LIFE IS NOT JUST BLACK AND WHITE: THE INFLUENCE OF SOCIAL CLASS CUES ON RACE IN AN AFFECT MISATTRIBUTION PROCEDURE

By

Alejandro Carrillo

A THESIS

Submitted to
Michigan State University
in partial fulfillment of the requirements
for the degree of

Psychology—Master of Arts

2022

ABSTRACT

LIFE IS NOT JUST BLACK AND WHITE: THE INFLUENCE OF SOCIAL CLASS CUES ON RACE IN AN AFFECT MISATTRIBUTION PROCEDURE

By

Alejandro Carrillo

Indirect measures of racial bias, such as the Affect Misattribution Procedure, Evaluative Priming Task, and the Implicit Association Task, have been used to provide evidence of stereotypical associations and valenced evaluations towards different racial categories. However, a common limitation shared across these tasks is the tendency to use simplistic racial stimuli that do not capture or account for the multiple categories people may belong to. That is, it is unlikely that people are perceived and evaluated along a single feature (i.e., race) but rather at the intersections of multiple categories (i.e., race, age, attractiveness, social class, etc.). Social class, in particular, is a strongly evaluated category and has been shown to share stereotypic associations with race (Moore-Berg & Karpinski, 2019). Thus, this thesis investigated the effects of social class on racial evaluations in an AMP task. Social class was manipulated using occupational clothing in Study 1 and residential areas in Study 2, while race was limited to Black and White men. Across two studies, participants demonstrated a consistent, unexpected pro-Black bias. In addition, an effect of social class was only found in Study 2 such that high-class primes were associated with positive responses. Regarding social class, the results suggest that the manipulations chosen may play an important role in categorization; however, future research is needed to examine just how different class representations impact evaluation.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

Introduction

Over the last 40 years, it has become a truism within social psychology that stored information in the form of stereotypes or evaluations can be automatically activated. It has been repeatedly and reliably shown that information associated with social categories can be activated "automatically" – that is, in an unintended, fast, unconscious, or uncontrollable manner. Although there are important differences in each of these features (Bargh, 1989), for the purposes of the present research, such nuances are not of central importance. Therefore, I refer to cognitive processes that have any of these features as "automatic" throughout this thesis. These automatic effects have often been referred to as implicit cognition or implicit bias (Greenwald & Banaji, 1995).

Several indirect measures have been developed to measure the automatic activation of evaluations and stereotypes. By "indirect measures," I refer to methods of obtaining responses from participants without directly asking them for the outcome of interest (Fazio & Olson, 2003). The most common of these measures include the Affect Misattribution Procedure (AMP; Payne, Cheng, Govorun, & Stewart, 2005), the Evaluative Priming Task (EPT; Fazio, 2001; Neely, 1977), and the Implicit Association Task (IAT; Greenwald, McGhee, & Schwartz, 1998). Less common variants of indirect measures include the Sequential Priming Task (SPT; Kramer, Jones, & Sharma, 2013), the go/no go association task, and approach-avoidance task (for a review, see Fazio & Olson, 2003; Gawronski, 2009). Given that the focus of the current work is the AMP, I describe this measure in the most detail next and then provide a brief overview of other measures following.

The AMP (Payne et al., 2005) uses a modified sequential priming procedure in which participants are asked to make judgments of a neutral target image after an affectively-arousing

prime image has been presented. In rendering a decision about how much they like the neutral target image, participants will use their feelings as information relevant to this judgement (Schwarz & Clore,1996). Because the prime image immediately preceding the neutral target was affectively-arousing, the feelings provoked by the prime will be misattributed to the neutral image and will therefore influence the judgments of the neutral target. The AMP task, then, is based on a misattribution effect such that the participant misattributes the feeling evoked by the prime image to the neutral image.

Figure 1 illustrates a specific example of an AMP intended to measure racial prejudice. Participants are first primed with a Black or White face. This prime face is then immediately replaced by a neutral Chinese character. The Chinese characters are chosen with the intent that they are neutral, novel targets for which participants have no pre-existing evaluations and which do not arouse any particular affective response. (To this end, participants who are familiar with Chinese characters are prevented from completing this task.) They are then asked to rate whether the neutral target image is more or less pleasant than other target images in the task. To the extent that the neutral images receive a greater proportion of pleasant responses following White faces rather than Black faces, it would be assumed that White faces induced more positive affect. This effect has been found to be replicable (Payne & Lundberg, 2014) and is observed even when participants are instructed that the prime may influence their responses (Payne et al., 2005).

Along the same vein, the EPT (Fazio, 2001) uses a sequential priming effect paradigm in which participants are asked to correctly categorize a valenced word as positive or negative after exposure to a prime image. The core difference between the AMP and the EPT is that the effect of interest is not a misattribution effect, but rather a semantic activation effect. For example, in

an EPT measuring racial prejudice, participants are first primed with an image of a Black or White face, which is immediately replaced by a positive or negative word. The participants must then categorize the target word as positive or negative as quickly as possible. To the extent that the preceding image facilitates categorization, the prime is assumed to be associated with positive or negative valence. For example, to the extent that participants hold negative evaluations of Blacks, presentation of a Black face should facilitate categorization of subsequent negative targets and interfere with categorization of subsequent positive targets.

The IAT (Greenwald et al., 1998) assesses the strength of associations between two categories and attributes. For instance, an IAT measuring racial prejudice would compare how quickly participants can make different categorizations when "Black" and "White" are paired with "good" or "bad" categories. In one block, "White" and "good" are paired together, and "Black" and "bad" are paired together. The speed at which participants categorize Black/White faces and positive/negative words is recorded. In a different block, these pairings are reversed, such that "White" and "bad" are paired together, and "Black" and "good" are paired together. Participants again categorize Black/White faces and positive/negative words under these pairings. To the extent that a participant associates one race or another with positive or negative evaluations, participants should be faster or more accurate to categorize under one or the other blocks. For example, if a participant is faster at categorizing Black faces when "Black" is paired with "bad," then this would indicate a negative association or evaluation with the category of Black (Greenwald et al., 1998). A commonality shared among these indirect measures is that each one is a variation of a rapid response task designed to assess associations that are activated in an automatic manner. In effect, these different indirect measures are used as a means of assessing one's immediate evaluations of different classes of stimuli. By and large, these

measures have reliably shown that they assess some associations that people have between categories and attributes and that such associations can be activated without the intention, control, or awareness of a person.

However, there are some significant limitations of these indirect measures, and these limitations are the focus of the current research proposal. Virtually all indirect measures, including measures assessing both evaluations and stereotypes, rely on the use of simplistic racial stimuli, typically close-cropped head-shots of Black and White faces. Over the last several decades, the use of this format has produced a body of literature that supports the notion that American participants hold stereotypical associations and valenced evaluations towards racial categories. This is to say, for example, that upon mere exposure to a Black face, participants experience some negative affect which influences subsequent evaluations or facilitates stereotype-congruent categorization.

However, generally speaking, humans do not encounter people who only have one identifiable characteristic (e.g., race), much less a static face. Humans are multiply categorizable and have intersecting features, such as race/social class combinations, race/sex combinations, race/physical attractiveness combinations, and so on. It is unclear whether the results obtained on indirect measures in response to simple faces, which only vary on a single category, remain the same when these faces might vary from each other on multiple dimensions, as they do in real life. To illustrate the importance of multiple categories, I conducted a pilot study with an AMP task in which race was crossed with gender and physical attractiveness. This study indicated that once attractiveness was added as a factor, the typical race effect of White primes leading to more pleasant ratings was no longer observed and instead there were just strong effects of attractiveness.

Moreover, a larger problem shared among indirect measures is not just the use of simplistic racial stimuli, but the assumption that people are only ever perceived in a singular, uninformative context. Stimuli on these tasks tend to be presented in neutral or uninformative backgrounds, but humans are only ever perceived in social and physical contexts. There is a growing body of literature that supports the notion that automatically-activated evaluations and stereotypes can change depending on a variety of contextual factors (for a review, see Gawronski, Rydell, Vervliet, & De Houwer, 2010). These contextual factors include the salience of a category (eg., Kühnen et al., 2001; Mitchell, Nosek, & Banaji, 2003), the social context of the target (eg., Dasgupta & Rivera, 2008; Maddux, Barden, Brewer, & Petty, 2005; Scherer & Lambert, 2009; Wittenbrink, Judd, & Park, 2001), the prototypicality of the target (eg., Livingston & Brewer, 2002; Ma & Correll, 2011), exposure to exemplars (eg., Dasgupta & Asgari, 2004; Dasgupta & Greenwald, 2001), the emotional state of the perceiver (eg., Dasgupta, DeSteno, Williams, & Hunsinger, 2009; DeSteno, Dasgupta, Bartlett, & Cajdric, 2004), and social role of the perceiver (eg., Richeson & Ambady, 2001, 2003).

Taking both of these points together, one important feature of both people and contexts is social class. People may be categorized as belonging to higher or lower social classes, and people may exist in and be perceived in physical spaces that convey higher or lower social classes. We might predict that social class is an important variable for automatic evaluations for both theoretical and empirical reasons. For theoretical reasons, researchers have argued that the human mind is attuned, perhaps for evolutionary reasons, to the ecological history of an individual (Williams, Sng, & Neuberg, 2016). Specifically, whether a person grew up in a harsh/desperate or safe/hopeful environment predicts important life outcomes, including the likelihood of adopting behaviors associated with fast or slow life-history strategies, and humans

have been shown to care about this feature. In some work, researchers have argued that race is a heuristic cue for ecology, and thus people care about and categorize others in terms of race because of the information it provides about ecological history (given the correlation between social class and race in the U.S.). In this instance, it is not so much that the category of race is encoded and activated, but instead, race is used to determine the probability that an individual comes from a resource-rich or resource-deprived environment. Williams et al. (2016) found that when participants were given information both about race and whether the person was raised in a harsh environment, participants treated Black and White targets equally and based their stereotypes entirely off the harsh environment cue. In other words, racial evaluations may be mistaken for evaluations of social class, and therefore indirect measures could reveal quite different automatic evaluations of Blacks and Whites when the additional cue of social class is included (Mattan, Kubota, & Cloutier, 2017).

Along these same lines, all humans are more than just a single category, and one of these categories is social class. That is, we do not see "Black" people, we see "ghetto Blacks;" we do not see "White" people, we see "White trash." Importantly, people can immediately judge social class, and so the automatic evaluations we experience may be influenced in important ways by social class (Mattan et al., 2017).

Empirically, we might predict social class is important because there are overlaps in the stereotypes of race and social class, with Black stereotypes overlapping with low class stereotypes and White stereotypes overlapping with high class stereotypes (Brown, Boniecki, & Walters, 2004; Bullock, 1999; Cozzarelli, Wilkinson, & Tagler, 2001; Devine, 1989; Devine & Elliot, 1995; Durante, Tablante, & Fiske, 2017; Fiske, Cuddy, Glick, & Xu, 2002; Katz & Braly, 1933; Klonis, Plant, & Devine, 2005; Moore-Berg & Karpinski, 2019; Niemann, Jennings,

Rozelle, Baxter, & Sullivan, 1994; Spencer, Fein, Wolfe, Fong, & Dunn, 1998; Tan, Zhang, Zhang, & Dalisay, 2009; Wu, Bai, & Fiske, 2018). Further, Lei and Bodenhausen (2017), using a reverse correlation design, found that participants with higher levels of economic prejudice had mental representations that were Blacker for the poor category than either the middle income or wealthy category. Moreover, mental representations were Whiter for the Wealthy category, than either the poor or middle-income category.

To conclude, we should be cautious about making strong inferences of mental representation from measures that simplify human stimuli in potentially misleading ways, given that those stimuli are never really perceived in actual life in the way that they are represented in the tasks (i.e., a close-cropped image with no other information present). Hence, this project aims to explore what happens when additional information is added and how race and social class may interact.

Current Research Program

Indirect measures of racial bias are an important tool to assess people's implicit attitudes; however, they are generally missing important information present in real social experiences, including the multiple dimensions of any given individual's demographics and the fact that people are perceived in more than one context. Thus, I aim to start this line of research inquiry by investigating the effects of social class indicators in a racial bias AMP task.

Although there is substantial evidence indicating that race and social class overlap in meaningful ways (Moore-Berg & Karpinski, 2019), there are relatively few empirical studies addressing this issue from an indirect measurement perspective (Mattan, Kubota, Li, Venezia, & Cloutier, 2019; Moore-Berg & Karpinski, 2019; Oh, Shafir, & Todorov, 2020). The studies conducted have either manipulated clothing as a contextual cue (Moore-Berg, Karpinski, &

Plant, 2017; Oh et al., 2020) or used color-coded backgrounds (Mattan et al., 2019). The design of these studies is based on clothing as a reliable indicator of social class. Oh et al. (2020) used photos of people digitally manipulated to be in richer or poorer clothing and found evidence for pictures of richer clothing people to be perceived as more competent than those in poorer clothing. This is taken to provide evidence towards the effect of subtle contextualization cues of clothing. On the other hand, Mattan et al. (2017, 2019) makes the argument that clothing may not provide reliable evidence of higher class because high-class clothing may be indicative of something else. An example is a waiter at an upscale restaurant; while they may be dressed "richly," they still hold a relatively low-class position. This distinction is why Mattan et al. (2019) decided to manipulate social class by assigning different color-coded backgrounds to represent different social classes. Thus, there is no agreed upon way to manipulate social class, and it is possible that the method chosen to convey social class in an indirect measurement task may have an important effect on person evaluation. I will turn briefly to how social class has been manipulated in indirect measurement tasks and how the type of manipulation influenced perception.

First, Moore-Berg et al. (2017) conducted a first-person shooter task in which the stimuli were manipulated to represent either high or low class. Participants were instructed that they would be shown a suspect holding either a gun or non-gun object and that they must choose to "shoot" or "not shoot" the suspect using corresponding keys on a keyboard. Typically, a race x object interaction is observed such that participants are slower to shoot an armed White suspect (vs. an armed Black suspect) or quicker to shoot an unarmed Black suspect (vs. an unarmed White suspect). Moore-Berg et al. (2017) studies provided evidence of race x object x class interaction such that participants were faster to respond "not shoot" to White high-class stimuli

rather than Black high-class stimuli. Further, there were no significant differences in response times between low-class Black, low-class White, and high-class Black stimuli; thus, the effect was driven by the high-class White stimuli. One possibility for the current study, then, is a race by social class interaction such that target images following high-class White stimuli will receive the greatest proportion of pleasant responses.

In a different manipulation, Mattan et al. (2019) conducted a set of four studies using an evaluative priming task in which social status was designated through the use of color-coded backgrounds and shirts. Mattan predicted that race and social class would either independently contribute to categorization or interact as in Moore-Berg et al. (2017). In the first two experiments, color-coded backgrounds were used. This design produced separate main effects for race and social status. That is, Blacks were categorized with negative words faster than Whites, and high status people were categorized with positive words faster than low status people. In Study 3 the number of training trials was reduced, and in Study 4, color-coded shirts replaced color-coded backgrounds. In Study 3 Mattan et al. (2019) found evidence of a race x class x word valence interaction such that high-class White primes were evaluated with positive associations faster than low-class Black, low class White, and high-class Black. This finding was similar in nature to the findings from Moore-Berg et al. (2017) in that high-class Black primes did not facilitate a more positive association compared to the two lower-class primes. In Study 4, Mattan et al. (2019) replicated the three-way interaction with the exception that high-class Black primes also facilitated more positive associations than either of the lower class primes. One possibility is that in the design of the first two studies, the perceptual salience of social class attenuated the influence of race (Gawronski, Cunningham, LeBel, & Deutsch, 2010). Whereas

with reduced training (Study 3) and incorporated social class (Study 4), the salience of this category was reduced and no longer overwhelmed the effect of race.

However, it is also possible that race is a heuristic cue towards ecology, and as such, the effect of clothing or social class will subsume the effect of race (Williams et al., 2016). In that case, clothing that denotes one's social status and potential availability of resources should be a stronger cue of home ecology and should overwhelm the race effect. Williams et al. (2016) provides this argument as their data suggests that when ecology information is presented, social inferences are influenced by ecology over race.

In sum, the current study is exploratory in nature. The effects of interest may either show up in a pro-White high-status interaction, or separate main effects for race and social class, or only in main effects of social class.

Study 1

Method

Participants

Student participants were recruited via Michigan State University's Department of Psychology HPR/SONA system to complete a "Cognition and Perception" task for a half credit towards the fulfillment of required credit hours in their introductory psychology class. The sample (N = 288; 46 male; 237 female, 3 NA) excluded 4 participants who either did not respond to all items in the Qualtrics survey or did not finish the study task. Most participants (202) self-identified as White, with additional representation for Asian (31), Black (27), Other (28). Further, members of the student participant pool who could read Chinese characters were prevented from participating.

Materials

36 unique faces from the Chicago Face Database were digitally manipulated to be in both high and low status occupational clothing. 18 occupations were selected (9 high, 9 low) and each occupation was used four times, twice for each race, creating a total of 72 prime images. The high status professions were architect, chemist, doctor, judge, Ph.D. holder, pilot, professor, and two business class suits. As it stands, many high status occupations are some variations of formal wear which led to multiple suit images. The low status professions were construction worker, delivery driver, grocer, industrial worker, janitor, maintenance worker, mechanic, and two store clerks (for examples see Fig.2). All data, syntax, and materials can be found at: OSF AMP Race x SES.

Manipulation Check

To verify whether the targets' clothing effectively manipulated perceived social class as intended, after the main task participants provided judgements on a subset of the faces. Specifically, there were four conditions with 18 faces each. Participants only saw one racial category and no repeated faces. They rated the faces on a scale of 0 (poor) to 10 (rich) and there were 72 ratings for each target. The high-status primes were consistently rated above the scale midpoint (M = 6.80), and low-status primes were consistently rated below the midpoint (M = 4.11); this difference was significant *(*$F(3,71) = 224.65$; $p < 0.001$) in a 2 (Race) x 2 (SES) ANOVA. Notably, we also see that ratings of Black primes (M = 5.75) were rated significantly higher on status than the White primes (M = 5.16) ($F(3,71) = 10.93$; $p = 0.002$).

Procedure

Participants completed an Affect Misattribution Procedure (AMP; Payne et al., 2005) online using Pavlovia. Pavlovia is an online platform that allows researchers to host experiments designed in PsychoPy. Participants received instructions to rate a Chinese character as either pleasant or unpleasant using the "e" and "i" keys, respectively. The participants received further instructions stating that throughout the experiment, a face would appear briefly to signal that a Chinese character would appear next and that this face should otherwise be ignored. Participants then completed five practice trials to prepare them for the main task. The AMP task consisted of 144 trials in two 72 image blocks in which each face and Chinese character were randomly paired and ordered. That is, each of the 36 faces were shown twice, once for each level of social class. Consistent with the standard AMP procedure, each face flashed for 50 milliseconds, after which a Chinese character displayed for 75 milliseconds. This was then replaced by a pattern mask that remained on the screen until participants responded. The task does not proceed to the

next trial until the participant has evaluated the character using the requisite keys. After the AMP task, participants were redirected to Qualtrics, where they were asked to provide an explicit rating on the perceived social class of the stimuli used in the AMP task. Participants also responded to a series of demographic questions regarding their ethnicity and gender, and subjective social status.

Analyses

Preregistered

Multilevel modeling (MLM) using maximum likelihood was used to predict the binary response to the target image as a function of the race of the prime and the social class of the prime, as well as their interaction. To account for non-independence across participants and targets (Judd, Westfall, & Kenny, 2012), the model specified as random: (1) the intercept, target race slope, target social class slope, and the interaction between race and social class for participants; (2) the intercept for primes; (3) the intercept for the target image.

Exploratory

Prior work has suggested that attractiveness may be an important predictor of participants response to the target image. Thus, a second model was specified that included prime attractiveness as a fixed effect. Otherwise, the fixed and random effect structure was the same as above. The ratings of attractiveness were taken from the Chicago Face Database, which means that the ratings were collected prior to my clothing manipulation. That is, I did not ask participants to rate the faces on attractiveness in this study.

Study 1 Results

To predict the proportion of pleasant responses, a multilevel logistic regression was estimated with fixed effects of prime race and prime social class. Random effects included random intercepts for participants, primes, and targets and random slopes for prime race and prime social class for participants. In this model, prime race predicted the proportion of pleasant responses (*OR* = 1.07, 95% *CI* [1.02, 1.11], *p* = 0.004; Table 2) such that Black primes generated more positive responses than White primes. The effect of prime social class was unrelated to the odds of a pleasant response (*OR* = 0.99, 95% *CI* [0.96, 1.03], *p* = 0.692), indicating that the social class of the primes did not influence participants' responses. Further, the main effect of prime race was not qualified by an interaction with social class (*OR* = 0.99, 95% *CI* [0.96, 1.02], *p* = 0.506).

This pattern was not expected, as prior work has consistently found that Black primes elicit more negative responses than White primes. One possibility is that the attractiveness of the primes played a role in this finding. Although the primes did not differ on attractiveness according to the ratings in the Chicago Face Database ($F(1, 34) = 0.33$, $p = 0.569$), there were differences in the distributions such that there were more attractive Black primes than White primes (see Fig. 3). Given that attractive faces elicit more positive responses, the greater positivity of Black primes may be an attractiveness effect.

To probe this possibility, a multilevel logistic regression model predicting the odds of a pleasant response from attractiveness, prime race, and prime social class was conducted to explore the effect of attractiveness. Random effects included random intercepts and slopes for participants, primes, and targets. Additionally, random slopes for prime race and prime social class were included for participants. In this model the proportion of pleasant responses was

predicted by prime race (*OR* = 1.06, 95% *CI* [1.09, 1.39], $p < 0.001$; Table 3) and prime attractiveness (*OR* = 1.07, 95% *CI* [1.01, 1.14], $p = 0.023$) such that Black primes and attractive primes increased the likelihood of a pleasant response.

However, this effect was qualified by a three-way interaction between prime race, prime class, and attractiveness (*OR* = 0.93, 95% *CI* [0.88, 0.99], $p = 0.022$; see Fig. 3). Follow-up analyses indicated that as attractiveness increases for Black low-class primes the odds of a pleasant response increases (*OR* = 1.14, 95% *CI* [0.99, 1.29], $p = 0.08$). As attractiveness increases for Black high-class primes, there was a weak decrease in the odds of pleasant response (*OR* = 0.94, 95% *CI* [0.82, 1.08], $p = 0.39$). While for White low-class primes increasing attractiveness was associated with a somewhat greater odds of a pleasant response (*OR* = 1.06, 95% *CI* [0.96, 1.17], $p = 0.24$). Finally, for White high-class primes, increasing attractiveness was associated with a significantly greater odds of a pleasant response (*OR* = 1.17, 95% *CI* [1.06, 1.29], $p = 0.002$).

Study 1 Discussion

The purpose of Study 1 was to gain a better understanding of how social class as denoted through occupational clothing would interact with race. Recall that the effects of interest could appear as a pro-high-class bias, separate pro-White and pro-high-class biases, or as a pro-White high-class interaction. However, participants primarily demonstrated a positive bias towards Black primes with no discernible difference in attitudes towards class. To investigate this further, prime attractiveness was added as a fixed effect. In this model, a three-way interaction emerged, suggesting that as attractiveness increased, Black low-class primes and White high-class primes elicited more positive responses. Interestingly, while not significant, the opposite was true for Black high-class primes.

Given that social class has been represented in different ways in the literature (color cues, backgrounds, and clothing), the goal of Study 2 was to extend the findings of this work to a different social class manipulation. In Study 2, social class was represented using pictures of low and high-class neighborhoods as backgrounds. Additionally, a new set of faces were selected from the Chicago face database that were more evenly matched on attractiveness.

Study 2

Method

Participants

Participant recruitment and exclusion was identical to that of Study 1. The sample (N = 293; 86 male; 197 female, 8 gender variant, 2 NA) excluded 5 participants who either did not respond to all items in the Qualtrics survey or did not finish the study task. Most participants (203) self-identified as White, with marginal representation for Asian (33), Black (21), Other (36).

Materials

36 unique faces from the Chicago Face Database were digitally manipulated to be in both high and low status neighborhoods. Faces were selected with approximately equal, moderate levels of attractiveness. Eighteen backgrounds were selected (9 high class, 9 low class) and each neighborhood was used four times, twice for each race, creating a total of 72 prime images (for examples see Fig. 5). The Chinese characters were the same as those used in Study 1.

Manipulation Check

To verify whether the primes' background setting effectively manipulated perceived social class as intended, after the main task participants provided judgements on a subset of the faces. Specifically, there were four conditions with 18 faces each. Participants only saw one racial category and no repeated faces. They rated the faces on a scale of 0 (poor) to 7 (rich) and there were 72 ratings for each target. The high-status primes were consistently rated above the scale midpoint ($M = 5.18$), and low-status primes were consistently rated below the midpoint ($M = 2.33$); this difference was significant ($F(1,68) = 1813.88$; $p < 0.001$) in a 2 (Race) x 2 (SES) ANOVA. Additionally, participants provided judgements of attractiveness on the same faces on

a scale of 0 (Unattractive) to 7 (Attractive) and there were 72 ratings for each target. Notably, we see that ratings of Black primes ($M = 2.87$) were significantly higher than the White primes ($M = 2.42$) ($F(1,68) = 10.93$; $p < 0.001$), despite the primes being chosen by their similar ratings in the Chicago Face Database.

Procedure

The procedure was identical to Study 1.

Analyses

Preregistered

Multilevel modeling (MLM) using maximum likelihood was used to predict the binary response to the target image as a function of the race of the prime and the social class of the prime, as well as their interaction. To account for non-independence across participants and targets (Judd et al., 2012), the model specified as random: (1) the intercept, target race slope, target social class slope, and the interaction between race and social class for participants; (2) the intercept for primes; (3) the intercept for the target image.

Exploratory

Once again, a second model was specified that included prime attractiveness as a fixed effect. However, in this model the ratings of attractiveness were collected from participants after the experiment and were not the Chicago Face Database ratings.

Study 2 Results

To predict the proportion of pleasant responses, a multilevel logistic regression was estimated with fixed effects of prime race and prime social class. Random effects included random intercepts for participants, primes, and targets. Additionally, random slopes for prime race and prime social class were included for participants. There was a main effect of prime race (*OR* = 1.05, 95% *CI* [1.02, 1.08], $p = 0.001$; Table 5) such that Black primes increased the likelihood of a pleasant response. Additionally, there was a main effect of prime social class (*OR* = 1.13, 95% *CI* [1.08, 1.18], $p < 0.001$) indicating that high-class primes increased the odds of a pleasant response. These effects were not qualified by an interaction (*OR* = 1.00, 95% *CI* [0.92, 1.02], $p = 0.890$).

Once again, the main effect of race with Black primes eliciting more pleasant responses was unexpected. Black primes were also rated as more attractive. Therefore, to explore the effect of attractiveness, a multilevel logistic regression model predicting the odds of a pleasant response from attractiveness, prime race, and prime social class was conducted. Random effects included random intercepts and slopes for participants, primes, and targets. Additionally, random slopes for prime race and prime social class by participant were included. In this model the proportion of pleasant responses was predicted by prime race (*OR* = 1.03, 95% *CI* [1.01, 1.06], $p = 0.004$; Table 6) and prime class (*OR* = 1.10, 95% *CI* [1.01, 1.18], $p = 0.02$). Additionally, there was a marginal effect of attractiveness (*OR* = 1.05, 95% *CI* [1.00, 1.10], $p = 0.060$).

However, these main effects were qualified by a race by attractiveness interaction (*OR* = 0.94, 95% *CI* [0.90, 0.99], $p = 0.018$; see Fig. 6). Attractiveness interacted with race such that for Black primes, higher attractiveness was unrelated to the odds of a pleasant response (*OR* = 0.99, 95% *CI* [0.91, 1.06], $p = 0.74$), whereas for White primes, higher attractiveness was associated

with a larger odds of a pleasant response (*OR* = 1.12, 95% *CI* [1.04, 1.19], *p* = 0.001). Given that this unexpected, exploratory effect was not observed in Study 1, it should be interpreted with caution.

Study 2 Discussion

The purpose of Study 2 was to investigate how social class represented by neighborhoods would interact with race. In this study, participants demonstrated both a pro-Black and a pro-high class bias. Given that participants rated Black primes as more attractive than White primes, we ran another model with attractiveness similar to Study 1. However, in this model, increases in attractiveness were associated with an increase of positive responses for White faces. While not significant, there is a slight negative association between increasing attractiveness and positive responses for Black faces. That is, it would appear that attractiveness does not explain the pro-Black bias among participants.

General Discussion

The current work aimed to explore how social status cues interact with race in an AMP task. In two studies, social status was measured using either occupational clothing (Study 1) or residential areas (Study 2). Across studies, participants consistently demonstrated a pro-Black preference, and ratings of prime attractiveness did not explain this effect. An effect of social class was observed in Study 2 as a pro-high-class preference, but there is no apparent effect in Study 1. It is possible that the social class effect is weak or noise; alternatively, it is possible that neighborhoods convey social class in more immediate and stronger ways.

Interpretation

The pro-Black bias was striking, and to investigate this further, an exploratory set of analyses that included prime attractiveness were conducted. While past work in the lab has shown that increasing attractiveness is related to an increased likelihood of a pleasant response, the findings were somewhat counter intuitive. That is, in Study 1, increases in attractiveness were positively associated with pleasant responses for high-class White and low-class Black primes but negatively associated with pleasant responses for high-class Black primes. In addition, in Study 2, increases in attractiveness were negatively associated with the odds of a pleasant response for Black faces in general. It would seem then that prime attractiveness does not answer the question as to why Black primes had a higher proportion of pleasant responses, and in fact, new questions are raised. It is not readily apparent why attractiveness would have different effects for different racial categories. There is some work showing that the role of the perceiver may influence perceptions of attractiveness; however, the effects studied were within sex rather than within race (Agthe, Spörrle, & Maner, 2011). Moreover, the variability in the exact nature of the effects suggests that these may be noise in the data rather than replicable,

important findings. Further work is necessary to investigate whether attractiveness is differentially affected by race.

Notably, Moore-Berg and Karpinski (2021) conducted a race by social class AMP task utilizing the stimuli from their set of studies in 2017. Across two studies Moore-Berg and Karpinski (2021) found that participants demonstrated a positive high-class bias and a positive Black bias, although the effect of race was non-significant in Study 1. This pro-Black bias has been found in other labs, which may point to other influencing factors. Perhaps the greater political climate of the last year and a half may have led to a shift in attitudes, or at the very least, participants may feel compelled to respond in a socially desirable manner (Huddy & Feldman, 2009). An interesting route forward may be to include measures of motivation to control prejudiced reactions (Dunton & Fazio, 1997).

Regarding social class, we see that the manipulations were differentially effective such that residential areas provided a robust and noticeable effect, whereas occupational clothing did not. There are several possible explanations for this outcome. First, participants may simply have stronger associations to neighborhood scenes than to people's occupations, especially since the neighborhoods selected represented extremes in terms of wealth or poverty. Second, it may be possible that different social class manipulations tap into different aspects of social status. Mattan et al. (2017) suggests that social status can be broken into three distinct categories (wealth, prestige, and dominance) and that evaluations can differ between categories and even contexts. For example, occupational stimuli may tap into concepts of prestige but not necessarily concepts of wealth, adding a level of complexity to the task. In comparison, in the AMP conducted by Moore-Berg and Karpinski (2021) the consistent pro-high-class bias may be explained by the stimuli set used, which essentially compared lay clothing to business attire.

Third, if we look at the explicit ratings of SES: the low status primes in Study 1 (M = 4.11) received higher ratings than the low status primes in Study 2 (M = 2.33). That is, the low status manipulation in Study 2 was perhaps a better indicator than occupational clothing. However, the high status primes in Study 1 (M = 6.80) received higher explicit ratings than the high status primes in Study 2 ((M = 5.18). This raises questions about how to effectively manipulate social status in photographic stimuli and whether the effects of interest are a "high social status is good" effect or a "low social status is bad" effect.

Class and Race

Note that these studies were exploratory in that it was unclear exactly how social class would impact participant decision-making. However, there were several possibilities: First, the effects of interest could have appeared as a pro-high-class bias with no effect of race. Second, there could have been separate pro-White and pro-high-class biases indicating that participants' responses reflect positive associations towards these separate categories but not towards the unique combinations. Third, there could have been a pro-White high-class interaction indicating that participants' responses demonstrate a positive association towards a unique intersecting category.

Notably, the pro-Black bias observed is inconsistent with these hypotheses and the broader literature, which would predict that White primes are associated with positivity. In particular, the first prediction is challenged given that it is based on the assumption of an evolutionary basis for prioritizing life history information over race information (Williams et al., 2016). This past research predicts that race would be used as a proxy in the absence of social class (or ecology), resulting in positive White associations. In the presence of social class (ecology) information, then the social class information will be used over that of race

information. However, even in the case of a main effect of social class in Study 2, there is still an observable pro-Black preference. While this does not disprove this position, it does raise questions about the conditions in which race provides valuable information.

Moving forward, the results from Study 2 provide moderate support for the second possibility. Participants' responses are influenced by their associations of the individual categories and not the unique combinations such as high-class Blacks or low-class Whites. These findings are somewhat consistent with the AMP studies conducted by Moore-Berg and Karpinski (2021), which indicate participants demonstrated a pro-high-class bias with a marginal but non-significant pro-Black bias in Study 1 and pro-high class and pro-Black bias in Study 2. At the same time, however, prior work has also shown combined category effects (Correll, Wittenbrink, Park, Judd, & Goyle, 2011; Moore-Berg et al., 2017) and has documented differences in attitudes towards intersecting identities (Moore-Berg & Karpinski, 2019).

With that in mind, several important limitations make interpretation difficult. The first mentioned earlier is that social class may be evaluated across several dimensions (Mattan et al., 2017). Thus, it may not always be clear how participants interpret different social class manipulations. For example, while occupational clothing may represent high or low social class, they may do so for different reasons. Perhaps professors are seen as holding some prestige, while a physician is seen to have both prestige and wealth. In addition, it may be informative for future work to consider whether the type of status representation influences the degree to which participants create or recognize the combined categories. An additional limitation is that comparisons across different study designs is made difficult because each emphasize different demand characteristics (categorization, shooting decisions, or misattribution task). It may be the case that the instructions set by the task change the accessibility of relevant associations; to test

this, multiple studies using the same set of social class manipulations across different designs would be necessary. Thus, the studies conducted in this thesis do not conclusively rule out any pattern of findings but do raise interesting avenues of new research in this domain.

Conclusion

Despite these limitations, the present study has enhanced our understanding of the relationship between race and social class in person evaluation. It is evident that social class is an important category to consider in tandem with race, but significantly more work is needed to understand how and under what conditions does it affect categorization and evaluation.

APPENDIX

Tables and Figures

Table 1 *Mean proportion of pleasant responses in Study 1 by Race and Social Class*

Prime Race	Prime Class	Mean
Black	High	55.29 (20.71)
Black	Low	55.99 (21.30)
White	High	53.20 (19.81)
White	Low	53.01 (20.18)

Note: Descriptive statistics for Study 1. Standard deviations are presented in parentheses.

Table 2 *Multilevel Logistic Regression Predicting Target Pleasantness from Race and Social Class in Study 1*

	Rating:			
Predictors	*Odds Ratios*	*CI*	*Z statistic*	*p*
Intercept	1.23	1.09 – 1.39	3.30	0.001
Prime Race	1.07	1.02 – 1.11	2.87	0.004
Prime Class	0.99	0.96 – 1.03	-0.40	0.692
Prime Race x Class	0.99	0.96 – 1.02	-0.66	0.506
Random Effects	*N*	*Variance*		
Participant	289	1.05		
Participant x Prime Race		0.25		
Participant x Prime Class		0.07		
Participant x Prime Race x Class		0.01		
Prime	72	0.00		
Target	144	0.02		
Observations	41,760			

Note: Race and Class were effects coded (1 = Black, -1 = White; 1 = High class, -1 = Low class).

Table 3 *Multilevel Logistic Regression Predicting Target Pleasantness from Race, Social Class, and Attractiveness in Study 1*

	Rating:			
Predictors	*Odds Ratios*	*CI*	*Z statistic*	*p*
Intercept	1.23	1.09 – 1.39	3.32	0.001
Prime Race	1.06	1.02 – 1.09	3.48	<0.001
Prime Class	1.00	0.97 – 1.03	-0.14	0.889
Prime Att.	1.07	1.01 – 1.14	2.27	0.023
Prime Race x Att.	0.96	0.91 – 1.02	-1.23	0.220
Prime Race x Class	0.99	0.96 – 1.02	-0.59	0.558
Prime Class x Att.	0.98	0.92 – 1.04	-0.72	0.469
Prime Race x Class x Att.	0.93	0.88 – 0.99	-2.29	0.022
Random Effects	*N*	*Variance*		
Participant	289	1.02		
Prime	72	0.00		
Target	144	0.02		
Observations	41,760			

Note: Race and Class were effects coded (1 = Black, -1 = White; 1 = High class, -1 = Low class). Attractiveness was mean centered.

Table 4 *Mean proportion of pleasant responses in Study 2 by Race and Social Class*

Prime Race	Prime Class	Mean
Black	High	59.90 (16.69)
Black	Low	54.69 (17.53)
White	High	57.73 (16.95)
White	Low	52.58 (17.63)

Note: Descriptive statistics for Study 2. Standard deviations are presented in parentheses.

Table 5 *Multilevel Logistic Regression Predicting Target Pleasantness from Race and Social Class in Study 2*

	Rating:			
Predictors	*Odds Ratios*	*CI*	*Z statistic*	*p*
Intercept	1.34	1.24 – 1.46	6.97	<0.001
Prime Race	1.05	1.02 – 1.08	3.29	0.001
Prime Class	1.13	1.08 – 1.18	4.99	<0.001
Prime Race x Class	1.00	0.98 – 1.02	0.14	0.890
Random Effects	*N*	*Variance*		
Participant	293	0.44		
Participant x Prime Race		0.03		
Participant x Prime Class		0.13		
Participant x Prime Race x Class		0.00		
Prime	72	0.00		
Target	144	0.02		
Observations	42,192			

Note: Race and Class were effects coded (1 = Black, -1 = White; 1 = High class, -1 = Low class).

Table 6 *Multilevel Logistic Regression Predicting Target Pleasantness from Race, Social Class, and Attractiveness in Study 2*

	Rating:			
Predictors	*Odds Ratios*	*CI*	*Z statistic*	*p*
Intercept	1.03	1.24 – 1.46	7.18	<0.001
Prime Race	1.03	1.01 – 1.06	2.88	0.004
Prime Class	1.11	1.09 – 1.14	9.27	<0.001
Prime Att.	1.11	1.09 – 1.14	1.88	0.060
Prime Race x Att.	0.94	0.90 – 0.99	-2.37	0.018
Prime Race x Class	1.00	0.97 – 1.02	-0.32	0.751
Prime Class x Att.	1.04	0.99 – 1.09	1.40	0.162
Prime Race x Class x Att.	1.02	0.97 – 1.07	0.71	0.478
Random Effects	*N*	*Variance*		
Participant	293	0.42		
Prime	72	0.02		
Target	144	0.00		
Observations	42,192			

Note: Race and Class were effects coded (1 = Black, -1 = White; 1 = High class, -1 = Low class). Attractiveness was mean centered.

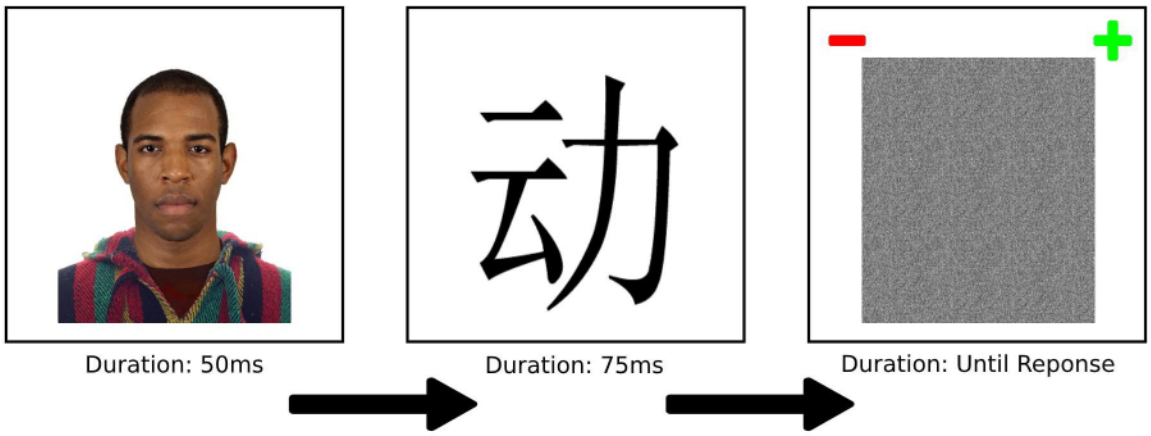

Figure 1 *A typical AMP trial designed to assess affective responses to racial groups.* Participants first see a picture of a Black or White prime (screen 1), then see a neutral Chinese symbol (screen 2). Participants must indicate whether the Chinese symbol is more pleasant or more unpleasant than the typical symbol using keyboard responses; responses are made during a pattern mask which remains visible until a response is registered (screen 3).

(a) Low (b) High

(c) Low (d) High

(e) Low (f) High

Figure 2 *Examples of low- and high-class occupations for three stimuli.*

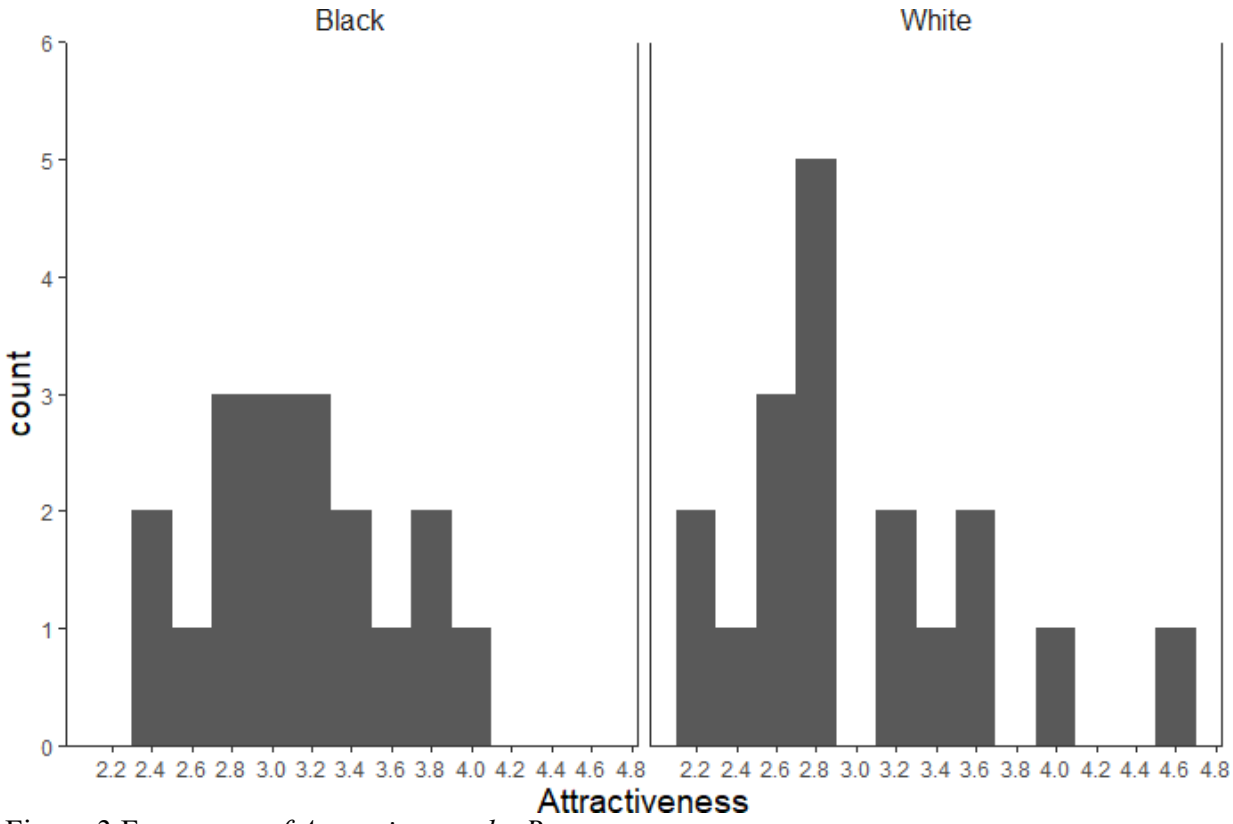

Figure 3 F*requency of Attractiveness by Race.*

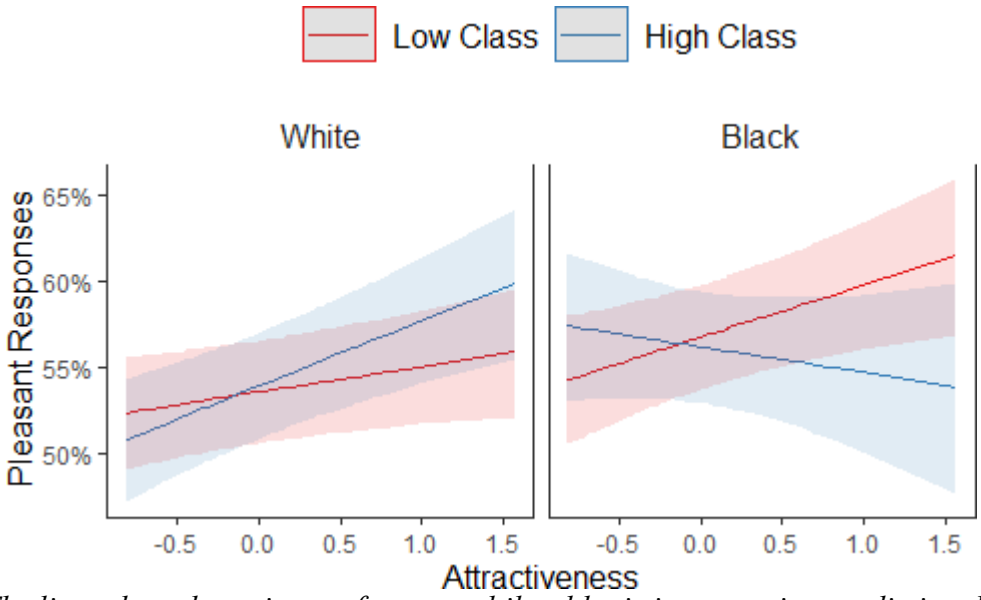

Figure 4 *The lines show the estimates from a multilevel logistic regression predicting the odds of pleasant response from race, social class, and attractiveness.* The shaded areas represent 95% CI.

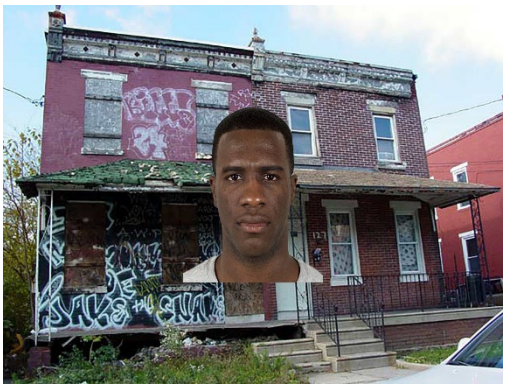

(a) Low

(b) High

(c) Low

(d) High

(e) Low

(f) High

Figure 5 *Examples of low- and high-class background conditions for three stimuli.*

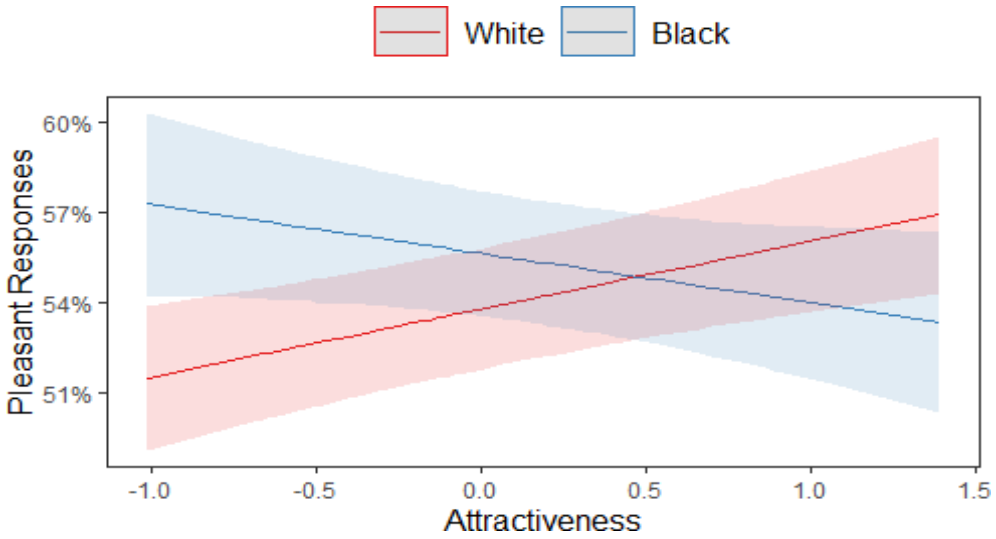

Figure 6 *The lines show the estimates from a multilevel logistic regression predicting the odds of pleasant response from race, social class, and attractiveness.* The shaded areas represent 95% CI. Note that social class isn't shown as it was not significant in the interaction.

REFERENCES

REFERENCES

Agthe, M., Spörrle, M., & Maner, J. K. (2011). Does being attractive always help? positive and negative effects of attractiveness on social decision making. Personality and Social Psychology Bulletin, 37(8), 1042–1054.

Bargh, J. A. (1989). Conditional automaticity: Varieties of automatic influence in social perception and cognition. Unintended thought, 3, 51–69.

Brown, L. M., Boniecki, K. A., & Walters, A. M. (2004). Intergroup flexibility and people's views of african americans. International Journal of Intercultural Relations, 28(5), 373–398.

Bullock, H. E. (1999). Attributions for poverty: A comparison of middle-class and welfare recipient attitudes 1. Journal of Applied Social Psychology, 29(10), 2059–2082.

Correll, J., Wittenbrink, B., Park, B., Judd, C. M., & Goyle, A. (2011). Dangerous enough: Moderating racial bias with contextual threat cues. Journal of experimental social psychology, 47(1), 184–189.

Cozzarelli, C., Wilkinson, A. V., & Tagler, M. J. (2001). Attitudes toward the poor and attributions for poverty. Journal of social issues, 57(2), 207–227.

Dasgupta, N., & Asgari, S. (2004). Seeing is believing: Exposure to counterstereotypic women leaders and its effect on the malleability of automatic gender stereotyping. Journal of experimental social psychology, 40(5), 642–658.

Dasgupta, N., DeSteno, D., Williams, L. A., & Hunsinger, M. (2009). Fanning the flames of prejudice: the influence of specific incidental emotions on implicit prejudice. Emotion, 9(4), 585.

Dasgupta, N., & Greenwald, A. G. (2001). On the malleability of automatic attitudes: combating automatic prejudice with images of admired and disliked individuals. Journal of personality and social psychology, 81(5), 800.

Dasgupta, N., & Rivera, L. M. (2008). When social context matters: The influence of long–term contact and short–term exposure to admired outgroup members on implicit attitudes and behavioral intentions. Social Cognition, 26(1), 112–123.

DeSteno, D., Dasgupta, N., Bartlett, M. Y., & Cajdric, A. (2004). Prejudice from thin air: The effect of emotion on automatic intergroup attitudes. Psychological Science, 15(5), 319–324.

Devine, P. G. (1989). Stereotypes and prejudice: Their automatic and controlled components. Journal of personality and social psychology, 56(1), 5.

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.

Dunton, B. C., & Fazio, R. H. (1997). An individual difference measure of motivation to control prejudiced reactions. Personality and Social Psychology Bulletin, 23(3), 316–326.

Durante, F., Tablante, C. B., & Fiske, S. T. (2017). Poor but warm, rich but cold (and competent): Social classes in the stereotype content model. Journal of Social Issues, 73(1), 138–157.

Fazio, R. H. (2001). On the automatic activation of associated evaluations: An overview. Cognition & Emotion, 15(2), 115–141.

Fazio, R. H., & Olson, M. A. (2003). Implicit measures in social cognition research: Their meaning and use. Annual review of psychology, 54(1), 297–327.

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.

Gawronski, B. (2009). Ten frequently asked questions about implicit measures and their frequently supposed, but not entirely correct answers. Canadian Psychology/Psychologie canadienne, 50(3), 141.

Gawronski, B., Cunningham, W. A., LeBel, E. P., & Deutsch, R. (2010). Attentional influences on affective priming: Does categorisation influence spontaneous evaluations of multiply categorisable objects? Cognition and Emotion, 24(6), 1008–1025.

Gawronski, B., Rydell, R. J., Vervliet, B., & De Houwer, J. (2010). Generalization versus contextualization in automatic evaluation. Journal of Experimental Psychology: General, 139(4), 683.

Greenwald, A. G., & Banaji, M. R. (1995). Implicit social cognition: attitudes, self-esteem, and stereotypes. Psychological review, 102(1), 4.

Greenwald, A. G., McGhee, D. E., & Schwartz, J. L. (1998). Measuring individual differences in implicit cognition: the implicit association test. Journal of personality and social psychology, 74(6), 1464.

Huddy, L., & Feldman, S. (2009). On assessing the political effects of racial prejudice. Annual Review of Political Science, 12, 423–447.

Judd, C. M., Westfall, J., & Kenny, D. A. (2012). Treating stimuli as a random factor in social psychology: A new and comprehensive solution to a pervasive but largely ignored problem. Journal of Personality and Social Psychology, 103(1), 54–69. doi: 10.1037/a0028347

Katz, D., & Braly, K. (1933). Racial stereotypes of one hundred college students. The Journal of Abnormal and Social Psychology, 28(3), 280.

Klonis, S. C., Plant, E. A., & Devine, P. G. (2005). Internal and external motivation to respond without sexism. Personality and Social Psychology Bulletin, 31(9), 1237–1249.

Kramer, R. S., Jones, A. L., & Sharma, D. (2013). Sequential effects in judgements of attractiveness: the influences of face race and sex. PLoS One, 8(12), e82226.

Kühnen, U., Schießl, M., Bauer, N., Paulig, N., Pöhlmann, C., & Schmidthals, K. (2001). How robust is the iat? measuring and manipulating implicit attitudes of east-and west-germans. Zeitschrift für Experimentelle Psychologie.

Lei, R. F., & Bodenhausen, G. V. (2017). Racial assumptions color the mental representation of social class. Frontiers in psychology, 8, 519.

Livingston, R. W., & Brewer, M. B. (2002). What are we really priming? cue-based versus category-based processing of facial stimuli. Journal of personality and social psychology, 82(1), 5.

Ma, D. S., & Correll, J. (2011). Target prototypicality moderates racial bias in the decision to shoot. Journal of Experimental Social Psychology, 47(2), 391–396.

Maddux, W. W., Barden, J., Brewer, M. B., & Petty, R. E. (2005). Saying no to negativity: The effects of context and motivation to control prejudice on automatic evaluative responses. Journal of Experimental Social Psychology, 41(1), 19–35.

Mattan, B. D., Kubota, J. T., & Cloutier, J. (2017). How social status shapes person perception and evaluation: A social neuroscience perspective. Perspectives on Psychological Science, 12(3), 468–507.

Mattan, B. D., Kubota, J. T., Li, T., Venezia, S. A., & Cloutier, J. (2019). Implicit evaluative biases toward targets varying in race and socioeconomic status. Personality and Social Psychology Bulletin, 45(10), 1512–1527.

Mitchell, J. P., Nosek, B. A., & Banaji, M. R. (2003). Contextual variations in implicit evaluation. Journal of Experimental Psychology: General, 132(3), 455.

Moore-Berg, & Karpinski, A. (2019). An intersectional approach to understanding how race and social class affect intergroup processes. Social and Personality Psychology Compass, 13(1), e12426.

Moore-Berg, & Karpinski, A. (2021). Race and social class as intersecting social categories: An analysis of implicit and explicit attitudes. Social Psychology.

Moore-Berg, Karpinski, A., & Plant, E. A. (2017). Quick to the draw: How suspect race and socioeconomic status influences shooting decisions. Journal of Applied Social Psychology, 47(9), 482–491.

Neely, J. H. (1977). Semantic priming and retrieval from lexical memory: Roles of inhibitionless spreading activation and limited-capacity attention. Journal of experimental psychology: general, 106(3), 226.

Niemann, Y. F., Jennings, L., Rozelle, R. M., Baxter, J. C., & Sullivan, E. (1994). Use of free responses and cluster analysis to determine stereotypes of eight groups. Personality and Social Psychology Bulletin, 20(4), 379–390.

Oh, D., Shafir, E., & Todorov, A. (2020). Economic status cues from clothes affect perceived competence from faces. Nature Human Behaviour, 4(3), 287–293.

Payne, B. K., Cheng, C. M., Govorun, O., & Stewart, B. D. (2005). An inkblot for attitudes: affect misattribution as implicit measurement. Journal of personality and social psychology, 89(3), 277.

Payne, B. K., & Lundberg, K. (2014). The affect misattribution procedure: Ten years of evidence on reliability, validity, and mechanisms. Social and Personality Psychology Compass, 8(12), 672–686.

Richeson, J. A., & Ambady, N. (2001). Who's in charge? effects of situational roles on automatic gender bias. Sex Roles, 44(9-10), 493–512.

Richeson, J. A., & Ambady, N. (2003). Effects of situational power on automatic racial prejudice. Journal of Experimental Social Psychology, 39(2), 177–183.

Scherer, L., & Lambert, A. J. (2009). Counterstereotypic exemplars in context: Evidence for intracategory differentiation using implicit measures. Social Cognition, 27(4), 522–549.

Schwarz, N., & Clore, G. L. (1996). Feelings and phenomenal experiences. In E. T. Higgins & A. Kruglanski (Eds.), Social psychology: Handbook of basic principles (p. 385-407). New York: Guilford Press.

Spencer, S., Fein, S., Wolfe, C., Fong, C., & Dunn, M. (1998). Stereotype activation under cognitive load: The moderating role of self-image threat. Personality and Social Psychology Bulletin, 24, 1139–1152.

Tan, A., Zhang, Y., Zhang, L., & Dalisay, F. (2009). Stereotypes of african americans and media use among chinese high school students. The Howard Journal of Communications, 20(3), 260–275.

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.

Wittenbrink, B., Judd, C. M., & Park, B. (2001). Spontaneous prejudice in context: Variability in automatically activated attitudes. Journal of personality and social psychology, 81(5), 815.

Wu, S. J., Bai, X., & Fiske, S. T. (2018). Admired rich or resented rich? how two cultures vary in envy. Journal of Cross-Cultural Psychology, 49(7), 1114–1143.