warm, or unlikely to want to harm others, and they are stereotyped to be incompetent, or unable to act on any positive or negative intentions they may have (Fiske, S. T., Cuddy, A. J., Glick, P., & Xu, J., 2002). This suggests that older adults may be liked and pitied, but not feared. Other work has similarly found that participants had a harder time identifying anger on the faces of elderly Black men, suggesting that old age may neutralize the threat commonly attributed to Black men (Kang & Chasteen, 2009). In fact, the prototypical older adult may be a slow, timid, and fragile person (Nelson, 2004). Most importantly, older adults are thought to be less agentic (Kite et al., 1991), suggesting that they may be seen as too weak to present a credible threat. A common portrayal of older adults is as the helpless victims of crime and violence. One might therefore expect that older adults should not be seen to pose a threat, regardless of the race of the person in question. On this account, older Black men should be seen as less threatening than younger Black men, because in latter case, their old age might outweigh the threat they would otherwise be perceived to pose as a Black man.

Will older adults of all races be seen as non-threatening? Black men are often stereotyped as being hostile, violent, and posing an imminent threat to perceivers' physical safety (Devine, 1989). Similarly, relative to White men, Black men are more likely to be seen as angry (Becker et al., 2010; Hugenberg & Bodenhausen, 2003). One consequence of these associations is that objects are more likely to be misidentified as threatening when held by a Black man or seen after an image of a Black man's face (Correll et al., 2002; Payne, 2001). Because older adults are stereotyped as non-threatening, racial biases may diminish for older adult targets. However, an alternative possibility is that racial bias will persist for older adult targets. Research conducted by Todd, Thiem, and Neel (2016) suggests that the stereotype of Blacks as violent and criminal may indeed be applied no matter what the target's age. Four experiments found that *young* age does not disarm the threat that Black men commonly evoke and that negative stereotypes of Black men extended even to young children. This work suggests that Black males may be viewed as hostile and violent regardless of their age.

The aim of the current research was to examine whether implicit stereotypes of Black men as violent and criminal also generalize to older Black men. To test this, we used variations of Payne's (2001) weapon identification task. In Experiment 1, participants categorized objects as weapons or tools after briefly being presented with faces of younger or older Black and White men. Experiment 2 was a conceptual replication, with two primary changes. First, instead of categorizing objects as guns or tools, participants categorized words as threatening or safe. Second, to avoid stimulus selection bias, we selected a new set of face images from online sources using the same criteria that were used to select the primes from Experiment 1. During both experiments we measured response times (RTs) and error rates. In addition, we used the process dissociation procedure (PDP) to estimate the degree to which automatic or controlled processes might be affecting participant's performance. We were particularly interested in answering two questions: (1) Is the racial bias smaller for older adults than young adults? (2) Does the effect hold for threat-related words as well as for objects?

Because previous research found that age (i.e., youth) does not diminish the application of racial stereotypes, we predicted that old age would not reduce racial bias in the identification of threatening objects and words. In Experiment 1, we predicted that participants would respond faster and with fewer errors in identifying guns after seeing images of Black men than after seeing images of White men, and that they would show the converse response pattern in identifying tools. In Experiment 2, we predicted that a similar pattern of racial bias would emerge when identifying threatening and safe words. For both experiments, we also predicted that age would not moderate racial bias – that is, the magnitude of the racial bias would be comparable across target age.

Experiment 1

Method

Undergraduates (N = 192) participated for course credit. A computer malfunction resulted in data loss for 1 participant. Data from 10 additional participants with below-chance accuracy (errors on >50% of trials) on the weapon identification task were excluded from analyses, leaving a final sample of 181 participants (114 women, 63 men, and 4 unreported; 140 White, 3 Black, 17 Latinx, 13 Asian, 4 who reported more than one race/ethnicity, and 4 unreported).

As part of a study investigating “rapid categorization judgments under distracting conditions,” participants saw two images flash in quick succession on the monitor. They were told to ignore the first image (prime), which always displayed a face, and to quickly and accurately categorize the second image (target object) as either a gun or a tool by pressing one of two response keys (counter-balanced across participants). The primes were 40 head-and-shoulder photos of younger and older White and Black men (10 of each category combination) taken from

Gawronski et al. (2010), all of which were standardized in size and facial expression. The target objects were 6 gun images and 6 tool images (e.g., screwdriver) taken from Payne (2001).

The sequence of each trial was as follows: blank screen (500 ms), face prime (200 ms), target object (200 ms), and pattern mask (remained on screen until participants responded). A message (*Please respond faster!*) appeared for 1 s if participants did not respond within a 500 ms response deadline, after which the next trial began. Participants completed a total of 288 randomly-ordered experimental trials. Sixteen practice trials preceded the experimental trials.

Analyses and Results

Descriptive statistics (mean RTs, error rates, and PDP estimates) for all prime–target combinations in Experiment 1 are displayed in Table 1. For both experiments, we report the results most pertinent to our primary hypotheses; additional results appear in the Appendix.

RTs. Prior to analysis, trials with errors (incorrect object identifications) and RTs <100 ms were excluded. RTs >500 ms were coded as errors, and 500 ms was used as the upper bound for RT inclusion. Although the RTs were log-transformed to reduce positive skew (Payne, 2001), raw RTs are reported in the Tables to improve the ease of interpretation.

A 2 (Prime Age) × 2 (Prime Race) × 2 (Target Object) repeated-measures ANOVA on the RTs yielded a significant Prime Race × Target Object interaction, $F(1, 180) = 110.39$, $p < .001$, $\eta_p^2 = .38$. Decomposing this interaction revealed that participants identified guns more quickly after Black primes than after White primes, $F(1, 180) = 54.90$, $p < .001$, $\eta_p^2 = .23$, whereas they identified tools more quickly after White primes than after Black primes, $F(1, 180) = 77.40$, $p < .001$, $\eta_p^2 = .30$. Neither the Prime Age × Target Object interaction, $F(1, 180) = 1.55$, $p = .215$, $\eta_p^2 = .01$, nor the Prime Age × Prime Race × Target interaction ($F < 1$, $p = .666$, $\eta_p^2 <$

.01), was significant, suggesting a comparable pattern of object identification across prime age and a comparable pattern of racial bias in object identification across prime age.

Error rates. A Prime Race × Target Object interaction emerged in an identical analysis on the error rates, $F(1, 180) = 58.74$, $p < .001$, $\eta_p^2 = .25$. Decomposing this interaction revealed that participants misidentified tools as guns more often after Black primes than after White primes, $F(1, 180) = 29.77$, $p < .001$, $\eta_p^2 = .14$, whereas they misidentified guns as tools more often after White primes than after Black primes, $F(1, 180) = 36.01$, $p < .001$, $\eta_p^2 = .17$. As was the case with the RT analyses, neither the Prime Age × Target Object interaction, $F(1, 180) = 1.43$, $p = .234$, $\eta_p^2 < .01$, nor the Prime Age × Prime Race × Target Object interaction ($F < 1$, $p = .511$, $\eta_p^2 < .01$), was significant, again indicating a comparable pattern of object identification across prime age and a comparable pattern of racial bias in object identification across prime age.

PDP estimates. We next conducted PDP analyses to estimate the unique contributions of automatic and controlled processes to task performance. Using the equations reported by Payne (2001), we computed estimates of automatic and controlled processing, separately for primes of each age–race combination. In the weapon identification task, controlled processing reflects the ability to distinguish guns from tools, independent of response biases. Automatic processing, in contrast, reflects the biasing influence of the primes when controlled processing fails. In cases of perfect performance (i.e., controlled processing = 1), automatic processing is undefined; thus, we applied an adjustment commonly used in signal-detection analyses (see Snodgrass & Corwin, 1988, for details). Because negative estimates of controlled processing violate assumptions of PDP that parameter estimates range from 0 to 1 (Jacoby, 1991), we replaced such instances with a value of 0 (Todd et al., 2016).

A 2 (Prime Age) × 2 (Prime Race) repeated-measures ANOVA revealed that estimates of automatic processing were greater for Black primes than for White primes, $F(1, 180) = 60.20$, $p < .001$, $\eta_p^2 = .25$. Neither the Prime Age main effect, $F(1, 180) = 1.30$, $p = .255$, $\eta_p^2 < .01$, nor the Prime Age × Prime Race interaction ($F < 1$, $p = .419$, $\eta_p^2 < .01$), was significant, indicating a comparable pattern of automatic processing across prime age and a comparable pattern of racial bias in automatic processing across prime age. An identical analysis on the estimates of controlled processing yielded no significant effects (Fs < 1.19, ps > .277, η_p^2s < .01). Together, these analyses indicate that the racial bias in weapon identification was driven entirely by differences in automatic processing.

Experiment 2

In line with our hypothesis, participants in Experiment 1 responded faster and identified guns with fewer errors following Black primes than following White primes. Similarly, participants responded faster and identified tools with fewer errors following White primes than following Black primes. This racial bias was comparable across target age and based on PDP estimates it was driven entirely by differences in automatic processing.

For Experiment 2, we were interested in testing whether we could conceptually replicate our findings using a similar sequential task. However, in order to test that the threat association was not just contingent upon objects, we replaced the object identification task with a word categorization task. Completely new face primes were also incorporated into the experiment to help ensure adequate stimulus sampling. We predicted that White primes would elicit faster and more accurate responses for safe words and that Black primes would elicit faster and more accurate responses for threatening words. Once again we did not think that the age of the target would have any significant effects.

Method

Undergraduates (*N* = 168) participated for course credit. Data from 8 participants with below-chance accuracy (errors on >50% of trials) on the weapon identification task were excluded from analyses, leaving a final sample of 160 participants (103 women, 56 men, and 1 unreported; 127 White, 5 Black, 9 Latinx, 9 Asian, 9 who reported more than one race/ethnicity, and 1 unreported).

Experiment 2 was identical to Experiment 1, but with three important changes. First, instead of a weapon identification task, participants completed a sequential priming task in which they categorized words as ‘threatening’ (*violent*, *dangerous*, *hostile*, *aggressive*, *criminal*, and *threatening*) or ‘safe’ (*innocent*, *harmless*, *friendly*, *trustworthy*, *peaceful*, and *safe*). Second, the response deadline was increased to 750 ms to account for the increased difficulty of categorizing words versus objects. Third, to ensure adequate stimulus sampling for our prime images (Wells & Windschitl, 1999), we selected a new sample of 40 photos of younger and older Black and White men from online sources using the same criteria as Gawronski et al. (2010). As in Experiment 1, participants completed a total of 288 randomly-ordered experimental trials, which were preceded by 16 practice trials.

Analyses and Results

Descriptive statistics (mean RTs, error rates, and PDP estimates) for all prime–target combinations in Experiment 2 are displayed in Table 2.

RTs. RT data were prepared as in Experiment 1. Prior to analysis, trials with errors (incorrect word identification) and RTs <100 ms were excluded. RTs >750 ms were coded as errors, and 750 ms was used as the upper bound for RT inclusion. Although the RTs were log-transformed to reduce positive skew, raw RTs are reported to improve the ease of interpretation.

A 2 (Prime Age) × 2 (Prime Race) × 2 (Target Word) repeated-measures ANOVA on the RTs yielded a significant Prime Race × Target Word interaction, $F(1, 159) = 19.04$, $p < .001$, $\eta_p^2 = .11$. Decomposing this interaction revealed that participants identified dangerous words more quickly after Black primes than after White primes, $F(1, 159) = 11.86$, $p = .001$, $\eta_p^2 = .07$, whereas they identified safe words more quickly after White primes than after Black primes, $F(1, 159) = 7.39$, $p =.007$, $\eta_p^2 = .04$. Neither the Prime Age × Target Word interaction, $F(1, 159) = 1.10$, $p = .296$, $\eta_p^2 < .01$, nor the Prime Age × Prime Race × Target Word interaction ($F < 1$, $p = .764$, $\eta_p^2 < .01$), was significant, suggesting a comparable pattern of word identification across prime age and a comparable pattern of racial bias in word identification across prime age.

Error rates. A Prime Race × Target Word interaction emerged in an identical analysis on the error rates, $F(1, 159) = 13.87$, $p < .001$, $\eta_p^2 = .08$. Decomposing this interaction revealed that participants misidentified safe words as threatening more often after Black primes than after White primes, $F(1, 159) = 6.62$, $p = .011$, $\eta_p^2 = .04$, whereas they misidentified threatening words as safe more often after White primes than after Black primes, $F(1, 159) = 11.35$, $p = .007$, $\eta_p^2 = .07$. As was the case with the RT analyses, neither the Prime Age × Target Word interaction, $F(1, 159) = 1.23$, $p = .269$, $\eta_p^2 < .01$, nor the Prime Age × Prime Race × Target Word interaction ($F < 1$, $p = .972$, $\eta_p^2 < .01$), was significant, again indicating a comparable pattern of word identification across prime age and a comparable pattern of racial bias in word identification across prime age.

PDP estimates. As in Experiment 1, we next conducted PDP analyses to estimate the unique contributions of automatic and controlled processes to task performance. A 2 (Prime Age) × 2 (Prime Race) repeated-measures ANOVA revealed that estimates of automatic processing were greater for Black primes than for White primes, $F(1, 159) = 12.38$, $p = .001$, $\eta_p^2 = .07$.

Neither the Prime Age main effect ($F < 1$, $p = .350$, $\eta_p^2 < .01$), nor the Prime Age × Prime Race interaction ($F < 1$, $p = .946$, $\eta_p^2 < .01$), was significant, indicating a comparable pattern of automatic processing across prime age and a comparable pattern of racial bias in automatic processing across prime age. An identical analysis on the estimates of controlled processing yielded no significant effects (*F*s < 1, *p*s > .499, η_p^2s < .01). Together, these analyses indicate that, much like in Experiment 1, the racial bias in word identification was driven entirely by differences in automatic processing.

Discussion

Results from two experiments confirmed our hypothesis that older Black men facilitate judgments of objects and words as threatening, relative to older White men. In conjunction with the findings of Todd, Simpson, et al. (2016) and Todd, Thiem, et al. (2016), these studies suggest that the implicit bias to associate Black males with danger extends across several age groups (young children, young adults, and older adults). These findings generalized across two different paradigms: Using both a weapon identification task (Experiment 1) and a word categorization task (Experiment 2), we found a similar pattern of threat association. Racial primes had a significant effect on participants' ability to respond accurately and quickly to threatening and non-threatening objects and words. Our analysis of the data using PDP estimates suggests that the racial bias displayed here was driven by automatic processes.

From Todd, Thiem, et al. (2016), we know that participants are able extract information about both age and race even when only briefly shown facial images (see also Ito & Urland, 2003). A potential limitation with these studies, however, is that the cropped facial images we used may be lacking in important cues. Future experiments may be well served by the use of more ecologically valid stimuli. For example, full-body images projected onto a screen at full size may allow participants to detect more useful information about the age of a target (Cesario, 2017) and provide a more ecologically valid test of racial bias in judgments of threat. For example, when participants are able to see the true stature of a young child or an older adult (who may appear slightly frail/hunched), they may be able to more accurately assess the potential threat this person might present to their safety. One might therefore predict that, when using such stimuli, age would moderate the effects of racial bias.

Providing participants with contextual cues or background information may also represent an important extension of this line of research. This could involve presenting images of different racial primes in varying settings (Correll et al., 2002), allowing participants to gather information about the situation's broader context. In reality, we rarely if ever make snap judgements about the threat an individual poses without using the contextual information available to us. Police officers, for example, rely on information provided by a dispatcher regarding a suspect to inform how they approach and respond to the suspect (Johnson, Cesario, & Pleskac, 2017). Providing participants with this type of information, through methods such as background images, information from third parties, and other contextual cues may help paint a more accurate picture of how participants would respond to people of different races in the real world.

Another goal could be to prime the kinds of environments a person grew up in. Past work shows that racial stereotypes often reflect whether someone comes from an environment where resources are scarce, rather than race per se (Williams et al., 2016). Stereotypes of people from resource-scarce environments suggest that they will try to acquire others' resources via coercion. Thus, if participants are given cues that suggest an outgroup individual (such as an older Black man for a White participant) has plenty of resources (the target is wearing a suit), the same racial stereotype of violence and criminality may no longer be applied.

In sum, consistent with previous research, we found racial bias in the identification of objects as weapons or tools and in the identification of words as threatening or safe. Also in line with past work, we found that this pattern emerged even when participants were presented with images of individuals that belong to age groups that are generally thought to be non-threatening (older adults). This suggests that old age, like youth, may be insufficient to disarm the threat associated with Black men. In conjunction with the findings of Todd et al. (2016), these results suggest that stereotypes of young Black men as dangerous generalize to Black males regardless of age. Building a better understanding of how racial bias can drive automatic processes, particularly those related to threat identification, is not just an important step in helping us to comprehend how rapid social categorization can shape behavior but perhaps also for the development of more effective interventions and training programs related to the management of potentially threatening situations.

References

Becker, D. V., Neel, R., & Anderson, U. S. (2010). Illusory conjunctions of angry facial expressions follow intergroup biases. *Psychological Science*, *21*(7), 938-940.

Cesario, J. (2017, April). Questioning the Use of Experimental Social Psychology Data to Understand Disparate Outcomes: The Case of Race Bias in Police Shootings. Paper presented at 89th Annual Meeting Midwestern Psychological Association, Chicago, IL.

Correll, J., Park, B., Judd, C. M., & Wittenbrink, B. (2002). The police officer's dilemma: Using ethnicity to disambiguate potentially threatening individuals. *Journal of Personality and Social Psychology*, *83*, 1314–1329.

Devine, P. G. (1989). Stereotypes and prejudice: Their automatic and controlled components. *Journal of Personality and Social Psychology*, *56*, 5–18.

Eberhardt, J. L., Goff, P. A., Purdie, V. J., & Davies, P. G. (2004). Seeing black: race, crime, and visual processing. *Journal of Personality and Social Psychology*, *87*(6), 876-893.

Fiske, S. T., Cuddy, A. J., Glick, P., & Xu, J. (2002). A model of (often mixed) stereotype content: competence and warmth respectively follow from perceived status and competition. *Journal of Personality and Social Psychology*, *82*(6), 878-902.

Hugenberg, K., & Bodenhausen, G. V. (2003). Facing prejudice implicit prejudice and the perception of facial threat. *Psychological Science*, *14*(6), 640-643.

Johnson, D.J., Cesario, J., & Pleskac, T.J. (2017). Dispatch information and the decision to shoot. Talk presented at the Society for Personality and Social Psychology, San Antonio, TX.

Kang, S. K., & Bodenhausen, G. V. (2015). Multiple identities in social perception and interaction: Challenges and opportunities. *Annual Review of Psychology*, *66*, 547–574.

Kite, M. E., Deaux, K., & Miele, M. (1991). Stereotypes of young and old: Does age outweigh gender? *Psychology and Aging*, *6*(1), 19-27.

Nelson, T. D. (2004). *Ageism: Stereotyping and prejudice against older persons*. Cambridge, MA: MIT press.

Payne, B. K. (2001). Prejudice and perception: The role of automatic and controlled processes in misperceiving a weapon. *Journal of Personality and Social Psychology*, *81*, 181–192.

Snodgrass, J. G., & Corwin, J. (1988). Pragmatics of measuring recognition memory: Applications to dementia and amnesia. *Journal of Experimental Psychology: General*, *117*, 34–50.

Todd, A. R., Thiem, K. C., & Neel, R. (2016). Does seeing faces of young black boys facilitate the identification of threatening stimuli?. *Psychological Science*, *27*(3), 384-393.

Todd, A. R., Simpson, A. J., Thiem, K. C., & Neel, R. (2016). The Generalization of Implicit Racial Bias to Young Black Boys: Automatic Stereotyping or Automatic Prejudice? *Social Cognition*, *34*(4), 306-323.

Wells, G. L., & Windschitl, P. D. (1999). Stimulus sampling and social psychological experimentation. *Personality and Social Psychology Bulletin*, *25*(9), 1115-1125.

Williams, K. E., Sng, O., & Neuberg, S. L. (2016). Ecology-driven stereotypes override race stereotypes. *Proceedings of the National Academy of Sciences*, *113*(2), 310–315.

Table 1

Mean RTs (in ms), Error Rates (in %), and PDP Estimates by Condition (Experiment 1)

	Prime age and race			
	Younger adult primes		Older adult primes	
Variable	Black primes	White primes	Black primes	White primes
RTs (in ms)				
Gun	259.9 (36.6)	270.0 (36.1)	262.4 (36.0)	270.5 (36.9)
Tool	291.0 (33.6)	281.5 (36.9)	296.1 (36.1)	285.0 (37.7)
Error rates (in %)				
Gun	14.9 (11.0)	17.5 (12.3)	13.8 (10.5)	17.4 (12.4)
Tool	18.3 (12.6)	15.4 (10.7)	18.4 (12.4)	15.6 (11.9)
PDP estimates				
Automatic	.55 (.17)	.47 (.18)	.58 (.17)	.47 (.17)
Control	.66 (.21)	.67 (.20)	.67 (.20)	.66 (.21)

Note. Standard deviations are in parentheses.

Table 2

Mean RTs (in ms), Error Rates (in %), and PDP Estimates by Condition (Experiment 2)

	Prime age and race			
	Younger adult primes		Older adult primes	
Variable	Black primes	White primes	Black primes	White primes
RTs (in ms)				
Threatening words	491.4 (49.9)	500.8 (47.1)	495.6 (49.4)	499.9 (51.2)
Safe words	499.1 (51.0)	495.4 (49.3)	504.5 (46.2)	496.5 (46.3)
Error rates (in %)				
Threatening words	18.5 (11.8)	20.5 (12.8)	19.1 (12.0)	21.2 (12.6)
Safe words	25.1 (12.6)	23.3 (11.9)	24.9 (12.6)	23.2 (11.5)
PDP estimates				
Automatic	.59 (.15)	.54 (.17)	.58 (.17)	.53 (.17)
Control	.56 (.20)	.56 (.20)	.56 (.20)	.56 (.19)

Note. Standard deviations are in parentheses.

Appendix

Here, we report additional results that, though significant, were not directly relevant to testing our main predictions.

Experiment 1

RTs. In addition to the Prime Race × Target Object interaction on the RTs reported in the main text, there was also a Prime Age effect, $F(1, 180) = 20.21$, $p < .001$, $\eta_p^2 = .10$, indicating that participants responded more quickly after younger adult primes than after older adult primes, and a Target Object main effect, $F(1, 180) = 215.60$, $p < .001$, $\eta_p^2 = .55$, indicating that participants identified guns more quickly than tools.

Error rates. In addition to the Prime Race × Target Object interaction on the error rates reported in the main text, there was also a Target Object main effect, $F(1, 180) = 3.93$, $p = .049$, $\eta_p^2 = .02$, indicating that participants were more likely to misidentify tools as guns than vice versa.

PDP estimates. All results of the PDP analyses were reported in the main text.

Experiment 2

RTs. In addition to the Prime Race × Target Word interaction reported in the main text, there was also a Prime Age effect, $F(1, 159) = 6.86$, $p = .010$, $\eta_p^2 = .04$, indicating that participants responded more quickly after younger adult primes than after older adult primes. A Prime Age × Prime Race interaction, $F(1, 159) = 4.37$, $p = .038$, $\eta_p^2 = .03$, further indicated that participants responded more quickly after younger Black men than after younger White men, whereas they responded equally quickly after older Black men and older White men.

Error rates. In addition to the Prime Race × Target Word interaction reported in the main text, there was also a Target Word main effect, $F(1, 159) = 27.90$, $p < .001$, $\eta_p^2 = .15$, indicating that participants were more likely to misidentify safe words as threatening than vice versa.

PDP estimates. All results of the PDP analyses were reported in the main text.