Received: 8 July 2021 | Accepted: 20 January 2022

DOI: 10.1002/ejsp.2840

EASP WILEY

RESEARCH ARTICLE

On the use of gender categories and emotion categories in threat-based person impressions

Heather R. Rees[1] | Jeffrey W. Sherman[1] | Karl Christoph Klauer[2] | Andrew R. Todd[1]

[1] Department of Psychology, University of California, Davis, California, USA

[2] Department of Social Psychology and Methodology, Albert-Ludwigs-University, Freiburg, Germany

Correspondence
Heather R. Rees or Andrew R. Todd, Department of Psychology, University of California, Davis, 1 Shields Avenue, Davis, CA 95616, USA.
Email: hrrees.1@gmail.com or atotdd@ucdavis.edu

Abstract

People often form impressions of others in contexts where both relatively static demographic cues (e.g., age, race, gender) and situationally flexible, dynamic cues (e.g., emotion expressions) are available. We examined whether and how attending to gender (male, female) versus emotion expression (neutral, smiling) affects threat-based person impressions. In three experiments, we heightened the salience of either gender categories or emotion categories in a sequential priming task. Category salience consistently moderated the use of both categories in threat impressions. The effect of gender (emotion) categories was stronger when attention was directed to gender (emotion) than when attention was directed to emotion (gender); however, category salience did not eliminate the use of either category. Implications for models of person construal and bias in person impressions are discussed.

KEYWORDS
gender bias, sequential priming task, social categorization, stereotyping, threat

1 | INTRODUCTION

Forming impressions of others is fundamental to social life. One important impression entails determining if another person poses a physical threat. Perceivers quickly process cues from multiple sources to inform such impressions. Included among these sources are situationally flexible, dynamic cues pertaining to emotion categories (e.g., anger, happiness) and comparatively more static cues pertaining to demographic categories (e.g., age, race, gender) that are often detectable in a person's face (Freeman et al., 2020; Haxby et al., 2000). Cues signalling approachability (e.g., a smile) typically evoke weaker threat impressions than do cues signalling hostility (e.g., a scowl) or even neutrality. Demographic categories with links to threat-based stereotypes also shape such impressions. For example, men are viewed as more aggressive than women (Archer, 2004), which can help explain why female targets typically evoke weaker threat-related responses than do male targets (Neel et al., 2012). In some cases, emotion categories and demographic categories have an interactive influence on person impressions (Freeman et al., 2020). For example, threat-based forms of racial bias are weaker when Black and White targets are smiling (Kubota & Ito, 2014; Richeson & Trawalter, 2008), suggesting that smiles may dampen threat signals stereotypically associated with race cues.

Although emotion categories and demographic categories are both processed quickly and early during person construal (Eimer et al., 2003; Ito & Urland, 2003), they may not affect threat impressions equally. Instead, cues pertaining to one category may ultimately guide impressions more than cues pertaining to other categories do (Petsko & Bodenhausen, 2020). Contextual factors, such as how salient (i.e., attention-eliciting) or goal-relevant different cues are, can determine which of several applicable categories has a greater relative impact on threat impressions. For example, Todd et al. (2021) found that weapon identification bias (i.e., mistaking a tool for a gun) in the presence of Black versus White men and boys was weaker when participants attended to the targets' age than when they attended to the targets' race (see also Gawronski et al., 2010; Macrae et al., 1995).

Ample evidence indicates that threat appraisal requires integrating multiple social cues, and that category salience can determine which of several concurrently activated cues have the biggest impact on such impressions (Freeman et al., 2020). Most research, however, has examined how different combinations of demographic categories, such as age and race (e.g., Gawronski et al., 2010; Jones & Fazio, 2010; Todd et al., 2021) or gender and race (e.g., Jones & Fazio, 2010; Macrae et al., 1995), jointly affect person impressions. Fewer studies have tested how combinations of demographic categories and other categories (e.g., emotion) jointly shape person construal (e.g., Hugenberg & Bodenhausen, 2003; Kubota & Ito, 2007; Raissi & Steele, 2021), even in the absence of category salience. Importantly, emotion categories may affect threat impressions differently than demographic categories do. For example, because emotion expressions are dynamic and provide an immediate signal of a target's approachability, they may provide a more bottom-up, "evidence based" route (Ames, 2005) to threat impressions than do demographic categories.

In three experiments, we investigated the joint effects of multiple social cues (i.e., those pertaining to gender and emotion categories) on threat impressions, as reflected in performance on the stereotype misperception task (SMT; Krieglmeyer & Sherman, 2012). The SMT is a well-validated sequential priming task that assesses the biasing impact of semantic content (e.g., threat appraisals) evoked by task-irrelevant prime faces on judgements of unrelated target faces. Across experiments, the prime faces varied on both gender (male, female)[1] and emotion expression (neutral, smiling). Examining judgements of target faces reveals the extent to which gender cues and expression cues detectable in the prime faces spontaneously influence threat impressions.

We were particularly interested in understanding the joint effects of gender categories and emotion categories in contexts where one category is especially salient. One factor affecting stimulus salience is its contextual distinctiveness (Taylor & Fiske, 1978). In Experiment 1, we altered how the prime faces appeared in the SMT. In the gender-salient condition, we varied the gender of the prime faces while holding emotion expression constant, thus ensuring that gender was more distinctive throughout the task. In the expression-salient condition, we varied the emotion expression of the prime faces while holding gender constant, thus ensuring that emotion expression was more distinctive. Experiment 2 was a close replication of Experiment 1, but it also included a control condition in which neither category was made salient. Inclusion of this condition afforded a baseline assessment of the default use of gender and emotion categories in threat impressions.

Finally, in Experiment 3, we replicated the procedure of Experiment 2 with a different category-salience manipulation. Participants completed a counting task (Olson & Fazio, 2003) in which they kept a mental tally of how many faces of a focal category (gender or emotion expression) appeared as primes in the SMT. In the gender-salient condition, participants counted male faces and female faces, thus directing their attention to gender. In the expression-salient condition, they counted smiling faces and neutral-expression faces, thus directing their attention to emotion expression. Like Experiment 2, Experiment 3 also included a control condition in which participants simply completed the SMT without a category-salience induction.

We expected that a particular category (e.g., gender) would affect threat impressions more when that category was salient than when a different category (e.g., emotion expression) was salient (e.g., Macrae et al., 1995; Todd et al., 2021). Although prior work has found that attending to a focal category amplifies use of that category and attenuates use of a non-salient category, questions remain about whether category salience *eliminates* use of the non-salient category, or whether the non-salient category is still used, albeit more weakly. Furthermore, because prior research on category salience has not examined emotion alongside a demographic category (e.g., gender), category salience might affect use of these categories differently than in research that has only compared demographic categories. Accordingly, our general prediction was that category salience would shape category use in threat impressions, resulting in category salience × prime gender and category salience × prime expression interactions. We also considered the possibility that a prime gender × prime expression interaction would emerge, whereby gender bias is weaker for smiling faces than for neutral-expression faces. That is, independent of category salience, smiling expressions might dampen the differential threat stereotypically associated with male faces versus female faces, much like how smiling expressions have been found to dampen the differential threat stereotypically associated with Black faces versus White faces (Kubota & Ito, 2014; Richeson & Trawalter, 2008; see also Raissi & Steele, 2021).

2 | EXPERIMENT 1

2.1 | Method

2.1.1 | Participants

In our own prior research, we have observed substantial variability in the magnitude of category-salience effects on bias in behavioural analyses of sequential priming tasks, ranging from small (η_p^2 = .011; Rees et al., 2020, Experiment 3) to large (η_p^2 = .152; Todd et al., 2021, Experiment 1). In this and all subsequent experiments, we set target sample sizes that would afford at least 80% power to detect medium-sized (η_p^2 = .06) category salience × prime gender/prime expression interactions in the SMT.

Undergraduates (N = 131) participated for course credit. Data were excluded from one participant who experienced a data-writing error. The final sample comprised 130 participants. A sensitivity power analysis using G*Power (Faul et al., 2007) indicated that a sample size of 130 participants in a 2 × 2 × 2 mixed design afforded 80% power (α = .05) to detect category salience × prime gender/prime expression interactions as small as η_p^2 = .059. Participant demographics for all experiments appear in Table 1.

2.1.2 | Procedure

In this and all subsequent experiments, participants arrived at the lab in groups of up to eight. They were led by an experimenter[2] to a private

[1] We recognize that gender is not a binary category. Because all the face stimuli used in this research belonged to people who self-identified as men or women, we use the descriptors "male" and "female" for these stimuli.

[2] Experimenter gender was not recorded; however, the majority of the experimenters were women.

TABLE 1 Sample Size and Participant Demographics in Each Experiment

		Age		Gender (%)			Race/Ethnicity (%)				
Experiment	*N*	*M*	*SD*	Male	Female	Nonbinary	W	B	A	L	M
1	130	19.8	2.4	14.6	85.4	0.0	13.8	0.8	57.7	20.0	5.4
2	253	20.0	1.7	21.3	75.9	0.0	14.5	1.6	49.4	27.3	7.2
3	297	19.5	2.5	14.5	84.5	0.4	17.8	1.7	44.8	29.0	5.7

Note. *N* = the number of participants included in the data analyses. Some participants did not report their gender or race/ethnicity. For race/ethnicity, W = White or European American, B = Black or African American, A = Asian American or Pacific Islander, L = Latinx or Hispanic, and M = reported other or more than one race/ethnicity.

cubicle where they completed an SMT (Krieglmeyer & Sherman, 2012) via computer. On each trial of the SMT, a prime face and target face appeared in quick succession. The prime faces were 12 self-identified White female faces (six smiling, six neutral-expression) and 12 self-identified White male faces (six smiling, six neutral-expression) from the Chicago Face Database (Ma et al., 2015).[3] We selected facial identities that (a) had both neutral-expression and smiling versions, (b) were reliably categorized as "male" or "female" (99% of all cases), and (c) did not have high unusualness ratings.

To determine whether the intended emotion expression was evident on the faces, we had 51 MTurk workers rate all 24 faces (order randomized) on emotion expression (1 = *extremely unhappy*, 4 = *neither unhappy nor happy*, 7 = *extremely happy*). Overall, faces with smiling expressions ($M = 6.07$, $SD = 0.43$) were rated as much happier than faces with neutral expressions ($M = 3.49$, $SD = 0.50$), $t(50) = 27.12$, $p < .001$, $d = 5.55$ $CI_{95\%}$ [4.39, 6.70]. This difference in perceived emotion expression also emerged separately for both male faces, $t(49) = 21.65$, $p < .001$, $d = 4.24$, $CI_{95\%}$ [3.32, 5.16], and female faces, $t(50) = 26.69$, $p < .001$, $d = 5.60$ $CI_{95\%}$ [4.42, 6.80].

The target faces, which have been used extensively in SMT research (Krieglmeyer & Sherman, 2012; Rees et al., 2019, 2020; Rivers et al., 2020), were blurred black-and-white drawings of 48 computer-generated facial morphs that varied in threat (Oosterhof & Todorov, 2008). They were created by morphing 24 unique facial identities with facial features that were +2 *SD* or −2 *SD* from the neutral point in threat, resulting in two groups of target faces that differed in their threat appearance (see Krieglmeyer & Sherman, 2012, for additional details).

Following two practice blocks (six trials each), participants completed four test blocks (72 trials each). Each trial had the following sequence: fixation cross (500 ms), prime face (150 ms), blank screen (50 ms), and target face (100 ms). Finally, a pattern mask appeared until participants judged the target face as "more threatening" or "less threatening" than the average target face in the task (see Figure 1).

Instructions urged participants to attend to the prime faces, but to avoid letting these faces affect their judgements of the target faces. Thus, any influence of the prime faces on judgements of the target faces can be described as unintentional. The proportion of "more threatening" responses after prime faces of each type serves as an indirect index of the threat impressions elicited by those faces. A greater proportion of "more threatening" responses after male versus female prime faces constitutes a *gender* bias, whereas a greater proportion of "more threatening" responses after neutral-expression versus smiling prime faces constitutes an *expression* bias.

To manipulate category salience, we structured the SMT so that either gender or emotion expression was more distinctive throughout the task (Macrae & Cloutier, 2009; Rees et al., 2020; Todd et al., 2021). Participants were randomly assigned to one of two category-salience conditions. In the *gender-salient* condition, faces of smiling men and smiling women appeared together as primes within some blocks of trials, and faces of neutral-expression men and neutral-expression women appeared together as primes within other blocks of trials. Varying gender while holding emotion expression constant within the same block of trials renders gender more salient than emotion expression in that block of trials (see Figure 2). In the *expression-salient* condition, faces of smiling men and neutral-expression men appeared together as primes within some blocks of trials, and faces of smiling women and neutral-expression women appeared together as primes within other blocks of trials. Varying emotion expression while holding gender constant within the same block of trials renders emotion expression more salient than gender in that block of trials. In both conditions, block order was counterbalanced across participants. Block order did not moderate gender bias or expression bias here or in Experiment 2.

2.2 | Results

Descriptive statistics for all experiments appear in Table 2. We submitted the proportion of "more threatening" responses to a 2 (Category Salience) × 2 (Prime Gender) × 2 (Prime Expression) mixed ANOVA, with repeated measures on the last two factors.[4] There were

[3] As is customary in the SMT, we also included a featureless face outline as a neutral prime on some trials. These trials are necessary for conducting multinomial processing tree models (i.e., the SMT model), the aim of which is to decompose the cognitive mechanisms underlying the pattern of behavioural responses. The SMT model is primarily concerned with measuring the extent to which stereotype activation and stereotype application independently contribute to behavioural performance on the SMT (Krieglmeyer & Sherman, 2012). We do not discuss the results of multinomial modelling analyses because model fit was unusually poor (all G^2s > 425, all *p*s < .001). In addition, the model failed to converge on a solution for parameter estimates in some experiments (details about model fit analyses are available from the first author on request). Consequently, we felt that presenting the results of the SMT model was inappropriate. Because they were not of theoretical interest, we do not include the trials with the featureless face outlines as primes in the analyses reported below.

[4] Gender (expression) bias on the SMT is indicated by a gender (expression) prime main effect, without regard to the "target" variable (Krieglmeyer & Sherman, 2012). Thus, target effects

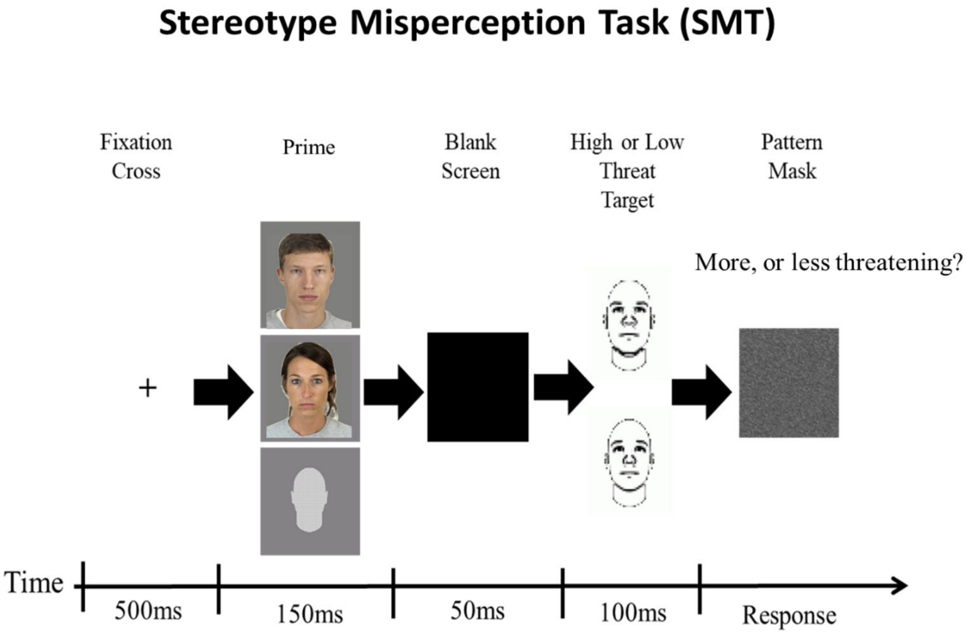

FIGURE 1 Depiction of the trial sequence in the stereotype misperception task

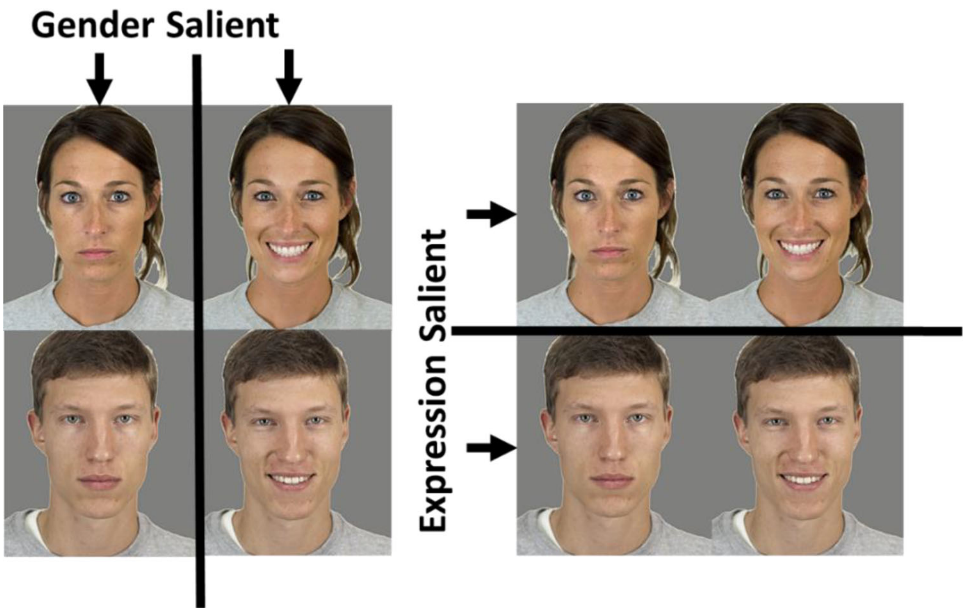

FIGURE 2 Depiction of the category-salience conditions

significant main effects of prime gender (male > female; gender bias), $F(1, 128) = 84.67$, $p < .001$, $\eta_p{}^2 = .398$, $CI_{90\%}$ [.290, .486],[5] and prime expression (neutral > smiling; expression bias), $F(1, 128) = 235.01$, $p < .001$, $\eta_p{}^2 = .647$, $CI_{90\%}$ [.566, .704].

Next, we examined whether category salience moderated gender bias or expression bias. The category salience × prime gender interaction was significant, $F(1, 128) = 19.29$, $p < .001$, $\eta_p{}^2 = .131$, $CI_{90\%}$ [.052, .222] (see Figure 3). Decomposing this interaction revealed that gender bias (i.e., the prime gender simple effect) was weaker in the expression-salient condition, $F(1, 63) = 19.17$, $p < .001$, $\eta_p{}^2 = .233$, $CI_{90\%}$ [.094, .366], than in the gender-salient condition, $F(1, 65) = 67.35$, $p < .001$, $\eta_p{}^2 = .509$, $CI_{90\%}$ [.361, .609].

The category salience × prime expression interaction was also significant, $F(1, 128) = 46.70$, $p < .001$, $\eta_p{}^2 = .267$, $CI_{90\%}$ [.163, .363]

were not of theoretical interest and were not included in the analyses. All conclusions remain unchanged when this variable is included in analyses. High-threat targets were judged as "more threatening" than were low-threat targets in all three experiments.

[5] Following Steiger (2004), we report 90% confidence intervals (CIs) for effect size estimates ($\eta_p{}^2$) accompanying all significant F tests.

TABLE 2 Proportion of "More Threatening" Judgements by Category Salience, Prime Gender, and Prime Expression (Experiments 1, 2, and 3)

	Prime expression and gender			
	Neutral-Expression		Smiling	
Category Salience	Male	Female	Male	Female
Experiment 1				
Gender salient ($n = 66$)	0.57 (0.29)	0.29 (0.26)	0.33 (0.27)	0.14 (0.18)
Expression salient ($n = 64$)	0.72 (0.26)	0.62 (0.26)	0.19 (0.18)	0.13 (0.16)
Experiment 2				
Gender salient ($n = 105$)	0.64 (0.24)	0.35 (0.24)	0.39 (0.28)	0.18 (0.22)
Control ($n = 51$)	0.71 (0.23)	0.57 (0.28)	0.23 (0.18)	0.15 (0.16)
Expression salient ($n = 97$)	0.71 (0.23)	0.67 (0.24)	0.23 (0.20)	0.17 (0.18)
Experiment 3				
Gender salient ($n = 98$)	0.72 (0.21)	0.52 (0.27)	0.24 (0.22)	0.13 (0.18)
Control ($n = 100$)	0.79 (0.19)	0.58 (0.28)	0.21 (0.22)	0.11 (0.18)
Expression salient ($n = 99$)	0.73 (0.23)	0.60 (0.28)	0.21 (0.24)	0.14 (0.20)

Note. n = the number of participants included in the data analyses for each category-salience condition. Standard deviations are in parentheses.

FIGURE 3 Proportion of "more threatening" judgements by category salience and prime gender in Experiment 1. Error bars depict standard errors

(see Figure 4). Decomposing this interaction revealed that expression bias (i.e., the prime expression simple effect) was weaker in the gender-salient condition, $F(1, 65) = 43.77$, $p < .001$, $\eta_p^2 = .402$, $CI_{90\%}$ [.247, .519], than in the expression-salient condition, $F(1, 63) = 207.17$, $p < .001$, $\eta_p^2 = .767$, $CI_{90\%}$ [.678, .817].

Finally, there was a significant prime gender × prime expression interaction, $F(1, 128) = 6.39$, $p = .013$, $\eta_p^2 = .048$, $CI_{90\%}$ [.006, .119] (see Figure 5). Decomposing this interaction revealed that gender bias (i.e., the prime gender simple effect) was weaker for smiling prime faces, $F(1, 129) = 46.52$, $p < .001$, $\eta_p^2 = .265$, $CI_{90\%}$ [.161, .361], than for neutral-expression prime faces, $F(1, 129) = 58.76$, $p < .001$, $\eta_p^2 = .313$, $CI_{90\%}$ [.206, .407].

2.3 | Discussion

In Experiment 1, category salience moderated the use of both gender categories and emotion categories in threat impressions: Gender bias was weaker when emotion expression was salient, and expression bias was weaker when gender was salient. These results align with prior evidence indicating that category salience can reduce bias in threat impressions (Rees et al., 2020; Todd et al., 2021) but with different categories. Although category salience altered how much either gender or emotion expression affected impressions, it did not eliminate gender bias or expression bias. Both categories continued to affect impressions (i.e., gender bias and expression bias were both still

FIGURE 4 Proportion of "more threatening" judgements by category salience and prime expression in Experiment 1. Error bars depict standard errors

FIGURE 5 Proportion of "more threatening" judgements by prime expression and prime gender in Experiment 1. Error bars depict standard errors

evident) regardless of which one was more salient. This latter finding contrasts with some prior evidence indicating that attending to one category can eliminate the effect of the non-salient category on person impressions (Gawronski et al., 2010; Macrae et al., 1995; but see Todd et al., 2021). Additionally, gender bias was weaker when the prime faces were smiling, which complements prior evidence indicating that smiling expressions can reduce *racial* biases in threat appraisals (Kubota & Ito, 2014; Richeson & Trawalter, 2008).

3 | EXPERIMENT 2

Because the category-salience manipulation in Experiment 1 drew attention to either gender or emotion expression, it is unclear if expression salience reduced gender bias, gender salience increased gender bias, or both. Similarly, it is unclear if gender salience reduced expression bias, expression salience increased expression bias, or both. In Experiment 2, we included a control condition wherein all prime types appeared randomly, thus affording a baseline assessment of the default use of gender and emotion categories in threat impressions.

3.1 | Method

3.1.1 | Participants

Undergraduates (N = 269) participated for course credit. Data were excluded from participants who experienced data-writing errors ($n = 9$)

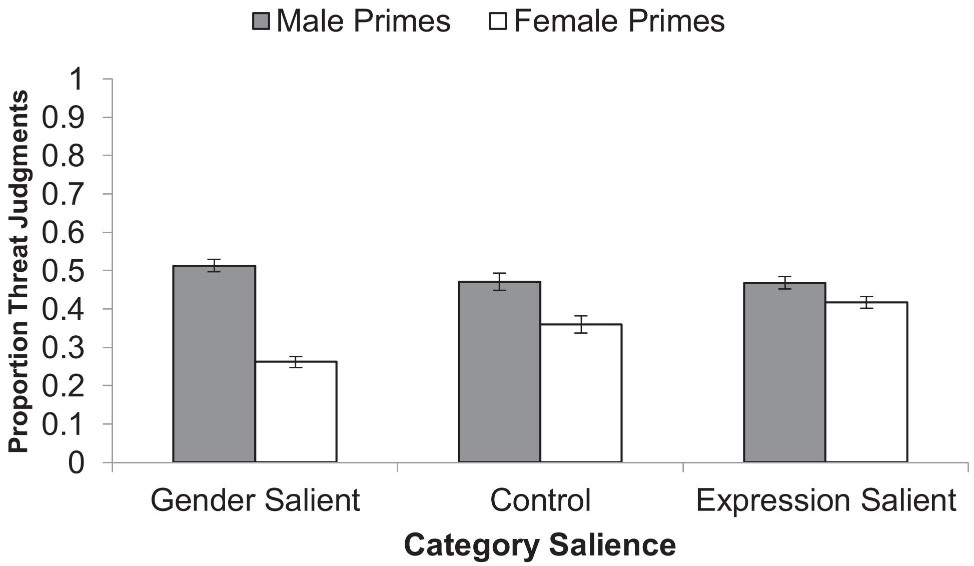

FIGURE 6 Proportion of "more threatening" judgements by category salience and prime gender in Experiment 2. Error bars depict standard errors

or whose SMT performance was 2.5 *SD* above their condition mean ($n = 2$). The final sample comprised 253 participants. A sensitivity power analysis using G*Power (Faul et al., 2007) indicated that a sample size of 253 participants in a 3 × 2 × 2 mixed design afforded 80% power ($\alpha = .05$) to detect category salience × prime gender/prime expression interactions as small as $\eta_p^2 = .038$. The preregistration for this experiment can be found at https://osf.io/xgszb.[6]

3.1.2 | Procedure

This experiment was nearly identical to Experiment 1, the only difference being the inclusion of a control condition[7] wherein all possible combinations of prime faces (male–neutral, female–neutral, male–smiling, female–smiling) appeared randomly, thus affording baseline assessments of gender bias and expression bias when neither category is especially salient. The order of prime expression and prime gender was counterbalanced in the gender-salient and expression-salient conditions.

[6] A typo in the target sample size of the preregistration stated that we planned to collect more data than intended. This sample estimate was intended for Experiment 3, which we powered more highly due to using a new category-salience manipulation. Both experiments were preregistered and conducted at similar times, accounting for the typo. In both Experiments 2 and 3, there is a typo in the "Study Design" section of the preregistration stating that we would manipulate prime race rather than prime gender. This was a copy-and-paste typo from preregistrations for other studies in which we have varied prime race in the SMT (i.e., studies assessing racial bias rather than gender bias in threat impressions).

[7] Due to a clerical error, the control condition ($n = 51$) had fewer participants than did the gender-salient ($n = 105$) and expression-salient conditions ($n = 97$). Although ANOVA is robust to unequal sample sizes, we conducted an analysis using a random sample of 51 participants from the gender-salient and expression-salient conditions as a robustness check. This analysis revealed virtually identical results, suggesting that the unequal sample sizes did not affect our conclusions.

3.2 | Results

We submitted the proportion of "more threatening" responses to a 3 (Category Salience) × 2 (Prime Gender) × 2 (Prime Expression) mixed ANOVA, with repeated measures on the last two factors. There were significant main effects of prime gender (male > female), $F(1, 250) = 144.08$, $p < .001$, $\eta_p^2 = .366$, $CI_{90\%}$ [.289, .433], prime expression (neutral > smiling), $F(1, 250) = 360.25$, $p < .001$, $\eta_p^2 = .590$, $CI_{90\%}$ [.529, .639], and category salience (expression salient > control > gender salient), $F(2, 250) = 4.40$, $p = .013$, $\eta_p^2 = .034$, $CI_{90\%}$ [.004, .074].

Next, we examined whether category salience moderated gender bias or expression bias. The category salience × prime gender interaction was significant, $F(2, 250) = 35.38$, $p < .001$, $\eta_p^2 = .221$, $CI_{90\%}$ [.145, .288] (see Figure 6). Decomposing this interaction revealed that gender bias (i.e., the prime gender simple effect) was weakest in the expression-salient condition, $F(1, 96) = 17.75$, $p < .001$, $\eta_p^2 = .156$, $CI_{90\%}$ [.059, .264], strongest in the gender-salient condition, $F(1, 104) = 131.28$, $p < .001$, $\eta_p^2 = .560$, $CI_{90\%}$ [.451, .634], with the control condition intermediate between the other two conditions, $F(1, 50) = 35.00$, $p < .001$, $\eta_p^2 = .412$, $CI_{90\%}$ [.234, .539].

Approaching this interaction differently, we created a gender-bias index (proportion of "more threatening" judgements on male prime trials minus proportion of "more threatening" judgements on female prime trials) and examined the effect of category salience on this index. The gender-bias index was weaker in the expression-salient condition than in either the gender-salient condition, $t(250) = 8.25$, $p < .001$, $d = 1.16$, $CI_{95\%}$ [0.87, 1.46], or the control condition, $t(250) = 2.01$, $p = .045$, $d = 0.35$, $CI_{95\%}$ [0.01, 0.69]. The gender-bias index was also stronger in the gender-salient condition than in the control condition, $t(250) = 4.77$, $p < .001$, $d = 0.81$, $CI_{95\%}$ [0.47, 1.16].

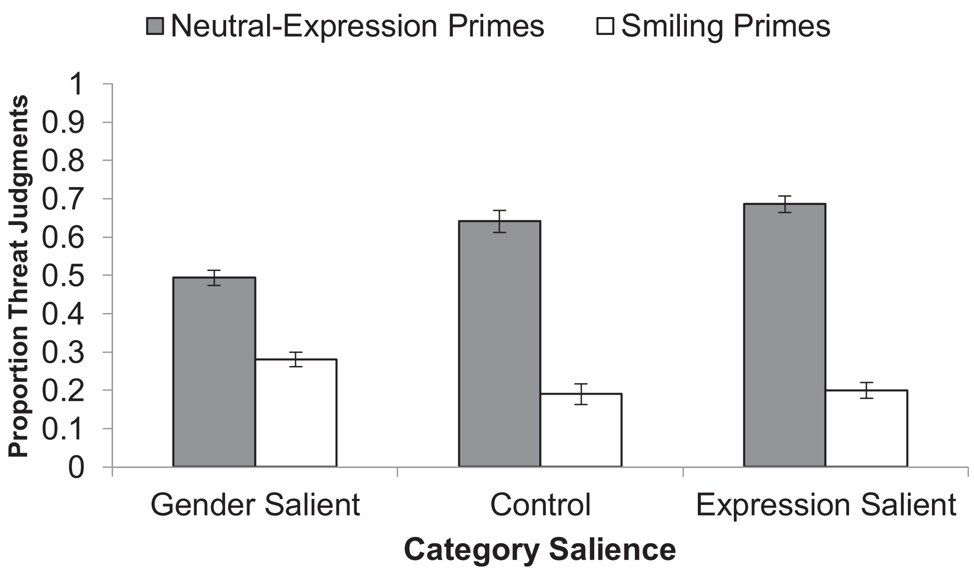

FIGURE 7 Proportion of "more threatening" judgements by category salience and prime expression in Experiment 2. Error bars depict standard errors

The category salience × prime expression interaction was significant, $F(2, 250) = 22.86$, $p < .001$, $\eta_p^2 = .155$, $CI_{90\%}$ [.088, .218] (see Figure 7). Decomposing this interaction revealed that expression bias (i.e., the prime expression simple effect) was weakest when gender was salient, $F(1, 104) = 62.58$, $p < .001$, $\eta_p^2 = .402$, $CI_{90\%}$ [.255, .474], strongest when expression was salient, $F(1, 96) = 218.56$, $p < .001$, $\eta_p^2 = .695$, $CI_{90\%}$ [.609, .750], with the control condition falling between the other two conditions, $F(1, 50) = 97.63$, $p < .001$, $\eta_p^2 = .661$, $CI_{90\%}$ [.522, .740].

Approaching this interaction differently, we created an expression-bias index (proportion of "more threatening" judgements on neutral-expression prime trials minus proportion of "more threatening" judgements on smiling prime trials) and examined the effect of category salience on this index. The expression-bias index was weaker in the gender-salient condition than in either the expression-salient condition, $t(250) = 6.38$, $p < .001$, $d = 0.90$, $CI_{95\%}$ [0.61, 1.19], or the control condition, $t(250) = 4.60$, $p < .001$, $d = 0.79$, $CI_{95\%}$ [0.44, 1.13]. The expression-bias index in the expression-salient and control conditions did not significantly differ, $t(250) = 0.67$, $p = .504$, $d = 0.12$, $CI_{95\%}$ [−0.22, 0.45].

Finally, there was a significant prime gender × prime expression interaction, $F(1, 250) = 5.60$, $p = .019$, $\eta_p^2 = .022$, $CI_{90\%}$ [.002, .060] (see Figure 8). Decomposing this interaction revealed that, contrary to Experiment 1, gender bias (i.e., the prime gender simple effect) was weaker for neutral-expression prime faces, $F(1, 252) = 100.82$, $p < .001$, $\eta_p^2 = .286$, $CI_{90\%}$ [.211, .355], than for smiling prime faces, $F(1, 252) = 104.92$, $p < .001$, $\eta_p^2 = .294$, $CI_{90\%}$ [.219, .363].[8]

3.3 | Discussion

In Experiment 2, category salience moderated the extent to which gender categories and emotion categories guided threat impressions, but again it did not eliminate either gender bias or expression bias. Regardless of which category was more salient, gender bias and expression bias were both evident. These results replicate findings from Experiment 1 and prior findings of category salience attenuating (but not eliminating) bias in threat impressions (Rees et al., 2020; Todd et al., 2021); however, they are inconsistent with prior findings of category salience eliminating bias (Gawronski et al., 2010; Macrae et al., 1995).

Novel to Experiment 2 was the inclusion of a control condition, which afforded an assessment of the default use of gender categories and emotion categories in threat impressions. Gender bias was weaker in the expression-salient condition and stronger in the gender-salient condition relative to control. Expression bias was weaker in the gender-salient condition relative to control; however, comparable expression bias emerged in the expression-salient and control conditions. This latter finding suggests that the default use of emotion categories may already be at ceiling in this task, which entailed appraising threat.

Unexpectedly, the pattern of responses underlying the prime expression × prime gender interaction was opposite to that in Experiment 1. That is, gender bias was *stronger* for smiling versus neutral-expression prime faces. We hesitate to speculate on the differing nature of this interaction pending replication.

[8] This reversal of the prime gender × prime expression interaction may best be considered in the context of an unexpected category salience × prime gender × prime expression interaction, $F(2, 250) = 5.24$, $p = .006$, $\eta_p^2 = .040$, $CI_{90\%}$ [.007, .083]. Follow-up analyses revealed that the expected prime gender × prime expression interaction (i.e., weaker gender bias for smiling versus neutral prime faces) emerged only in the gender-salient condition. The prime gender × prime expression interaction was not significant in the expression-salient condition and was weakly reversed in the control condition. The category salience × prime gender × prime expression interaction was not significant in either of the other experiments. We have included a full decomposition of this interaction in the supplementary materials.

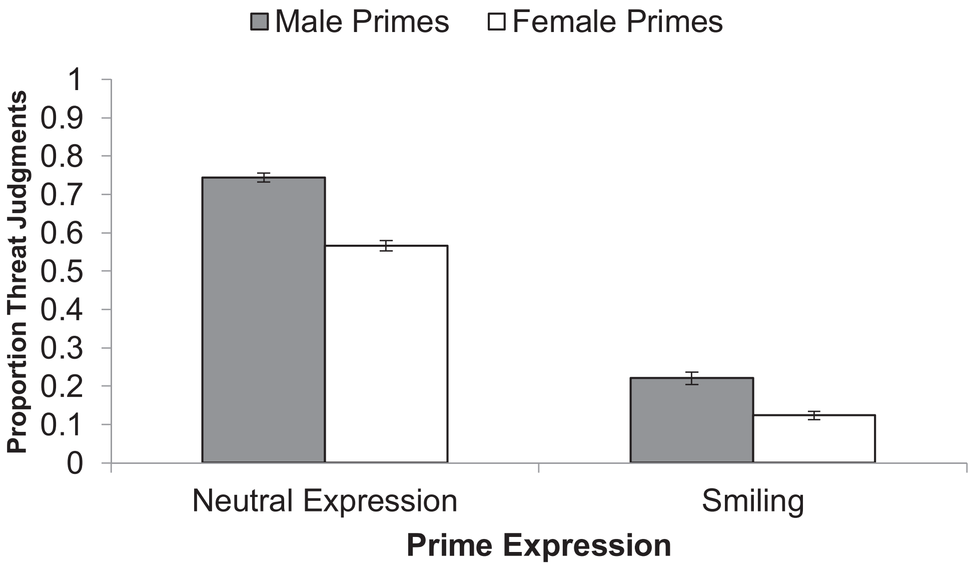

FIGURE 8 Proportion of "more threatening" judgements by prime expression and prime gender in Experiment 2. Error bars depict standard errors

4 | EXPERIMENT 3

Our first two experiments manipulated category salience by varying the structure of the SMT. Although we had no reason to expect that our findings are unique to this manipulation, we nevertheless conducted a conceptual replication with a different category-salience manipulation. In Experiment 3, we used a counting task developed by Olson and Fazio (2003) wherein participants counted the number of male faces and female faces (gender-salient condition) or the number of smiling faces and neutral-expression faces (expression-salient condition) that appeared as primes during the SMT. Such counting manipulations have been effective in directing attention to different applicable categories of prime faces in sequential priming tasks (Gawronski et al., 2010; Jones & Fazio, 2010).

4.1 | Method

4.1.1 | Participants

Undergraduates (N = 319) participated for course credit. Data were excluded from participants who experienced data-writing errors (n = 1), whose SMT performance was 2.5 *SD* above their condition mean (n = 2), and who incorrectly reported their instructions condition or did not answer the count question on all trials (n = 19). All data exclusions based on the manipulation check were preregistered at https://osf.io/jkmbs. The final sample comprised 297 participants. A sensitivity power analysis using G*Power (Faul et al., 2007) indicated that a sample size of 297 participants in a 3 × 2 × 2 mixed design afforded 80% power ($\alpha = .05$) to detect category salience × prime gender/prime expression interactions as small as $\eta_p^2 = .033$.

4.1.2 | Procedure

Experiment 3 was nearly identical to Experiment 2, the only difference being the category-salience manipulation. Participants were randomly assigned to count prime faces according to either their gender or their emotion expression (instructions adapted from Gawronski et al., 2010). In the *gender-salient* [*expression-salient*] condition, participants kept "a running mental tally of the number of appearances of both men and women [of faces with neutral and smiling expressions] as they appear during each block". In the *control* condition, they kept "a running mental tally of the faces that appear BEFORE the target faces during each block". All participants completed an SMT in which all possible combinations of prime gender and prime expression appeared randomly across four blocks. After each block, participants reported how many faces they counted according to their instructions.

4.2 | Results

We submitted the proportion of "more threatening" responses to a 3 (Category Salience) × 2 (Prime Gender) × 2 (Prime Expression) mixed ANOVA, with repeated measures on the last two factors. There were significant main effects of prime gender (male > female), $F(1, 294) = 202.98$, $p < .001$, $\eta_p^2 = .408$, $CI_{90\%}$ [.339, .468], and prime expression (neutral > smiling), $F(1, 294) = 631.60$, $p < .001$, $\eta_p^2 = .682$, $CI_{90\%}$ [.636, .719].

Next, we examined whether category salience moderated gender bias. The category salience × prime gender interaction was significant, $F(2, 294) = 3.36$, $p = .036$, $\eta_p^2 = .022$, $CI_{90\%}$ [.0008, .053] (see Figure 9). Decomposing this interaction revealed that gender bias (i.e., the prime gender simple effect) was weaker in the expression-salient condition,

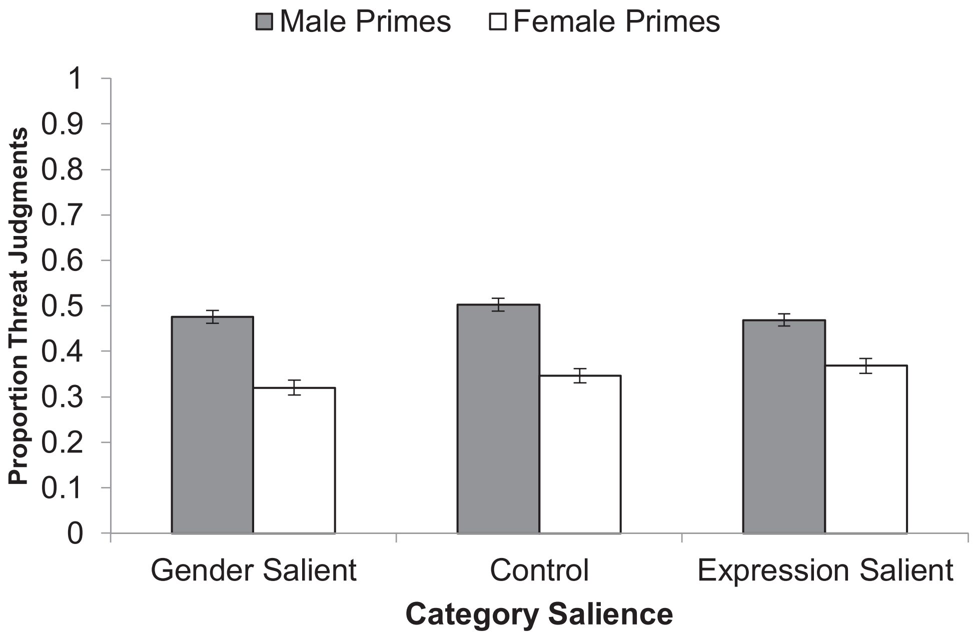

FIGURE 9 Proportion of "more threatening" judgements by category salience and prime gender in Experiment 3. Error bars depict standard errors

$F(1, 105) = 61.06$, $p < .001$, $\eta_p^2 = .368$, $CI_{90\%}$ [.247, .466], than in either the control condition, $F(1, 101) = 74.37$, $p < .001$, $\eta_p^2 = .424$, $CI_{90\%}$ [.302, .518], or the gender-salient condition, $F(1, 103) = 72.60$, $p < .001$, $\eta_p^2 = .413$, $CI_{90\%}$ [.293, .508].

Approaching this interaction differently, we examined the effect of category salience on the same gender-bias index from Experiment 2. This gender-bias index was weaker in the expression-salient condition than in either the gender-salient condition, $t(294) = 2.20$, $p = .029$, $d = 0.31$, $CI_{95\%}$ [0.03, 0.59], or the control condition, $t(294) = 2.30$, $p = .022$, $d = 0.33$, $CI_{95\%}$ [0.05, 0.60]. The gender-bias index in the control and gender-salient conditions did not significantly differ, $t(294) = 0.08$, $p = .934$, $d = 0.01$, $CI_{95\%}$ [−0.27, 0.29].

We next tested whether category salience moderated expression bias. Unlike Experiments 1 and 2, the category salience × prime expression interaction was not significant, $F(1, 294) = 1.81$, $p = .180$, $\eta_p^2 = .006$ (see Figure 10).

There was also a prime gender × prime expression interaction, $F(1, 294) = 60.18$, $p < .001$, $\eta_p^2 = .170$, $CI_{90\%}$ [.109, .233] (see Figure 11). As in Experiment 1 but unlike Experiment 2, decomposing this interaction revealed that gender bias (i.e., the prime gender simple effect) was weaker for smiling prime faces, $F(1, 296) = 117.16$, $p < .001$, $\eta_p^2 = .284$, $CI_{90\%}$ [.214, .348], than for neutral-expression prime faces, $F(1, 296) = 193.59$, $p < .001$, $\eta_p^2 = .395$, $CI_{90\%}$ [.326, .456].

4.3 | Discussion

These results largely replicated those from the first two experiments. Category salience moderated the extent to which gender categories affected threat impressions: Attending to emotion categories weakened gender bias, relative to both the control and gender-salient conditions. Unlike Experiments 1 and 2, however, category salience did not moderate the effect of emotion expression on threat impressions. Rather, there was evidence of expression bias in all conditions. Once again, a prime expression × prime gender interaction indicated that, as in Experiment 1 but unlike Experiment 2, gender bias was weaker when prime faces were smiling than when they had a neutral expression. Finally, we continued to observe effects of prime gender and prime expression on threat impressions regardless of category salience.

5 | INTERNAL META-ANALYSIS

We used Goh et al.'s (2016) mini meta-analysis template, which implements a fixed-effects procedure and inverse variance weighting, to quantify the cumulative effects of category salience on gender bias and expression bias. The experiments reported here are the only ones we have conducted examining category-salience effects on gender bias and expression bias in threat impressions (i.e., there is no file drawer).

First, we compared the three category-salience conditions on the gender-bias index (see Experiment 2). Gender bias was weaker in the expression-salient condition than in the gender-salient condition ($d = 0.71$, $CI_{95\%}$ [0.54, 0.89], $z = 7.91$, $p < .001$), and in the control condition ($d = 0.40$, $CI_{95\%}$ [0.18, 0.61], $z = 3.59$, $p = .003$). Gender bias was also stronger in the gender-salient condition than in the control condition ($d = 0.27$, $CI_{95\%}$ [0.06, 0.49], $z = 2.47$, $p = .014$).

Next, we compared the three category-salience conditions on the expression-bias index (see Experiment 2). Expression bias was weaker in the gender-salient condition than in the expression-salient condition

FIGURE 10 Proportion of "more threatening" judgements by category salience and prime expression in Experiment 3. Error bars depict standard errors

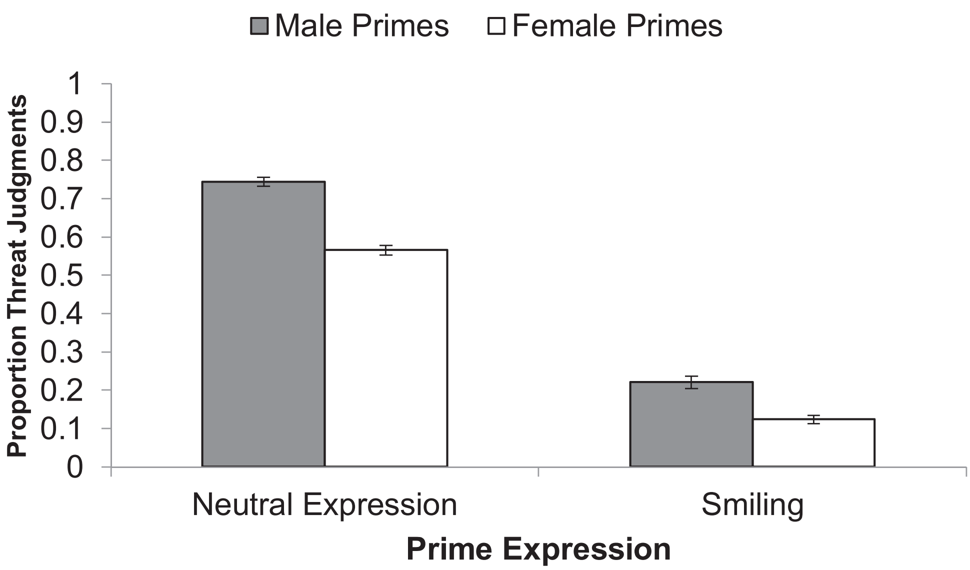

FIGURE 11 Proportion of "more threatening" judgements by prime expression and prime gender in Experiment 3. Error bars depict standard errors

($d = 0.68$, $CI_{95\%}$ [0.50, 0.85], $z = 7.48$, $p < .001$), and in the control condition ($d = 0.49$, $CI_{95\%}$ [0.28, 0.71], $z = 4.44$, $p < .001$). Expression bias in the latter two conditions did not significantly differ ($d = 0.02$, $CI_{95\%}$ [−0.20, 0.23], $z = 0.16$, $p = .874$).

These meta-analytic tests indicate that category salience had pronounced effects on both gender bias and expression bias. Attending to gender categories strengthened gender bias, but it weakened expression bias, relative to control. Attending to emotion categories, by contrast, weakened gender bias, but it did not significantly affect expression bias, relative to control.

6 | GENERAL DISCUSSION

Three experiments investigated threat impressions in contexts with multiple informative social cues. We manipulated the salience of two categories—gender and emotion expression—and found that category salience shaped the extent of both gender bias and expression bias. Meta-analytic results indicated that making a target category (e.g., emotion expression) salient reduced bias deriving from another category (e.g., gender) relative to control. Contrary to some prior work (Gawronski et al., 2010; Macrae et al., 1995) but consistent with other work (Todd et al., 2021), category salience did not eliminate the use of either category on threat impressions: Participants were influenced by both gender cues and emotion expression cues, judging target faces as less threatening when they were preceded by prime faces that were female or smiling versus male or neutral in expression. In addition, emotion categories and gender categories had an interactive effect on threat impressions; in two of the three experiments, gender bias was weaker when prime faces were smiling than when they had a neutral expression.

Overall, these results generally replicated other findings indicating that category salience moderates the use of social categories in person impressions (e.g., Gawronski et al., 2010; Jones & Fazio, 2010; Macrae et al., 1995; Rees et al., 2020; Todd et al., 2021). Unlike much of the previous research, we included a control condition, which allowed us to determine whether category salience increased the influence of the focal category, decreased the influence of another category, or both. We found that making a category salient reduced the effect of the non-salient category on threat impressions relative to baseline, but that the non-salient category still exerted an influence on threat impressions.

Our experiments also addressed the question of whether category-salience effects on threat impressions differ depending on the type of category in question—specifically, whether a more dynamic cue like emotion expression affects impressions differently than a demographic category like gender. Although both cues consistently affected threat impressions, effect sizes for expression bias were larger than were those for gender bias across experiments, suggesting that emotion expression was more impactful than gender in guiding threat impressions. Furthermore, meta-analytic results revealed that gender salience increased gender bias relative to control, whereas expression salience did not increase expression bias. Considered alongside the findings that expression bias was stronger than gender bias, expression bias might have been at ceiling in the control condition. Altogether, these findings suggest that emotion categories might have a greater impact than gender categories on threat impressions. Such results are compatible with the idea that emotion expression may be a more diagnostic cue for determining a target person's approachability during threat appraisal (Ames, 2005). That emotion expression had such a powerful effect on threat impressions underscores the importance of including non-demographic cues in research examining the joint effects of multiple social cues on person construal.

Notably, the category-salience effects on the SMT we observed here differ from category-salience effects on the affect misattribution procedure (AMP; Payne et al., 2005) observed elsewhere. For example, Gawronski et al. (2010) found that directing attention to the age (race) of prime faces that varied in both age (younger adult, older adult) and race (Black, White) had no discernible effect on racial (age) bias on the AMP. Although the SMT and the AMP are both sequential priming tasks, different mechanisms underlie performance on the two tasks (Sherman & Rivers, 2021). Whereas AMP performance is driven by a misattribution mechanism (Payne et al., 2010), SMT performance is driven by a response-conflict mechanism (Krieglmeyer & Sherman, 2012). Indeed, when using a response-conflict task (evaluative priming task; Fazio et al., 1995), Gawronski et al. (2010) found that racial bias and age bias were both weaker when the non-focal category was salient. This finding aligns with the current results and with other research documenting category-salience effects on response-conflict tasks (Jones & Fazio, 2010; Todd et al., 2021). Given the differences between our findings and that of previous work, future experiments should use multiply categorizable stimuli with other sequential priming tasks to better understand the extent to which our findings are generalizable.

That we used a misattribution-driven task to examine gender bias and expression bias may have influenced the relative effectiveness of our different category-salience manipulations. In our first two experiments, we grouped the prime stimuli in a way that made either gender or emotion expression the more salient category, which is a more stimulus-driven method for manipulating category salience. In Experiment 3, we had participants count the number of primes of each type, which is a more goal-driven method for manipulating category salience. The category-salience effects in Experiments 1 and 2 were stronger than they were in Experiment 3. It is possible that performance on the SMT, which measures unintentional reactions to prime stimuli, is more responsive to stimulus-driven manipulations than to more overt, goal-driven manipulations. Although future research that directly tests this possibility is needed, these findings are consistent with some prior research using the SMT. Specifically, whereas cognitive load has little effect on SMT bias, stimulus-driven manipulations and cognitively efficient interventions (e.g., forming implementation intentions) can have pronounced effects on SMT bias (Rees et al., 2019, 2020).

Multiple potential cognitive mechanisms may underlie the current findings. We found that salient categories increased bias relative to a control condition, which could suggest that category salience increases the *activation* or the *application* of information pertaining to that specific category (e.g., gender) relative to non-salient categories (e.g., emotion). Our multinomial modelling analyses, which aim to disentangle activation and application processes, resulted in poor model fit (see footnote 3); thus, we hesitate to speculate on the processes at play here. Future research will be needed to make more conclusive claims about the specific mechanism(s) underlying the findings reported here.

Our findings align theoretically with integration models of person construal (Freeman et al., 2020), which propose that although category salience increases the weighting of particular social cues, multiple applicable categories are concurrently considered at above-baseline levels. Accordingly, making one category salient (e.g., emotion expression) should reduce the use of a different category (e.g., gender), but it should not eliminate the use of either category on person impressions. Integration models also predict that categories may interact in driving impressions. Our findings comport with such models: Although category salience reliably moderated bias strength, both gender bias and expression bias persisted in all conditions. Furthermore, we often observed an interaction between prime gender and prime expression on threat impressions, suggesting that these social cues were not used independently.

The current work indicates that attending to non-demographic cues (e.g., emotion expression) can reduce the use of demographic categories (e.g., gender) in person impressions. Our findings may be relevant to bias-reduction research, which is often concerned with discouraging perceivers from using demographic categories and their associated stereotypes when forming impressions, and instead encouraging them to use other relevant information about a person. We found that although attending to emotion categories dampened the use of gender cues, gender cues continued to bias threat impressions. These results suggest that it may be difficult to eliminate gender bias simply by structuring the environment in a way that encourages perceivers to attend to other applicable social cues, such as the emotion expression on a person's face. Such findings may differ from past research because we used a task that measures the threat spontaneously evoked by prime faces, rather than a task that measures explicit threat judgements of those same faces (Radeke & Stahelski, 2020).

We conclude by noting several limitations of the current work, which suggest fruitful directions for future research. First, because we examined only two social categories (gender and emotion), both of which may be diagnostic for threat impressions, our specific findings might not apply to all category combinations or judgements. Furthermore, because we only used two exemplars of emotion expression (i.e., neutral expressions and smiling expressions indicative of happiness), it is unclear if other emotion expressions (e.g., scowling/frowning expressions indicative of anger) respond differently to category salience or if they potentially interact with demographic categories (e.g., gender) in a similar way to smiling expressions. In addition, the strength of gender bias observed here may have been affected by our participant sample. The majority of our participants were women, who may have evaluated their ingroup more positively and thus as less threatening. Finally, the emotion expressions we examined were unambiguous (i.e., the smiling expressions were clearly smiling), which raises questions about whether and to what extent less clearly expressed emotions likewise affect the use of gender categories. With many open questions remaining, future studies should continue to use multiply categorizable person stimuli to improve our understanding of biases in person impressions.

CONFLICT OF INTEREST

The authors have no conflict of interest to declare.

ETHICS STATEMENT

The manuscript adheres to ethical guidelines specified in the APA Code of Conduct as well as the authors' local ethics guidelines. Additional information and materials are available from the corresponding author upon request. This research was facilitated by National Science Foundation Grant BCS-1764097 (awarded to ART).

DATA AVAILABILITY STATEMENT

The data and syntax for these experiments are available at: https://osf.io/p4f8w/?view_only=7e8a8ea4a3d24519b48e9800235f7dd4.

ORCID

Heather R. Rees https://orcid.org/0000-0001-8755-3297

REFERENCES

Ames, D. R. Everyday solutions to the problem of other minds. In B. F. Malle & S. D. Hodges (Eds.), (2005). *Other minds: How humans bridge the divide between self and others* (pp. 158–173). Guilford Publications.

Archer, J. (2004). Sex differences in aggression in real-world settings: A meta-analytic review. *Review of General Psychology*, *8*, 291–322. https://doi.org/10.1037/1089-2680.8.4.291

Eimer, M., Holmes, A., & McGlone, F. P. (2003). The role of spatial attention in the processing of facial expression: An ERP study of rapid brain responses to six basic emotions. *Cognitive, Affective, & Behavioral Neuroscience*, *3*, 97–110.

Faul, F., Erdfelder, E., Lang, A. G., & Buchner, A. (2007). G* Power 3: A flexible statistical power analysis program for the social, behavioral, and biomedical sciences. *Behavior Research Methods*, *39*, 175–191. https://doi.org/10.3758/BF03193146

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*, 1013–1027. https://doi.org/10.1037/0022-3514.69.6.1013

Freeman, J. B., Stolier, R. M., & Brooks, J. A. (2020). Dynamic interactive theory as a domain-general account of social perception. *Advances in Experimental Social Psychology*, *61*, 237–287. https://doi.org/10.1016/bs.aesp.2019.09.005

Gawronski, B., Cunningham, W. A., LeBel, E. P., & Deutsch, R. (2010). Attentional influences on affective priming: Does categorisation influence spontaneous evaluations of multiply categorisable objects? *Cognition and Emotion*, *24*, 1008–1025. https://doi.org/10.1080/02699930903112712

Goh, J. X., Hall, J. A., & Rosenthal, R. (2016). Mini meta-analysis of your own studies: Some arguments on why and a primer on how. *Social and Personality Psychology Compass*, *10*, 535–549. https://doi.org/10.1111/spc3.12267

Haxby, J. V., Hoffman, E. A., & Gobbini, M. I. (2000). The distributed human neural system for face perception. *Trends in Cognitive Sciences*, *4*, 223–233. https://doi.org/10.1016/S1364-6613(00)01482-0

Hugenberg, K., & Bodenhausen, G. V. (2003). Facing prejudice: Implicit prejudice and the perception of facial threat. *Psychological Science*, *14*, 640–643. https://doi.org/10.1046/j.0956-7976.2003.psci_1478.x

Ito, T. A., & Urland, G. R. (2003). Race and gender on the brain: Electrocortical measures of attention to the race and gender of multiply categorizable individuals. *Journal of Personality and Social Psychology*, *85*, 616–626. https://doi.org/10.1037/0022-3514.85.4.616

Jones, C. R., & Fazio, R. H. (2010). Person categorization and automatic racial stereotyping effects on weapon identification. *Personality and Social Psychology Bulletin*, *36*, 1073–1085. https://doi.org/10.1177/0146167210375817

Krieglmeyer, R., & Sherman, J. W. (2012). Disentangling stereotype activation and stereotype application in the stereotype misperception task. *Journal of Personality and Social Psychology*, *103*, 205–224. https://doi.org/10.1037/a0028764

Kubota, J. T., & Ito, T. A. (2007). Multiple cues in social perception: The time course of processing race and facial expression. *Journal of Experimental Social Psychology*, *43*, 738–752. https://doi.org/10.1016/j.jesp.2006.10.023

Kubota, J. T., & Ito, T. A. (2014). The role of expression and race in weapons identification. *Emotion*, *14*, 1115–1124. https://doi.org/10.1037/a0038214

Ma, D. S., Correll, J., & Wittenbrink, B. (2015). The Chicago face database: A free stimulus set of faces and norming data. *Behavior Research Methods*, *47*, 1122–1135. https://doi.org/10.3758/s13428-014-0532-5

Macrae, C. N., Bodenhausen, G. V., & Milne, A. B. (1995). The dissection of selection in person perception: Inhibitory processes in social stereotyping. *Journal of Personality and Social Psychology*, *69*, 397–407. https://doi.org/10.1037/0022-3514.69.3.397

Macrae, C. N., & Cloutier, J. (2009). A matter of design: Priming context and person perception. *Journal of Experimental Social Psychology*, *45*, 1012–1015. https://doi.org/10.1016/j.jesp.2009.04.021

Neel, R., Becker, D. V., Neuberg, S. L., & Kenrick, D. T. (2012). Who expressed what emotion? Men grab anger, women grab happiness. *Journal of Experimental Social Psychology*, *48*, 583–586. https://doi.org/10.1016/j.jesp.2011.11.009

Olson, M. A., & Fazio, R. H. (2003). Relations between implicit measures of prejudice: What are we measuring? *Psychological Science*, *14*, 636–639. https://doi.org/10.1046/j.0956-7976.2003.psci_1477.x

Oosterhof, N. N., & Todorov, A. (2008). The functional basis of face evaluation. *Proceedings of the National Academy of Sciences*, *105*, 11087–11092. https://doi.org/10.1073/pnas.0805664105

Payne, B. K., Cheng, C. M., Govorun, O., & Stewart, B. D. (2005). An inkblot for attitudes: Affect misattribution as implicit measurement. *Journal of Personality and Social Psychology*, *89*, 277–293. https://doi.org/10.1037/0022-3514.89.3.277

Payne, B. K., Hall, D. L., Cameron, C. D., & Bishara, A. J. (2010). A process model of affect misattribution. *Personality and Social Psychology Bulletin*, *36*, 1397–1408. https://doi.org/10.1177/0146167210383440

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*, e12518. https://doi.org/10.1111/spc3.12518

Radeke, M. K., & Stahelski, A. J. (2020). Altering age and gender stereotypes by creating the Halo and Horns Effects with facial expressions. *Humanities and Social Sciences Communications*, *7*, 1–11. https://doi.org/10.1057/s41599-020-0504-6

Raissi, A., & Steele, J. R. (2021). Does emotional expression moderate implicit racial bias? Examining bias following smiling and angry primes. *Social Cognition*, *39*, 570–590. https://doi.org/10.1521/soco.2021.39.5.570

Rees, H. R., Rivers, A. M., & Sherman, J. W. (2019). Implementation intentions reduce implicit stereotype activation and application. *Personality and Social Psychology Bulletin*, *45*, 37–53. https://doi.org/10.1177/0146167218775695

Rees, H. R., Ma, D. S., & Sherman, J. W. (2020). Examining the relationships among categorization, stereotype activation, and stereotype application. *Personality and Social Psychology Bulletin*, *46*, 499–513. https://doi.org/10.1177/0146167219861431

Richeson, J. A., & Trawalter, S. (2008). The threat of appearing prejudiced and race-based attentional biases. *Psychological Science*, *19*, 98–102. https://doi.org/10.1111/j.1467-9280.2008.02052.x

Rivers, A. M., Sherman, J. W., Rees, H. R., Reichardt, R., & Klauer, K. C. (2020). On the roles of stereotype activation and application in diminishing implicit bias. *Personality and Social Psychology Bulletin*, *46*, 349–364. https://doi.org/10.1177/0146167219853842

Sherman, J. W., & Rivers, A. M. (2021). There's nothing social about social priming: Derailing the "train wreck. *Psychological Inquiry*, https://doi.org/10.1080/1047840X.2021.1889312

Steiger, J. H. (2004). Beyond the *F* test: Effect size confidence intervals and tests of close fit in the analysis of variance and contrast analysis. *Psychological Methods*, *9*, 164–182. https://doi.org/10.1037/1082-989X.9.2.164

Taylor, S. E., & Fiske, S. T. (1978). Salience, attention, and attribution: Top of the head phenomena. *Advances in Experimental Social Psychology*, *11*, 249–288. https://doi.org/10.1016/S0065-2601(08)60009-X

Todd, A. R., Johnson, D. J., Lassetter, B., Neel, R., Simpson, A. J., & Cesario, J. (2021). Category salience and racial bias in weapon identification: A diffusion modeling approach. *Journal of Personality and Social Psychology*, *120*, 672–693. https://doi.org/10.1037/pspi0000279

SUPPORTING INFORMATION

Additional supporting information may be found in the online version of the article at the publisher's website.

How to cite this article: Rees, H. R., Sherman, J. W., Klauer, K. C., & Todd, A. R. (2022). On the use of gender categories and emotion categories in threat-based person impressions. *European Journal of Social Psychology*, *52*, 597–610. https://doi.org/10.1002/ejsp.2840