Journal of Experimental Social Psychology 123 (2026) 104870

Contents lists available at ScienceDirect

Journal of Experimental Social Psychology

journal homepage: www.elsevier.com/locate/jesp

Registered Report Stage 2: Full Article

Intersectional prejudice in context: Do naturalistic background images moderate implicit bias?☆,☆☆

Christopher D. Petsko [a,1,*], Paul Connor [b,1]

[a] *The University of North Carolina at Chapel Hill*
[b] *Meta Platforms, Inc.*

A R T I C L E I N F O

Keywords:
Intersectional stereotyping
Prejudice
Implicit bias
Normative fit
Reproducibility

A B S T R A C T

Intersectional stereotyping is a context-dependent phenomenon. For example, in some contexts, perceivers stereotype Black women specifically as *Black women.* In other contexts, perceivers stereotype Black women as *Black* more generally, or as *women* more generally. According to one recent model of stereotyping—the lens model—a factor that ought to sharpen perceivers' attention on race (more than gender) or on gender (more than race) is normative fit, or the extent to which a context itself is stereotypically associated with one social identity over another. The purpose of the present analysis was to test whether normative fit cues, operationalized in the form of background images against which targets are depicted, can indeed moderate patterns of implicit racial and gender prejudice. Toward this end, we conducted three experiments (two preliminary, one registered; total $N =$ 1269) in which participants encountered intersectional targets either against background images that normatively fit the concept of race (e.g., a graffitied street corner), or against background images that did not (e.g., a gendered bathroom entryway). Surprisingly, we found no evidence that normative fit cues moderated the degree to which participants exhibited racial prejudice or gender prejudice. Collectively, these experiments suggest that perhaps normative fit cues in and of themselves are not psychologically powerful enough to augment lens salience. More practically, these experiments suggest that previously reported empirical findings—findings suggesting that implicit attitudes vary as a function of the background images against which targets are depicted—may no longer hold in contemporary samples of U.S. adults.

1. Introduction

Everyone we encounter can be categorized in multiple ways: by their gender groups, their age groups, their sexual orientation groups, and the like. Yet since its inception, psychological research on stereotyping has largely overlooked this complexity (Cole, 2009; Goff & Kahn, 2013; Plaut, 2010). For example, research on racial stereotyping has tended to focus on stereotypes directed toward Black men, and research on gender stereotyping has tended to focus on stereotypes directed toward White women. As a result, there has been a relative scarcity of research on how Black women are stereotyped (Purdie-Vaughns & Eibach, 2008). In an attempt to correct for oversights such as these—and indeed, to garner a more complete understanding of stereotyping and prejudice—psychologists have increasingly taken an *intersectional* look at the topic of stereotyping (for reviews, see: Hudson et al., 2024; Kang & Bodenhausen, 2015; Nicolas et al., 2017; Lei et al., 2023; Remedios & Vinluan, 2024; Rosette et al., 2025).

The term intersectional was born out of Black feminist scholarship. In its original formulation, the concept of intersectionality was invoked as a tool for characterizing the ways in which systems of oppression mutually reinforce one another (e.g., Crenshaw, 1991; King, 1988), and for questioning whether positivist systems of knowledge could ever

☆ This research was supported by a National Science Foundation Postdoctoral Research Fellowship (SRF-2104594), which was awarded to Paul Connor. The authors would like to thank Will Maddux for providing helpful comments on a draft of this manuscript. Pre-registration documentation, data files, codebooks, survey materials, and R scripts associated with this article are available on the Open Science Framework (OSF) website: https://osf.io/43fbp/.

☆☆ This paper has been recommended for acceptance by Evava Pietri

* Corresponding author at: Kenan-Flagler Business School, University of North Carolina at Chapel Hill, 300 Kenan Center Drive, Chapel Hill, NC 27599, United States.

E-mail address: christopher_petsko@kenan-flagler.unc.edu (C.D. Petsko).

[1] Both authors contributed equally to the development of this manuscript and should be considered joint first authors.

https://doi.org/10.1016/j.jesp.2025.104870
Received 13 November 2025; Received in revised form 18 December 2025; Accepted 19 December 2025
Available online 31 December 2025

adequately capture the experiences of multiply marginalized individuals (Else-Quest & Hyde, 2016a, 2016b). In psychological science, intersectionality often refers more narrowly to an epistemological stance in which one questions whether basic psychological phenomena, like stereotyping, depend on the multiple social groups to which individuals belong (Cole, 2009). In accordance with this approach, the present analysis defines intersectional stereotyping as the study of how perceivers stereotype targets as a function of targets' multiple social identities (Ghavami & Peplau, 2013; Hall et al., 2019; Petsko and Bodenhausen, 2019b).

The study of intersectional stereotyping has produced a variety of novel scientific insights. One key insight is that any given person can face distinct sets of stereotypes depending on how they are construed by perceivers in a given social context (Petsko & Bodenhausen, 2020; Remedios & Vinluan, 2024). For example, there appear to be some contexts in which gay Black men are categorized as *Black* and face racial stereotyping (Petsko & Bodenhausen, 2019a), other contexts in which gay Black men are categorized as *gay* and face sexual stereotyping (Petsko & Bodenhausen, 2019b), and still other contexts in which gay Black men are categorized not as *gay* or as *Black* per se, but as *gay Black men* in particular (Pedulla, 2014; Preddie & Biernat, 2021). In these latter contexts, gay Black men appear to face not generalized sexual or racial stereotypes, but stereotypes that are specific to their intersecting identities. For example, when gay Black men are construed as *gay Black men* (rather than as *gay* more generally or as *Black* more generally), they are evaluated as having higher leadership potential (Wilson et al., 2017). This finding is at odds with the observation that leadership is typically associated with Whiteness more than with Blackness (Petsko & Rosette, 2023), and also with being straight more than with being gay (Pellegrini et al., 2020).

The purpose of this registered report is to critically examine one theoretical factor that is implicated in governing how perceivers construe intersectional targets: the visual contexts in which targets are embedded. We begin by briefly describing a contemporary model of intersectional stereotyping—the lens model—which invokes the importance of social contexts in shaping how stereotyping unfolds (Petsko et al., 2022). We then discuss data from two initial experiments (Experiments 1a and 1b) that test whether manipulating the visual contexts in which intersectional targets are embedded impacts the magnitude with which perceivers exhibit racial prejudice and gender prejudice against these targets, respectively. After describing what appear to be surprising results—results that raise the question of whether visual contexts have any bearing at all on implicit prejudice—we turn our attention to a registered experiment (Experiment 2). This registered experiment, like the initial experiments, sought to examine the question of whether visual contexts influence perceivers' prejudices, but did so using methods that were more closely aligned with those used in past published work on this topic.

1.1. The lens model of intersectional stereotyping

The lens model posits that people possess a repertoire of different *lenses* in their minds that they can use as bases for stereotyping intersectional targets, and critically, that people use just one lens at a time in a given social environment (Petsko et al., 2022). In this perspective, lenses are construed as situationally activated, identity-specific schemas for categorizing and forming impressions of others. The idea is that perceivers have a lens for the concept of race, a lens for the concept of gender, and so forth. And when one lens comes into focus, perceivers are expected to sharpen their focus on social identities that are relevant to that lens, but lose focus, at least momentarily, on identities irrelevant to that lens (see also: Bodenhausen & Macrae, 1998; Sinclair & Kunda, 1999 and Turner et al., 1987). Importantly, the lens model also argues that lenses can be specifically intersectional: that perceivers can have a lens for the concept of race and gender in tandem, for example.

If perceivers indeed focus on one identity—or one intersection of identities—at a time when construing intersectional targets, a relevant question concerns what it is that causes perceivers to use one lens over others. In their review of the literature on intersectional person perception, Petsko and Bodenhausen (2020) identified four theoretical factors that purportedly impact lens salience: (1) *accessibility*, or the ease with which a social lens can be retrieved from memory (Bruner, 1957; Higgins, 1996); (2) *fit*, or the extent to which a lens "explains," normatively or comparatively, patterns of intergroup behavior in a context (Oakes et al., 1991); (3) *perceiver goals*, or desired end states that motivate the use of some lenses over alternatives (e.g., Neel & Lassetter, 2019; Sinclair & Kunda, 1999); and (4) *distinctiveness*, or the extent to which a lens-associated category is rare and thus attention-grabbing in a social context (Hamilton & Gifford, 1976; McGuire et al., 1978). In the present analysis, we focus our attention on the concept of fit—specifically on the concept of normative fit.

1.2. Normative fit and its operationalization in visual context

According to the principle of *normative fit*, perceivers are more likely to use a given lens for construing others when the social category implied by that lens is normatively (i.e., stereotypically) aligned with a given social context (Oakes, 1987; Oakes et al., 1991; Turner et al., 1987). A key reason why perceivers are expected to do this is because lenses are thought to serve epistemic goals (Petsko et al., 2022). For example, in the United States, the criminal justice system has been argued to be a stereotypically racialized social context (e.g., Eberhardt et al., 2004). As such, perceivers in criminal justice contexts are expected to use the lens of race more than alternative lenses as a basis for construing social targets. This is because in criminal justice contexts, race (more than, say, gender) is expected to *feel* epistemically useful to perceivers—to give perceivers a subjective sense that they can understand and predict the behavior of others.

Past research is indeed suggestive of the possibility that criminal sentencing contexts (which normatively fit the concept of race) sharpen perceivers' focus on targets' racial groups. For example, when making judgments about the culpability of Black vs. White defendants who have been accused of stereotypically Black crimes, mock jurors exhibit detectible levels of anti-Black bias, but not—in these moments—detectible levels of anti-gay bias (Petsko & Bodenhausen, 2019a, Experiment 1). Likewise, in Weapons Identification Tasks (Payne, 2001), which make the concept of crime situationally salient, perceivers have been shown to exhibit detectable levels of anti-Black bias, but not detectable levels of age-related bias (Todd et al., 2016a; 2016b). In sum, existing research suggests that when the lens of race normatively fits a social environment, perceivers indeed seem to focus their attention on targets' race, and to ignore targets' other social identities (i.e., whether targets are gay vs. straight, or old vs. young).

One method of experimentally manipulating normative fit—which we will focus on throughout the present manuscript—is via the visual context against which targets are depicted. Indeed, past research has revealed, in at least a few notable examples, that visual contexts can moderate patterns of implicit prejudice. In one study, the magnitude with which perceivers expressed anti-Black bias was observed to be greater when targets were depicted against the backdrop of a graffitied street corner than against the backdrop of a church (Wittenbrink et al., 2001). Likewise, anti-Black bias has been found to be greater when targets are depicted against the backdrop of a jail cell rather than against the backdrop of church (Allen et al., 2010; Barden et al., 2004).[2] To be clear, the concept of normative fit was not explicitly named in these prior experiments. Nevertheless, it is plausible that background images depicting graffitied street corners and jail cells, respectively, provided greater normative fit to the concept of race than did churches—and in turn, that they came to activate the lens of race in the minds of perceivers.

1.3. Past explanations for contextual moderation of implicit bias

If the experiments cited in the prior section did not explicitly name the factor of normative fit, then what explanatory factors did these experiments implicate? Researchers have offered varying accounts for why visual contexts might influence the magnitude of implicit bias. According to Wittenbrink et al. (2001), visual contexts influence the magnitude of implicit anti-Black bias because people engage in *gestalt-based processing of social categories*. The idea in this analysis is that certain contexts (e.g., graffitied street corners) bring to mind negatively valenced aspects of a social category (i.e., the category *Black people*), whereas other contexts (e.g., churches) instead bring to mind positively valenced aspects of that same social category. According to Barden et al. (2004), visual contexts influence the magnitude of implicit anti-Black bias not because of gestalt-based processing of social categories, but instead because people engage in *role-based processing of social categories.* In this latter point of view, visual contexts are theorized to bring to mind specific social roles (e.g., the role of "prisoner" when one sees a jail cell; or the role of "church-goer" when one sees a church), which combines in perceivers' minds with a social category exemplars (causing perceivers to think specifically about *Black prisoners* vs. *White prisoners*, for example, or instead about *Black church-goers* vs. *White church-goers*). Although these prior investigations each offer different explanations for why visual contexts might moderate implicit anti-Black bias, they share in common a key assumption: that social contexts modulate how perceivers think about race itself. The lens model, of course, advocates for a different assumption—namely, that contexts do not modulate how perceivers think about race so much as whether perceivers think about race at all.

An issue we wish to raise before we describe the present experiments, which will expand upon in the general discussion, is that the present experiments are designed only for the purposes of testing whether visual contexts moderate patterns of implicit bias, as the lens model predicts. These experiments are not designed to adjudicate between different low-level explanations as to *why* visual contexts might do so. Suppose, for example, that depicting targets against a racialized background (e.g., a graffitied street corner) were to increase participants' racial bias relative to depicting those same targets against a gendered background (e.g., a bathroom at a mall). In this circumstance, it would be unclear what might account for such an increase in racial bias: do graffitied street corners bring to mind different social roles than bathrooms at the mall do (Barden et al., 2004)? Do graffitied street corners (relative to bathrooms) activate different aspects of racial representations in memory (Wittenbrink et al., 2001)? Or do people just use the lens of race more when confronted with graffitied street corners than they do when confronted with bathrooms at the mall (Petsko et al., 2022)? The fact that all three perspectives align with the possibility that implicit anti-Black bias may be elevated against the backdrop of a graffitied street corner than against the backdrop of a bathroom creates something of an explanatory confound.[3] Nevertheless, we wish to note that examining whether or not contextual moderation of implicit bias occurs at all is still a worthy scientific endeavor. Indeed, examining whether implicit biases are moderated by visual contexts enables a clear test of an important assumption of the lens model (if also, by extension, downstream implications of other models), and it likewise informs present-day researchers about the reproducibility of past research.

2. The present experiments

The present experiments test a key assumption of the lens model: that normative fit cues—operationalized in the form of background images against which targets are depicted—function to "switch on" the use of certain lenses (and in turn, certain forms of implicit prejudice), but to "switch off" the use of other lenses others (and in turn, other types of implicit prejudice). We begin our analysis with a discussion of two initial experiments (total $N = 804$) that tested, empirically, whether manipulating background images impacted the magnitude with which perceivers exhibited implicit racial prejudice and implicit gender prejudice, respectively. After describing the results of these experiments, which cast doubt on the notion that visual contexts indeed moderate implicit prejudice, we discuss a registered experiment (total $N = 465$). This registered experiment again examined the question of whether background images influence perceivers' implicit prejudices against intersectional targets. However, it did so using methods that were more closely aligned with past research on this topic.

Overall, then, this registered report enabled us to weigh in on three interrelated questions: (1) whether implicit evaluations are indeed moderated by the background images against which intersectional targets are depicted; (2) if so, whether activating specific lenses (e.g., the lens of race) reduces the activation of others (e.g., the lens of gender), as predicted by the lens model; and (3) whether an established moderator of implicit prejudice—at least according to three prior studies (Allen et al., 2010; Barden et al., 2004; Wittenbrink et al., 2001)—holds up to scrutiny when examined in a contemporary, highly powered set of experiments. Whereas the former two questions have implications for theories of intersectional stereotyping, the latter has implications for the reproducibility of past research.

3. Statement of transparency and openness

We report all manipulations, exclusions, and measured variables. Across experiments, raw effect sizes are reported in terms of unstandardized units (in this analysis, milliseconds), and standardized effect sizes are reported in terms of standard betas (βs).[4] Experiments 1a and 1b were both pre-registered in advance of data collection, but in light of their highly similar methods and results, we have aggregated the data from these two experiments into a single analysis (see the online supplement for a test demonstrating a lack of meaningful differences between these experiments). Given that the presentation of these experiments' results differs from what we initially pre-registered,

[2] To our knowledge, the articles mentioned here are the only three published articles that have yielded evidence for the idea that visual backdrops can moderate the extent to which perceivers exhibit anti-Black prejudice (Allen et al., 2010; Barden et al., 2004; Wittenbrink et al., 2001). A fourth paper (Maddux et al., 2005), which is conceptually related, also found evidence for the idea that visual backdrops of prisons (vs. churches) increase perceivers' anti-Black prejudice—but only on the condition that participants are low in the motivation to respond without prejudice.

[3] This all said, the lens model actually does make more straightforward predictions than past explanations do about what ought to happen to perceivers' levels of *gender* bias. In particular, the lens model—more parsimoniously than past explanations—argues not just that racial bias ought to become attenuated against the backdrop of a bathroom (vs. a graffitied street corner), but that gender bias ought to become exacerbated against this same backdrop.

[4] Throughout this manuscript, standardized effect sizes (βs) are computed by regressing z-standardized outcomes onto orthogonal condition contrasts (which sum to zero and always have a range of one; Judd et al., 2017). Under this approach, standard betas for main effect tests can be interpreted similarly to Cohen's *ds*.

readers should regard the reported results from these experiments as exploratory.[5] Experiment 2—the experiment that lies at the heart of the registered report—was confirmatory, pre-registered, and peer-reviewed in advance of data collection. Data files, materials, analysis scripts, and pre-registration documentation associated with this manuscript are available on the Open Science Framework (OSF) website: https://osf.io/43fbp/. All experiments were conducted under the oversight of the Institutional Review Board at UNC Chapel Hill (IRB Protocol Number: 22–2262).

4. Experiments 1a and 1b (single-category IAT)

We conducted two experiments (total $N = 804$) on the topic of whether implicit racial bias and gender bias, respectively, depend on whether intersectional targets are depicted in environments that normatively fit the lenses of race or of gender. According to the lens model, background contexts that normatively fit the lens of race ought to "switch on" racial bias, but also to "switch off" gender bias. Conversely, background contexts that normatively fit the lens of gender ought to do the reverse—that is, "switch on" gender bias, but "switch off" racial bias.

4.1. Method

Participants completed four iterations of the Single-Category Implicit Association Test (i.e., four SC-IATs) in which we measured the relative speed with which they associated targets (in every case, a mix of Black and White women and men) with positive vs. negative words. In SC-IATs, implicit racial bias is measured by the relative speed with which participants are faster to associate White (vs. Black) targets with positive (vs. negative) words. Likewise, implicit gender bias is measured by the relative speed with which participants are faster to associate women (vs. men) with positive (vs. negative) words.[6] As an aside, pro-woman biases on implicit measures are commonplace. Indeed, there are several prior papers documenting this type of gender bias (e.g., Aidman & Carroll, 2003; Connor et al., 2021).

Each SC-IAT included one of two sets of target people (always containing a mix of Black and White men and women, see online supplement), as well as a visual background context that either normatively fit the lens of race or the lens of gender (Fig. 1). Thus, both Experiments 1a and 1b contained within-person manipulations of target race (Black, White), target gender (women, men), and critically, visual context (normative fit with race, normative fit with gender).

4.1.1. Participants

For Experiment 1a, we recruited 401 participants from Prolific, but excluded seven participants for responding faster than 300 ms on 10 % or more of SC-IAT trials, and one for self-acknowledging low data quality (final $N = 393$; $M_{age} = 38.52$, $SD_{age} = 12.32$; 235 men, 150 women, 6 non-binary gender, 2 gender unreported; 304 White, 28 Latino, 25 Black, 14 East Asian, 5 South Asian, 17 other race). For Experiment 1b, we combined 212 undergraduates from the University of North Carolina at Chapel Hill (UNC) and 213 participants from Prolific, but excluded ten participants for responding faster than 300 ms on 10 % or more of SC-IAT trials, and four for self-acknowledging low data quality (final $N = 411$; $M_{age} = 30.50$, $SD_{age} = 13.76$; 210 women, 195 men, 4 non-binary gender, 2 gender unreported; 274 White, 37 Latino, 29 Black, 19 East Asian, 15 South Asian, 15 South-East Asian, 22 other race).

4.1.2. Procedure

All participants completed four Single-Category Implicit Association Tests (SC-IATs; Karpinski & Steinman, 2006) that were embedded into an online survey experiment. Experiments were completed via desktop computers using Inquisit Web software. Participants were told that the purpose of these online experiments was "to better understand the mechanisms underlying how we respond to others." Participants who were recruited from Prolific completed the experiments at a time and place of their choosing; student participants completed the study in laboratory data collection sessions (usually in groups of four to ten students at a time).

Target Photographs. We selected 40 full-body color photographs of targets at various intersections of race and gender (10 Black women, 10 Black men, 10 White women, and 10 White men) from a publicly available database containing target photos and norming data (Connor et al., 2021).[7] According to norming data, all targets were reliably categorized as belonging to their intended racial groups at least 88 % of the time.[8] The total pool of 40 targets was divided into two groups of 20 targets (Groups A and B, respectively). Groups A and B each contained 5 targets from each race-by-gender intersection (i.e., 5 Black women, 5 Black men, etc.). Targets in each of the race-by-gender intersections were matched as closely as possible with one another on prior ratings of perceived age, social class, and attractiveness (mean raters per target = 33.4; see the online supplement for more detail).

Visual Contexts. Experiments 1a and 1b employed two different background images that were selected as normatively fitting the concepts of gender or race, respectively. The background image that was selected to normatively fit the concept of gender was of a bathroom entryway with clearly marked men's and women's restrooms (see Fig. 1, left-hand side). The background image that was selected to normatively fit the concept of race was of a graffitied street corner similar to that used in Wittenbrink et al. (2001; see Fig. 1 right-hand side). Participants were told that the visual contexts had been added to the study to "make the experiment more like real life," and were told: "please keep the context in mind as you perform the task, and imagine the sights, sounds, and smells that would be present in that context." Of note, these instructions closely followed those that have been used in prior work (e.g., Allen et al., 2010; Barden et al., 2004).

SC-IATs. Participants completed four Single-Category Implicit Association Tasks (SC-IATs; Karpinski & Steinman, 2006) in which they categorized positively and negatively valenced words (e.g., "Marvelous," "Tragic") and target photographs as quickly as possible via timed computer key presses. In all blocks, positively valenced words were

[5] In addition, although this manuscript discusses Experiments 1a and 1b in terms of their implications for the lens model alone, these experiments were actually designed to test competing predictions. One prediction was that social contexts might augment whether race or gender becomes situationally activated in the minds of perceivers (as the lens model argues: Petsko et al., 2022); the other prediction was that gender bias might inevitably be greater than racial bias, regardless of social context (as gender-dominance findings suggest: Connor et al., 2023). See OSF for all original pre-registration documentation: https://osf.io/43fbp/.

[6] Experiments 1a and 1b used SC-IATs rather than regular IATs because, as noted in Footnote 4, these experiments were originally designed to test whether findings related to gender-dominance (which were established using SC-IATs: Connor et al., 2023) might be moderated by visual backdrops (as the lens model suggests: Petsko et al., 2022).

[7] Of note, prior experiments (e.g., Barden et al., 2004) used only male targets as stimuli, whereas our experiments used male and female targets. Exploratory analyses, reported in full in the online supplement, suggest that findings from Experiment 1a and 1b related to racial bias hold even when looking at trials that feature male stimuli only.

[8] In the online supplement, we report additional analyses in which we examined whether main findings from Experiment 1a and 1b might depend on how race-typical—according to norming data—targets were rated, on average. In brief, we found no evidence that targets who were rated as more race-typical (e.g., categorized as their intended race 100 % of the time) elicited significantly different response patterns than targets who were rated as less-race typical (e.g., categorized as their intended race 88 % of the time).

Fig. 1. *Visual Contexts and Target Presentations for Experiments 1a and 1b.*
Note. SC-IATs in Experiments 1a and 1b used the gendered visual context (restrooms) depicted in the left panel and the racialized visual context (street corner) depicted in the right panel, as well as the full-body target presentation (positively valenced blocks shown).

categorized as 'Good' to the left of screens and negatively valenced words were categorized as 'Bad' to the right of screens. In positively valenced blocks, targets were categorized as a 'Person' to the left of screens using the same key as positive words, and in negatively valenced blocks, targets were categorized as a 'Person' to the right of screens using the same key as negative words. In each SC-IAT, participants were randomly assigned to either complete two positively valenced blocks of 20 then 40 trials followed by two negatively valenced blocks of 20 and 40 trials, or vice versa, with negative blocks preceding positive blocks. All three categories of stimuli (positive words, negative words, and targets) were presented in equal proportions.[9]

The SC-IATs comprised each possible combination of the two visual contexts (graffitied street corner and gendered bathroom, respectively) and the two target groups (Groups A and B, respectively), presented in randomized order. The experiments, therefore, used a fully crossed 2 (visual context: race-fit vs. gender-fit) × 2 (target race: Black vs. White) × 2 (target gender: female vs. male) × 2 (block valence: positive vs. negative) within-subjects factorial design. The only difference between Experiments 1a and 1b was that Experiment 1a's word stimuli were presented approximately at targets' eyelines (vs. approximately at the middle of targets' torsos), whereas Experiment 1b's word stimuli were presented approximately at the middle of targets' torsos (vs. approximately at targets' eyelines).[10]

4.1.3. Power analysis

Estimates of power came from a bootstrapping technique in which we determined the smallest effects of interest that we could detect 90 % of the time given our analytic approach, final sample size, and an alpha level of 0.05. According to this technique, we had 90 % power to detect two-way interactions between target race (or gender) and block valence—the key test for identifying implicit bias—as small as $\beta = 0.11$; we had 90 % power to detect three-way interactions between target race (or gender), block valence, and visual context—the key test for identifying moderation of implicit bias by background image—as small as $\beta = 0.08$; and we had 90 % power to detect simple effects of block valence within race (or gender)—the key test for determining if a subgroup (e.g., Black targets) elicited an evaluative preference—as small as $\beta = 0.09$. Details on precise power simulation methods can be found in the online supplement.

4.2. Results

To analyze the data from Experiments 1a and 1b, response latencies (in milliseconds) for target photographs were regressed, in a multilevel model, onto within-person contrast codes that represented the full 2 (visual context: race vs. gender) × 2 (target race: Black vs. White) × 2 (target gender: female vs. male) × 2 (block valence: positive vs. negative) factorial design. To account for the nested data structure, this model included random intercepts for both targets and participants; to account for random variation in evaluation positivity of both targets and participants, this model included random slopes of block valence within each target and within each participant. Based on the recommendations of Connor et al. (2021), response times below 100 ms (0.03 % of trials) and above 4000 ms were excluded (0.2 % of trials), and trials in which participants hit the incorrect key were penalized by replacing the latency on these trials with participants' individual mean response latencies in the corresponding trial type (positive or negative valence) plus 600 milliseconds (4.6 % of trials).

4.2.1. Did implicit racial bias depend on background Context?

According to the lens model, implicit racial bias should emerge in contexts that normatively fit the concept of race, but not in contexts that normatively fit the concept of gender. In order to examine whether this was the case, participants' response latencies were subjected to the 2 (visual context) × 2 (target race) × 2 (target gender) × 2 (block valence) factorial analysis described above. As one might expect, this analysis revealed that participants, on average, exhibited a racial bias. In particular, this analysis yielded a significant two-way interaction between target race and block valence: $\beta = 0.11$, $b = 37.21$, 95 % CI[28.13, 46.29], $t(39.50) = 8.03$, $p < .001$,[11] the nature of which was that participants were significantly faster to associate White targets with good (vs. bad) concepts [$\beta = -0.09$, $M_{\text{diff}} = -29.37$ ms, 95 % CI [-40.57,

[9] It is common for SC-IATs to alter presentation ratios of each category of stimuli to equalize the number of stimuli categorized to the right and left of screens (Karpinski & Steinman, 2006). However, following Connor et al. (2021) we opted to present each category of stimuli equally to increase the overall number of target photograph trials, which are the sole focus of our analyses.

[10] This change was made to test if presenting words at targets' eye line in Experiment 1a had caused its results to differ from those of Connor et al. (2021), where target words were presented approximately at the middle of targets' torsos. We found no credible evidence that where the stimulus words were presented influenced results.

[11] Analyses throughout this manuscript were conducted using the "lme4" and "lmerTest" packages in R (Bates et al., 2015; Kuznetsova et al., 2017), which use the Satterthwaite approximation to estimate degrees of freedom. Degrees of freedom are all estimated to two decimal places. Fluctuations in degrees of freedom are attributable to approximation variability rather than to missing data. The reason why some t-tests report degrees of freedom as high as 128,113 is because this analysis was conducted at the level of trials rather than at the level of participants.

−18.17], t(77.76) = −5.14, $p < .001$] than they were to associate Black targets with good (vs. bad) concepts [β = 0.02, M_{diff} = 7.63 ms, 95 % CI [−1.98 ms, 17.23 ms], t(239.86) = 1.56, p = .121]. Surprisingly, however, this racial bias—that is, the magnitude of the two-way interaction between target race and block valence—was *not* significantly moderated by the visual context in which targets were embedded. There was no evidence of a three-way interaction between target race, block valence, and visual context: β = 0.01, b = 2.60, 95 % CI[−10.53, 15.73], t(128,112.50) = 0.39, p = .698. Thus, in contrast to the propositions of the lens model, visual contexts that were designed to normatively fit the concept of race made racial bias no greater than visual contexts designed to normatively fit the concept of gender (see Fig. 2, dark bars).

4.2.2. Did implicit gender bias depend on background Context?

According to the lens model, implicit gender bias should emerge in contexts normatively fitting the concept of gender, but not in contexts normatively fitting the concept of race. Tests of whether this ended up being the case came from the same 2 (visual context) × 2 (target race) × 2 (target gender) × 2 (block valence) factorial analysis described in the previous section. Overall, this analysis revealed evidence of a gender bias—that is, evidence of a significant two-way interaction between target gender and word valence: β = −0.10, b = −32.38, 95 % CI [−41.46, −23.30], t(39.49) = −6.99, $p < .001$. The nature of this interaction was that participants were significantly faster to associate women with good (vs. bad) concepts [β = −0.08, M_{diff} = −26.93 ms, 95 % CI [−38.46 ms, −15.40 ms], t(79.83) = −4.58, $p < .001$] than they were to associate men with good (vs. bad) concepts [β = 0.02, M_{diff} = 5.28 ms, 95 % CI [−4.28 ms, 14.84 ms], t(395.07) = 1.08, p = .280]. However, as was the case with racial bias, gender bias was not significantly moderated by visual context. There was no evidence of a three-way interaction between target gender, block valence, and visual context: β = 0.02, b = 5.86, 95 % CI[−7.27, 18.99], t(128,112.52) = 0.88, p = .382. Thus, once again, there was no support for the propositions of the lens model. Visual contexts that were designed to normatively fit the concept of gender made gender bias no greater than visual contexts that were designed to normatively fit the concept of race (see Fig. 2, light bars).

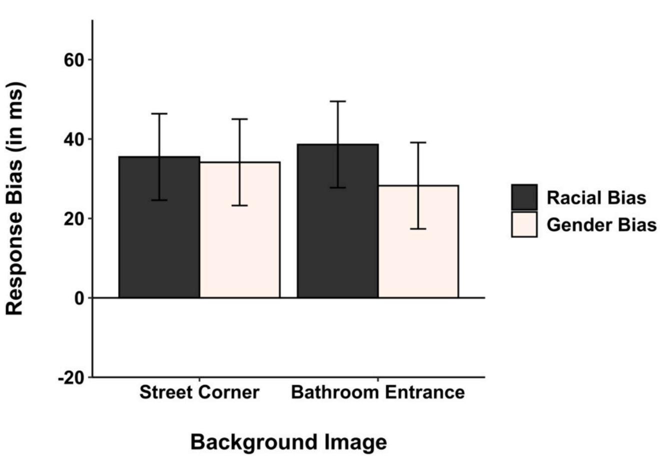

Fig. 2. *Implicit Racial and Gender Bias on SC-IATs (Experiments 1a and 1b).*
Note. Higher racial bias scores indicate the extent to which participants were relatively faster in racial stereotype-compatible trials (White targets categorized in positively-valenced blocks, Black targets categorized in negatively-valenced blocks) compared with racial stereotype-incompatible trials (the inverse pairings). Higher gender bias scores indicate the extent to which participants were relatively faster in gender stereotype-compatible trials (target women categorized in positively-valenced blocks, target men categorized in negatively-valenced blocks) compared with gender stereotype-incompatible trials (the inverse pairings). Effect size estimates are encompassed by 95 % confidence intervals.

4.3. Discussion

In Experiments 1a and 1b, normative fit cues did not significantly moderate implicit racial bias or implicit gender bias against intersectional targets. This was unexpected not just in light of the predictions the lens model, but likewise in light of prior studies demonstrating that racialized background images—including graffitied street corners (Wittenbrink et al., 2001)—can exacerbate racial bias on implicit measures (see also: Allen et al., 2010; Barden et al., 2004). However, while the results of Experiments 1a and 1b might be considered a failed conceptual replication of previous work in this area, it is important to note that their methods differed from earlier research in a number of potentially important ways. In light of this, we designed an additional experiment, which was proposed in the present registered report, to rigorously examine the extent to which visual contexts moderate the expression of implicit bias using methods more closely aligned with prior work. This higher-fidelity replication attempt was meant to be not only more definitive toward the end of establishing how reproducible earlier findings may be (e.g., those of Allen et al., 2010), but likewise toward the end of establishing whether normative fit, a theoretical factor that is implicated in lens use (Petsko et al., 2022), can indeed moderate implicit bias.

5. Experiment 2: A registered experiment (evaluative priming task)

Given the results of Experiments 1a and 1b, we ran an additional registered experiment to further examine the question of whether normative fit—operationalized in the form of background images against which targets are depicted—influences perceivers' implicit racial and gender biases against intersectional targets. However, unlike Experiments 1a and 1b, this registered experiment incorporated methods and images that were more closely aligned with past research on this topic. Incorporating these methods afforded a principled test of whether, as the lens model argues, normative fit cues can indeed “switch on” certain forms of bias and simultaneously “switch off” other forms of bias.[12]

Experiment 2 was meant to be closely modeled after prior experiments in a number of ways. First, instead of using SC-IATs, Experiment 2 measured implicit evaluations via evaluative priming tasks (EPTs; Fazio et al., 1995), which have been used in all prior studies demonstrating moderation of implicit bias by visual contexts (Allen et al., 2010; Barden et al., 2004; Wittenbrink et al., 2001). Second, instead of using images of novel street corners vs. restrooms (which we used in Experiments 1a and 1b), Experiment 2 used *actual* background images that have been used in prior studies (specifically, the jail vs. church background images originally used by Barden et al., 2004, and subsequently by Allen et al., 2010). Third, instead of using the full-body images of targets used in Experiments 1a and 1b, Experiment 2 used close-up headshots of Black and White women and men, matching how photographic stimuli were presented in prior studies (Allen et al., 2010; Barden et al., 2004; Wittenbrink et al., 2001).

The primary goal of the registered experiment was to attempt to replicate a finding from past research that would provide evidence for the lens model: specifically, the finding that anti-Black implicit racial bias is greater when targets are encountered in a jail compared to a church. This particular effect was demonstrated by both Allen et al. (2010) and Barden et al. (2004). By contrast, Wittenbrink et al. (2001) showed implicit racial bias to be absent when targets were presented in a

[12] Although it is true that the registered experiment enabled not just a test of whether the lens model holds up to scrutiny, but likewise a test of whether findings of past research could be conceptually replicated, the *primary* purpose of the registered experiment was to scrutinize the assumptions of the lens model.

church, but did not test biases using a jail context. We therefore considered this visual context contrast—jail vs. church—to be the most established effect in this research area, albeit one that was previously observed in relatively small samples.[13]

If we could replicate the moderating effect of visual context (specifically, jail vs. church) on implicit racial bias, a second goal of the registered study was to again test the lens model's prediction regarding lens trade-offs in intersectional stereotyping. If we could demonstrate that jails amplify implicit racial bias relative to churches, then in theory we ought also to have established the psychological conditions under which a lens trade-off would occur. That is, if it really is the case that an image of a jail cell "switches on" the use of a race lens, so too should it "switch off" the use of a gender lens, and in turn, participants' exhibition of a gender bias.

5.1. Method

Participants completed an online survey experiment containing two Evaluative Priming Tasks (EPTs; Fazio et al., 1995). In EPTs, racial bias is measured by the relative speed with which participants are faster to recognize "Good" (vs. "Bad") words when primed with the faces of White (vs. Black) targets. Likewise, gender bias is measured by the relative speed with which participants are faster to recognize "Good" (vs. "Bad") words when primed with the faces of women (vs. men).

Each EPT contained one of two stimulus sets of targets (in all cases a mix of Black and White women and men), and used one of two different background images: a background image that normatively fit the lens of race (i.e., a jail cell), or a background image that did not (i.e., a church; see Fig. 3). Thus, Experiment 2 was a within-person factorial experiment that contained manipulations of target race (Black, White), target gender (women, men), and critically, visual context (jail cell, church).

5.1.1. Participants

We aimed to recruit a final analytic sample of $N = 450$ participants from Prolific. As in Experiments 1a and 1b, we planned to exclude participants who self-acknowledged low data quality, as measured by a single item administered at the end of the experiment. In line with previous recommendations, we also planned to exclude any participant who responded incorrectly on 40 % or more of critical trials on the EPTs (Bar-Anan & Nosek, 2014). Based on our previous exclusion rates (8 of 401 participants in Experiment 1a and 14 of 425 participants in Experiment 1b), we intended to recruit a total sample of $N = 470$ participants, so as to have approximately $N = 450$ participants after exclusions.

5.1.2. Procedure

As in Experiments 1a and 1b, participants in Experiment 2 completed a study online via desktop computers using Inquisit Web software, and they were instructed that the purpose of the study was "to examine how people process and recall information about different concepts and social groups." After providing informed consent, participants reported their age, their gender, and their race/ethnicity.

Target Photographs. We used the same 40 photographs of Black and White women and men used in Experiments 1a and 1b, but presented targets via close-up headshots, similar to how they have been presented in prior work (see Fig. 3). We also supplemented these previously used stimuli with 40 novel headshots of Black and White women and men for use in training and testing blocks (see procedure below).

Visual Contexts. Experiment 2 used the same images of a jail and of a church, respectively, that were used originally by Allen et al. (2010) and Barden and colleagues (2004; see Fig. 3).

EPTs. Evaluative Priming Tasks (EPTs) were administered as part of a modified eight-block version of Fazio et al.' (1995) original procedure (see Table 1). In block 1 of this procedure, which was described as a practice block, the labels "Good" and "Bad" appeared at the top left and top right of screens, and participants were asked to categorize words (e.g., "amazing," "depressing") as either good or bad as quickly as possible using their E and I keys (the labels Good and Bad appeared at the top left and top right of screens throughout the block). In each of 10 trials, fixation crosses "+++" appeared in the center of screens for 500 ms, followed by a 100 ms interstimulus gap, followed by a word that remained on screen until categorized. A 3000 ms pause followed each trial. In block 2, described as a test block, participants completed the same task, but with a new set of positive and negative adjectives. Block 2 was 20 trials in length (see Table 1).

In block 3, described as a memorization task, no labels appeared on screen, and participants observed without responding as images of 12 novel Black and White female and male faces were displayed for 215 ms, each followed by a 3000 ms pause. Each face appeared twice, for a total of 24 trials. Participants were told to try to memorize the faces for a subsequent memory test. In block 4, described as a memory test, the labels "Familiar" and "Unfamiliar" appeared at the top left and top right of screens, and participants were asked to categorize faces as familiar or unfamiliar, viewing in succession the 12 faces presented during block 3 and 8 novel faces, for a total of 20 trials.

Block 5 was the first critical block, which contained the first complete EPT. In this block, participants were told that the previous tasks were being combined, and that they would now try to categorize words and memorize faces simultaneously. Additionally, participants were told that "in order to make the experiment more like real life, a picture of an everyday context will be.

added on the screen," and were asked to "please try to keep this context in mind as you perform the combined task, and imagine the sights, sounds, and smells that would be present." A background image then filled the screen for the remainder of the block (either an image of a jail or church). Trials proceeded as in blocks 1 and 2, with the exception that following the fixation crosses, a blank screen appeared for 100 ms,[14] followed by a face prime superimposed on each.

background for 215 ms (see Fig. 3). Twenty faces appeared (either from Group A or B) six times each, with each face preceding 3 positive and 3 negative words, for a total of 120 trials.

Block 6 was another memory test block, proceeding similarly to block 4, with 12 faces from block 5 and 8 novel faces categorized as familiar or unfamiliar. Block 7 was the second critical block, and thus contained the second complete EPT. Block 7 proceeded the same as block 5, but used the alternate visual context and the alternate target group that was not used in block 5. The order of both the visual contexts and whether participants saw Group A or B faces was randomized for each participant. Block 8 was a final memory test block, with 12 faces from block 7 and 8 novel faces categorized as familiar or unfamiliar.

Motivation to control prejudice. Following the EPTs, participants completed: (1) the Motivation to Control Prejudiced Responses scale (MCPR; Dunton & Fazio, 1997); and (2) the Internal and External Motivation to Respond without Prejudice Scales (IMS; EMS; Plant & Devine, 1998). We included these measures for exploratory purposes only. Prior research studies have suggested that motivations to control prejudice—as indexed by the MCPR scale—moderate the effects of visual contexts on implicit biases (Allen et al., 2010; Maddux et al., 2005). However, previous studies obtained conflicting results regarding the direction of this moderation pattern. Specifically, one study found that individuals high in MCPR showed relatively stronger effects of visual context in the expected direction (higher racial ingroup bias amidst

[13] Sample sizes (after exclusions) for previous studies demonstrating moderation of implicit bias by visual contexts in evaluative priming tasks were $n = 49$ (Allen et al., 2010), $n = 44$ (Wittenbrink et al., 2001, Study 2) and ns of 87, 47, and 60 (Barden et al., 2004, Studies 1, 2, and 3, respectively).

[14] This 100 ms of blank screen was included to mimic precisely the procedures of previous work (Allen et al., 2010; Barden et al., 2004).

Fig. 3. Visual Contexts and Target Presentations for (Registered) Experiment 2.

Table 1
Registered Evaluative Priming Task Procedure for Experiment 2.

Block No.	Block Description	Visual Context	Fixation (500 ms)	Prime (215 ms)	Target Types	"E" Key	"I" Key	Trial Count
1	Practice	None	+++	None	Good and Bad Words[a]	Good	Bad	10
2	Test	None	+++	None	Good and Bad Words[b]	Good	Bad	20
3	Memory Training	None	None	Face Stimuli	None	N/A	N/A	24
4	Memory Test	None	None	None	Face Stimuli[c]	Familiar	Unfamiliar	20
5[d]	Combined Test/Training	First Context[e]	+++[f]	Face Stimuli	Good and Bad Words[b]	Good	Bad	120
6	Memory Test	None	None	None	Face Stimuli[c]	Familiar	Unfamiliar	20
7[d]	Combined Test/Training	Second Context[e]	+++[f]	Face Stimuli	Good and Bad Words[b]	Good	Bad	120
8	Memory Test	None	None	None	Face Stimuli[c]	Familiar	Unfamiliar	20

Note. [a] Good and Bad words used in block 1 were *Good*: euphoric, amazing, creative; *Bad*: painful, depressing, wasted. [b] Good and Bad words used in blocks 2, 5, and 7 were *Good*: friend, smiling, adore, joyful, pleasure, friendship, happy, attractive, wonderful, kindness; *Bad*: bothersome, poison, pain, nasty, dirty, hatred, rotten, horrific, terrible, evil. [c] Each memory test presented 12 familiar faces used during training and 8 unfamiliar faces. [d] Blocks 5 and 7 were critical blocks measuring response times to Good and Bad words following face stimuli primes. [e] Context A and B refer to images of either the jail or the church (see Fig. 3), respectively. For each participant, one of these two contexts was randomly assigned to appear as Context A or B. [f] In blocks 5 and 7, 100 ms of (black) blank screen was inserted between the fixation cross and face stimuli to mimic prior studies.

threatening contexts; Allen et al., 2010), whereas the other study found the opposite: that individuals high in MCPR displayed outgroup-favoring implicit bias amidst threatening contexts (e.g., jails) compared with non-threatening contexts (e.g., church; Maddux et al., 2005). Consequently, we included the original MCPR scale, as well as two conceptually related scales (i.e., IMS and EMS), to examine whether we might also see evidence that motivations to control prejudice moderate the effects of visual contexts on implicit biases.

5.1.3. *Power analysis*

Estimates of power came from a bootstrapping technique in which we determined the smallest effects of interest that we could detect 90 % of the time given our analytic approach, proposed sample size of 450, and an alpha level of 0.05. This analysis revealed that we could expect to have 90 % power to detect two-way interactions between target race (or gender) and word valence—the key test for identifying implicit bias—as small as $\beta = 0.05$; we could expect to have 90 % power to detect three-way interactions between target race (or gender), word valence, and visual context—the key test for identifying moderation of implicit bias by background image—as small as $\beta = 0.09$; and we could expect to have 90 % power to detect simple effects of word valence within race (or gender)—the key test for determining if a subgroup (e.g., Black targets) elicited an evaluative preference—as small as $\beta = 0.05$. Details on precise power simulation methods can be found in the online supplement.

5.2. *Analysis plan*

Analyses focused on response times from critical blocks (5 and 7), which contained the EPTs against the church backdrop and the jail backdrop, respectively. To pre-process data, we removed incorrect responses (as was done by each of the previous studies in the area), and any response times below 300 ms or above 1000 ms [as was done by both Allen et al., 2010 and Maddux et al., 2005, and as recommended by Koppehele-Gossel et al., 2020].

5.2.1. *Key predictions*

This registered report made the following predictions: (1) Participants would, on average, exhibit evidence of a racial bias; (2) The magnitude of participants' racial bias would be moderated by background image, such that racial bias would be greater when the background image normatively fit the lens of race (vs. when the background image did not normatively fit the lens of race); (3) Participants would, on average, exhibit evidence of a gender bias; (4) The magnitude of participants' gender bias would be moderated by background image, such that gender bias would be weaker when the background image normatively fit the lens of race (vs. when the background image did not normatively fit the lens of race).

5.2.2. *Exploratory questions*

After examining whether participants' racial bias is indeed greater when the background image normatively fits the lens of race (vs. not), we planned to follow up by examining whether this tendency would be moderated by participants' motivations to control prejudiced responding. In particular, we planned to examine whether this bias would be starker among participants who score higher on MCPR (as well as on IMS and EMS, as some research suggests: Allen et al., 2010) or whether it would instead be weaker among participants who score higher on MCPR (as well as on IMS and EMS, as other research suggests: Maddux et al., 2005).

After examining whether participants' gender bias is indeed weaker when the background image normatively fits the lens of race (vs. not), we planned to follow up by examining whether gender bias would be merely attenuated under these conditions, or whether it would become "switched off" entirely (as the lens model suggests: Petsko et al., 2022).

5.2.3. Statistical model for testing key predictions and exploratory questions

Participants' response latencies (in milliseconds) were regressed, in a multilevel model, onto orthogonal contrast codes that represented the 2 (visual context: jail vs. church) × 2 (target race: Black vs. White) × 2 (target gender: female vs. male) × 2 (word valence: positive vs. negative) within-person factorial design of this experiment. Main effect contrasts were centered on zero and had a range of one (visual context: *jail* = ½; *church* = −½; target race: *Black* = ½; *White* = −½; target gender: *female* = ½; *male* = −½; word valence: *positive* = ½; *negative* = −½). Two-, three-, and four-way interaction contrasts were derived from the products of these main effect contrasts. To account for the nested data structure, this model included random intercepts for both targets and participants.[15] This model was run using the *lme4* package in R. Degrees of freedom (and in turn, standard error estimates, *t*-values, and *p*-values) were derived from the *lmerTest* package, which employs the Satterthwaite method of approximating degrees of freedom. Overall, the statistical model that was deployed in Experiment 2 was highly similar to the model run in Experiments 1a and 1b.

5.2.4. Inference criteria—main predictions

Prediction (1) was that participants would, on average, exhibit evidence of a racial bias. Evidence in favor of this prediction was expected to take the form of a significant (target race) × (word valence) interaction,[16] the nature of which was expected to be that participants would be faster to recognize "Good" (vs. "Bad") words when primed with the faces of White targets relative to when primed with the faces of Black targets.[17]

Prediction (2) was that the magnitude of participants' racial bias would be moderated by background image, such that racial bias would be greater when the background image normatively fit the lens of race (vs. when it did not). Evidence in favor of this prediction was expected to take the form of a significant (visual context) × (target race) × (word valence) interaction. The nature of this three-way interaction was expected to be such that racial bias—operationalized as a tendency to be significantly faster to recognize "Good" (vs. "Bad") words when primed with the faces of White (vs. Black) targets—would be significantly greater when the visual context was of a jail cell (vs. when the visual context was of a church).

Prediction (3) was that participants would, on average, exhibit evidence of a gender bias. Evidence in favor of this prediction was expected to take the form of a significant (target gender) × (word valence) interaction, the nature of which was expected to be that participants would be faster to recognize "Good" (vs. "Bad") words when primed with the faces of women relative to when primed with the faces of men.

Prediction (4) was that the magnitude of participants' gender bias would be moderated by background image, such that gender bias would be weaker when the background image normatively fit the lens of race (vs. when it did not). Evidence in favor of this prediction was expected to take the form of a significant (visual context) × (target gender) × (word valence) interaction. The nature of this three-way interaction was expected to be such that gender bias—operationalized as a tendency to be significantly faster to recognize "Good" (vs. "Bad") words when primed with the faces of women (vs. men)—would be significantly weaker when the visual context was of a jail cell (vs. when the visual context was of a church).

5.2.5. Inference criteria—exploratory questions

Our first exploratory question concerned whether the motivation to respond without prejudice (as indexed by IMS, EMS, and the MCPR scale, respectively) would play a role in moderating the magnitude by which background context influenced participants' expression of racial bias. To examine whether this was the case, the previously-described 2 (visual context) × 2 (target race) × 2 (target gender) × 2 (word valence) model was run three additional times—once each to examine whether participants' IMS, EMS, or MCPR scores, respectively, moderated the findings observed in the basic model described previously. In each of these three follow-up models, a given moderator (e.g., IMS scores) was z-standardized and entered as a predictor, and was likewise allowed to interact with each other predictor in the model. Evidence that motivation to respond without prejudice influences the magnitude by which visual context moderates racial bias was expected to take the form of a significant four-way interaction between any of the potential moderators, visual context, target race, and word valence. If any four-way interactions were to be significant, we planned to conduct follow-up spotlight analyses to examine whether it was the case that jail cells (vs. churches) exacerbate racial bias more strongly among those higher in motivation(s) to respond without prejudice (as some research suggests: Allen et al., 2010), or whether jail cells (vs. churches) instead exacerbate racial bias less strongly among those higher in motivation(s) to respond without prejudice (as other research suggests: Maddux et al., 2005).[18]

Our second exploratory question concerned whether gender bias would be merely attenuated among participants who saw a background image that normatively fits the lens of race (i.e., a jail cell), or whether it would be "switched off" entirely (as the lens model suggests: Petsko et al., 2022). If we were indeed to see evidence that jail cells (vs. churches) attenuate gender bias, we planned to follow up by examining to what degree gender bias becomes attenuated. Strong evidence in favor of the lens model—that is, evidence that the race lens "switches off" gender bias entirely—was expected to take the form of a *non-significant* (target gender) × (word valence) interaction in the jail cell condition. Weak evidence in favor of the lens model was instead expected to take the form of a significant—if attenuated, relative to its magnitude in the church condition—(target gender) × (word valence)

[15] In Phase 1 of our registered report, and also in our pre-registration documentation, we had stated as well that we wanted our model to also include random slopes of word valence within each target and within each participant. However, we noted in Phase 1 (and in our pre-registration documentation) that if the specified model did not converge, we would remove one or both of the random slopes, as needed, in order to achieve model convergence. Ultimately, the inclusion of one or both random effects of slope did not allow for model convergence, which caused us to simplify our random-effects structure, as intended.

[16] Statistical significance was determined on the basis of whether any given fixed effect is associated with a (Satterthwaite approximated) *p*-value below the 0.05 threshold.

[17] Spotlight analyses were, in all cases, conducted by manipulating the zero point on particular contrast codes. In this procedure, all contrasts remained orthogonal to each other; only the relevant zero-point(s) were manipulated (for more on this approach, see Judd et al., 2017). As an example, to estimate the speed with which participants recognized "Good" (vs. "Bad") words as a function of target race, the main statistical model was run twice: once wherein the target race contrast was set to *Black* = *0* and *White* = *−1*, and once wherein the target race contrast was set to *Black* = *1* and *White* = 0. The former variation estimated the simple effect of word valence among trials with Black targets; the latter variation estimated the simple effect of word valence among trials with White targets.

[18] In order to conduct spotlights among those who are "high" (i.e., +1 *SD*) or "low" (i.e., −1 *SD*) on an individual difference variable (e.g., MCPR), we added or subtracted 1 from participants' z-standardized variable scores. For example, to examine effects among those who were "high" in MCPR, we subtracted 1 from participants' z-standardized MCPR scores (which made those who were one standard deviation above the mean on MCPR the new zero point in the model among which effects were being estimated).

interaction in the jail cell condition.

5.3. Timeline

After Stage 1 acceptance, we planned to complete data collection, analysis, and a write up of our results by the end of October 2025.

5.4. Results

We received complete, non-duplicated survey responses from a total of N = 481 participants, all of whom were recruited from Prolific. Following our pre-registered exclusion criteria, we eliminated 15 participants for self-acknowledging low data quality, and we eliminated one participant for responding incorrectly on more than 40 % of critical trials on the EPTs. After exclusions, our final analytic sample comprised a total of $N = 465$ participants ($M_{\text{age}} = 43.05$, $SD_{\text{age}} = 13.17$; 230 men, 226 women, 9 non-binary gender; 305 White, 76 Black, 30 Latino, 21 East Asian, 12 South-East Asian, 5 South Asian, 3 Alaska Native, 10 other race, and 3 race unreported). Finally, and as pre-registered, we removed 2.62 % of trials for having incorrect responses, as well as 12.17 % of trials for being faster than 300 ms or slower than 1000 ms.

5.4.1. Key findings: did implicit racial bias depend on background context?

Prediction (1) was that participants would, on average, exhibit evidence of a racial bias. As noted, we investigated whether there was support for this prediction by subjecting participants' response latencies to the 2 (visual context: jail vs. church) × 2 (target race: Black vs. White) × 2 (target gender: female vs. male) × 2 (word valence: positive vs. negative) factorial analysis described previously. This analysis indeed yielded evidence of a racial bias, which took the form of a 2 (target race) × 2 (word valence) interaction: β = 0.03, b = 3.85, 95 % CI [1.08, 6.63], t(94,457.41) = 2.72, p = .006. The nature of this two-way interaction was that as expected, participants were significantly faster to categorize good (vs. bad) concepts when primed with White faces [β = −0.09, M_{diff} = −11.87 ms, 95 % CI [−13.84 ms, −9.91 ms], t(94,458.10) = −11.85, p < .001] relative to when primed with Black faces [β = −0.06, M_{diff} = −8.02 ms, 95 % CI [−9.98 ms, −6.06 ms], t(94,458.57) = −8.01, p < .001]. Thus, we found support for Prediction (1). Participants exhibited detectable levels of a racial bias.

Prediction (2) was that the magnitude of participants' racial bias would be moderated by background image, such that racial bias would be greater when the background image normatively fit the lens of race (vs. when it did not). Did we find evidence in support of this prediction? According to the same factorial analysis described above, we did not. That is, the magnitude of the racial bias described previously was not moderated by background image (three-way interaction between target race, word valence, and visual context: β = 0.01, b = 1.00, 95 % CI [−4.55, 6.55], t(94,455.85) = 0.35, p = .725). This is all to say that in contrast to the predictions of the lens model, and in contrast as well to previously published findings (e.g., Allen et al., 2010; Barden et al., 2004), implicit racial bias was no weaker against the backdrop of a church than it was against the backdrop of a jail cell (see Fig. 4, dark bars).

5.4.2. Key findings: did implicit gender bias depend on background context?

Prediction (3) was that participants would, on average, exhibit evidence of a gender bias. Supporting this prediction, the 2 (visual context) × 2 (target race) × 2 (target gender) × 2 (word valence) analysis described previously revealed evidence of a two-way interaction between target gender and word valence: β = −0.09, b = −12.09, 95 % CI [−14.87, −9.32], t(94,456.60) = −8.54, p < .001. The nature of this two-way interaction was that as expected, participants were significantly faster to categorize good (vs. bad) concepts when primed with women's faces [β = −0.12, M_{diff} = −15.99 ms, 95 % CI [−17.95 ms, −14.03 ms], t(94,457.74) = −15.97, p < .001] relative to when primed with men's faces [β = −0.03, M_{diff} = −3.90 ms, 95 % CI [−5.86 ms, −1.93 ms], t(94,458.12) = −3.89, p < .001]. Thus, participants indeed exhibited detectable levels of a gender bias, evaluating women more positively than men (see also: Connor et al., 2021).

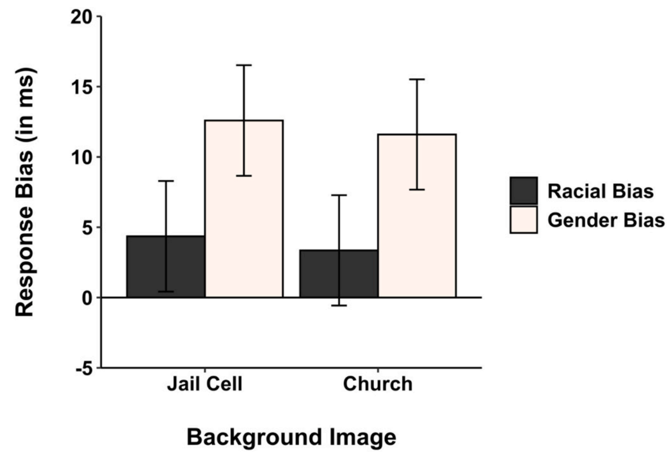

Fig. 4. Implicit Racial and Gender Bias on Evaluative Priming Tasks (Registered Experiment 2).
Note. Higher racial bias scores indicate the extent to which participants were relatively faster to categorize “Good” (vs. “Bad”) words when primed with White faces (vs. Black faces). Higher gender bias scores indicate the extent to which participants were relatively faster to categorize “Good” (vs. “Bad”) words when primed with women's faces (vs. men's faces). Effect size estimates are encompassed by 95 % confidence intervals.

Finally, prediction (4) was that the magnitude of participants' gender bias would be moderated by background image, such that gender bias would be weaker when the background image normatively fit the lens of race (vs. when it did not). However, in contrast to this prediction—and in line with what we have been documenting throughout this manuscript—there was no credible evidence that background image moderated the magnitude of gender bias. Indeed, the same factorial analysis referenced above yielded no evidence that the two-way interaction between target gender and word valence was further moderated by background image (three-way interaction: β = −0.01, b = −0.99, 95 % CI [−6.54, 4.56], t(94,455.81) = −0.35, p = .726). This is all to say that in contrast to the prediction that normative fit—operationalized in the form of background images against which targets are depicted—might influence perceivers' implicit gender biases, we instead found that gender biases were statistically similar regardless of whether participants encountered intersectional targets against the backdrop of a jail cell vs. a church (see Fig. 4, light bars).

5.4.3. Exploratory findings

As noted previously, we were interested in whether the motivation to respond without prejudice (as indexed by IMS, EMS, and the MCPR scale, respectively) would play a role in moderating the magnitude by which background image influenced participants' expression of racial bias.[19] To examine whether this was the case, we followed the analytic plan described previously. These exploratory analyses yielded no evidence of a four-way interaction between visual context, target race, word valence, and any of our potential moderating variables. This is all to say that there was no evidence that contextual moderation of racial bias depended on participants' IMS scores (four-way interaction: β = 0.03, b = 4.34, 95 % CI [−1.23, 9.91], t(93,765.33) = 1.53, p = .127), no

[19] In our pre-registration documentation, we had also mentioned taking an exploratory interest in whether gender bias might be “switched off” entirely against the backdrop of a jail cell (vs. the backdrop of a church). Given that participants' gender bias ultimately ended up being unaffected by background image, we had no reason to conduct exploratory analyses that looked into this question.

evidence that contextual moderation of racial bias depended on participants' EMS scores (four-way interaction: β = −0.02, $b = -2.86$, 95 % CI [−8.44, 2.72], $t(93{,}356.91) = -1.01$, $p = .315$), and there was no evidence that contextual moderation of racial bias depended on targets' MCPR scores (four-way interaction: β = 0.02, $b = 3.07$, 95 % CI [−2.48, 8.62], $t(94{,}326.67) = 1.09$, $p = .278$). Thus, in contrast to earlier findings (Allen et al., 2010; Maddux et al., 2005), the magnitude by which participants' racial bias depended on background image was not contingent on their underlying motivations to respond without prejudice—at least not to a degree that we could detect, given our sample size and analytic approach.

5.5. Discussion

The purpose of Experiment 2 was to further examine whether the theoretical factor of normative fit—operationalized in the form of background images against which targets are depicted—influences perceivers' implicit racial and gender biases against intersectional targets, respectively. Unlike Experiments 1a and 1b, however, Experiment 2 was (a) fully registered in advance of data collection, and was (b) designed to closely approximate methods that have been used to document contextual moderation of implicit bias in prior research (Barden et al., 2004). Consistent with Experiments 1a and 1b, we again found no evidence in favor of the hypothesis that the background images against which targets are depicted moderates the magnitude with which perceivers' exhibit racial bias or gender bias against those targets. As we will suggest below, the lack of moderation of implicit bias by background image raises questions about whether normative fit cues in and of themselves play a role moderating lens use (as the lens model argues: Petsko et al., 2022), and it likewise has implications for the reproducibility of past research. It is to these issues that we turn our attention in the general discussion.

6. General discussion

According to the lens model, how perceivers stereotype intersectional targets is likely to depend on the social environment in which targets are embedded (see also: Bodenhausen & Macrae, 1998; Sinclair & Kunda, 1999 and Turner et al., 1987). If a social environment makes the lens of race salient, for example, perceivers are expected to exhibit racial patterns of stereotyping, but not so much gendered patterns of stereotyping. If a social environment makes the lens of gender salient, perceivers are expected to do the reverse: that is, to exhibit gendered patterns of stereotyping, but not so much racial patterns of stereotyping. And finally, if an environment makes an intersectional lens salient—for example, a lens that sharpens perceivers' focus on both race and gender simultaneously—perceivers are expected to exhibit specifically intersectional patterns of stereotyping. However, the theoretical factors that are implicated in lens salience, and that in turn ought to shape how intersectional stereotyping unfolds in a particular context, are each in need of greater scientific scrutiny (for a review, see: Petsko & Bodenhausen, 2020). It was against this backdrop that the present experiments were conducted.

The central goal of the present manuscript was to examine whether one theoretical factor that is implicated in lens salience—the factor of normative fit—can itself influence whether perceivers use the lens of gender or race, respectively, when stereotyping intersectional targets. Broadly speaking, the principle of normative fit suggests that perceivers ought to be more likely to use a given lens for construing others when the social category implied by that lens is stereotypically aligned with a given social environment (Oakes, 1987; Oakes et al., 1991; Turner et al., 1987). This is to say that if a social environment feels stereotypically aligned with the concept of race, it ought to facilitate perceivers' likelihood of stereotyping others through the lens of race, and to decrease perceivers' likelihood of stereotyping others through the lens of alternative social categories (like gender). However, across three experiments wherein we manipulated the visual backdrops against which targets were depicted ($N = 1269$), we found no support for this hypothesis. That is, we found no evidence that depicting intersectional targets against backdrops that are stereotypically aligned with the concept of race—backdrops of graffitied street corners or prison cells, respectively—increased the magnitude by which perceivers exhibited racial stereotyping against these targets. In fact, and in contrast to prior empirical demonstrations (e.g., Allen et al., 2010; Wittenbrink et al., 2001), we found no evidence that visual backdrops exhibited any influence on participants' implicit prejudice at all.

6.1. Theoretical and practical implications

A key theoretical implication of the present work is that perhaps an assumption of the lens model—that normative fit cues can, in and of themselves, determine which lenses perceivers use in particular social environments—needs to be revised. Indeed, the present experiments, which contained high-fidelity manipulations of normative fit (but see also our discussion on limitations, below), suggest that perhaps normative fit cues are not powerful enough *on their own* to influence lens salience. Perhaps instead normative fit cues need to be combined with other psychological factors if they are to augment which lenses become situationally salient to perceivers. Something notable about this theoretical implication is that it bears consistency with other perspectives on identity salience. For example, according to self-categorization theory, racial categorization (and in turn, racial stereotyping) is not expected to result from fit cues by themselves, but is instead expected to result from fit cues *in tandem* with category accessibility (Oakes et al., 1991; Turner et al., 1987). This is to say that according to self-categorization theory, simply seeing a person against the backdrop of a prison cell is not expected to be psychologically powerful enough to cause race to become salient to perceivers. Instead, this perspective suggests that perceivers would need to see a prison cell (or some other stimulus that normatively fits the concept of race) at a moment in time when race has already been made readily retrievable from memory (i.e., accessible; Bruner, 1957; Higgins, 1996). Perhaps the lens model, which has previously suggested that its salience factors operate independently from one another, would benefit from considering the question of whether certain factors need to co-occur with other factors in order to meaningfully impact lens salience.

A key practical implication of the present work concerns the reproducibility of prior research on this topic. When we began conducting the research studies reported in this manuscript, we expected to find evidence for a phenomenon that has been repeatedly reported in the past: that changing the configural backdrop against which targets are depicted can change the magnitude by which participants exhibit implicit prejudice against those targets (Allen et al., 2010; Barden et al., 2004; Wittenbrink et al., 2001). Nevertheless, and despite our attempts to conduct well-powered replications of this phenomenon, we found no credible evidence that these earlier findings were able to be reproduced in contemporary samples of U.S. adults. In light of all this, one practical implication of the present work is that the scientific consensus on whether background images moderate degrees of implicit prejudice may need to be updated. It may well have been the case that seeing Black (vs. White) individuals against the backdrop of graffitied street corners facilitated implicit anti-Black bias 25 years ago (Wittenbrink et al., 2001). But among participants of today (whose intergroup attitudes, to be sure, differ considerably from those of participants from two decades ago: Charlesworth & Banaji, 2022), degrees of anti-Black bias appear to be relatively immune to the contextual backdrops against which targets are depicted—at least in our data.

6.2. Limitations of the present research

One limitation of the present research is that the methods of—and perhaps even the realism of—the experiments reported here were not isomorphic with those of the experiments we sought to replicate (Allen et al., 2010; Barden et al., 2004; Wittenbrink et al., 2001). As such, it is not possible to know whether our failure to detect moderation of implicit bias by background image was attributable to the lack of this effect's existence, or instead to differences between our experiments and those reported in prior research. For example, all three of the prior empirical demonstrations that we have been citing (Allen et al., 2010; Barden et al., 2004; Wittenbrink et al., 2001) relied on undergraduate samples, and all three collected their data on human subjects in the lab (i.e., in-person, under the supervision of research associates). Our experiments, in contrast, included broader samples of U.S. adults, and the majority of our data were collected via online survey platforms. In addition, all three of these prior demonstrations used exclusively male stimulus targets, whereas we used male as well as female stimulus targets. Finally, all three of these prior demonstrations were conducted at a time when implicit measurement tasks like the ones used here might have *felt* more cutting-edge or engaging to participants (i.e., at a time when these tasks might have elicited greater experimental realism; Wilson et al., 2010). In light of all this, it is possible that a more direct replication of past work—or alternatively, a conceptual replication of past work that is higher in experimental realism—would indeed yield evidence for past findings' reproducibility. Future research should investigate this possibility.

A second limitation of the present research is that whereas Experiments 1a and 1b used social contexts that were designed to normatively fit either the lenses of race or gender (i.e., a graffitied street corner vs. a gendered bathroom entryway, respectively), Experiment 2 did not. Indeed, Experiment 2 used social contexts that either were or were not expected to normatively fit the lens of race alone (i.e., a prison cell vs. a church, respectively). This particular design choice of Experiment 2 complicates the conclusions that can be drawn from Experiment 2, both because a) this design choice may be expected to decrease any perceptual tension there might have been between race *and* gender in the minds of perceivers, and because b) churches, like prisons, can be said to be social contexts that normatively fit the concept of race.[20] Given these design choices of Experiment 2, it is plausible that implicit prejudices were unmoderated by contextual images in Experiment 2, at least in part, because the contrast between race and gender was relatively weak in this experiment, and likewise because both contexts were arguably associated with race. Future research might benefit from investigating whether race is indeed more readily associated with the concept of crime in the minds of perceivers than it is with the concept of religion, as we have been arguing (see also: Devine & Elliot, 1995; Eberhardt et al., 2004). Such testing may be necessary toward the end of interpreting Experiment 2 as being a high-fidelity test of whether normative fit cues moderate lens salience. For example, if present-day Americans do indeed associate race more strongly with the concept of crime than with the concept of religion, it is reasonable to interpret Experiment 2 as being a high-fidelity test of the lens model's assumptions. In contrast, if present-day Americans do *not* associate of race with the concept of crime more than with the concept of religion, it may be best to interpret Experiment 2 as a well-powered replication attempt of an established phenomenon (e.g., Barden et al., 2004), but not necessarily as a high-fidelity test of the lens model's assumptions.

A third limitation of the present research is that its key theoretical implication—that perhaps lens salience results not from normative fit cues by themselves, but from normative fit cues in tandem with lens accessibility—is, at this point, speculative. Unfortunately, the present experiments did not measure any kind of variable that might serve as a proxy for lens accessibility, and they likewise did not include manipulations of accessibility itself. As such, although the hypothesis that salience results from a two-way interaction between normative fit and accessibility is sensible—and is aligned with other perspectives of identity salience (Oakes et al., 1991; Turner et al., 1987)—we unfortunately have no way of directly testing this hypothesis here. Future research should seek to further clarify whether this theoretical implication holds merit. For example, a series of experiments that orthogonally manipulate whether race is a) accessible to perceivers, and b) normatively fitting of the social environment, would provide insight into whether normative fit cues indeed need to be combined with lens accessibility, as we have been suggesting, or whether instead normative fit cues have no influence on lens salience, even when combined with other factors.

A fourth limitation of the present research is that it does little to inform present-day social scientists about where contextual malleability in implicit attitudes—when it does occur—might come from. As mentioned previously, a variety of theoretical accounts have been offered for why implicit attitudes might be contextually malleable. Gestalt-based processing accounts suggest that perhaps different contexts activate different aspects of racial representations in working memory (Wittenbrink et al., 2001). Role-based processing accounts suggest that perhaps different contexts bring to perceivers' minds different social roles, which, when combined with racial category information, can change the meaning perceivers make of race itself (Barden et al., 2004). And then of course lens-based processing accounts suggest that social contexts modulate not so much how perceivers think about race, but whether perceivers think about race at all (Petsko et al., 2022). Assuming that future research is able to answer the question of *when* implicit prejudice might vary as a function of social context, a relevant domain of inquiry will be to answer the question of *why* it may do so. If, hypothetically, gestalt- or role-based processing accounts were to outperform lens-based processing accounts as explanations for contextual variation in implicit prejudice, it may be unwise to recommend lens-based interventions as tools for mitigating intergroup discrimination. For example, if it is not the case that focusing perceivers' attention on the lens of profession attenuates perceivers' likelihood of exhibiting a racial bias, it may be unwise for researchers to recommend such an intervention as a tool for reducing the influence of racial bias on workplace decision making.

7. Conclusion

A central goal of the present manuscript was to examine whether one theoretical factor that is implicated in lens salience—the factor of normative fit—can itself influence whether perceivers use the lens of gender or race, respectively, when stereotyping intersectional targets (Oakes et al., 1991; Petsko et al., 2022). Toward this end, we conducted three experiments (two of which were preliminary, one of which was registered; total $N = 1269$) in which participants encountered intersectional targets—in all cases, a mix of Black and White men and women—either against background images that normatively fit the concept of race (e.g., a graffitied street corner), or against background images that did not (e.g., a gendered bathroom entryway). Contrary to the predictions of the lens model and prior research, we found no evidence that normative fit cues moderated the degree to which participants exhibited implicit racial or gender prejudice against targets. In light of these findings, we suggest that perhaps normative fit cues in and of themselves are not psychologically powerful enough to augment lens salience. In addition, we suggest that perhaps previously reported empirical findings—findings suggesting that implicit attitudes vary as a function of the background images against which targets are depicted

[20] In support of this latter point, sociological accounts have repeatedly emphasized the importance of Black churches as hubs for community organizing and information exchange among Black Americans living in the U.S (Allen, 2019; Patillo-McCoy, 1998). In addition, the concept of religion may not only be associated with Blackness, but may as well be increasingly associated with Whiteness in contemporary U.S. society. Indeed, over the last 50 years, White evangelicals as a group have become a more visible force in U.S. politics (e.g., Du Mez, 2020), which may have led to an increase in Americans' proclivity to associate the concept of religion with Whiteness.

(Allen et al., 2010; Barden et al., 2004; Wittenbrink et al., 2001)—no longer hold up to scrutiny in contemporary samples of U.S. adults. Of course, verifying each of these implications requires further empirical testing. Our hope is that others join us in conducting such testing, as doing so could help to answer important theoretical questions about the psychological factors that give rise to lens salience, as well as important practical questions about the reproducibility of research on implicit bias.

Open practices

This manuscript includes Open Data and Open Materials. Preregistration documentation, data files, codebooks, survey materials, and R scripts associated with this article are available on the Open Science Framework (OSF) website: https://osf.io/43fbp/.

CRediT authorship contribution statement

Christopher D. Petsko: Writing – review & editing, Writing – original draft, Methodology, Funding acquisition, Formal analysis, Data curation, Conceptualization. **Paul Connor:** Writing – review & editing, Software, Methodology, Funding acquisition, Formal analysis, Data curation, Conceptualization.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

References

Aidman, E. V., & Carroll, S. M. (2003). Implicit individual differences: Relationships between implicit self-esteem, gender identity, and gender attitudes. *European Journal of Personality, 17*(1), 19–36. https://doi.org/10.1002/per.465

Allen, S. E. (2019). Doing Black Christianity: Reframing Black church scholarship. *Sociology Compass, 13*, Article e12731. https://doi.org/10.1111/soc4.12731

Allen, T. J., Sherman, J. W., & Klauer, K. C. (2010). Social context and the self-regulation of implicit bias. *Group Processes & Intergroup Relations, 13*(2), 137–149. https://doi.org/10.1177/1368430209353635

Bar-Anan, Y., & Nosek, B. A. (2014). A comparative investigation of seven indirect attitude measures. *Behavior Research Methods, 46*, 668–688. https://doi.org/10.2139/ssrn.2074556

Barden, J., Maddux, W. W., Petty, R. E., & Brewer, M. B. (2004). Contextual moderation of racial bias: The impact of social roles on controlled and automatically activated attitudes. *Journal of Personality and Social Psychology, 87*(1), 5–22. https://doi.org/10.1037/0022-3514.87.1.5

Bates, D., Maechler, M., Bolker, B., & Walker, S. (2015). Fitting linear mixed-effects models using lme4. *Journal of Statistical Software, 67*(1), 1–48. https://doi.org/10.18637/jss.v067.i01

Bodenhausen, G. V., & Macrae, C. N. (1998). Stereotype activation and inhibition: Advances in social cognition, volume XI. In R. S. Wyer, Jr. (Ed.), *Stereotype activation and inhibition: Advances in social cognition, volume XI* (pp. 1–52).

Bruner, J. S. (1957). On perceptual readiness. *Psychological Review, 64*, 123–152. https://doi.org/10.1037/h0043805

Charlesworth, T. E. S., & Banaji, M. R. (2022). Patterns of implicit and explicit attitudes: IV. Change and stability from 2007 to 2020. *Psychological Science, 33*, 1347–1371. https://doi.org/10.1177/09567976221084257

Cole, E. R. (2009). Intersectionality and research in psychology. *American Psychologist, 64*, 170–180. https://doi.org/10.1037/a0014564

Connor, P., Varney, J., Keltner, D., & Chen, S. (2021). Social class competence stereotypes are amplified by socially signaled economic inequality. *Personality and Social Psychology Bulletin, 47*(1), 89–105. https://doi.org/10.1177/0146167220916640

Connor, P., Weeks, M., Glaser, J., Chen, S., & Keltner, D. (2023). Intersectional implicit bias: Evidence for asymmetrically compounding bias and the predominance of target gender. *Journal of Personality and Social Psychology, 124*(1), 22–48. https://doi.org/10.1037/pspa0000314

Crenshaw, K. (1991). Mapping the margins: Intersectionality, identity politics, and violence against women of color. *Stanford Law Review, 43*, 1241–1299. https://doi.org/10.2307/1229039

Devine, P. G., & Elliot, A. J. (1995). Are racial stereotypes really fading? The Princeton trilogy revisited. *Personality and Social Psychology, 21*, 1139–1150. https://doi.org/10.1177/01461672952111002

Du Mez, K. K. (2020). *Jesus and John Wayne: How white evangelicals corrupted a faith and fractured a nation*. New York, NY: W. W. Norton.

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. https://doi.org/10.1177/0146167297233009

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*, 876–893. https://doi.org/10.1037/0022-3514.87.6.876

Else-Quest, N. M., & Hyde, J. S. (2016a). Intersectionality in quantitative psychological research: I. Theoretical and epistemological issues. *Psychology of Women Quarterly, 40*(2), 155–170. https://doi.org/10.1177/0361684316629797

Else-Quest, N. M., & Hyde, J. S. (2016b). Intersectionality in quantitative psychological research: II. Methods and techniques. *Psychology of Women Quarterly, 40*(3), 319–336. https://doi.org/10.1177/0361684316647953

Fazio, R. H., Jackson, J. R., Dunton, B. C., & Williams, C. J. (1995). Variability in automatic activation as an unobtrusive measure of racial attitudes: A bona fide pipeline? *Journal of Personality and Social Psychology, 69*(6), 1013–1027. https://doi.org/10.1037/0022-3514.69.6.1013

Ghavami, N., & Peplau, L. A. (2013). An intersectional analysis of gender and ethnic stereotypes: Testing three hypotheses. *Psychology of Women Quarterly, 37*, 113–127. https://doi.org/10.1177/0361684312464203

Goff, P. A., & Kahn, K. B. (2013). How psychological science impedes intersectional thinking. *Du Bois Review: Social Science Research on Race, 10*, 365–384. https://doi.org/10.1017/S1742058X13000313

Hall, E. V., Hall, A. V., Galinsky, A. D., & Phillips, K. W. (2019). MOSAIC: A model of stereotyping through associated and intersectional categories. *Academy of Management Review, 44*, 643–672. https://doi.org/10.5465/amr.2017.0109

Hamilton, D. L., & Gifford, R. K. (1976). Illusory correlation in interpersonal perception: A cognitive basis of stereotypic judgments. *Journal of Experimental Social Psychology, 12*, 392–407. https://doi.org/10.1016/S0022-1031(76)80006-6

Higgins, E. T. (1996). Knowledge activation: Accessibility, applicability, and salience. In *Social psychology: Handbook of basic principles* (pp. 133–168). New York, NY: Guilford Press.

Hudson, S. T. J., Myer, A., & Berney, E. C. (2024). Stereotyping, prejudice, and discrimination at the intersection of race and gender: An intersectional theory primer. *Social and Personality Psychology Compass*. https://doi.org/10.1111/spc3.12939. e12939.

Judd, C. M., McClelland, G. H., & Ryan, C. S. (2017). *Data analysis: A model comparison approach* (3rd ed.). Routledge. https://doi.org/10.4324/9781315744131

Kang, S. K., & Bodenhausen, G. V. (2015). Multiple identities in social perception and interaction: Challenges and opportunities. *Annual Review of Psychology, 66*, 547–574. https://doi.org/10.1146/annurev-psych-010814-015025

Karpinski, A., & Steinman, R. B. (2006). The single category implicit association test as a measure of implicit social cognition. *Journal of Personality and Social Psychology, 91*(1), 16–32. https://doi.org/10.1037/0022-3514.91.1.16

King, D. K. (1988). Multiple jeopardy, multiple consciousness: The context of a Black feminist ideology. *Signs: Journal of Women in Culture and Society, 14*(1), 42–72. https://doi.org/10.1086/494491

Koppehele-Gossel, J., Hoffmann, L., Banse, R., & Gawronski, B. (2020). Evaluative priming as an implicit measure of evaluation: An examination of outlier-treatments for evaluative priming scores. *Journal of Experimental Social Psychology, 87*, Article 103905. https://doi.org/10.1016/j.jesp.2019.103905

Kuznetsova, A., Brockhoff, P. B., & Christensen, R. H. B. (2017). lmerTest package: Tests in linear mixed effects models. *Journal of Statistical Software, 82*(13), 1–26. https://doi.org/10.18637/jss.v082.i13

Lei, R. F., Foster-Hanson, E., & Goh, J. X. (2023). A sociohistorical model of intersectional social category prototypes. *Nature Reviews Psychology, 2*(5), 297–308. https://doi.org/10.1038/s44159-023-00165-0

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. https://doi.org/10.1016/j.jesp.2004.05.002

McGuire, W. J., McGuire, C. V., Child, P., & Fujioka, T. (1978). Salience of ethnicity in the spontaneous self-concept as a function of one's ethnic distinctiveness in the social environment. *Journal of Personality and Social Psychology, 36*, 511–520. https://doi.org/10.1037/0022-3514.36.5.511

Neel, R., & Lassetter, B. (2019). The stigma of perceived irrelevance: An affordance-management theory of interpersonal invisibility. *Psychological Review, 126*, 634–659. https://doi.org/10.1037/rev0000143

Nicolas, G., de la Fuente, M., & Fiske, S. T. (2017). Mind the overlap in multiple categorization: A review of crossed categorization, intersectionality, and multiracial perception. *Group Pro- cesses & Intergroup Relations, 20*, 621–631. https://doi.org/10.1177/1368430217708862

Oakes, P. J. (1987). The salience of social categories. In J. C. Turner, M. A. Hogg, P. J. Oakes, S. D. Reicher, & M. S. Wetherell (Eds.), *Rediscovering the social group* (pp. 117–141). Oxford, UK: Blackwell.

Oakes, P. J., Turner, J. C., & Haslam, S. A. (1991). Perceiving people as group members: The role of fit in the salience of social categorizations. *British Journal of Social Psychology, 30*, 125–144. https://doi.org/10.1111/j.2044-8309.1991.tb00930.x

Patillo-McCoy, M. (1998). Church culture as a strategy of action in the Black community. *American Sociological Review, 63*, 767–784. https://doi.org/10.2307/2657500

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. https://doi.org/10.1037/0022-3514.81.2.181

Pedulla, D. S. (2014). The positive consequences of negative stereotypes: Race, sexual orientation, and the job application process. *Social Psychology Quarterly, 77*, 75–94. https://doi.org/10.1177/0190272513506229

Pellegrini, V., De Cristofaro, V. D., Giacomantonio, M., & Salvati, M. (2020). Why are gay leaders perceived as ineffective? The role of the type of organization, sexual prejudice and gender stereotypes. *Personality and Individual Differences, 157*, 1–11. https://doi.org/10.1016/j.paid.2020.109817

Petsko, C. D., & Bodenhausen, G. V. (2019a). Race-crime congruency effects revisited: Do we take defendants' sexual orientation into account? *Social Psychological and Personality Science, 10*, 73–81. https://doi.org/10.1177/1948550617736111

Petsko, C. D., & Bodenhausen, G. V. (2019b). Racial stereotyping of gay men: Can a minority sexual orientation erase race? *Journal of Experimental Social Psychology, 83*, 37–54. https://doi.org/10.1016/j.jesp.2019.03.002

Petsko, C. D., & Bodenhausen, G. V. (2020). Multifarious person perception: How social perceivers manage the complexity of intersectional targets. *Social and Personality Psychology Compass, 14*, Article e12518. https://doi.org/10.1111/spc3.12518

Petsko, C. D., & Rosette, A. S. (2023). Are leaders still presumed White by default? Racial bias in leader categorization revisited. *Journal of Applied Psychology, 108*, 330–340. https://doi.org/10.1037/apl0001020

Petsko, C. D., Rosette, A. S., & Bodenhausen, G. V. (2022). Through the looking glass: A lens-based account of intersectional stereotyping. *Journal of Personality and Social Psychology, 123*, 763–787. https://doi.org/10.1037/pspi0000382

Plant, E. A., & Devine, P. G. (1998). Internal and external motivation to respond without prejudice. *Journal of Personality and Social Psychology, 75*, 811–832. https://doi.org/10.1037/0022-3514.75.3.811

Plaut, V. C. (2010). Diversity science: Why and how difference makes a difference. *Psychological Inquiry, 21*, 77–99. https://doi.org/10.1080/10478401003676501

Preddie, J. P., & Biernat, M. (2021). More than the sum of its parts: Intersections of sexual orientation and race as they influence perceptions of group similarity and stereotype content. *Sex Roles, 84*(9–10), 554–573. https://doi.org/10.1007/s11199-020-01185-3

Purdie-Vaughns, V., & Eibach, R. P. (2008). Intersectional invisibility: The distinctive advantages and disadvantages of multiple subordinate-group identities. *Sex Roles: A Journal of Research, 59*, 377–391. https://doi.org/10.1007/s11199-008-9424-4

Remedios, J. D., & Vinluan, A. C. (2024). Intersectionality in social cognition: Stereotyping across interlocking social categories. In D. E. Carlston, K. Johnson, & K. Hugenberg (Eds.), *Oxford handbook of social cognition* (2nd ed.). https://doi.org/10.1093/oxfordhb/9780197763414.013.18

Rosette, A. S., Li, X., Samuel, N., & Petsko, C. D. (2025). From primary to pluralistic: A typology of intersectionality. *Current Opinion in Psychology, 61*, Article 101944. https://doi.org/10.1016/j.copsyc.2024.101944

Sinclair, J., & Kunda, Z. (1999). Reactions to a Black professional: Motivated inhibition and activation of conflicting stereotypes. *Journal of Personality and Social Psychology, 77*, 885–904. https://doi.org/10.1037/0022-3514.77.5.885

Turner, J. C., Hogg, M. A., Oakes, P. J., Reicher, S. D., & Wetherell, M. S. (1987). *Rediscovering the social group: A self-categorization theory*. Cambridge, MA, US: Basil Blackwell. https://doi.org/10.1111/j.2044-8309.1987.tb00799.x

Wilson, J. P., Remedios, J. D., & Rule, N. O. (2017). Interactive effects of obvious and ambiguous social categories on perceptions of leadership: When double-minority status may be beneficial. *Personality and Social Psychology Bulletin, 43*, 888–900. https://doi.org/10.1177/0146167217702373

Wilson, T. D., Aronson, E., & Carlsmith, K. (2010). The art of laboratory experimentation. In S. T. Fiske, D. T. Gilbert, & G. Lindzey (Eds.), *Handbook of social psychology* (5th ed., pp. 51–81). John Wiley & Sons, Inc.. https://doi.org/10.1002/9780470561119.socpsy001002

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), 815–827. https://doi.org/10.1037//0022-3514.81.5.815