Journal of Experimental Social Psychology 53 (2014) 185–192

Contents lists available at ScienceDirect

Journal of Experimental Social Psychology

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

The influence of target group status on the perception of the offensiveness of group-based slurs ☆

P.J. Henry [a,*], Sarah E. Butler [b], Mark J. Brandt [c]

[a] Department of Psychology, New York University – Abu Dhabi, United Arab Emirates
[b] Department of Psychology, The Sage Colleges, United States
[c] Department of Social Psychology, Tilburg University, Netherlands

HIGHLIGHTS

- We report two studies, one with experimental methodology.
- Slurs are seen as more offensive when directed at lower v. higher status groups.
- The effect is mediated by the expected emotional reaction of the target.

ARTICLE INFO

Article history:
Received 24 June 2013
Revised 19 March 2014
Available online 1 April 2014

Keywords:
Intergroup relations
Group-based status
Taboos
Group-based insults
Prejudice

ABSTRACT

Two studies investigate the effects of target group status on perceptions of the offensiveness of group-based slurs. Using real-world groups as targets, Study 1 showed that the perception that a group is of lower status in society is associated with the perceived offensiveness of insults targeting that group. Experimental methods in Study 2 showed that people perceive slurs against a low status group as especially offensive, a pattern that was mediated by the expectation that low-status targets would be emotionally reactive to the insult. The results suggest that cultural taboos emerge concerning insults against low-status groups that may be due in part to how those target groups are expected to respond emotionally to those insults.

Introduction

On July 1, 2013, CNN aired a brief panel discussion with the headline "N-word vs 'Cracker': Which is worse?" in which the panel debated the offensiveness of one word used derisively toward Blacks compared to the other used against Whites. The panel concluded the obvious, that the "n-word" was more offensive than "cracker," but they struggled to explain why. Throughout the past decades, high profile cases have emerged of the use of ethnic slurs against Blacks by celebrities who have been censured or even lost their jobs or sponsorship for the act, including sportscaster Howard Cosell, radio personality Don Imus, comedian Michael Richards, and most recently celebrity chef Paula Deen (see Kurylo, 2013, for other vivid examples). In contrast, however, there are relatively few cases of public controversy over ethnic slurs lobbed against Whites. These examples point to an "offensiveness gap" when determining the offensiveness of slurs across different groups. What makes one group-based slur offensive but another innocuous?

The present paper is a response to a call for more research on taboo language, including understanding the forces that encourage and discourage their use (Jay, 2009). We consider the group-based status of the target as a basis for the offensiveness gap. This factor can explain the Black–White offensiveness gap described in the CNN example above, but can also be generalized across all target groups that vary in status.

Considering qualities of the target group: group-based status

We focus our attention on aspects of the target group, a component that to this point has received little attention. Qualities of target groups that have been studied include whether the target of an insult belongs to an ascribed group (e.g., ethnicity) versus an acquired group (e.g., obesity), with ethnicity-based slurs perceived as more socially harmful and deserving of greater punishment than obesity slurs (Boeckmann & Liew,

☆ We thank Pat Shrout, John Jost, Erin Godfrey, Jojanneke van der Toorn, Irina Feygina, Margarita Krochik, Anesu Mandisodza, Avital Mentovich, Christine Reyna, Jen Zimmerman, Jim Davis, other members of the LSSR lab at DePaul University, and several anonymous reviewers for comments on earlier versions of this manuscript. Special thanks to Geoff Wetherell for assistance with data collection for Study 2b.

* Corresponding author at: Department of Psychology, New York University – Abu Dhabi, PO Box 129188, Abu Dhabi, United Arab Emirates.
E-mail address: pj.henry@nyu.edu (P.J. Henry).

http://dx.doi.org/10.1016/j.jesp.2014.03.012
0022-1031/

2002). Other target factors include characteristics of specifically ethnic groups, such that slurs directed against ethnic groups that are smaller in number and less familiar are perceived as less complex and more negative compared to slurs directed against ethnic groups that are larger and more familiar (Mullen, Rozell, & Johnson, 2000, 2001).

We extend this research to consider target group status, a factor relevant to all social groups. Although this factor has not been considered in the literature on group-based slurs and the offensiveness gap, there are clues in the existing literature as to its importance. Recent research has shown that dehumanizing words are more offensive when used against women, a lower status group compared to men (Haslam, Loughnan, & Sun, 2011). Furthermore, the finding that ethnic slurs are perceived as especially negative when directed at smaller groups that are less familiar (Mullen et al., 2001) may be a function of the lower status of the smaller groups, given that foreigners and ethnic minorities are typically lower in status in a country.[1] Thus we propose that one possible contextual influence on the offensiveness of a group-based slur is the group-based status of the target of the slur.

The causal direction could be the opposite, however, such that offensive group-based slurs could lead people to perceive target groups as lower in status. For example, a series of studies showed that when harsh ethnic slurs were directed against a hypothetical Black trial lawyer (Kirkland, Greenberg, & Pyszczynski, 1987) or an actual Black interaction partner (Greenberg & Pyszczynski, 1985), participants were more likely to negatively evaluate the target (although the effect may be restricted to participants already holding anti-Black attitudes; Simon & Greenberg, 1996). Other studies have shown that priming derogatory slurs against gays leads to greater activation of negative associations with gays (Carnaghi & Maass, 2007), and the use of slurs against fictional groups leads to preferences of exclusion and greater social distance from those groups (Leader, Mullen, & Rice, 2009). These kinds of reactions may lead to a perception of lower social status of the groups being disparaged.

Why group-based status?

We predict that group-based slurs directed against members of a low-status group are especially likely to be perceived as offensive, and there are plausible reasons concerning participants' expectations of the reactions of the target that lead us to expect this pattern of results. First, participants may expect low-status targets to be emotionally reactive to an insult. Allport (1954) speculated about the "vigilance and hypersensitiveness" (p. 145) among members of oppressed social groups, and a growing literature has documented that members of low-status groups may be especially emotionally reactive to insults (Henry, 2009) and signs of disrespect (Henry, 2011) because they are more vigilant to threats and rejection (Henry, 2009; Kraus, Horberg, Goetz, & Keltner, 2011; Mendoza-Denton, Downey, Purdie, Davis, & Pietrzak, 2002) due to a lifetime of threatening experiences associated with prejudice and discrimination (Brandt & Henry, 2012; Henry, 2009). Vignette studies where participants read about a case of sexual harassment (Hunter & McClelland, 1991) and racial harassment (McClelland & Hunter, 1992) show that the inclusion of an emotional reaction from the target increases the judgment of the seriousness of the harassment. If others in society even anticipate the possibility of a greater emotional reaction to an insult from those of lower status groups, and see that such emotional reactions signal that an insult is offensive, then taboos could form surrounding the use of those insulting words.

A second possible mechanism concerning expectations of target reactions is the concern that the target of the insult will engage in activist behaviors in response to an insult. Members of low-status groups under certain circumstances may take individual or collective action against prejudice (e.g., Gurr, 1970; Runciman, 1966), and the broader society may recognize the likelihood that members of low-status groups will protest or even press charges in response to group-based insults. Establishing taboos against expressing group-based slurs may help to avoid such protest and activism.

We explore the influence of both mechanisms in Study 2.

The importance of context

No given word in any language is inherently offensive. A word becomes an offensive slur because of the context in which it is used. For example, group-based slurs are seen as more offensive when directed against members of an outgroup compared to an ingroup (Asim, 2007; Haslam et al., 2011), with some offensive words even seen as endearing and a source of connection to others when spoken by one ingroup member to another ingroup member (Croom, 2011; Galinsky et al., 2013), as in the case for slurs relevant to gays (Bronski, 2011) and Blacks (Kennedy, 2002). Other work shows that the offensiveness of a group-based slur changes depending on the historical use of a word within a given context (e.g., in the evolution of the word "faggot" from meaning a bundle of twigs in Shakespeare's day to its modern usage as a slur towards gay men; Cresswell, 2009). Finally, cultural norms establish the injustice of targeting any group as the brunt of derisive slurs, such that within contexts that celebrate diversity and equality, taboos may form concerning group-based slurs more generally (e.g., Jeffries, Hornsey, Sutton, Douglas, & Bain, 2012). To focus on how the specific features of status of the target group plays a role in the offensiveness gap, we control for these contextual features in the two studies we present.

The present studies

This paper is one of a few empirical studies to consider how qualities of the target group itself determine the offensiveness gap, and the first to test the role of perceived group status of the targets of offensive words as a causal agent. In Study 1, we tested whether target group status matters at all in determining the perceived offensiveness of participant-generated slurs. In Study 2 we experimentally manipulated the group-based status of a target of a novel group-based slur to determine its causal influence on the perceived offensiveness of the slur, and tested two possible mediating mechanisms, expected emotional and behavioral reactions of the target.

In these studies our methods controlled for a number of alternative contextual variables that could influence the perceived offensiveness of a word. First, whereas in Study 1 we consider words that are used in the everyday lexicon and could have historical and learned influences in the meanings given to them, we control for this influence in Study 2 by using a novel offensive slur directed toward a novel group that would have no historical context for the participants to interpret. Second, we were specifically interested in what makes a group-based slur offensive when it has the *intention* of being offensive (as opposed to serving as a source of connection between ingroup members). Therefore, the instructions and manipulations in our studies were explicit about the offensive use of the words. Third, given that the emotional response to an insult can influence the judgment of the seriousness of an insult (Hunter & McClelland, 1991; McClelland & Hunter, 1992), we provided the participant with no such clues to the target's response to the insult, so that we could assess instead the participant's anticipation of both emotional and behavioral reactions.

Study 1

The first study was designed to test the prediction that slurs targeting groups that participants perceived as having lower social status would be seen as more offensive than slurs targeting groups participants perceived as having higher social status.

[1] There are exceptions, for example in some African countries.

Method

Participants

Participants were recruited through the undergraduate psychology research participation system at a private Midwestern university to participate in a study for partial course credit. The sample was composed of 214 college students, ranging in age from 18 to 30 years (median = 19), including 49 men and 165 women. Ethnic/racial identification was mostly White (145) followed by Latino/a (26), Asian (21), and Black (9), with 13 indicating "other" or providing no indication.

Materials and procedure

Participants completed all materials on a central online website. Participants were asked a series of questions concerning 15 target groups, including groups representing different races (*Black/African-American, Latino(a), White/European-American, Asian/Asian-American, Arab/Arab-American*), genders (*male, female*), sexual orientations (*gay people, straight people*), mental status (*highly intelligent people, mentally ill people, mentally disabled people*), religious affiliation (*Jewish people*), age (*elderly people*), and physical status (*obese people*).

To hold constant the intended use of the slur as offensive, for each target group participants generated the most offensive word they could think of to name that group.[2] Following the word generation, participants rated the offensiveness of the word, "regardless of who says it," on an 11-point scale, anchored at 1 = "not at all offensive" and 11 = "extremely offensive." Next, participants completed ratings of the relative status of each target group in American society using 9-point scales anchored at 1 = "lowest status" and 9 = "highest status."[3]

Results and discussion

Fig. 1 plots the means for each target group on their perceived social status (*x*-axis) and the perceived offensiveness of the self-generated slur (*y*-axis). A generalized estimating equation (GEE) analysis was conducted to determine if participants' perceptions of the status of the target groups were related to the perception of the offensiveness of the most offensive word they could generate for each of the 15 target groups, while managing the dependence of the responses within each participant. This technique is conceptually similar to computing a correlation for each participant and testing if the correlation across participants is on average statistically different from zero.

The results showed a significant negative relationship between the perceived status of the target group and the perceived offensiveness of the self-generated word directed at the group, $b = -.39$, $SE\ (b) = .03$, Wald $\chi^2 = 206.4$, $p < .001$ (see Fig. 1). To get a better sense of the size and consistency of the association between perceived offensiveness and target group status we computed a correlation for each participant in the sample. The average within-participant correlation across the sample was $M\ r = -.30$ ($SD\ r = .29$), which was significantly different from zero, $t(209) = 15.26$, $p < .001$, Cohen's $d = -1.03$. Further exploration of these results revealed that 85.7% of the sample recorded correlations less than zero, indicating that the negative association between offensiveness and target group status was consistent across participants. In short, the results show that the higher a participant perceived the status of the target group, the less offensive the participant rated the self-generated group-based slurs.

Exploratory results: participant gender and ethnicity

To examine whether the status of the participant was related to perceptions of offensiveness, we explored two dimensions of status that are relevant to our student population, gender and ethnic group membership, to determine if these may have influenced the perception of the offensiveness of slurs. For example, Cowan and Hodge (1996) found gender and gender by ethnicity interaction effects on judgments of the offensiveness of hate speech messages targeted against minority groups, specifically showing that men, and especially white men, perceived hate speech to be less offensive than other groups.

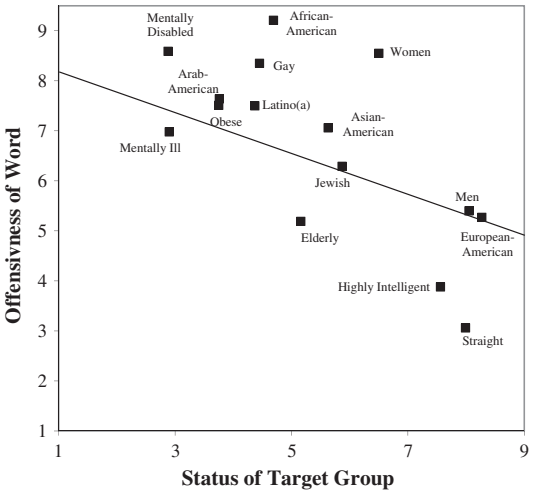

Fig. 1. Study 1: Predicting the perceived offensiveness of words based on the perceived status of the target group. Note: The data points represent the location for the average of target group status and offensiveness of the word in the sample. The regression shown here is not derived from the data points shown in the figure (which would give a steeper slope), but is instead derived from the analyses using the general estimating equations (GEE) function given in the text.

Again using a GEE analysis, we tested two separate two-way interactions, one between perceived target group status and participant gender, and the other between perceived target group status and participant ethnicity.[4] The two-way interaction between perceived target group status (continuous, mean-centered within participants) × participant gender (0.5 = male, −0.5 = female) was statistically significant, $b = -.25$, $SE\ (b) = .07$, Wald $\chi^2 = 11.9$, $p < .001$. Probing this interaction, simple slopes revealed that the effect of perceived target group status on the perceived offensiveness of a word was more strongly negative for women ($b = -.45$, $SE\ (b) = .03$, Wald $\chi^2 = 266.2$, $p < .001$), than for men ($b = -.20$, $SE\ (b) = .07$, Wald $\chi^2 = 229.7$, $p = .002$).

The two-way interaction between the perceived target group status (continuous, mean-centered within participants) × participant ethnicity (−1 = White, 1 = non-White ethnic minority) was not statistically significant, $b = .00$, $SE\ (b) = .03$, Wald $\chi^2 = .00$, $p = .99$. However, there was a main effect of participant ethnicity on overall perceptions of the offensiveness of group-based slurs, $b = .45$, $SE\ (b) = .16$, Wald $\chi^2 = .7.8$, $p = .005$, such that ethnic minorities rated slurs overall to be more offensive than their White counterparts irrespective of the target group status.

Exploratory results: target group size

We did not measure participants' perception of the size of the target group in our study, but this feature of groups is important to consider for determining the offensiveness of slurs, at least those directed toward ethnic groups (Mullen et al., 2001). As an exploratory substitute we created a measure of target group size derived from the actual size of the group according to demographic statistics available from multiple sources, including the U.S. census and from nationally representative

[2] The list of (often very offensive) slurs, plus their frequency of generation by the participants, is available in the supplemental materials.

[3] The selection of an 11-point scale for perceived offensiveness but a 9-point scale for perceived status was accidental and arbitrary, not based on past methodology or theory.

[4] The sample size prevented running a meaningful three-way interaction between perceived target group status, participant gender, and participant ethnicity, as one of the four cells had a very small *n* (male, nonWhite, $n = 14$) relative to the other cells.

data sets.[5] We re-estimated the model where we regressed offensiveness on target group status, this time including the measure of target group size.

First, we found that the higher the target group size, the higher its perceived status ($b = -.044$, $SE\ (b) = .0015$, Wald $\chi^2 = 875.7$, $p < .001$). Group size also revealed an independent significant main effect predicting the offensiveness of the slurs ($b = -.027$, $SE\ (b) = .0021$, Wald $\chi^2 = 164.4$, $p < .001$) confirming prior research by showing that as the size of the group decreases, the perception of the offensiveness of slurs directed toward those groups increases. (These results reduced somewhat the negative relationship between perceived target status and perceived slur offensiveness, but it remained statistically significant, $b = -.23$, $SE\ (b) = .03$, Wald $\chi^2 = 69.5$, $p < .001$.) Although this was not the main focus of this study, we were able to extend research by Mullen et al. (2001) by showing that the negative relationship between group size and the perceived offensiveness of words functions beyond just ethnic groups (the effect remains when removing the ethnicity-based Black, White, Latino, Asian, Jewish, and Arab targets).

Study 2

Study 1 had strengths insofar as it tested participants' perceptions using real-world groups and the generation of real insults directed against those groups. It showed the expected relationship between the perceived status of a group and the perceived offensiveness of slurs directed against that group.

However, this analysis was limited in three ways. First, we could not rule out the influence of other important factors influencing the perception of the offensiveness of the words, such as social learning, the historical use of the particular words that the participants generated, and the historical treatment of the groups considered. Second, we were not able to test for causal relationships, showing that the status of a group can cause the perceived offensiveness of a word directed at that group. The direction theoretically could be reversed such that words that are considered offensive could lead a person to perceive the target group as lower in status. A third limitation of Study 1 is that we did not address the question of *why* the perceived status of a group would influence the offensiveness of slurs directed against them. Here we explore the two possible mechanisms discussed in the introduction, those of the expected emotional reactions and the expected behavioral reactions of the target of the slur.

To address these points we exposed participants to a target who was insulted based on the target's experimentally manipulated group-based status. We predicted that a group-based insult would be perceived as particularly offensive when directed toward a target manipulated to be part of a low-status versus high-status group. We were interested in testing for the offensiveness of a group-based slur without the possible confounding effects of social learning or the cultural or historical uses of the word. In this effort, we invented for this study the group the target belonged to and the group-based insult directed at the target. It was also made clear that the use of the slur had hostile intent, or was meant to be offensive, to clear any ambiguity concerning the flexibility of a word's meaning depending on context (see also Leader et al., 2009, for the advantages of using novel groups and novel insults in research). Furthermore, we assessed participants' expected reactions from the target to determine if such expectations might be driving the predicted offensiveness gap in insults directed at high- versus low-status groups. Study 2 encompasses two separate studies, 2a and 2b. Study 2b replicates Study 2a exactly, but adds a manipulation of the status of the perpetrator as well as that of the status of the target of the insult to remove any participant assumptions that a perpetrator of a slur would have higher status than a target.

[5] We thank an anonymous reviewer for this suggestion. The sources used to gather demographic percentages are available in the supplementary materials.

Method

Participants

Participants were recruited from the introductory psychology research participation system at a private Midwestern university. All participants received partial course credit for their participation. Study 2a had 265 participants, ranging in age from 18 to 33 years (median = 19), including 202 women and 61 men, with 2 participants not reporting gender.[6] Ethnic/racial identification was mostly White (174) followed by Latino/a (41), Black (18), and Asian (15), with 17 indicating "other" or providing no indication. Study 2b had 365 participants, ranging in age from 18 to 56 years (median = 19), including 261 women and 102 men, with 2 participants not reporting gender. Ethnic/racial identification was mostly White (245) followed by Latino/a (57), Asian (25), and Black (18), with 20 indicating "other" or providing no indication.

Materials and procedure

Participants completed all materials on a central online experiment website.

Target group-based status manipulation. Participants were randomly assigned to evaluate one of two brief scenarios, shown below, which describe an incident that takes place in a job setting where the participant witnesses a target being insulted. The target was a member of a group, "creative developers," manipulated to be high- or low-status in the company. The target was male so as to focus on their manipulated group-based status as the reason for the insult, rather than other factors such as perceived sexism of the perpetrator. The insult was clearly directed at the individual based on his group membership. The high-status target manipulation is presented here, with the low-status version indicated in brackets. The places where the changes occur across conditions are indicated in boldface.

> You are going through an information session to learn to use the new payroll system at your job. During the session one of the creative developers asks a number of questions. Creative Developers are the people who design the products that your company manufactures, and the position is one of the **most [least]** important and influential in the company. They make **very good [little]** money, have **very good [no]** benefits, and frequently **get three-day weekends [have to work on weekends]**. After his fifth question you hear one of the people from payroll whisper to his neighbor, just loud enough so that the Creative Developer asking the questions will hear, "Can you believe this guy! How does he not get this stuff? Though really, what else can you expect from a Crappo!" You know that "Crappo" is a combination of the words "creative" and "poser" and is a derogatory, negative term that some people in the company use to refer to creative developers.

The key experimental manipulation was whether the target was a member of a high-status group (creative developers who have influence in the workplace and good pay) or a low-status group (creative developers who have minimal influence and low pay). Care was taken to keep equal across both conditions the name of the target's department (group) in the company, the nature of the job responsibilities, the word used as the insult, and its description as derogatory.

In Study 2b we added immediately after the description of the status of the target the extra manipulation of the status of the perpetrator of the insult, creating four between-subjects conditions based on the 2 (target status: high v. low) × 2 (perpetrator status: high v. low) design, to which participants were randomly assigned. When the target was

[6] Two participants who indicated an age of 17 were removed to maintain a homogenous adult sample, as well as two who did not indicate their age. Their removal did not change the results.

Table 1
Studies 2a and 2b: mean differences across experimental conditions.

	Study 2A				Study 2B			
	Experimental condition				Experimental condition			
	Low-status target	High-status target			Low-status target	High-status target		
	M (*SD*)	*M* (*SD*)	*F*	η^2	*M* (*SD*)	*M* (*SD*)	*F*	η^2
Status of the target (manipulation check)	2.3 (1.8)	5.5 (2.4)	160.5***	.38	4.3 (2.2)	6.0 (2.0)	51.4***	.13
Offensiveness of the word	7.4 (1.7)	6.5 (2.2)	14.4***	.05	6.4 (1.8)	5.9 (2.0)	5.7*	.02
Expected negative emotions	7.3 (1.4)	6.4 (1.5)	28.5***	.10	6.3 (1.7)	5.9 (1.5)	5.0*	.02
Expected activist behaviors	4.0 (1.5)	4.2 (1.6)	0.4^{ns}	.00	4.3 (1.4)	4.5 (1.4)	1.1^{ns}	.00

Note: ^{ns}Not statistically significant ($p > .30$). Study 2b accounts for the non-significant manipulation (all *p*s > .16) of the status of the perpetrator of the insult, and the non-significant interaction between the status of the target and the status of the perpetrator (all *p*s > .17).
*** $p < .001$.
* $p < .05$.

high status, the high-status perpetrator condition was manipulated the following way, with the low-status perpetrator version indicated in brackets. The places where the changes occur across conditions are indicated in boldface.

> People in payroll are responsible for the distribution of paychecks to the employees and are **also [in contrast]** of **high [low]** status **in the company, receiving a similar kind of treatment, salary, and benefits. [in the company. They make very little money, have no benefits, and frequently have to work on weekends.]**

When the target was low status, the low-status perpetrator condition was manipulated the following way, with the high-status perpetrator version indicated in brackets. The places where the changes occur across conditions are indicated in boldface.

> People in payroll are responsible for the distribution of paychecks to the employees and are **also [in contrast]** of **low [high]** status **in the company, receiving a similar kind of treatment, salary, and benefits. [in the company. They make very good money, have very good benefits, and frequently get three day weekends.]**

Measures. A manipulation check was administered to test if the participants perceived the differential status of the groups across the conditions, asking simply "What is the status of Creative Developers at this company?" For Study 2b we also added, "What is the status of people in Payroll (the person who delivered the insult) at this company?" Both items were anchored at 1 = lowest status and 9 = highest status.

Similarly to Study 1, the dependent variable was the *offensiveness* of the insulting word, with an item that read "How offensive is the use of the term 'Crappo' in this situation," anchored at 1 = "not at all" and 9 = "very much so."[7]

Three items measured the *expected target's feelings* towards the insult. The first asked "How insulted do you think the Creative Developer would feel because of that comment?" anchored at 1 = "not at all insulted" and 9 = "very insulted." The second asked "How much do you think that comment would make the Creative Developer feel bad about himself?" anchored at 1 = "not at all bad" and 9 = "very bad." The third asked "How angry do you think the Creative Developer would be about that comment?" anchored at 1 = "not at all angry" and 9 = "very angry." These items were averaged into a scale, Cronbach's $\alpha = .86$ for Study 2a, .85 for Study 2b.

Four items measured the *expected target's behaviors* in reaction to the insult, or the perceived probability that the target would take action in response to the insult. All four items were measured on a 9-point scale anchored at 1 = "not at all likely" and 9 = "very likely." The questions asked, (1) "How likely do you think the Creative Developer would be to report the comment to a supervisor?" (2) "How likely do you think the Creative Developer would be to confront the individual in a private setting and tell the person that they don't appreciate being called names?" (3) "How likely do you think the Creative Developer would be to write a letter to the company newsletter about how using certain names for groups is inappropriate?" and (4) "How likely do you think the Creative Developer would be to say nothing and ignore the comment?" (reverse coded). The four items were averaged into a scale, Cronbach's $\alpha = .77$ for Study 2a, .71 for Study 2b.

Results and discussion

Manipulation check

The results of the manipulation check for the status of the *target*, given in the first row of Table 1, showed that participants saw the target in the low-status condition as lower in status compared to the target in the high-status condition. For Study 2b, the result for the manipulation of the status of the *perpetrator* is that participants saw the perpetrator in the high-status condition as higher in status ($M = 5.8$, $SD = 2.0$) compared to the perpetrator in the low-status condition ($M = 4.4$, $SD = 2.2$; $F(1, 333) = 34.7$, $p < .001$, $\eta^2 = .095$) confirming the effectiveness of the manipulation. There were no significant main effects or interactions for either manipulation check (all *p*s > .15).

The effect of target status on the perceived offensiveness of the word and expected target reactions

We next considered the influence of the target's status on the perceived offensiveness of the slur. The results, shown in the second row of Table 1, confirmed the prediction that those who read about the coworker who was a member of a low-status group perceived the word as being more offensive than those who read about the coworker who was a member of a high-status group, despite the fact that the word itself and its use in the context did not change across conditions. Study 2b replicated this main effect of target status. The status of the perpetrator did not significantly predict the offensiveness of the slur, either as a main effect ($F(1, 332) = 0.00$, $p = .97$) or in an interaction with the status of the target ($F(1, 332) = 0.39$, $p = .53$).

We tested whether the manipulated status of the target would influence the expected emotional reactions and activist behaviors of the target to determine if either (or both) play a role as an underlying mechanism. The third row of Table 1 confirms that the low-status target was expected to have a more negative emotional reaction to the insult (be insulted, feel bad, feel angry) compared to the high-status person. Study 2b replicated this main effect of target status. Again, perpetrator status did not factor as a predictor of emotional reactions, either as a main effect ($F(1, 333) = 0.07$, $p = .80$) or as it interacted with target status ($F(1, 333) = 0.25$, $p = .62$).

The fourth row of Table 1 shows that there were no differences in terms of expected activist behaviors: Both low-status and high-status targets were perceived to be equally likely to take action against the

[7] To be clear, the item asks about the *participants'* perceived offensiveness of a word, not the participants' belief that a word is offensive to the target. Nevertheless, the latter is probably represented in the expected emotional and behavioral reactions to the slur.

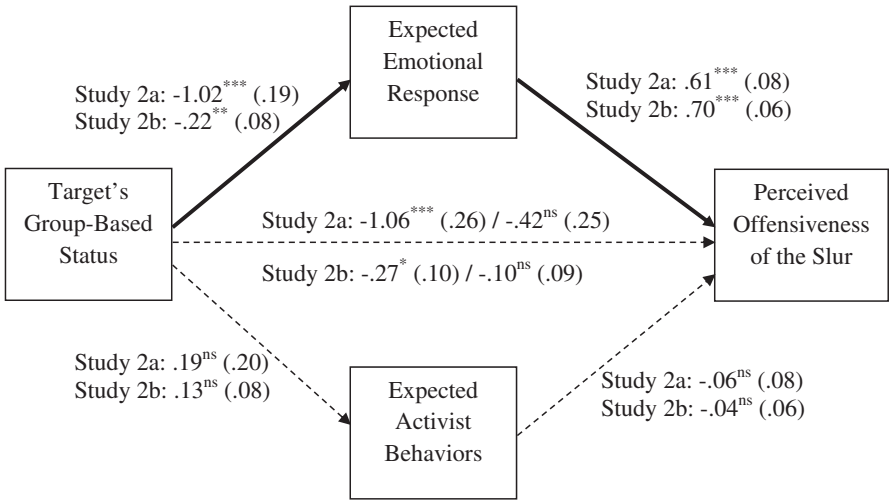

Fig. 2. Study 2: Mediations between target group-based status and perceived offensiveness of the slur. Note: Numbers indicate unstandardized coefficients (with standard errors in parentheses). Numbers to the left of the slash indicate the direct effect of the target's group-based status on the perceived offensiveness of the slur; numbers to the right of the slash indicate the relationship after including the mediators in the model. All models included participant gender and ethnicity as covariates.

slur. Furthermore, in Study 2b there were no effects of perpetrator status (F (1, 333) = 1.9, p = .17) or the interaction between perpetrator status and target status (F (1, 333) = 1.8, p = .18).

Testing the mediators between target status and offensiveness of the slur

We explored whether the effect of status on the perceived offensiveness of the word would be due in part to the expected negative emotional response of members of low-status groups or to the expected activist behaviors taken as a result of the insult, or both. Fig. 2 shows the results of the multiple mediation analysis using 5000 bootstrapped resamples with 95% bias and corrected confidence intervals (Preacher & Hayes, 2004). The indirect effect through expected negative emotions was significant: for Study 2a, the indirect effect = −.62, SE = .14, 95%CI (−.94, −.37); for Study 2b, the indirect effect = −.16, SE = .06, 95%CI (−.28, −.04). The indirect effect through behavioral expectations, however, was not significant: for Study 2a, the indirect effect = −.01, SE = .02, 95%CI (−.10, .01); for Study 2b, the indirect effect = −.01, SE = .01, 95%CI (−.04, .01). Furthermore, the indirect effect through behavioral expectations was significantly smaller than the indirect effect through expected emotions: for Study 2a, the indirect effect difference = .61, SE = .15, 95%CI (.35, .92); for Study 2b, the indirect effect difference = −.15, SE = .06, 95%CI (−.28, −.03). Status seems to determine the offensiveness of a slur based on the expected *emotional reaction* to the use of the word on the part of the target, rather than the expected *actions* the target would take in response to the use of the word.

Exploratory results: participant gender and ethnicity

As with Study 1 we considered whether or not characteristics of the participants influenced the judgments of the offensiveness of the slur, looking at gender and ethnicity of the participants. For Study 2a we considered two separate two-way interactions, one between target group status and participant gender, and the other between target group status and participant ethnicity. For Study 2b we looked at two separate three-way interactions, identical to the Study 2a two-way interactions but adding the extra factor of perpetrator status.[8]

The analyses do not change the main findings of the effect of the manipulation of the target status on perceptions of the offensiveness of the slur reported in Table 1. The only exploratory effect that replicated across the Study 2a and 2b results was the main effect of participant gender: Women across the board rated the slurs as more offensive than men, in Study 2a (women: M = 7.1, SD = 2.0, men: M = 6.5, SD = 1.9; F (1, 253) = 3.8, p = .052, η^2 = .015), and in Study 2b (women: M = 6.4, SD = 1.8, men: M = 5.5, SD = 2.0; F (1, 320) = 16.5, p < .001, η^2 = .049).

The interaction between participant gender and target status found in Study 1 did not replicate in Study 2a, but was marginally significant in Study 2b, however, in the opposite direction compared to the Study 1 results. In Study 2b there was a larger difference for men in judging the offensiveness of the slur directed at the low-status target (M = 6.0, SD = 1.8) compared to the high-status target (M = 5.0, SD = 2.1), relative to women who showed a smaller difference between the low-status target (M = 6.6, SD = 1.8) compared to the high-status target (M = 6.2, SD = 1.9); F (1, 320) = 3.6, p = .059, η^2 = .011.

No other effects predicting the perceived offensiveness of the slur were statistically significant (all ps > .12). The details of these analyses can be found in the online supplemental materials.

General discussion

These studies test the social psychological mechanisms underlying the perceived offensiveness of group-based slurs, with a focus on perceptions of the characteristics of the target group. The results from Study 1, using participant-generated offensive insults toward real-world groups, showed a relationship between the perceived status of the target group and the perceived offensiveness of the slur directed against that group. Study 2 established that this relationship is causal, and showed that the relationship is mediated by the expected negative emotional reaction of the low-status target.

Study 2 also showed that the offensiveness of a slur against a member of a low-status group is probably not driven by concerns that a target will take social action against such prejudices. However, one should not conclude that our participants expected no behavioral action from our target. While the average score fell somewhat below the midpoint (between 4.0 and 4.4, depending on the condition, on a 9-point scale), the participant response still indicates at least some activist behavior is expected. Although the action of members of low-status groups to repair inequities is certainly a real phenomenon and is perceived by our participants to be at least somewhat possible, it does not appear to be the source determining whether or not a group-based slur is

[8] As with Study 1, sample size issues prevented us from running interactions of participant gender X participant ethnicity, as several cells in both studies were too small (ns < 10) for a meaningful interpretation of results.

perceived to be offensive in our study. We believe our participants do not establish the offensiveness of a term based on concerns about social activism, which can take a long time to create meaningful change, but instead rely on more immediate expected cues such as emotional reactions to insults.

Interestingly, Study 2b showed there was no effect of the status of the *perpetrator* of the insult on the perception of the insult as offensive. Although this was the finding in the setting of our controlled laboratory experiment, one can imagine that the group-based status of a perpetrator probably should matter in real-world settings. For example, the celebrity scandals of group-based slurs raised in the introduction all involve high-status Whites slandering lower-status Blacks. However, even in these real-world situations it still may not be the perpetrator status per se that is driving the offensiveness of the slurs, but other factors that we controlled for in our experiment, such as the history of relations between groups in society. Because groups that are historically oppressive (e.g., Whites in the United States) tend to be also higher in status, these factors almost always become conflated in real-world examples. Therefore, it may be these historical factors rather than perpetrator status per se that is behind the perception that slurs from dominant groups are more offensive. However, this is an empirical question that requires further investigation of the phenomenon outside the controlled setting of the lab.

Exploratory analyses

Interestingly, Study 2 with its novel slur directed against a novel group revealed some different patterns of results compared to Study 1, with its evaluations of actual slurs directed against real-world groups. Most notably, in predicting the offensiveness of a slur only one of the two studies in Study 2 replicated the Study 1 interaction between target status and participant gender, but it was in the opposite direction (with men in Study 2 showing a greater sensitivity to target status compared to women). Nor was there a main effect of participant ethnicity in Study 2 as there was in Study 1. These findings may point to the importance of knowing the historical treatment of real groups, and the reduced sensitivity that men may have toward lower-status concerns and the increased sensitivity ethnic minorities may have to all group-based slander in real-world contexts. These findings do not diminish the importance of the Study 2 findings that used novel groups, but instead illustrate further the added complexity that historical forces bring to individual judgments. When those historical factors are stripped away as in Study 2, we can see more clearly how target status operates in influencing perceptions of slurs, and in particular that male participants, too, are sensitive to target status cues.

Cultural contexts

To understand the offensiveness of any group-based slur requires an understanding of the broader cultural context in which it is uttered. In particular, we have focused on a context that champions equality and encourages sensitivity toward members of other groups. The location of the present studies, a university context, helps ensure this kind of egalitarian value environment (cf. Henry, 2008). With that said, the results here are not intended to generalize to societies that have more clear norms of intolerance of diversity and endorsement of group-based hierarchies (cf. Jeffries et al., 2012), such as the pre-civil rights era in the United States or Nazi Germany. Imagine, for example, in the years leading up to WWII in Germany, insults directed against Christian Germans versus insults directed against Jews. In such situations, it is likely that the lower status Jews found such slurs towards their group extremely offensive, but in this context the intolerance of diversity and acceptance of the hierarchical system would preclude a recognition of this offensiveness by society at large. It is also possible that perpetrators would know that such slurs would be offensive to low status groups, but they would be used nevertheless to help maintain the hierarchy. The role of these historical and cultural influences in places of intolerance describes a wholly different context and probably a different set of mechanisms, one that would be better informed and fruitfully explored by other intergroup theories such as social dominance theory (Sidanius & Pratto, 1999).

Instead, we deliberately sought a context of intergroup sensitivity because we believe that this is the trajectory that the United States has taken, as can be witnessed by the shift in prejudicial attitudes from blatant to more subtle forms for many groups (e.g., Morrison & Morrison, 2002; Sears & Henry, 2005; Swim, Aikin, & Hall, 1995), a trend that is also true for other modern, Western societies (e.g., Pedersen & Walker, 1997; Pettigrew & Meertens, 1995). This kind of study in this kind of context will be useful for regulators such as the Federal Communications Commission (FCC) for understanding the processes underlying the public's judgment of the offensiveness of speech. The scandals surrounding group-based slurs that the FCC has to address have erupted, not coincidentally, within this egalitarian value context.

Limitations and future directions

Part of the tradeoff of studying real-world phenomena such as offensive slurs using laboratory and survey methods is that the scientific requirements of precision and control remove the participants' experience from how slurs are used and perceived in real-world situations. For example, the participants evaluated slurs in their written format and, in Study 2, a hypothetical scenario. These slurs may not have the same impact compared to their more oral and live counterparts. For example, there are discrepancies between people's anticipated emotional reactions to an ethnic slur versus their actual emotional reactions when they witness a live slur (Kawakami, Dunn, Karmali, & Dovidio, 2009). People may also rely on other sources for determining the offensiveness of a slur (e.g., direct and actual instead of imagined reactions of the target, reactions of other observers, real-time nonverbal emotional cues, inferences of intent). Furthermore, the target's group-based status was told to participants in Study 2, whereas judgments of status in the real world happen through other cues. Of course, field studies of offensive slurs would pose their own practical and ethical challenges, but future research could incorporate more ecologically valid and realistic components (e.g., a video recorded insult rather than a written description of an insult) that would help tie these psychological components of analysis to their real-world counterparts.

We also proposed that the expected emotional reactions of a target influence the perceived offensiveness of a slur, and while we indeed found that participants expected greater emotional reactivity from the lower-status targets of an insult, we do not know from where our participants get this information for these expectations. In the introduction we based this prediction on evidence that members of lower-status groups tend to be more emotionally reactive to insults and more sensitive to social threats. The inference we make, then, is that third-party observers may be relying on accuracy or personal experience in predicting greater emotional reactivity from lower-status group members. However, we do not test this idea. We found that the offensiveness of a word is caused by the perceived lower status of a target, and found that the expectation of the emotional reactivity of that lower status target is probably driving (at least in part) the offensiveness of the word, but the question remains open as to why there is a greater expectation of emotional reactivity from members of a lower-status group, including a novel group.

Additionally, we did not comprehensively test all plausible mediating mechanisms other than expected emotional and behavioral reactions of the targets. For example, the perception that low-status groups are victims of prejudice and discrimination could trigger feelings of sympathy toward those groups, which could in turn render judgments of slurs against them as offensive. Other such plausible alternative mechanisms could exist. However, what we were able to demonstrate is that the expectation of a negative emotional reaction from a target is sufficient

to explain the relationship between perceived status of a group and perceived offensiveness of a slur against them. The mediation in Study 2 rendered the status-offensiveness effect non-significant, even with a relatively large sample. These analyses are only the first step in identifying the mechanisms driving the relationship between group status and offensiveness, and future work will help document the multiple likely mediating processes.

Future research could also extend the analyses we conducted in broader ways. For example, would perceptions of acts of discrimination work in a similar way to perceptions of group-based slurs? If group-based slurs can be perceived as a kind of discriminatory act (e.g., Rosette, Carton, Bowes-Sperry, & Hewlin, 2013), then broader forms of discrimination may fall under the same principles. For example, academics virtually never use the harsher and loaded term "racism" to describe discrimination against Whites (Henry & Pratto, 2010), and while Whites in the general public may make claims of "reverse racism" such claims are often associated with White beliefs that discrimination against Blacks no longer exists (Fraser & Kick, 2000) and are therefore equal in status to Whites.

Further exploration could consider the relative importance of target status with other aspects of perceptions of groups that could influence the offensiveness of words. For example, other influencing factors could include the ingroup–outgroup membership of the target, negative feelings toward the target group, other more complex emotional reactions toward target groups (e.g., Fiske, Cuddy, & Glick, 2002), the perceived justification of the group as a recipient of a slur, among innumerable other potential influencing factors.

Conclusion

Language is a form of communication whose power depends on cultural contexts. Broadcasters, public figures, teachers, comedians, and anyone making public statements about groups in society have known, for the most part, to be especially careful about how to use language. The present set of studies contributes to our understanding of why, and proposes one causal mechanism, that of the group-based status of the targets of a slur. Future research can and should deepen this understanding through consideration of the multiplicity of influences that result in the complexity of reactions to real-world examples of offensive slurs.

Appendix A. Supplementary data

Supplementary data to this article can be found online at http://dx.doi.org/10.1016/j.jesp.2014.03.012.

References

Allport, G. W. (1954/1979). *The nature of prejudice (25th anniversary edition)*. Reading. MA: Addison-Wesley.

Asim, J. (2007). *The N word: Who can say it, who shouldn't, and why*. New York: Houghton Mifflin.

Boeckmann, R. J., & Liew, J. (2002). Hate speech: Asian-American students' justice judgments and psychological responses. *Journal of Social Issues*, *58*, 363–381.

Brandt, M. J., & Henry, P. J. (2012). Psychological defensiveness as a mechanism explaining the relationship between low socioeconomic status and religiosity. *International Journal for the Psychology of Religion*, *22*, 321–332.

Bronski, M. (2011). *A queer history of the United States*. Boston: Beacon Press.

Carnaghi, A., & Maass, A. (2007). In-group and out-group perspectives in the use of derogatory group labels: Gay versus fag. *Journal of Language and Social Psychology*, *26*, 142–156.

Cowan, G., & Hodge, C. (1996). Judgments of hate speech: The effects of target group, publicness, and behavioral responses of the target. *Journal of Applied Social Psychology*, *26*, 355–374.

Cresswell, J. (Ed.). (2009). *Oxford dictionary of word origins* (2nd ed.). New York: Oxford University Press.

Croom, A.M. (2011). Slurs. *Language Sciences*, *33*, 343–358.

Fiske, S. T., Cuddy, A. J. C., & Glick, P. (2002). Emotions up and down: Intergroup emotions result from perceived status and competition. In D.M. Mackie, & E. R. Smith (Eds.), *From prejudice to intergroup emotions: Differentiated reactions to social groups* (pp. 247–264). Philadelphia, PA: Psychology Press.

Fraser, J., & Kick, E. (2000). The interpretive repertoires of whites on race-targeted policies: Claims making of reverse discrimination. *Sociological Perspectives*, *43*, 13–28.

Galinsky, A.D., Wang, C. S., Whitson, J. A., Anicich, E. M., Hugenberg, K., & Bodenhausen, G. V. (2013). The reappropriation of stigmatizing labels: The reciprocal relationship between power and self-labeling. *Psychological Science*, *24*, 2020–2029.

Greenberg, J., & Pyszczynski, T. (1985). The effect of an overheard ethnic slur on evaluations of the target: How to spread a social disease. *Journal of Experimental Social Psychology*, *21*, 61–72.

Gurr, T. R. (1970). *Why men rebel*. Princeton, NJ: Princeton University Press.

Haslam, N., Loughnan, S., & Sun, P. (2011). Beastly: What makes animal metaphors offensive? *Journal of Language and Social Psychology*, *30*, 311–325.

Henry, P. J. (2008). College sophomores in the laboratory redux: Influences of a narrow data base on social psychology's view of the nature of prejudice. *Psychological Inquiry*, *19*, 49–71.

Henry, P. J. (2009). Low-status compensation: A theory for understanding the role of status in cultures of honor. *Journal of Personality and Social Psychology*, *97*, 451–466.

Henry, P. J. (2011). The role of group-based status in job satisfaction: Workplace respect matters more for the stigmatized. *Social Justice Research*, *24*, 231–238.

Henry, P. J., & Pratto, F. (2010). Power and racism. In A. Guinote, & T. Vescio (Eds.), *The social psychology of power* (pp. 341–362). New York: Guilford.

Hunter, C., & McClelland, K. (1991). Honoring accounts for sexual harassment: A factorial survey analysis. *Sex Roles*, *24*, 725–752.

Jay, T. (2009). The utility and ubiquity of taboo words. *Perspectives on Psychological Science*, *4*, 153–161.

Jeffries, C. H., Hornsey, M. J., Sutton, R. M., Douglas, K. M., & Bain, P. G. (2012). The David and Goliath principle: Cultural, ideological, and attitudinal underpinnings of the normative protection of low-status groups from criticism. *Personality and Social Psychology Bulletin*, *38*, 1053–1065.

Kawakami, K., Dunn, E., Karmali, F., & Dovidio, J. F. (2009). Mispredicting affective and behavioral responses to racism. *Science*, *323*, 276–278.

Kennedy, R. (2002). *Nigger: The strange career of a troublesome word*. New York: Random House.

Kirkland, S. L., Greenberg, J., & Pyszczynski, T. (1987). Further evidence of the deleterious effects of overheard ethnic slurs: Derogation beyond the target. *Personality and Social Psychology Bulletin*, *13*, 216–227.

Kraus, M. W., Horberg, E. J., Goetz, J. L., & Keltner, D. (2011). Social class rank, threat vigilance, and hostile reactivity. *Personality and Social Psychology Bulletin*, *37*, 1376–1388.

Kurylo, A. (2013). *The communicated stereotype: From celebrity vilification to everyday talk*. Plymouth, UK: Lexington Books.

Leader, T., Mullen, B., & Rice, D. (2009). Complexity and valence in ethnophaulisms and exclusion of ethnic out-groups: What puts the "hate" into hate speech? *Journal of Personality and Social Psychology*, *96*, 170–182.

McClelland, K., & Hunter, C. (1992). The perceived seriousness of racial harassment. *Social Problems*, *39*, 92–107.

Mendoza-Denton, R., Downey, G., Purdie, V. J., Davis, A., & Pietrzak, J. (2002). Sensitivity to status-based rejection: Implications for African American students' college experience. *Journal of Personality and Social Psychology*, *83*, 896–918.

Morrison, M.A., & Morrison, T. G. (2002). Development and validation of a scale measuring modern prejudice toward gay men and lesbian women. *Journal of Homosexuality*, *43*, 15–37.

Mullen, B., Rozell, D., & Johnson, C. (2000). Ethnophaulisms for ethnic immigrant groups: Cognitive representation of 'the minority' and 'the foreigner. *Group Processes and Intergroup Relations*, *3*, 5–24.

Mullen, B., Rozell, D., & Johnson, C. (2001). Ethnophaulisms for ethnic immigrant groups: The contributions of group size and familiarity. *European Journal of Social Psychology*, *31*, 231–246.

Pedersen, A., & Walker, I. (1997). Prejudice against Australian aborigines: Old-fashioned and modern forms. *European Journal of Social Psychology*, *27*, 561–587.

Pettigrew, T. F., & Meertens, R. W. (1995). Subtle and blatant prejudice in Western Europe. *European Journal of Social Psychology*, *25*, 57–75.

Preacher, K. J., & Hayes, A. F. (2004). SPSS and SAS procedures for estimating indirect effects in simple mediation models. *Behavior Research Methods, Instruments, & Computers*, *36*, 717–731.

Rosette, A. S., Carton, A.M., Bowes-Sperry, L., & Hewlin, P. F. (2013). Why do racial slurs remain prevalent in the workplace? Integrating theory on intergroup behavior. *Organization Science*, *24*, 1402–1421.

Runciman, W. G. (1966). *Relative deprivation and social justice*. Berkeley: University of California Press.

Sears, D. O., & Henry, P. J. (2005). Over thirty years later: A contemporary look at symbolic racism and its critics. *Advances in Experimental Social Psychology*, *37*, 95–150.

Sidanius, J, & Pratto, F. (1999). Social dominance: An intergroup theory of social hierarchy and oppression. New York: Cambridge University Press.

Simon, L., & Greenberg, J. (1996). Further progress in understanding the effects of derogatory ethnic labels: The role of preexisting attitudes toward the targeted group. *Personality and Social Psychology Bulletin*, *22*, 1195–1204.

Swim, J. K., Aikin, K. J., & Hall, W. S. (1995). Sexism and racism: Old-fashioned and modern prejudices. *Journal of Personality and Social Psychology*, *68*, 199–214.