Quantifying shared and idiosyncratic judgments of racism in social discourse

Joel Eduardo Martinez^{1,2} & Elizabeth Levy Paluck^{1,2}

¹ Department of Psychology, Princeton University, Princeton, NJ, USA

² Princeton School of Public and International Affairs, Princeton University, Princeton, NJ, USA

Abstract

Rising nationalism on the one hand and anti-racist protests on the other have called into question norms surrounding what can be considered racist. How do individuals decide whether anti-immigrant statements are “racist”? This project sought to identify determinants of agreement about what kind of speech is “racist”. Across two studies and one replication ($N = 1211$) from a nationally representative panel, perceivers rated how “racist” they found social media statements about immigration. The statements were drawn following the announcement of immigration policies in the United States. We introduce variance component analyses to quantify the extent to which judgments of racism are shared or idiosyncratic. We calculated the contributions of perceiver demographics and attitudes (e.g., racial category, age, gender, political leaning, attitudes toward hierarchy and race), statement attributes (e.g., aggressiveness or target of the statement), and norm cues (e.g., statement popularity) to measure sources of agreement. While traditional aggregate analyses suggested that judgments about what constitutes racism differed on average between racial categories, the variance analyses identified substantial heterogeneity, even within racial categories and when norm cues varied (i.e., number of likes). Some agreement occurred when perceivers shared political and racial ideology, and when the statements contained aggressive messages. Judgments of racism may not be shared within predicted social cleavages (e.g., racial categories) in this particular case, because anti-immigrant sentiment is highly politicized. Mapping variance, compared to only analyzing differences in means, can improve theoretical inferences and avoid “group” essentializing within the psychological study of racism, group-based beliefs, and anti-prejudice endeavors.

Keywords: racism, immigrants, agreement, political ideology, variance component analysis

Introduction

Marginalization and exclusion are prominent topics in politics today. With the rise of nationalism in response to large-scale migration patterns (Kaufmann, 2019), public discourse increasingly demarcates who does not belong (e.g., refugees, asylum seekers, “illegals”). New waves of immigrants are often racialized through nationalistic projects where they are ascribed a salient group identity and interpersonally and institutionally treated as inferior and a threat to Americans (Abrajano & Hajnal, 2015; Brown et al., 2018; Inda, 2006; Ngai, 2004; Provine & Doty, 2011; Santa Ana, 1999; Warner, 2005). In the context of these racializing events and anti-racist mobilization fueling widespread discussions about racism, to what extent is there a public consensus that targeting immigrants is “racist”?

The language we have available in the U.S. for describing prejudice or discrimination is moralized such that a person or event labeled as “racist” is perceived as negative and socially unacceptable (Doane, 2006). In this way, charges of racism are an effective form of social condemnation. Yet, this condemnation of racism contrasts with increasing support for nationalism as a positive or necessary political force to secure national sovereignty. What determines interpersonal or collective agreement on whether anti-immigrant sentiment is “racist”? Without a shared understanding of what constitutes racism, discussing the problem or mobilizing solutions to address racism is intractable. Pew data suggest U.S. adults *believe* that other adults disagree about what constitutes racist language (Pew Research Center, 2019). Whether judgments of racism are actually idiosyncratic is an open question.

Social media provides an arena where race and racism as concepts are debated, (re)constructed, and reified (Chaudhry, 2015, 2016). In this study, we consider public statements made on social media in response to two Trump administration announcements: of the travel ban

from seven Middle Eastern countries, and of the building of a wall along the Mexico-U.S. border in early 2017 (Amatulli, 2017). These two announcements sparked a wide range of public statements on social media such as Twitter, which varied from enthusiastic support for these endeavors (González-Ramírez, 2017) to campaigns countering the anti-immigrant spirit of the policy announcements (Martinez, Feldman, & Cikara, 2017).

While previous research has sought to understand people's theories about the nature of racism and racist people (e.g., Hudson & Esses, 2005; Sommers & Norton, 2006), the goal of our endeavor is not to arrive at a particular definition of racism, but is rather an inquiry into the determinants of consensus about appropriate labeling of "racist". The aims of this project are twofold: 1) use a variance decomposition technique unfamiliar to this area of research which quantifies the degree to which racism judgments of these statements are shared or idiosyncratic, and 2) empirically quantify the contributions of characteristics of the perceivers, statements, and contexts to agreement or disagreement about judgments of racism. We highlight throughout how these endeavors can bring insights to the study of racism and group-based beliefs.

Theoretical characteristics shaping lay perceptions of racism

Whether racism can or will be used to describe anti-immigrant statements depends on how racism itself has been defined. Generally speaking, scholars across disciplines have identified racism at different levels of analysis (i.e., systems, institutions, ideologies, groups, and individuals) with varying implications for what racism is or is not (e.g., Blumer, 1958; Bonilla-Silva, 1997; Cazenave, 2016; Feagin, 2006; Hoyt, 2012; Richeson & Sommers, 2016; Salter et al., 2017; Sears & Henry, 2005). The lack of scholarly consensus on the nature of racism reveals how disagreements about the racist nature of a statement cannot be easily resolved by referring to a single definition or expert judgment. Absent a singular standard, to what extent will people's

racism judgments be shared or idiosyncratic? And, what influences the extent to which these judgments are shared or idiosyncratic? Quantifying the relative contributions of characteristics of the person (“perceiver”), the stimulus (a statement), and the situation (“context”) to agreement or disagreement can help adjudicate among theorized sources that contribute to individuals’ stances on racism.

Characteristics of the individual person (“The Perceiver”) that shape judgments of racism

We focus on an individual’s racial classification as the main characteristic of interest for its centrality in racism, but additionally investigate relevant ideological characteristics: political orientation, endorsement of symbolic racism, blatant dehumanization, and social dominance orientation.

Race is considered the fundamental cleavage of society (Bonilla-Silva, 1997) and the ideological core of racism (Cazenave, 2016; Fields & Fields, 2012; Golash-Boza, 2016). Individuals classified as different racial categories are thought to be motivated to perceive or deny racism depending on whether they are targeted as victims or implicated as perpetrators (Adams et al., 2006; Carter & Murphy, 2015). Correspondingly, group-based theorizing suggests that different racial categories consider different phenomena as racist (Sommers & Norton, 2006) and perceive discrimination at different rates (Norton & Sommers, 2011), conditional on individuals’ identification with their classification (Kaiser & Wilkins, 2010; Operario & Fiske, 2001; Sellers & Shelton, 2003).

A starting assumption for research on judgments of racism is that individuals classified as the same racial category will share similar judgments (Blauner, 1992; Blumer, 1958; Bobo, 1999), but disagree across racial categories due to different understandings of, experiences with, and relationships to racism (Adams et al., 2006; Carter & Murphy, 2015). For example, Latinxs

are not uniformly thought of as a race, but are commonly racialized through illegality in the U.S. (Flores & Schachter, 2018). Latinxs could be considered as targeted by tweets discussing undocumented immigrants coming from the Mexico-U.S. border (Warner, 2005) and may want to defend their group against these statements or signal a shared status as a discriminated non-white group by judging them as racist (Brown, 2013; Vargas, 2017). Whites are often thought of the default representative group of the United States (Devos & Banaji, 2005), have developed a language or frame for denying the centrality of racism to life experience (Bonilla-Silva, 2006; Feagin, 2006), and could be considered the beneficiaries of anti-immigrant sentiment with interests in defending anti-immigrant sentiment (Abrajano & Hajnal, 2015).

However, racialized individuals often encounter other forms of classification that facilitate competing judgments about racism. For example, ideological cleavages also establish shared goals and shared means for interpreting reality (Jost et al., 2008; Ondish & Stern, 2017; Stern et al., 2014; Stern & Ondish, 2018). The combination of shared reality within an endorsed political orientation and a sense of conflict between “right” and “left” political groups (Sides et al., 2019) has the potential to shape judgments of racism. Theories of political ideology as motivated cognition highlight the association between conservatism and the justification of inequality (Jost et al., 2003). Individuals who endorse more conservative stances therefore may learn to judge racism in a way that legitimizes their partisan membership and views of racial difference (e.g., by making fewer attributions of racism).

Shared judgments may also, or instead, come from individuals’ endorsement of ideologies that are not necessarily based on institutionally established identities. These ideologies justify inequality and anti-immigrant sentiment, and represent a perceptual denial of racism: endorsement of symbolic racism, social dominance orientations (SDO), or blatant

dehumanization (Berg, 2013; Bruneau et al., 2017; Costello & Hodson, 2011; Craig & Richeson, 2014; Esses et al., 2008). Endorsers of symbolic racism reframe racism as the inevitable consequence of targets' work ethic (Bonilla-Silva, 2006; Sears & Henry, 2003; cf. Sniderman et al., 1991). Individuals high in SDO prefer group-based hierarchy as the hegemonic form of social organization and endorse narratives about their societies that legitimize hierarchies (Liao et al., 2016; Pratto et al., 2013). Individuals who blatantly dehumanize others on a monkey-to-man scale also endorse discriminatory policies, show less compassionate responses to injustices towards racial outgroups (Bruneau et al., 2017; Kteily et al., 2015), and willingly disregard social desirability concerns about appearing racist to experimenters (Axt, 2017).

Characteristics of the Stimulus (Statement) that shape judgments of racism

A “stimulus” in this analysis is any statement that is under consideration as possibly racist (for our study purposes the stimuli are tweeted statements about US immigration). We focus on two characteristics of the stimulus: the group that is targeted by the stimulus and the extent to which the message is aggressive. The extent to which these characteristics contribute to agreement or disagreement depends on whether perceivers share understandings of racism that map unto these characteristics.

People may be more likely to judge stimuli as racist when the statements reference groups that have historically been targeted by racist antipathy (such as black people in the U.S. context). An equally logical counter-prediction is that these stimuli might be judged as less racist if this historical antipathy has decreased sensitivity to perceptions of racism against that group, or if the group is not commonly perceived as a race. These predictions are quickly complicated by considering the interaction between the beliefs of the perceiver and the group targeted by the stimulus. For example, some individuals are more likely to perceive racism against whites, even

though whites have not historically been the targets of racism, because of their attitudes about the changing demographics of the United States (Craig et al., 2018).

Aggressive or subtle antagonistic language in the content of a message can also influence judgments of racism. If a statement conveys the same racist message using subtle or coded language, it can be seen as reasonable discourse (Huddy & Feldman, 2009; Williams et al., 2016). This is exemplified, for example, by common yet dehumanizing metaphors in media and political rhetoric comparing migrants to animals or dangerous flood waters (Santa Ana, 1997, 1999). Statements with aggressive language are more likely to be flagged as racist, but they might also not due to recent shifts in norms of political speech that increased public acceptance of explicit racial and racist messaging (Crandall et al., 2018; Portice & Reicher, 2018; Schaffner, 2018; Valentino et al., 2018).

Characteristics of the context that shape judgments of racism

We consider times of political conflicts and social norms as two important contextual factors that should influence how we evaluate racism by serving as interpretative frames. If a speech act or behavior is encountered in a context of intergroup conflict, it is more likely to be regarded through a racial lens. For example, in times of economic downturn elite politicians often frame intergroup relations as competitive to shift blame (Green et al., 1998). These statements can shift public perceptions of racism and the degree of intergroup conflict (T. E. Nelson et al., 2007; Schaffner, 2018).

Relatedly, judgments of racism may also be guided by social norms. Normative information – e.g., how many others in a group or society also do or believe something – implies whether a judgment is shared (Crandall & Stangor, 2005; Tankard & Paluck, 2016). For example, a “liked” feature of Twitter communicates how socially popular a tweet has been on the

website. Thus, if a tweet is liked or shared many times, a perceiver can infer that many other people know about the tweet and also may extrapolate that many people endorse the tweet's content. This kind of normative information may shape impressions of whether that tweet is racist (Sechrist & Stangor, 2001; Stangor et al., 2001). If one believes that society is racist as a whole and a tweet is extremely popular, the tweet would also be seen as racist. Conversely, to the extent that many have endorsed a tweet and one believes that it is socially unacceptable to be racist, one may conclude the tweet is not racist. The perception of normative information, may in this way, account for agreement or disagreement in judgments of racism.

In sum, the psychological literature illustrates how characteristics of the perceiver, stimulus, and context can affect judgments of racism. Most of the previously listed characteristics have been linked to anti-immigrant acts or attitudes, but not to judgments of anti-immigrant statements. We investigate anti-immigrant statements made during a particularly racializing and politically turbulent moment for immigrants in The United States, a time in which we would expect many of the perceiver, stimulus, and context characteristics mentioned above to relate to these judgments. Our interest lies in quantifying *how much* agreement or disagreement about these statements is explained by each characteristic, rather than *whether* these characteristics relate on average to people's judgment of a statement. In this way, our approach speaks to questions about the determinants of a public consensus (or lack of consensus), which is a different kind of question than typical questions about the average determinants of individual attitudes, behaviors, or judgments. Below we describe how we empirically quantify these contributions using the variance component analysis (VCA).

Quantifying shared and idiosyncratic definitions of racism: The variance component analysis

The study of collective schemas and social representations has been a longstanding concern of psychologists (Echterhoff et al., 2009; Hardin & Higgins, 1996; Moscovici, 2000). A common practice in this kind of investigation is to draw inferences about shared beliefs from differences-in-means between groups or stimulus manipulations. As we will show, these average comparisons can obscure and shift focus away from important (and often large) judgment heterogeneity, which misleads inferences made about individuals' or groups' beliefs and stimulus effects. By contrast, the VCA revisits these concerns with new tools for understanding. VCA quantifies the *variance* of the judgments that can be attributed to differences across perceivers, stimuli, and to the interaction between perceivers and stimuli. By shifting the analytic focus to variance, VCA allows a researcher to 1), quantify the *extent* of agreement and 2), partition the contributions of various factors to that agreement.

Running a VCA requires ratings from multiple people who rate multiple stimuli multiple times. Repeated measures are methodologically crucial for investigations of shared or idiosyncratic judgments for they allow investigators to estimate the reliability of judgments and disentangle meaningful sources of agreement and disagreement from random error (Martinez et al., 2020). Collecting repeated measures is uncommon in psychology, but even rarer in the study of racism perception. In all of these ways, VCA represents an advance to the study of the collective psychological concept of “racism”.

VCA follows the social relations model (Kenny, 1994) for assessing accuracy in interpersonal perception, and generalizability theory, where error variance is decomposed into multiple variance components (i.e., sources of non-independence in the data; Shavelson et al., 1989). The VCA estimates three components of interest: the variance between participants' averages, between stimuli's averages, and between averages from each participant-stimulus pair

(i.e., their interaction) as derived from random-effects regression modeling (Judd, Westfall, & Kenny, 2012).

Figure 1A provides a schematic of the specific ratings that make up each variance component. **Figure 1B** shows the pattern of variance component sizes to expect given mostly shared or idiosyncratic judgments. One source of judgment disagreement comes from differences between perceivers' averages (*perceiver* σ^2) which originate from disagreements between groups of perceivers or from general rating tendencies of individuals. However, another less-investigated source of judgment disagreement comes from perceivers' idiosyncratic and *reliable* stimulus rankings that are not shared with other perceivers (*perceiver x stimulus interaction* σ^2). For example, a perceiver consistently finds tweet A as more racist than B while the opposite is true for other perceivers. Crucially, the interaction component requires repeated measures to be estimated otherwise it is confounded with residuals and labeled as error. The interaction component represents purely idiosyncratic contributions, while the perceiver component represents an "impure" idiosyncrasy. Perceivers might disagree via differences in means but can belong to clusters of perceivers with similar means and can still share stimulus rankings if there are no perceiver x stimulus interactions (e.g., all judge stimulus A as relatively more racist than B). Lastly, the stimulus variance component (*stimulus* σ^2) purely reflects shared judgments because agreement about distinctions between stimuli allows stimulus means to vary (e.g., a main effect of a stimulus characteristic). Stimulus A can only have a higher mean than stimulus B if most perceivers agree on that difference.

Figure 1. Variance component analysis. This figure provides an intuition about what the variance components represent, actual variance calculations involve maximum likelihood estimators rather than simple averaging. A) A simplified schematic on deriving the main variance components from multiple perceivers repeatedly rating multiple stimuli. The average of each perceiver (e.g., $P1$, $P2$, $P3$) represents the average of all ratings given by each perceiver (red); the variance of the differences between these averages and the grand mean (the average of all ratings) constitutes the perceiver component. The average of each stimulus (e.g., $S1$, $S2$, $S3$) represents the average of all the ratings given to each stimulus (blue); the variance of the differences between these averages and the grand mean constitutes the stimulus component. Lastly, the average of each stimulus-perceiver pairing (e.g., $P1S1$, ..., $P3S3$) represents the average of the repeated ratings given to each stimulus by each perceiver (purple). The variance of the differences between these averages and the participant and stimulus means (e.g., $P3S1$ from $P3$ and $S1$) constitutes the perceiver x stimulus interaction component. B) The left panel provides visual examples of the relative variances we might expect for each component given largely idiosyncratic judgments, the right panel represents expectations for shared judgments.

The partitioning of variance into meaningful components of agreement and disagreement make the VCA perfectly suited to answer questions about shared judgments. However, an underdeveloped application of this model is comparing within-group agreement about the racism of a statement and mapping the extent of agreement onto characteristics of perceivers or of stimuli. Some reasons for this underdevelopment are that the VCA does not automatically provide summaries for different groups of perceivers, and there is no established method for mapping variables' contributions to the different variance components (e.g., how much racial categories explain perceiver variance, or tweet aggressiveness explains stimulus variance).

We attempt to address both issues here. The first can be addressed by partitioning the data by racial, or any other, category and computing a VCA per category. With this, one can examine whether agreement, as indexed by variance components, varies according to category (e.g., Xie et al., 2018). The second can be addressed by directly estimating the random effects

that constitute the different variance components, and then calculating how much variance in those random effects is explained by perceiver or stimulus characteristics – thus quantifying their contributions to agreement or disagreement. VCA, as presented here, accomplishes the following: 1) provides a rigorous statistical assessment of the homogeneity of evaluations within a category, 2) identifies the key social cleavages that exhibit agreement or stimulus characteristics that elicit agreement, and 3) quantifies the relative importance of competing classifications or ideologies for guiding an individual's judgment.

Overview of studies

The main study design consists of participants from nationally representative online panels who rated statements (i.e., anti-Muslim and anti-undocumented immigrant tweets) on how racist they thought each statement was. Study 1 investigated levels of agreement within and between racial categories regarding judgments of racism, when participants were only presented with the text from the tweet (i.e., no information about social engagement). Study 2 directly replicated Study 1 with a different sample. Study 3 probed the impact of manipulated social norm cues, i.e., the popularity of the statements as reflected by number of likes attached to the statements.

Based on previous literature, we hypothesized that perceivers' shared racial classification should lead to shared racism judgments with each category, but disagreement between categories. However, we also tested the possibility that agreement might also, or instead, come from other measured perceiver and stimulus characteristics. In terms of averages, for Study 1 and 2, we predicted that whites would display on average lower racism ratings than black or Latinx perceivers. For Study 3, we predicted that whites would on average see popular tweets as less racist, but that black and Latinx participants might use tweet popularity to infer more or less

racism. Lastly, we predicted that black and Latinx participants would see unpopular tweets as more racist than whites on average. Our predictions for the studies were all pre-registered:

<https://osf.io/wz83n/register/565fb3678c5e4a66b5582f67>.¹

Study 1 - Quantifying Shared or Idiosyncratic Judgments of Racism

Methods

All data, scripts, and study materials can be found at the following archive:

https://osf.io/v2cey/?view_only=69a1bebeaf684236ad3a6ddc08d263ef.

Sample

368 participants living in the United States were recruited from a nationally-representative Qualtrics panel with the goal of obtaining 100 participants per each of three racial categories: Latinx, black, and white. Latinx represents a category targeted by the statements, white represents beneficiaries of the statements, and black represents a category also targeted by immigration enforcement yet is often erased in public discourse about immigration according to stereotypes of undocumented immigrants and Muslims as “brown” (Rivera, 2014). While, sampling individuals who identify solely with one of these categories ignores category intersections and over-simplifies the relationship between classification and group-based motivations, that is the point. Race is considered the fundamental cleavage in the organization of society that shapes psychology and racial discourse (Bobo, 1999; Bonilla-Silva, 1997; Feagin, 2006; Mills, 1997). All data collection for this and the rest of the studies occurred in June of 2018.

¹ The preregistration includes further manipulations that were not completed due to project constraints and may be assessed in future studies. For a list of all deviations from preregistration please see Appendix A.

Our exclusion criteria required participants show test-retest reliability on judgments of racism across the three blocks: specifically, they needed to demonstrate an intraclass correlation (ICC) of judgments greater than zero. This criterion excluded individuals whose responses were completely inconsistent across repetitions and individuals who pressed the same response for all the stimuli as there was no variance to calculate reliability.

After exclusions ($N = 62$), the final sample included 306 participants: 103 black, 103 Latinx, and 100 white. Our sample was skewed towards males ($N = 218$). The average age was 43, $SD = 16$, and ranged between 18 and 82. The ICCs ranged from .01 to .94 (average = .45, median = .50, $SD = .28$). For a full composition breakdown of all measured participant factors by racial category, see **Supplementary Table S1**.

A sample-size analysis via a simulation study showed that many of the components computed in the VCA begin to stabilize in terms of their estimation when the number of participants were 60, the number of stimuli 30, and the number of repeated measures greater than 2; diminishing improvements on precision was found with greater numbers (Martinez et al., 2020). Our studies have approximately 100 participants in each racial category (total N is around 300 per study), 50 stimuli, and 3 repeated measures to help optimize estimate precision.

Stimuli

We collected tweets from Twitter from February 6 to February 13, 2017. Tweets were found by scraping the hashtags #AmericaFirst and #MuslimBan,² which developed in reaction to these events. The set was deliberately large to avoid the confounds of other studies that present participants with a few stimuli, but also pared down to a reasonable number that could be rated

² Hashtags organize tweets that contain the relevant hashtag into a centralized public space where people collectively discuss a topic.

multiple times within one session. In order to narrow down to fifty tweets, two judges (one of whom is the first author) chose tweets from this collection that antagonistically targeted Muslims, undocumented immigrants, or immigrants in general. They also eliminated tweets that were not comprehensible, that only featured hashtags or images without a personal comment, and that had exceedingly offensive content (e.g., n-word or images of murder). Judges also excluded similar or repetitive statements, and specifically sought out comments addressing a range of social, economic, and political perspectives. This process was aimed at creating a stimulus set that incorporated relevant theoretical dimensions, rather than a representative set of tweets on immigration.³ This design choice reflects practice in social perception research. For example, when people evaluate the beauty of abstract objects, rather than assemble a representative set of all possible objects (which is impossible), researchers assemble a set of objects that vary on a few dimensions. In this way, intractable questions assessing agreement about the beauty of all possible objects get transformed into tractable questions that attempt to identify factors related to agreement about a smaller set of objects – whether from the perceivers or the chosen object dimensions. This approach then lends itself to meta-analytically testing

³ For example, by including only tweets that antagonistically discuss a group, we precluded the ability to understand how people distinguish antagonistic vs. innocuous tweets in terms of racism. However, that kind of agreement is not of interest here. Instead, we are focused on statements that deride groups as interpersonal discrimination is the most popular understanding of racism (Adams et al., 2008). Likewise, by removing the most egregious of tweets, we also precluded our ability to capture judgments of “obviously” racist tweets. One could argue this omission could omit another form of agreement. There are ethical considerations to be made about the value of exposing participants to *extreme* disturbing content, but just as important we believe many of the chosen tweets were egregious enough (**Supplementary Table S2**). It is an open question whether the kind of tweets left out of the set would elicit shared judgments (e.g., racist murders by police or military are often defended and non-black people still comfortably and publicly say the n-word despite pushback on both). While acknowledging this possible limitation, we direct readers to **Supplementary Figure S2** for additional discussion regarding the tweets’ rating ranges and textual content to better situate the results.

generalizability and boundary conditions across studies that study the same phenomena using VCA.

The final group of tweets about the Muslim ban and the wall varied along two dimensions: the aggressiveness of the content and the groups that are targeted, as subjectively judged in the pairing-down process. To the first point, we included a variety of tweets ranging from blatantly aggressive statements to more coded language that linked particular groups to economic or political worries. The more aggressive tweets contained statements that dehumanized others or called for violence. The political tweets contained desires for safety from threats, financial responsibility, or helping Americans first. With respect to targeting, we considered two measures of group targeting. The first is the hashtag from where the tweet originated: *#AmericaFirst* or *#MuslimBan*. The second codes what specific group the tweet discussed: undocumented immigrants, Muslims, both, or all immigrants.

For this study, we did not display any information about social engagement. We obscured the tweeter's picture and username and the number of likes, favorites and retweets and the hashtag to anonymize the tweet for the study.

Procedure

Participants first completed a demographic survey where we asked about age, racial or ethnic background (Latinx or Hispanic, black or African American, Asian/Pacific Islander or Asian American, white or European-American, Indigenous, or other), gender identity (woman, man, nonbinary, other, prefer not to say), and twitter usage (daily, several times a week, weekly, less often, never). Participants who did not identify as one of the three main racial categories of interest could not advance to complete the study.

Participants were instructed that they would read a series of tweets related to the two U.S. policies about the wall and the ban in three repeated blocks that randomly displayed all the stimuli. In between blocks there were instructions that all the same tweets would be shown again for the same task. Their task was to rate “how racist is this tweet?” on a scale of 1 (not at all) to 9 (extremely). Participants were asked to use their gut judgment when rating the tweets.

The task was a fully crossed design (Westfall et al., 2014): every participant rated all fifty tweets which were repeatedly presented in three blocks, randomized within block. Finally, participants completed four different measures: the social dominance orientation 8 item scale (Ho et al., 2015), the symbolic racism 8 item scale (Henry & Sears, 2002), the blatant dehumanization scale (Kteily et al., 2015), and political orientation. There was no time pressure for completing each rating and for completing the measures, but given the number of tweets, participants moved expeditiously through the questions (average time = 35 mins).

Measures

The order of the first three scales listed were randomized across participants. Political orientation was always collected last. A correlation matrix displaying the relationship among all the collected perceiver measures can be found in **Supplementary Figure S1** and average scores of these measures across the sample and within racial categories can be found in **Supplementary Table S1**.

Demographics. Racial or ethnic background, age, gender identity, and Twitter use were measured. Twitter use is treated as a continuous measure, where the lowest number, 1, represents daily use and the highest number, 5, represents never having used Twitter before.

Blatant Dehumanization. The blatant dehumanization scale asked participants to rate four groups (white, black, Latinx/o/a, Hispanic, or Muslim) using the ascent of man diagram, a

continuous rating underneath a visual scale featuring a monkey to man continuum of evolution (Kteily et al., 2015). The participants were instructed to use the sliders to indicate how evolved each group is. No numeric values were shown on the slider, however the range varied from 0 (least human) to 100 (most human).

We derived five blatant dehumanization scores to understand the relative importance of general outgroup vs. specific group dehumanization in judgments of racism. The first is outgroup dehumanization (BD-Outgroup) which takes each participant's ingroup rating and subtracts the average rating of the other three groups. The other four scores subtracted each participant's rating of the target outgroup from their ingroup rating, so that the score is relative to their dehumanization of their own group (Black: BD-Black, Latinx: BD-Latinx, Muslim: BD-Muslim, or White: BD-White). Positive scores represent the magnitude of dehumanization while negative scores provide evidence of ingroup dehumanization.

Social Dominance Orientation. Participants responded whether they strongly oppose or strongly favor eight items on a 7-point scale, half of which were relevant to the egalitarian component of SDO (*SDOe*, e.g. *"It is unjust to try to make groups equal."*) and the other half to the dominance component (*SDOd*, e.g. *"Some groups of people are simply inferior to other groups."*) (Ho et al., 2015). After reverse coding some contrasting items, items were averaged to compute component scores. Higher scores indicated a higher level of *SDOe* or *SDOd*.

Symbolic Racism. We used the eight-item version of the Symbolic Racism scale (Henry & Sears, 2002). For example, participants would respond their level of agreement with the following statement, *"It's really a matter of some people not trying hard enough; if blacks would only try harder they could be just as well off as whites."* After reverse scoring, all the ratings

were summed together following convention to create the composite score that could range from 8 to 31, where higher scores indicate more endorsement of symbolic racism.

Political Orientation. Political orientation was assessed using two questions: social and cultural stance and economic stance from extremely liberal to extremely conservative on a 7-point scale (e.g., “*In terms of {social and cultural, economic} issues, how liberal or conservative are you in an American context?*”). The two were correlated at $r = .88$ and were averaged into one political orientation score.

Analysis

This analysis takes a multi-step approach to understanding the factors that contribute to shared judgments of racism. First, we report a linear regression analysis to estimate average directional associations between characteristics of the perceivers or of the stimuli and racism judgments. Second, we calculate a rough summary estimate between 0 and 1 of the agreement in judgments about racism in the tweets (the “beholder” index). Third, we use variance partitioning to describe the sources of variance, specifically the perceivers, stimuli, and their interaction (“VPC” or variance partitioning coefficients). Lastly, we use characteristics of the perceivers and stimuli to quantify how much variance they explain in each component. For example, when there is high variance between perceivers, perceiver characteristics like political orientation may contribute to disagreement. When there is more variance between the stimuli, this suggests that there are properties of the stimuli driving agreement.

Aggregate judgment analyses

To estimate the relationship between characteristics and judgments, we used mixed effect linear models in lme4 package version 1.1.13 (Bates et al., 2015) with maximum likelihood estimation. The following constituted the model formula: $\text{rating} \sim \text{race} + \text{political orientation} +$

age + twitter use + gender + SDOd + SDOe + blatant dehumanization outgroup + symbolic racism + tweet hashtag + targeted group + tweet aggressiveness + (tweet aggressiveness|participant) + (1|stimulus) + (1|participant:stimulus). To help with convergence, models were optimized with “bobyqa”, and allowed to run for a maximum of 500,000 iterations, and the tweet hashtag and group were removed as random slopes.

Confidence intervals, slope estimates, and multiple comparisons were computed using the “emmeans” package version 2.26.3 in R (Lenth, 2014) with Satterthwaite approximations of the degrees of freedom. Multiple comparisons were corrected using the Benjamini & Hochberg method for controlling false discovery rate. All continuous variables were mean centered, except twitter use which was centered on daily use.

Agreement analyses

Variance Component Analysis. We calculated two measures of agreement: variance partitioning coefficients (VPC) and beholder indices (BI). These measures are both derived from variance components analyses (Hönekopp, 2006; Martinez et al., 2020), mainly from separate components estimating variance between perceivers, stimuli, and their interaction around a grand mean (Judd et al., 2012). The model equation, combining the two levels of the model, is:

$$y_{ij} = \mu + p_i + s_j + ps_{ij} + e_{ij}. \quad (1),$$

where:

$$p_i \sim N(0, \sigma_p^2) \quad (2)$$

$$s_i \sim N(0, \sigma_s^2) \quad (3)$$

$$ps_{ij} \sim N(0, \sigma_{ps}^2) \quad (4)$$

Here, y_{ij} represents the judgment given by i th perceiver to the j th tweet, μ represents the grand mean over all judgments, p_i represents the difference of the i th perceiver's mean from the grand

mean, s_j represents the difference of the j th tweet's mean from the grand mean, ps_{ij} is the divergence of i th perceiver's average rating of j th tweet from their own mean and the tweet's mean, and e_{ij} represents random error in i th perceiver's judgment of j th tweet. The variance components of the perceiver (σ_p^2), stimulus (σ_s^2), and perceiver x stimulus interaction (σ_{ps}^2) random effects are the focus of this analysis and are assumed to be normally distributed and have a mean of zero (Shavelson et al., 1989; Shoukri, 2011, p. 44).⁴ These components were estimated on the raw data with random-effects regression models from the lme4 package version 1.1.13 in R, with restricted maximum likelihood.

Variance Partitioning Coefficients. The VPCs we report are computed by taking each variance component and dividing it by the total variance, including residuals (Goldstein et al., 2002). Confidence intervals for the VPCs are computed using 1,000 bootstrapped samples of the random effect models. Differences in VPCs between studies or conditions were computed by first subtracting the bootstrapped sample traces, calculating confidence intervals of the differences, and checking if they contain zero.

Beholder Indices. The beholder indices (BIs) represent two ratios of idiosyncratic to shared variance from the variance components (Hönekopp, 2006): idiosyncratic σ^2 (shared σ^2 + idiosyncratic σ^2)⁻¹. The $b1$ index assumes that the interaction component is the only true measure of idiosyncrasy and disregards the perceiver component: interaction σ^2 (stimulus σ^2 + interaction σ^2)⁻¹. The $b2$ index takes into account the perceiver component as idiosyncrasy: perceiver σ^2 + interaction σ^2 (stimulus σ^2 + perceiver σ^2 + interaction σ^2)⁻¹. These indexes provide more binary

⁴ We also compute block, block x stimulus, and block x perceiver variance components. These components do not factor into the agreement estimates, but can be important in interpreting the impact of repeated presentations on judgments. We find very little impact in this study, but we describe their purpose in **Appendix B**. For the complete set of estimated variance components for each study, see **Supplementary Figure S3**.

conclusions than the VPC. If the index is above a .5, evaluations are more idiosyncratic in nature and if the index is below a .5 the evaluations are more shared. An index at .5 would suggest equal shared and idiosyncratic contributions.

We used all of the data to obtain overall estimates across all perceivers as the main reference point. We then partitioned the data into racial categories and fed them into the model to obtain category-specific estimates (i.e., black, Latinx, white).

Variance contributions. As previously mentioned, there is no established method to map variable contributions to specific variance components. In a typical linear regression, one would compare a variance explained metric (R^2) from a model with and without the variable of interest, the additional variance explained is considered its contribution to the dependent variable. In a multi-level regression, a similar additive analysis can be done by, for example, comparing the overall log-likelihood of the models. However, maximum likelihood estimations of *variance components* are non-linear and dependent on model specification, therefore this additive strategy for quantifying and comparing variable contributions to variance components can lead to unexpected results (e.g., a variance component can *increase* with the added variable) (Hox, 2010, pp. 69–78).

Here, we used a two-step analysis to address this issue. The first step obtained empirical Bayes estimates of all of the random effects (i.e., the adjustments of each perceiver or stimulus or interaction mean), which resulted in a distribution per random effect estimate, instead of point estimates. The estimates were calculated using 1,000 simulations from the posterior distribution of the model with the merTools R package version 0.3.0 (Knowles & Frederick, 2016). In step two, the modes from every random effect's distribution were then submitted to ordinary least-squares regressions where each characteristic was independently added as a regressor and a

variance explained metric (R^2) was estimated. This R^2 metric was transformed to a percentage and for our purposes represents the independent variance contribution of each regressor.

Lastly, a model with all of the measures (excluding the group-specific blatant dehumanization scores because of collinearity) was estimated for each component to examine the total variance contributed simultaneously. This process was implemented separately for each variance component. Importantly, perceiver-level regressors only explained the perceiver component, stimulus-level regressors only explained the stimulus component, and interactions between every perceiver and stimulus variables were used to explain the interaction component. All continuous variables were mean centered to provide interpretable intercepts and variance components (Bryk & Raudenbush, 2002).

Results

Regression coefficients. **Table 1** displays the average directional effects of regressors on racism judgments. The grand mean of the model was 5.12, 95% CI[3.93, 6.32], traversing the midpoint of the scale. *Symbolic racism*, *political orientation*, and the *aggressiveness of tweets* were significantly associated with racism ratings. Lower racism ratings were given to politically framed tweets and by individuals high in symbolic racism or with a conservative political orientation. *Race* also predicted the evaluations whereby Latinxs on average (6.35, 95% CI[5.44, 7.26]) gave higher racism ratings than white (5.66, 95% CI[4.73, 6.59]) and black (5.58, 95% CI[4.70, 6.46]) participants. All other factors were not significant.

The tweet with the most disagreement was “We have enough pedophiles and rapists in America, we don't need anymore.” ($M_{\text{rating}} = 5.4$, $SD = 3.2$), the tweet with the least disagreement was “We're very close to having Pakis banned from America. Love it!” ($M_{\text{rating}} = 6.9$, $SD = 2.6$).

Table 1. Estimated marginal means and slopes from racism ratings in Study 1.

Regressor	Coefficient <i>b</i>	SE	95% CI	P value
<i>Age</i>	.002	.007	-.01, .02	.719
<i>Twitter use</i>	-.11	.07	-.23, .02	.098
<i>Social Dominance</i>	.04	.11	-.17, .25	.684
<i>Orientation egalitarian</i>				
<i>Social Dominance</i>	-.17	.11	-.38, .04	.12
<i>Orientation dominance</i>				
<i>Blatant</i>	.002	.005	-.008, .01	.733
<i>Dehumanization</i>				
<i>Outgroup</i>				
<i>Symbolic Racism</i>	-.18	.02	-.23, -.13	<.0001
<i>Political Orientation</i>	-.33	.07	-.47, -.19	<.0001
<i>Race</i>				
Latinx	6.35	.37	5.44, 7.26	-
Black	-.77	.26	-1.38, -.15	.009
White	-.68	.26	-1.32, -.05	.014
<i>Gender</i>				
Woman	6.28	.37	5.39, 7.17	-
Man	.06	.24	-.51, .63	.989
Nonbinary	-.02	1.7	-4.22, 4.18	.989
<i>Hashtag</i>				
#MuslimBan	6.5	.39	5.6, 7.4	-
#AmericaFirst	-.15	.22	-.60, .29	.477
<i>Tweet Type</i>				
Political	5.22	.35	4.41, 6.03	-
Aggressive	1.12	.19	.73, 1.51	<.0001
<i>Target group</i>				
All immigrants	6.34	.37	5.39, 7.29	-
Both	.039	.65	-1.76, 1.84	.952
Middle East	.12	.26	-.62, .85	.952
Undocumented	.26	.27	-.5, 1.03	.952

Beholder indices. Both beholder indices, *b1* and *b2*, suggest that perceivers' evaluations were mostly idiosyncratic when estimated within each racial category and across the full sample. Across all partitions, *b1* values hovered around .7 (All = .73; Latinx = .79; black = .70; white = .68) and *b2* values around .9 (All = .91; Latinx = .93; black = .88; white = .88). These are all above the .5 threshold that represents equal shared and idiosyncratic contributions. To put these values into perspective, beholder indices derived from beauty evaluations of faces were .55 (*b1*)

and .79 (*b2*) (Martinez et al., 2020) and .74 (*b1*) and .79 (*b2*) for the likability of artwork (Leder et al., 2016). This stands in stark contrast to beholder indices of highly shared evaluations of the dangerousness of objects with sharp shapes, .10 (*b1*) and .28 (*b2*) (Martinez et al., 2020). This suggests that despite average differences in ratings between racial categories, within-categories judgments contain a lot of disagreement.

Figure 2. Variance partitioning coefficient estimates and their contributing factors in racism judgments. A) Columns represent the category partitioning of the data for estimating the rater (red), stimulus (blue), and interaction (purple) VPCs. Error bars represent 95% bootstrapped confidence intervals. B) Columns represent variance components, y axis lists the regressors, and x axis represents the amount of variance explained by each regressor for each component. The specific amount is written next to the bars. The interaction component depicts the first 16 regressors that explained the most variance.

Variance partitioning coefficients. The relative sizes of the variance components can explain why the beholder indices indicate highly idiosyncratic judgments. The perceiver and interaction VPCs explain the most variance across all partitions of the data (**Figure 2A**). Conversely, the stimulus VPC is relatively small, which can occur if the range of tweets was homogenously racist (**Appendix B**). Examining the within-stimulus variance suggests otherwise. While the range of tweet averages hovered around the middle of the scale (around 4 - 7), they also had large standard deviations indicating each tweet was given ratings across the full scale (**Supplementary Figure S2**). Therefore, the relatively small range of tweet averages is due to

rater disagreements in the ratings of each tweet, as reflected in the large perceiver and interaction components.

Variance contributions. The stimulus component mainly consisted of the aggressiveness of the tweets which explained 48% of the stimulus variance. Together, all the stimulus factors explained 49% of the stimulus component. The largest contributors to perceiver variance were symbolic racism, political orientation, and SDOd explaining 27%, 17%, and 9% of the variance, respectively. The rest of the measures contributed 7% or less, however, altogether the measures explained 37% of the perceiver component. Lastly, none of the two-way interactions explained more than 1% of the interaction component.

Discussion

The results exhibited how racism judgments of anti-immigrant tweets are influenced by multiple sources. In line with theories of group-motivated perceptions, racial categories differed in their average racism ratings; Latinx perceivers rated the tweets as more racist than other categories. However, these average differences were overshadowed by a large amount of perceiver-related variance, a level of disagreement comparable with aesthetic evaluations of artwork (Leder et al., 2016). There was more within-category disagreement about racism than would be expected on the basis that individuals classified as a racial category have a shared psychology due to their historical relationship to racism (Adams et al., 2006; Blauner, 1992; Carter & Murphy, 2015). Perceiver variance, or disagreement, was more associated with ideological variables like symbolic racism and political orientation; lower average judgments of racism were given by conservatives and endorsers of symbolic racism. Agreement therefore occurred more within opposing ideological divides. To put these patterns into perspective, while all three characteristics were significant, race categories only explained 7% of the perceiver

variance (3.5% of the total variance) while symbolic racism and political orientation explained 17% and 27% of the perceiver variance (8.5% to 13.5% of the total variance).

Other sources of agreement were more cross-cutting. The aggressiveness of tweets, which explained about half of the stimulus variance, elicited shared judgments from the sample such that aggressive tweets were perceived as more racist than politically framed tweets. Even so, the overall impact of this source of agreement was minimal as it only explained about 5% of the total variance. Lastly, there was a sizable interaction variance component, indicative of truly idiosyncratic judgments where a perceiver's racism ranking of the tweets was unique, reliable, and not shared by other perceivers. The sources of these idiosyncrasies remain elusive despite our analytic attempts at understanding them.

Overall, the VCA revealed insights that would remain hidden by analyses that solely focus on aggregate results. The first is a reminder that we cannot reliably predict an individual's stance on racism based on their racial classification or from differences in means between racial categories. The second is that while one major source of disagreement came from mean differences between perceivers, another major contributor (20% of the total variance) was disagreements between perceivers' unique and reliable stimulus rankings, which were only estimable because of the repeated measures. The large interaction variance indicates that we cannot assume the stimulus set is perceived in the same way across perceivers, even when they share the same mean judgment (e.g., a pair of liberal Latinxs). Lastly, after accounting for variables theorized to be strongly related to judgments of racism, a majority of the meaningful variance remained unexplained. These insights warn researchers to be cautious about discussing categories as psychologically-bounded groups (Loveman, 1999), flattening evaluative

heterogeneity, and overclaiming variable impact when discussing or making claims about collective psychology from aggregate analyses.

Study 2 – Replication of Study 1

The aim of this study was to replicate Study 1 using the same stimuli and a new sample from a nationally representative Qualtrics panel of the United States. Study 2's sample was collected simultaneously alongside Study 1. In addition to the benefits of replication generally, Study 2 provides a variance-based benchmark for assessing replication (Gantman et al., 2018).

Methods

Sample

From 370 participants, 69 were excluded for not meeting reliability criteria. The final sample ($n = 301$) included 99 black, 101 Latinx, and 101 white participants. The participants were collected from the nationally representative Qualtrics panels. The ages ranged between 18 and 81 (average = 43, $SD = 16$). The ICCs ranged from .01 to .94 (average = .46, median = .49, $SD = .28$). For a full description of the sample see **Supplementary Table S1**. A correlation matrix of perceiver measures can be seen in **Supplementary Figure S4**, which show congruent patterns to Study 1's matrix.

Stimuli

The 50 tweets were the same exact set as used in Study 1.

Procedure

Procedures for this study mirrored those of Study 1.

Perceiver measures

The same measures were collected as Study 1.

Analysis

Aggregate judgment analysis

The same mixed effect regression model specifications of Study 1 were used for this study.

Agreement analyses

The goal is to obtain variance estimates and compare them to Study 1. Changes in estimates across both studies can help identify if the disagreement found in Study 1 can be replicated. Beholder indices and VPCs are calculated across the full sample and within racial categories as done in Study 1. For the difference scores in beholder indices and VPCs, Study 1 bootstrapped estimates are subtracted from this study's bootstrapped estimates. For this replication, we also conduct the variance contribution analysis to examine if measures' contributions are consistent.

Results

Regression coefficients. The grand mean of the model was 5.66, 95% CI[5.3, 6.0]. Similar to Study 1, *symbolic racism*, *political orientation*, and the *aggressiveness of tweets* significantly predicted racism ratings. Scoring high on symbolic racism, being more conservative, and blatantly antagonistic tweets lead to lower racism ratings. This time however, *race* only showed the predicted patterns descriptively ($ps > .12$). Latinxs rated the tweets as more racist (6.32, 95% CI[5.5, 7.2]) than white (6.1, 95% CI[5.3, 7]) and black (5.8, 95% CI[4.9, 6.7]) participants. Other regressors were not significant (**Table 2**).

The tweet with the most disagreement was “And to me, it's not even America First. it's simply protecting us from #terrorists who wish to kill us and our children” ($M_{\text{rating}} = 5.0$, $SD = 3.1$), the tweet with the least disagreement was “We're very close to having Pakis banned from America. Love it!” ($M_{\text{rating}} = 7.1$, $SD = 2.5$).

Table 2. Estimated marginal means and slopes from racism ratings in Study 2.

Regressor	Coefficient	SE	95% CI	P value
<i>Age</i>	-.004	.007	-.02, .01	.587
<i>Twitter use</i>	-.03	.06	-.16, .09	.592
<i>Social Dominance Orientation egalitarian</i>	-.19	.11	-.42, .02	.072
<i>Social Dominance Orientation dominance</i>	-.02	.11	-.24, .20	.852
<i>Blatant Dehumanization Outgroup</i>	.001	.005	-.008, .01	.775
<i>Symbolic Racism</i>	-.15	.03	-.20, -.10	<.0001
<i>Political Orientation</i>	-.34	.07	-.48, -.19	<.0001
<i>Race</i>				
Latinx	6.32	.35	5.48, 7.17	-
Black	-.51	.25	-1.12, .09	.125
White	-.21	.26	-.84, .41	.412
<i>Gender</i>				
Woman	6.41	.36	5.6, 7.2	-
Man	-.08	.23	-.52, .36	.714
<i>Hashtag</i>				
#MuslimBan	6.5	.37	5.7, 7.3	-
#AmericaFirst	-.16	.21	-.59, .27	.449
<i>Tweet Type</i>				
Political	5.26	.33	4.51, 6.02	-
Aggressive	1.06	.19	.68, 1.43	<.0001
<i>Target group</i>				
All immigrants	6.33	.35	5.44, 7.21	-
Both	-.09	.63	-1.84, 1.64	.882
Middle East	.05	.26	-.65, .77	.882
Undocumented	.28	.27	-.46, 1.02	.882

Beholder indices. As in Study 1, both beholder indices suggest that evaluations of racism are largely idiosyncratic across the full sample and within racial categories, contrary to expectations of group-based agreement. Once again, *b1* is consistently around .7 (All = .73, Latinx = .75, black = .72, white = .70) and *b2* around .9 (All = .91, Latinx = .93, black = .88, white = .91). Differences from Study 1 estimates were very small in magnitude and do not constitute meaningful changes in conclusion about the idiosyncratic nature of racism judgments.

Variance partitioning coefficients. The VPCs suggest that the idiosyncratic beholder indices are driven by large perceiver and perceiver x stimulus variance (**Figure 3A**). The stimulus VPC is relatively smaller driven by disagreements about the rating of each tweet (**Supplementary Figure S2**), consistent with Study 1. The VPC differences between Study 1 and 2 are very small, even within the racial categories, and do not constitute meaningful changes from Study 1 (**Supplementary Figure S5A**). Taken together, the pattern of results suggests that the overall level of disagreement in Study 1 was replicated.

Figure 3. Variance partitioning coefficients and their contributing factors in racism judgments. A) Columns represent the category partitioning of the data for estimating the rater (red), stimulus (blue), and interaction (purple) VPCs. Error bars represent 95% bootstrapped confidence intervals. B) Columns represent variance components, y axis lists the regressors, and x axis represents the amount of variance explained by each regressor for each component. The specific amount is written next to the bars. The interaction component only depicts the first 16 regressors that explained the most variance.

Variance contributions. The stimulus component was mainly explained by tweet aggressiveness with 45% of the variance (**Figure 3B**). The tweets' hashtag and the group discussed explained a minimal amount of the interaction variance, but all together 47% of the stimulus variance was contributed by our measured stimulus factors. The perceiver variance had the largest contributions by the following four perceiver factors: symbolic racism (24%), political orientation (21%), and social dominance orientation egalitarian (9%). The rest contributed less than 6% of the variance. Together, all perceiver factors contributed 33% of the

variance. Lastly, none of the two-way interactions explained more than 1% of the interaction component. These results are consistent with Study 1.

Discussion

We replicated the main finding of the first study that racism evaluations of anti-immigrant sentiments are largely idiosyncratic within racial categories and across them. Evaluations showed consistently large perceiver and perceiver x stimulus variance components and a smaller stimulus variance. Most of the same factors (i.e., symbolic racism, political orientation, tweet aggressiveness) identified in Study 1 contributed to these components in similar proportions. Symbolic racism and political orientation contributed to differences between perceivers whereby higher scores were associated with lower racism ratings. Likewise, more politically framed tweets were seen as less racist than antagonistic tweets which contributed to between-perceiver agreement. Most of the meaningful variance again remained unexplained.

This replication contained an important point of departure from Study 1 that is worth highlighting. Here, race was not a significant regressor of racism judgments. We have now shown large within-category disagreement about racism in the presence and absence of significant between-category differences in racism ratings. Thus, perceiver and perceiver x stimulus sources of disagreement beyond race may be a persistent quality of racism judgments of these anti-immigrant statements. Despite this consistency, future investigations should build on these findings and explore whether greater perceiver influence applies more generally to racism judgments in other contexts and with other stimuli.

Study 3 - The Role of Social Norm Cues in Racism Judgments

The previous studies concluded that judgments of racism are largely idiosyncratic when evaluating anti-immigrant tweets that contained no information regarding how others engaged

with the tweet. In this study, we examine whether displaying the number of responses to each tweet as a cue to social agreement with the tweet may set the standards for what is acceptable discourse. The impact of social norm cues is indexed by comparing levels of agreement or disagreement relative to the combined data from Study 1 and Study 2 whose tweets contained no social engagement information. We predicted that tweets that exhibited lots of social support (i.e., popular) might increase the level of agreement within and across racial categories as indexed by low beholder indices and relatively large stimulus variances and smaller perceiver-related variances. Conversely, tweets that showed a low level of social support (i.e., unpopular) would overall lead to idiosyncratic evaluations as the ambiguity of the tweets' acceptability among others may lead people to stick to prior beliefs.

Methods

Sample

The samples were collected from the nationally representative Qualtrics panels. After exclusions for reliability ($n = 69$), the final sample in the *unpopular* condition included 301 participants: 102 black, 102 Latinx, and 97 white. The ages ranged between 18 and 90 (average = 42, $SD = 16$). The ICCs ranged from .01 to .94 (average = .44, median = .49, $SD = .28$).

For the *popular* condition, after exclusions ($n = 75$) there were 303 participants: 103 black, 98 Latinx, and 102 white. The ages ranged between 18 and 89 (average = 41, $SD = 16$). The ICCs ranged from .01 to 1 (average = .43, median = .46, $SD = .27$).

For full demographic characteristics and mean scores of perceiver measures, see **Supplementary Table S3**. The correlation matrices of perceiver measures for both conditions showed similar patterns (**Supplementary Figures S6 and S7**).

Stimuli

The 50 tweets were the same as the previous studies. However, we removed the social information censors and added information about the tweets' number of likes, retweets, and comments as cues to social support. These tweet properties were assigned randomly generated values. In the unpopular condition, retweets could range between 1 to 20, likes between 1 and 20, and comments between 0 and 5. In the popular condition, retweets could range between 1,000 to 10,000, likes between 1,000 to 10,000, and comments between 100 to 400. After randomization, the ratio of retweets to likes was adjusted so it was not impossibly high or low as there are usually less retweets than likes ($M = .71$, $SD = .45$). These ranges were chosen to signal heavy and light engagement with tweets.

Procedure

Procedures are consistent with the previous studies.

Measures

The blatant dehumanization, social dominance orientation, symbolic racism, and political orientation scales and race, age, gender, twitter use were measured and computed the same as the previous studies.

Analysis

Aggregate judgment analyses

We constructed two full mixed effect regression models with all the same measures and model specifications as Study 1 and 2, one for each condition.

Agreement analyses

The agreement estimates, beholder indices and VPCs, were computed as in the previous studies. This time, the estimates from the popular and unpopular conditions are compared to each other and to the estimates from the combined data from Study 1 and Study 2. The former

examines the effect of level of social support while the latter examines the impact of social norm information on variance components and, by implication, agreement.

We also compute variance contributions of perceiver and stimulus variables as in previous studies. Variance contributions are computed separately for the popular and unpopular conditions with the goal of examining which explanatory measures are consistently relevant across conditions and which measures are unique to a condition.

Results

Regression Coefficients

As some evidence that our social norm manipulations functioned as intended, averages across conditions fit the expected pattern. Specifically, all racial categories rated the popular tweets as less racist than unpopular tweets. This could arise either because the popularity information signals acceptable discourse or because unpopular tweets signal unacceptable discourse. A descriptive comparison to Study 1 and 2 where no social norm information was given suggests both may be at play. Average racism ratings in all categories were the highest in the unpopular condition (white = 6.17, black = 6.15, Latinx = 6.84), followed by the Study 2 (white = 6.11, black = 5.81, Latinx = 6.33) and Study 1 averages (white = 5.66, black = 5.58, Latinx = 6.34), and the popular tweets yielded the lowest averages (white = 5.52, black = 5.56, Latinx = 5.52).

Unpopular condition. The grand mean was 5.8, 95% CI[5.1, 6.5]. The two dominant regressors were *symbolic racism* and the *aggressiveness of the tweets*. Social dominance orientation *egalitarianism* and *dominance* were also significant. Lastly, *race* was significant in the predicted manner. Latinxs (6.61, 95% CI[5.73, 7.50]) rated the tweets as more racist than

white (5.97, 95% CI[5.10, 6.84]) and black participants (5.97, 95% CI[5.10, 6.83]). Other regressors were not significant (**Table 3**).

The tweet with the most disagreement was “American lives everywhere are being saved by deporting ILLEGAL CRIMINAL ALIENS. Drug dealers, gang members and sex offenders ALL OUT!” ($M_{\text{rating}} = 5.6$, $SD = 3.2$), the tweet with the least disagreement was “We're very close to having Pakis banned from America. Love it!” ($M_{\text{rating}} = 7.2$, $SD = 2.4$).

Popular condition. The grand mean was 5.97, 95% CI[4.7, 7.3]. *Symbolic racism*, *tweet aggressiveness*, social dominance orientation *egalitarianism*, *political orientation*, and *race* (Latinx = 6.26, 95% CI[5.5, 7], black = 5.56, 95% CI[4.8, 6.3], white = 5.52, 95% CI[4.77, 6.27]) were significant. Other regressors were not significant (**Table 3**).

The tweet with the most disagreement was “We have enough pedophiles and rapists in America, we don't need anymore.” ($M_{\text{rating}} = 5.7$, $SD = 3.2$), the tweet with the least disagreement was “We're very close to having Pakis banned from America. Love it!” ($M_{\text{rating}} = 7.1$, $SD = 2.5$).

Table 3. Estimated marginal means and slopes from racism ratings in Study 3.

Regressor	Coefficient	Unpopular Tweets			Coefficient	Popular Tweets		
		SE	95% CI	P		SE	95% CI	P
<i>Age</i>	-.001	.007	-.01, .01	.939	.0008	.007	-.01, .02	.906
<i>Twitter use</i>	-.11	.06	-.23, .02	.092	-.06	.07	-.19, .07	.327
<i>Social Dominance Orientation egalitarian</i>	-.27	.11	-.47, -.06	.011	-.21	.10	-.42, -.007	.042
<i>Social Dominance Orientation dominance</i>	.21	.11	.002, .42	.048	-.003	.11	-.21, .20	.977
<i>Blatant Dehumanization Outgroup</i>	.001	.005	-.009, .01	.794	-.007	.004	-.02, .002	.132

<i>Symbolic Racism</i>		-.22	.03	-.27, -.17	<.0001	-.19	.03	-.24, -.14	<.0001
<i>Political Orientation</i>		-.11	.07	-.25, .02	.094	-.19	.07	-.32, -.05	.007
<i>Race</i>									
	Latinx	6.61	.36	5.73, 7.50	-	6.26	.31	5.5, 7.01	-
	Black	-.65	.25	-1.25, -.05	.021	-.69	.27	-1.33, -.05	.015
	White	-.64	.26	-1.27, -.01	.021	-.74	.27	-1.39, -.08	.015
<i>Gender</i>									
	Woman	6.25	.37	5.35, 7.15	-	6.44	.35	5.55, 7.33	-
	Man	-.36	.23	-.92, -.19	.341	-.19	.24	-.84, .46	.661
	Nonbinary	-.68	.99	-3.09, 1.73	.745	1.92	1.8	-2.9, 6.7	.661
<i>Hashtag</i>									
	#MuslimBan	6.7	.39	5.86, 7.61	-	6.5	.33	5.7, 7.2	-
	#AmericaFirst	-.12	.22	-.56, .32	.583	-.19	.20	-.60, .21	.339
<i>Tweet Type</i>									
	Political	5.58	.35	4.79, 6.37	-	5.26	.29	4.59, 5.92	-
	Aggressive	1.03	.19	.65, 1.42	<.0001	.99	.18	.64, 1.35	<.0001
<i>Target group</i>									
	All immigrants	6.61	.37	5.69, 7.54	-	6.26	.31	5.47, 7.04	-
	Both	-.17	.65	-1.95, -1.61	.979	-.03	.59	-1.67, 1.61	.958
	Middle East	-.15	.27	-.88, .57	.979	.08	.24	-.58, .76	.958
	Undocumented	-.28	.28	-1.04, .47	.979	.31	.25	-.38, 1.01	.958

Beholder indices. Beholder indices in the unpopular condition were high, around .7 for *b1* (All = .73, Latinx = .76, black = .71, white = .71) and .9 for *b2* (All = .91, Latinx = .91, black = .85, white = .93). Beholder indices for the popular condition were also around .7 for *b1* (All = .75, Latinx = .79, black = .71, white = .71) and .9 for *b2* (All = .92, Latinx = .93, black = .87,

white = .92), suggesting evaluations were largely idiosyncratic in both conditions. Differences from Study 1 and 2 and between conditions were very small in magnitude and do not meaningfully change conclusions.

Figure 4. VPCs and beholder indices for the unpopular and popular tweet conditions. A) VPC estimates for the unpopular condition. B) The amount of variance contributed (x axis) by regressors (y axis). C) VPC estimates for the unpopular condition. D) The amount of variance contributed (x axis) by regressors (y axis). In plots A & C the columns represent different data partitions, in plots B & D they represent the difference variance components. Across the plots, red represents the perceiver component, blue the stimulus component, and purple the perceiver x stimulus interaction component. Error bars represent bootstrapped 95% confidence intervals.

Variance partitioning coefficients. The perceiver and perceiver x stimulus interaction

VPCs were relatively larger than the stimulus VPC in both conditions (**Figure 4A & 4C**). The small stimulus variance represents perceivers' disagreements in the rating of each tweet

(**Supplementary Figure S2**). There were minor differences between all the VPCs of the popular

and unpopular conditions compared to Study 1 and 2 and between the themselves

(**Supplementary Figure S5B-D**).

Variance contributions. Consistent with Study 1 and 2, both conditions show that the perceiver variance is largely explained by symbolic racism (26% - 30%) and the egalitarian dimension of social dominance orientation (7% - 11%). All measures together contributed 33% to 38% of the perceiver component. The aggressiveness of tweets also contributed the most stimulus variance (45%) and altogether the measures contributed 46% - 47% of the variance. None of the two-way interactions contribute more than 1% of the interaction component variance. However, it is notable that the interaction between SDOe or SDOd and tweet aggressiveness has consistently explained the largest interaction variance across all studies. A difference that arises from Study 1 and 2 is that the rest of the measures contribute less than 8% of the variance, including less contributions from political orientation (6% - 7%).

Discussion

The regression results in both models show that symbolic racism and the aggressiveness of tweets continue to be significant measures, such that greater endorsement of symbolic racism and more politically framed tweets on average led to less racist evaluations. This time, SDO-egalitarian showed more consistent significance than political orientation in predicting judgments of racism. Still, both of their effects were in the predicted direction - lower racism ratings.

The variance analyses showed minimal differences in agreement regarding judgments of racism between the popular and unpopular tweet conditions. In fact, they were consistent with each other and with estimates from the tweets without social engagement information (i.e., Study 1 and 2 combined). Evaluations are largely idiosyncratic due to large perceiver and interaction variances. Symbolic racism and race explained more of the between-perceiver variance and are

therefore implicated in increasing disagreement as reflected in the aggregate results. The contributions to perceiver x stimulus interaction variance remain elusive. The small amount of cross-cutting agreement was elicited by the aggressiveness of the tweets, which explained about half of the stimulus variance. Most of the meaningful variance once again remains unexplained. Overall, these results suggest that both social norm manipulations had minimal effects on disagreement levels such that they remained as high as they were in Study 1 and 2.

General Discussion

We sought to quantify agreement in lay judgments of racism of anti-immigrant statements, particularly in a political context with online Twitter discussions about policies that racialize immigrants: the travel ban on Middle Eastern countries and the wall along the Mexico-U.S. border. The context of these two policies were expected to align the perceptions of racial categories in The United States in ways that reflect varied motivational relationships to anti-immigrant sentiment. To study this, we introduced VCA as a novel method for assessing collective beliefs and representations. However, across three studies that compared judgments within and between racial categories and that varied cues regarding the social approval of the tweets, the results consistently showed more disagreement than agreement. One consistent source of agreement was within ideological cleavages: individuals who shared a political ideology or views on how black people are being treated in American society today (i.e., symbolic racism). The other source came from the aggressiveness of the tweets such that more aggressive tweets were rated as more racist.

We first discuss how different characteristics contributed to agreement and disagreement and the inferences we can draw from the results. Then we discuss the implications for the scientific investigation of shared beliefs in racial categories and contexts, the function of

agreement within ideological cleavages, implications for theorizing “racism”, and practical applications.

Shared and idiosyncratic contributions to judgments of racism

What factors influence individuals’ stances on racism? Based on the regression analyses, ideological cleavages were consistently significantly associated with differences in racism judgments, racial classifications less so. The scales that we used to capture ideological positions on race covered a range of different stances that individuals could take—from the more blatant (blatant dehumanization and social dominance) to the more coded and subtle (symbolic racism). Interestingly, we find that those who endorsed symbolic racism on average rated the tweets as less racist. Similarly, political orientation was also significantly associated with differences in racism judgments whereby conservatives judged the tweets as less racist on average than liberals. Aggressive tweets were also rated as more racist than the politically framed tweets, likely due to shared distinctions of racism as an explicitly aggressive phenomenon (Huddy & Feldman, 2009). Other variables (e.g., demographic variables, twitter use, targeted group) were either inconsistently or not significantly associated with racism judgments.

The previous paragraph is an example of the typical summary that follows from aggregate analyses that identify significant measures. Theoretical inferences usually focus on directional effects or lack of differences between means. Here, the VCA grants the opportunity to ground inferences in a deeper understanding of the measures. While we show that some variables influenced racism judgments in the directions theorized, focusing on how much variance they contributed paints a different picture of their impact. The perceiver variance explained about 50% of the total variance while symbolic racism independently explained about 24% to 30% and political orientation about 6% to 21% of that 50%. All perceiver characteristics together

explained only a third of the perceiver component (they are not additive due to collinearity), which is about 16% of the total variance. The stimulus variance explained less than 10% of the total variance and tweet aggressiveness consistently explained about 50% of that 10%, which is 5% of the total variance. While these variables reached statistical significance as regressors, the VCA reveals that a huge portion of meaningful variance remains unexplained. Contributions to about 66% of the perceiver variance (i.e., disagreement) and 50% of the stimulus variance (i.e., agreement) are consistently unknown. This result demonstrates why effect size, and not statistical significance, is a welcome emphasis of VCA.

Characteristics explaining the perceiver variance explain people's tendency to give more or less racism judgments, and characteristics explaining the stimulus variance explain why some tweets are judged as more or less racist than other tweets. We can then ask: why do *specific* individuals consistently judge *specific* tweets as more or less racist than others? Disagreement from perceiver x stimulus interactions was a consistently large factor in judgments of racism, around 20% of the total variance. For perspective, this is larger than the 16% total variance contribution from all perceiver characteristics. In an exploratory analysis where we interacted our perceiver and stimulus measures, the largest variance explained was less than 1% in all four studies and thus 99% of this variance is unexplained (**Figures 2-4**). Combining the components, around 80% of the total variance constitutes meaningful variance, 29% of it (23% of total variance) is explained by measures theorized to be strong influences on judgments of racism, 71% (57% of total variance) is unexplained. These relative patterns likely exist in many judgment studies, but without the proper analysis and study design to estimate components and contributions, they remain hidden and can lead to overclaiming and overtheorizing measures' impact.

Knowing the directional effects of variables is one way to inform theory. Knowing much they contribute to different variance components and how much variance is left unexplained is an underappreciated tool in theory building (Gantman et al., 2018; Yarkoni, 2019). In this case, one could try to find more important influences on individuals' stances on racism by explaining more perceiver or stimulus variance through additional measures of group characteristics (e.g., system justification; Jost & Banaji, 1994), unique characteristics (e.g., historical knowledge of racism; Nelson et al., 2013), design and measurement characteristics (e.g., differences in construal of or general use of scales; Paluck & Shafir, 2017) or hidden dimensions of the stimulus set that perceivers are detecting. For the interaction variance, a nomothetic approach provides a severely incomplete understanding of stable judgments arising from the dynamic between a specific perceiver and stimulus, suggesting a need to collect exhaustive measures with complex interactions (e.g., situational, attentional, construal, affect, social networks, knowledge) or to consider an idiographic approach that centers subjective experiences or unique identities.

The role of social norm perception in racism judgments

Popular tweets (potentially signaling social norms of greater approval) were rated on average as less racist compared to unpopular tweets. However, those ratings still showed consistent disagreement about the racist quality of tweets. Manipulating tweet engagement as a social norm cue did not bring judgments of racism into general alignment. We identify two possibilities for why this occurred.

The first involves the possibility that summary statistics of tweet engagement may not signal approval from a personally-relevant "referent group" that might color the participants' judgment of the tweet (Tankard & Paluck, 2016). In support of this interpretation are separate data showing people are more likely to share tweets from a salient in-group member (Rudat &

Buder, 2015). Twitter has even come to recognize that users seek to make sense of tweets by assessing who has shared what, and why. Recently they have added features (after this study's data collection) that provide information about whether sharing comes from one's own connections, one's extended network, and whether sharing involves a simple retweet or additional comments that imply critique.

Second and related, people may start with different assumptions about the more general norms of Twitter—specifically, whether the norms are racist or anti-racist. When aggressively anti-immigrant tweets receive a lot of engagement, this may serve as confirmatory evidence to some people that the tweet is racist (because of their impression of rampant racism on Twitter). However, widespread positive reception of same tweet may suggest to other people that the tweet is not as racist as their own personal impression, because they have relatively weaker impressions of Twitter as a racist or hostile environment. Thus, idiosyncratic construal of the relationship between different sources of norm information within the context of Twitter may facilitate disagreement in judgments.

What is a racial “group”? The role of evaluative agreement.

Despite a vast literature showing that perceptions of discrimination differ enormously, and that these differences can be predicted by individuals' racial classification (Carter & Murphy, 2015; Kraus et al., 2017; Norton & Sommers, 2011), we consistently find large amounts of disagreement about judgments of racism within racial “groups”. Why is this the case? Race and racial categories have been theorized as social constructions driven by shifting racial projects (Omi & Winant, 2015) and as complex multidimensional constructs (Roth, 2016) akin to a “bundle of sticks” (Sen & Wasow, 2016). Evaluative agreement about racism could potentially be a stick given research on the importance of consensus in topics consequential to

group identity (Haslam et al., 1999; Reicher et al., 1997). However, extracting a group's beliefs from their identity, called the identity-to-politics link, is a practice that has been criticized for overlooking many factors when treating individuals by a common category label (Lee, 2008). This is considered the *fallacy of groupism*, which is “a tendency to treat categories of people as if they were internally homogenous and externally bounded... and to take ethnic and racial groups as basic constituents of social life, chief protagonists of social conflict, and fundamental units of social analysis” (Brubaker, 2009).

Still, structural conditions in society assure that individuals classified as particular race categories within a hierarchical system share many common experiences, even across lines of class or gender—such as white advantages. For this reason, average differences in attitudes and judgments might be expected. This assertion is supported by much psychological evidence (Blumer, 1958; Bobo, 1999; Carter & Murphy, 2015; Kraus et al., 2017; Norton & Sommers, 2011; Sidanius & Pratto, 1999) and our results from Study 1 and 3. So why do our results also consistently highlight larger within-category disagreement? We suggest one reason is psychology's overwhelming focus on group averages, which proceeds at the expense of parsing within-group heterogeneity. Since aggregate group analyses are not very representative of individual-level data (Fisher et al., 2018), interrogating variance should become common practice when studying shared beliefs. Researchers that only use one approach could arrive at misleading conclusions regarding perceptions of racism.

Another reason is that investigations begin and end with “race” as the assumed protagonist. It is important from a normative and an empirical standpoint to acknowledge that racial categories: are unstably imposed on or adopted by people (Loveman, 1999; Sen & Wasow, 2016), analytically confound dimensions such as identity, classification, or ancestry (Fields,

2001; Roth, 2016; Sen & Wasow, 2016), are internally heterogeneous (Crenshaw, 1991; Sellers & Shelton, 2003; Wimmer, 2015; Zuckerman, 1990), and are produced and maintained *by* racism as its core ideology for manufacturing oppressive double standards (Fields & Fields, 2012; Golash-Boza, 2016). In using them as predictor variables to make “group” inferences, researchers *essentially* transform differences in means into races with different beliefs and behaviors. This analytic race-talk naturalizes “race” as an agentic entity that itself organizes beliefs and behaviors (Brubaker & Cooper, 2000; Cazenave, 2016), a logic that conjures the myth of “race” as causal or explanatory instead of as the *outcome* of racist practices (Fields & Fields, 2012; Zuberi, 2000). From this perspective, clarity would be better gained from placing race instead as a dependent variable (Benjamin, 2014). The VCA can also help researchers minimize race-talk by 1) forcing inferences about racial categories to foreground heterogeneity and 2) identifying more influential sources of agreement that may not conform to racial cleavages.

Functions of ideological disagreement about racism

While theories predict that a sense of racial group position (within material and historical social hierarchies) would factor more strongly in how individuals view the world (Blumer, 1958; Bobo, 1999; Bonilla-Silva, 1997), these data indicate that ideological cleavages can be more informative about whether individuals will agree on judgments of racism (though not comprehensively as indicated by the large unexplained variance). These results are consistent with a view of “racism” as a discursive battle site for definitions that are politically beneficial (Doane, 2006). In this view, racism as a concept is understood as a violation, so charges of “racism” are deployed or deflected in a strategic manner to manifest or defend personal or political interests.

A perceiver-focused account of ideology-based agreement could argue that using charges of “racism” in particular ways can signal or construct identity and group membership (Bucholtz & Hall, 2005). The act of labeling anti-immigrant statements as racist or not performs an ideological or partisan group membership, a way to share reality (Jost et al., 2008; Stern & Ondish, 2018). A stimulus-focused account could argue that the statements targeted immigrant groups whose status as a “race” of people who can experience “racism” is ambiguous (Roth, 2016). Both accounts allow political elites to influence racism perception by setting clear ideological stances to follow (Nelson et al., 2007), stances that represent a struggle for maintaining or upending existing racial structures and racist practices (Doane, 2006). Connecting these ideas to the current study, the tweets were made in response to immigration policies enacted by the Trump administration, an association that participants likely made when reading the statements. Perceivers’ responses could then signal their partisanship and stance on Donald Trump, who functions as a signifier of various politicized content domains (e.g., moral, religion, threat, harm) associated with inter-ideological disagreement (Kubin & Brandt, 2020). Anti-immigrant statements that are not tied to a particular political elite might by contrast elicit agreement about racism within different ideological cleavages or may not elicit agreement within ideological cleavages at all (although it is difficult to imagine statements about immigration that are not politicized or connected to a political elite; see **Footnote 3**).

Evaluative agreement is therefore not completely driven by individuals’ psychology or some characteristic inherent to the targeted group; political elites and projects can mobilize beliefs and discursive norms about groups and racism (Jost & Banaji, 1994; Paluck, 2012; Portice & Reicher, 2018). This point emphasizes the need for future work to continue monitoring how racism is socially understood across perceivers, stimulus sets, and contexts since belief

patterns can change with political zeitgeist. For example, perceivers during World War II showed *greater* consensus in labeling statements against people from other nations as “prejudiced” (Dyer, 1945).

Tensions in using “racism” as an analytical concept

Our results identify a tension between lay perceptions of racism (i.e., a highly politicized term with little agreement except within ideological camps) and the manner in which academics try to reach a working analytic definition based on objective material conditions. This tension has implications for racism scholarship: what should be the relationship between scholarly and folk epistemologies on racism? While striving for an academic solution to the question “what is racism?” can risk focusing efforts on taxonomy debates at the expense of political mobilization (Reed Jr, 2009), it can also mitigate issues associated with a flexible use of “racism” (e.g., collapsing systemic vs. interpersonal levels of analysis; Wacquant, 1997). Likewise, raising consciousness about the systemic scope of racism can challenge racist miseducation (Mills, 1997; Salter et al., 2017) and insisting on confrontational labels like “racism” can challenge power relations sustained by linguistic rules about (in)appropriate racial discourse (Cazenave, 2016).

Yet, academic discourse that does not engage with lay usage may become all the more siloed and politically unrecognizable (Banton, 1979). Like the mismatch that occurs when importing academic forms of racial categorization to a society whose folk experience rejects those labels (Loveman, 1999), defining racism independently of lay usage may limit reception of that definition. Learning about the way people use, accept, or reject terms heavily theorized by scholars can inform us about the phenomenon in question.

Practical implications

The political context that generated these tweets is ongoing and will continue to spur online public discourse, much of it racist. This discourse can have serious consequences: legitimizing and propagating power structures (Santa Ana, 1999; van Dijk, 1993), inciting violence against targets (Müller & Schwarz, 2018), or even spurring violent backlash from those targeted (Kteily & Bruneau, 2017). Idiosyncratic ideas about whether anti-immigrant statements are racist, complicates predictions about collective and even individual responses to these events. For example, the “reporting” function on many social media sites allows individual users to report content violating community standards. These standards are typically nonspecific—defined as standards against “bullying” or “harassment.” There are two sources of subjectivity in this process - individuals’ perception of the content and the moderators’ perception of what is considered to fit the general list of violations. This subjectivity has led to protests against the unprincipled manner in which media companies enforce their standards (Tobin et al., 2017). To the extent that there are high levels of disagreement in the judgment process about violations of normative standards, companies will need to implement concrete anti-racist norms, sanctions, or systems that account for these social dynamics (Crawford & Gillespie, 2016).

Conclusions

When it comes to labeling anti-immigrant statements as racist, there is more disagreement than agreement, even within racial categories. However, this does not mean that racism is not real, or that there is not a very real material structural racism that impacts the daily lives of minoritized individuals (Bonilla-Silva, 1997; Feagin, 2006). Consequently, their capacity for *knowing* the racism they experience should be respected, especially as shifting and accessible conceptualizations of racism provide opportunities for naming one’s experience and connecting it to others’ (Fricker, 2007). These results also do not suggest that perceivers are

interpreting the literal meaning of the tweet content in such idiosyncratic ways that linguistic communication about them impossible. Disagreement about judgments of racism in this context should be taken to mean two things: 1) that judgments can be shared among other social cleavages - in our data it is ideological divides and 2) definitions have a political function, especially when charges of racism (sometimes) hold power to socially sanction. Thus, labeling a statement as racist or not mainly reflects competing motivations and power struggles – e.g., defending or mobilizing against a discriminatory social environment or performing group membership.

Acknowledgments

We would like to thank Hannah Alarian, Jason Chin, Chelsey Clark, Ajua Duker, Robin Gomila, Michael Kraus, Aaron Kurosu, Nathan Matias, Vivek Nemana, John Henry Pezzuto, Roni Porat, Jordan Starck, Ting Qian, Sherry Wu, Janet Xu, Nechumi Yaffe, Yang-Yang Zhou, and the Intergroup Relations Workshop for their much appreciated help and advice.

References

- Abrajano, M., & Hajnal, Z. (2015). *White backlash: Immigration, race, and American politics*. Princeton University Press.
- Adams, G., Edkins, V., Lacka, D., Pickett, K. M., & Cheryan, S. (2008). Teaching About Racism: Pernicious Implications of the Standard Portrayal. *Basic and Applied Social Psychology*, 30(4), 349–361. <https://doi.org/10.1080/01973530802502309>
- Adams, G., Tormala, T. T., & O'Brien, L. T. (2006). The effect of self-affirmation on perception of racism. *Journal of Experimental Social Psychology*, 42(5), 616–626. <https://doi.org/10.1016/j.jesp.2005.11.001>
- Amatulli, J. (2017). *Americans respond to Trump's plans for country with #NoBanNoWall*. https://www.huffingtonpost.com/entry/trump-plans-nobannowall-hashtag_us_588a20f8e4b0024605fe3866
- Axt, J. R. (2017). The best way to measure explicit racial attitudes is to ask about them. *Social Psychological and Personality Science*, 194855061772899–194855061772899. <https://doi.org/10.1177/1948550617728995>
- Banton, M. (1979). Analytical and folk concepts of race and ethnicity. *Ethnic and Racial Studies*, 2(2), 127–138. <https://doi.org/10.1080/01419870.1979.9993258>
- Bates, D. M., Mächler, M., Bolker, B., & Walker, S. (2015). Fitting linear mixed-effects models using lme4. *Journal of Statistical Software*, 67(1), 1–48. <https://doi.org/10.18637/jss.v067.i01>
- Benjamin, R. (2014). Conjuring difference, concealing inequality: A brief tour of Racecraft: A Karen E. Fields and Barbara J. Fields, Racecraft: The Soul of Inequality in American Life. New York: Verso, 2012. *Theory and Society*, 43(6), 683–688. <https://doi.org/10.1007/s11186-014-9238-z>
- Berg, J. A. (2013). Opposition to pro-immigrant public policy: Symbolic racism and group threat. *Sociological Inquiry*, 83(1), 1–31. <https://doi.org/10.1111/j.1475-682x.2012.00437.x>
- Blauner, B. (1992). Talking past each other: Black and white languages of race. *American Prospect*, 56–60.

- Blumer, H. (1958). Race prejudice as a sense of group position. *The Pacific Sociological Review*, 1(1), 3–7. <https://doi.org/10.2307/1388607>
- Bobo, L. D. (1999). Prejudice as group position: Microfoundations of a sociological approach to racism and race relations. *Journal of Social Issues*, 55(3), 445–472. <https://doi.org/10.1111/0022-4537.00127>
- Bonilla-Silva, E. (1997). Rethinking racism: Toward a structural interpretation. *American Sociological Review*, 62(3), 465–465. <https://doi.org/10.2307/2657316>
- Bonilla-Silva, E. (2006). *Racism without racists: Color-blind racism and the persistence of racial inequality in the United States* (Second). Rowman & Littlefield.
- Brown, H. E. (2013). Race, legality, and the social policy consequences of anti-immigration mobilization. *American Sociological Review*, 78(2), 290–314. <https://doi.org/10.1177/0003122413476712>
- Brown, H. E., Jones, J. A., Hollingsworth, A., Lawrence, C., Reece, R., & Dame, N. (2018). *The Racialization of Latino Immigrants in New Destinations: Criminality, Ascription, and Countermobilization*. <https://doi.org/10.7758/RSF.2018.4.5.06>
- Brubaker, R. (2009). Ethnicity, race, and nationalism. *Annual Review of Sociology*, 35(1), 21–42. <https://doi.org/10.1146/annurev-soc-070308-115916>
- Brubaker, R., & Cooper, F. (2000). Beyond “identity.” *Theory and Society*, 29, 1–47.
- Bruneau, E., Kteily, N., & Laustsen, L. (2017). The unique effects of blatant dehumanization on attitudes and behavior towards Muslim refugees during the European ‘Refugee Crisis’ across four countries. *European Journal of Social Psychology*. <https://doi.org/10.1002/ejsp.2357>
- Bryk, A., & Raudenbush, S. (2002). *Hierarchical linear models: Applications and data analysis methods*.
- Bucholtz, M., & Hall, K. (2005). Identity and interaction: A sociocultural linguistic approach. *Discourse Studies*, 7(4–5), 585–614. <https://doi.org/10.1177/1461445605054407>
- Carter, E. R., & Murphy, M. C. (2015). Group-based differences in perceptions of racism: What counts, to whom, and why? *Social and Personality Psychology Compass*, 9(6), 269–280. <https://doi.org/10.1111/spc3.12181>

- Cazenave, N. A. (2016). *Conceptualizing racism: Breaking the chains of racially accommodative language*. Rowman & Littlefield.
- Chaudhry, I. (2015). # Hashtagging hate: Using Twitter to track racism online. *First Monday*, 1–13.
- Chaudhry, I. (2016). “Not so black and white”: Discussions of race on Twitter in the aftermath of #Ferguson and the shooting death of Mike Brown. *Cultural Studies ↔ Critical Methodologies*, 16(3), 296–304. <https://doi.org/10.1177/1532708616634814>
- Costello, K., & Hodson, G. (2011). Social dominance-based threat reactions to immigrants in need of assistance. *European Journal of Social Psychology*, 41(2), 220–231. <https://doi.org/10.1002/ejsp.769>
- Craig, M. A., & Richeson, J. A. (2014). Not in my backyard! Authoritarianism, social dominance orientation, and support for strict immigration policies at home and abroad. *Political Psychology*, 35(3), 417–429. <https://doi.org/10.1111/pops.12078>
- Craig, M. A., Rucker, J. M., & Richeson, J. A. (2018). Racial and political Dynamics of an Approaching “Majority-Minority” United States. *Annals of the American Academy of Political and Social Science*, 677(1), 204–214. <https://doi.org/10.1177/0002716218766269>
- Crandall, C. S., Miller, J. M., & White, M. H. (2018). Changing Norms Following the 2016 U.S. Presidential Election: The Trump Effect on Prejudice. *Social Psychological and Personality Science*, 9(2), 186–192. <https://doi.org/10.1177/1948550617750735>
- Crandall, C. S., & Stangor, C. (2005). Conformity and prejudice. In J. F. Dovidio, P. Glick, & L. A. Rudman (Eds.), *On the Nature of Prejudice: Fifty years after Allport* (pp. 295–309). Blackwell Publishers.
- Crawford, K., & Gillespie, T. (2016). What is a flag for? Social media reporting tools and the vocabulary of complaint. *New Media & Society*, 18(3), 410–428. <https://doi.org/10.1177/1461444814543163>
- Crenshaw, K. (1991). Mapping the margins: Intersectionality, identity politics, and violence against women of color. *Stanford Law Review*, 43(6), 1241–1299. <https://doi.org/10.2307/1229039>

- Devos, T., & Banaji, M. R. (2005). American = White? *Journal of Personality and Social Psychology*, 88(3), 447–466. <https://doi.org/10.1037/0022-3514.88.3.447>
- Doane, A. (2006). What is racism? Racial discourse and racial politics. *Critical Sociology*, 32(2–3), 255–274. <https://doi.org/10.1163/156916306777835303>
- Dyer, H. S. (1945). The usability of the concept of “prejudice.” *Psychometrika*, 10(3), 219–224.
- Echterhoff, G., Higgins, E. T., & Levine, J. M. (2009). Shared reality: Experiencing commonality with others’ inner states about the world. *Perspectives on Psychological Science*, 4(5), 496–521. <https://doi.org/10.1111/j.1745-6924.2009.01161.x>
- Esses, V. M., Veenvliet, S., Hodson, G., & Mihic, L. (2008). Justice, Morality, and the Dehumanization of Refugees. *Social Justice Research*, 21(1), 4–25. <https://doi.org/10.1007/s11211-007-0058-4>
- Feagin, J. R. (2006). *Systemic racism: A theory of oppression*. Taylor & Francis.
- Fields, B. J. (2001). Whiteness, Racism, and Identity. *International Labor and Working-Class History*, 60, 48–56.
- Fields, K. E., & Fields, B. J. (2012). *Racecraft: The soul of inequality in American life*. Verso.
- Fisher, A. J., Medaglia, J. D., & Jeronimus, B. F. (2018). Lack of group-to-individual generalizability is a threat to human subjects research. *Proceedings of the National Academy of Sciences*, 201711978–201711978. <https://doi.org/10.1073/pnas.1711978115>
- Flores, R. D., & Schachter, A. (2018). Who are the “Illegals”? The Social Construction of Illegality in the United States. *American Sociological Review*, 83(5), 839–868.
- Fricke, M. (2007). *Epistemic injustice: Power and the ethics of knowing*. Oxford University Press. <https://doi.org/10.1093/acprof:oso/9780198237907.001.0001>
- Gantman, A., Gomila, R., Martinez, J. E., Matias, J. N., Paluck, E. L., Starck, J., Wu, S., & Yaffe, N. (2018). A pragmatist philosophy of psychological science and its implications for replication. *Behavioral and Brain Sciences*, 41, e127–e127. <https://doi.org/10.1017/S0140525X18000626>
- Golash-Boza, T. (2016). A critical and comprehensive sociological theory of race and racism. *Sociology of Race and Ethnicity*, 2(2), 129–141. <https://doi.org/10.1177/2332649216632242>

- Goldstein, H., Browne, W., & Rasbash, J. (2002). Partitioning variation in multilevel models. *Understanding Statistics*, 1(4), 223–231. https://doi.org/10.1207/S15328031US0104_02
- González-Ramírez, A. (2017). *Why is Twitter promoting a hashtag from an anti-immigrant hate group?* <https://www.refinery29.com/2017/04/150702/twitter-build-the-wall-hashtag-federation-for-american-immigration-reform>
- Green, D. P., Glaser, J., & Rich, A. (1998). From Lynching to Gay Bashing: The Elusive Connection Between Economic Conditions and Hate Crime. *Journal of Personality & Social Psychology*, 75(1), 82–92.
- Hardin, C. D., & Higgins, T. E. (1996). Shared Reality: How social verification makes the subjective objective. In R. M. Sorrentino & T. E. Higgins (Eds.), *Handbook of motivation and cognition: The interpersonal context* (3rd ed., pp. 28–77). Guilford.
- Haslam, S. A., Oakes, P. J., Reynolds, K. J., & Turner, J. C. (1999). Consensus, social identity salience, and the emergence of stereotypes. *Personality and Social Psychology Bulletin*, 25(7), 809–818.
- Henry, P. J., & Sears, D. O. (2002). The symbolic racism 2000 scale. *Political Psychology*, 23(2), 253–283. <https://doi.org/10.1111/0162-895X.00281>
- Ho, A. K., Sidanius, J., Kteily, N., Sheehy-Skeffington, J., Pratto, F., Henkel, K. E., Foels, R., & Stewart, A. L. (2015). The nature of social dominance orientation: Theorizing and measuring preferences for intergroup inequality using the new SDO₇ scale. *Journal of Personality and Social Psychology*, 109(6), 1003–1028. <https://doi.org/10.1037/pspi0000033>
- Hönekopp, J. (2006). Once more: Is beauty in the eye of the beholder? Relative contributions of private and shared taste to judgments of facial attractiveness. *Journal of Experimental Psychology. Human Perception and Performance*, 32(2), 199–209. <https://doi.org/10.1037/0096-1523.32.2.199>
- Hox, J. J. (2010). *Multilevel Analysis: Techniques and Applications*. <https://doi.org/10.1002/1521-3773>
- Hoyt, C. (2012). The pedagogy of the meaning of racism: Reconciling a discordant discourse. *Social Work*, 57(3), 225–234. <https://doi.org/10.1093/sw/sws009>

- Huddy, L., & Feldman, S. (2009). On assessing the political effects of racial prejudice. *Annual Review of Political Science*, 12(1), 423–447. <https://doi.org/10.1146/annurev.polisci.11.062906.070752>
- Hudson, G., & Esses, V. M. (2005). Lay perceptions of ethnic prejudice: Causes, solutions, and individual differences. *European Journal of Social Psychology*, 35(3), 329–344.
<https://doi.org/10.1002/ejsp.251>
- Inda, J. X. (2006). *Targeting immigrants: Government, technology, and ethics*. Blackwell Pub.
- Jost, J. T., & Banaji, M. R. (1994). The role of stereotyping in system-justification and the production of false consciousness. *British Journal of Social Psychology*, 33, 1–27.
<https://doi.org/10.1111/j.2044-8309.1994.tb01008.x>
- Jost, J. T., Glaser, J., Kruglanski, A. W., & Sulloway, F. J. (2003). Political conservatism as motivated social cognition. *Psychological Bulletin*, 129(3), 339–375. <https://doi.org/10.1037/0033-2909.129.3.339>
- Jost, J. T., Ledgerwood, A., & Hardin, C. D. (2008). Shared reality, system justification, and the relational basis of ideological beliefs. *Social and Personality Psychology Compass*, 2(1), 171–186.
<https://doi.org/10.1111/j.1751-9004.2007.00056.x>
- Judd, C. M., Westfall, J., & Kenny, D. A. (2012). Treating stimuli as a random factor in social psychology: A new and comprehensive solution to a pervasive but largely ignored problem. *Journal of Personality and Social Psychology*, 103(1), 54–69. <https://doi.org/10.1037/a0028347>
- Kaiser, C. R., & Wilkins, C. L. (2010). Group identification and prejudice: Theoretical and empirical advances and implications. *Journal of Social Issues*, 66(3), 461–476.
<https://doi.org/10.1111/j.1540-4560.2010.01656.x>
- Kaufmann, E. (2019). Ethno-traditional nationalism and the challenge of immigration. *Nations and Nationalism*, 25(2), 435–448. <https://doi.org/10.1111/nana.12516>
- Kenny, D. A. (1994). *Interpersonal Perception*. The Guilford Press.
- Knowles, J. E., & Frederick, C. (2016). *merTools: Tools for Analyzing Mixed Effect Regression Models* (0.3.0) [Computer software]. <https://CRAN.R-project.org/package=merTools>

- Kraus, M. W., Rucker, J. M., & Richeson, J. A. (2017). Americans misperceive racial economic equality. *Proceedings of the National Academy of Sciences*, 201707719–201707719.
<https://doi.org/10.1073/pnas.1707719114>
- Kteily, N., & Bruneau, E. (2017). Backlash: The politics and real-world consequences of minority group dehumanization. *Personality and Social Psychology Bulletin*, 43(1), 87–104.
<https://doi.org/10.1177/0146167216675334>
- Kteily, N., Bruneau, E., Waytz, A., Cotterill, S., Bruneau, E., & Waytz, A. (2015). The ascent of man: Theoretical and empirical evidence for blatant dehumanization. *Journal of Personality & Social Psychology*, 108(7), 1–31.
- Kubin, E., & Brandt, M. J. (2020). Identifying the domains of ideological similarities and differences in attitudes. *Comprehensive Results in Social Psychology*, 4(1), 53–77.
<https://doi.org/10.1080/23743603.2020.1756242>
- Leder, H., Goller, J., Rigotti, T., & Forster, M. (2016). Private and shared taste in art and face appreciation. *Frontiers in Human Neuroscience*, 10, 1–7.
<https://doi.org/10.3389/fnhum.2016.00155>
- Lee, T. (2008). Race, immigration, and the identity-to-politics link. *Annual Review of Political Science*, 11(1), 457–478. <https://doi.org/10.1146/annurev.polisci.11.051707.122615>
- Lenth, R. V. (2014). *Using lsmeans*. <https://cran.r-project.org/web/packages/lsmeans/vignettes/using-lsmeans.pdf>
- Liao, H.-Y., Hong, Y., & Rounds, J. (2016). Perception of Subtle Racism: The Role of Group Status and Legitimizing Ideologies. *The Counseling Psychologist*, 44(2), 237–266.
<https://doi.org/10.1177/0011000015625329>
- Loveman, M. (1999). Is “Race” Essential? *American Sociological Review*, 64(6), 891–898.
<https://doi.org/10.2307/2657409>
- Martinez, J. E., Feldman, L. A., & Cikara, M. (2017). Showcasing immigrant excellence. *Scientific American*. <https://blogs.scientificamerican.com/guest-blog/showcasing-immigrant-excellence/>

- Martinez, J. E., Funk, F., & Todorov, A. (2020). Quantifying idiosyncratic and shared contributions to judgment. *Behavior Research Methods*. <https://doi.org/10.3758/s13428-019-01323-0>
- Mills, C. W. (1997). *The racial contract*. Cornell University Press.
- Moscovici, S. (2000). The phenomenon of social representations. In S. Moscovici & G. Duveen (Eds.), *Social representations: Explorations in social psychology* (pp. 18–77). Wiley.
- Müller, K., & Schwarz, C. (2018). Making America hate again? Twitter and hate crime under Trump. *Social Science Research Network*. <https://doi.org/10.2139/ssrn.3149103>
- Nelson, J. C., Adams, G., & Salter, P. S. (2013). The Marley Hypothesis: Denial of Racism Reflects Ignorance of History. *Psychological Science*, 24(2), 213–218.
<https://doi.org/10.1177/0956797612451466>
- Nelson, T. E., Sanbonmatsu, K., & McClerking, H. K. (2007). Playing a different race card: Examining the limits of elite influence on perceptions of racism. *Journal of Politics*, 69(2), 416–429.
<https://doi.org/10.1111/j.1468-2508.2007.00540.x>
- Ngai, M. M. (2004). *Impossible subjects: Illegal aliens and the making of modern America*.
<http://public.eblib.com/choice/publicfullrecord.aspx?p=1584940>
- Norton, M. I., & Sommers, S. R. (2011). Whites See Racism as a Zero-Sum Game That They Are Now Losing. *Perspectives on Psychological Science*, 6(3), 215–218.
<https://doi.org/10.1177/1745691611406922>
- Omi, M., & Winant, H. (2015). *Racial formation in the United States* (Third Edition). Routledge.
- Ondish, P., & Stern, C. (2017). Liberals possess more national consensus on political attitudes in the United States: An examination across 40 years. *Social Psychological and Personality Science*.
<https://doi.org/10.1177/1948550617729410>
- Operario, D., & Fiske, S. T. (2001). Ethnic identity moderates perceptions of prejudice: Judgments of personal versus group discrimination and subtle versus blatant bias. *Personality and Social Psychology Bulletin*, 27(5), 550–561. <https://doi.org/10.1177/0146167201275004>

- Paluck, E. L. (2012). The dominance of the individual in intergroup relations research: Understanding social change requires psychological theories of collective and structural phenomena. *The Behavioral and Brain Sciences*, 35, 443–444. <https://doi.org/10.1017/S0140525X12001288>
- Paluck, E. L., & Shafir, E. (2017). The psychology of construal in the design of field experiments. In *Handbook of Field Experiments* (pp. 245–268). <https://doi.org/10.1016/bs.hefe.2016.12.001>
- Pew Research Center. (2019). *Public highly critical of state of political discourse in the U.S.* <https://www.people-press.org/2019/06/19/public-highly-critical-of-state-of-political-discourse-in-the-u-s/>
- Portice, J., & Reicher, S. (2018). Arguments for European disintegration: A mobilization analysis of anti-immigration speeches by U.K. political leaders. *Political Psychology*, 39(6), 1357–1372. <https://doi.org/10.1111/pops.12551>
- Pratto, F., Stewart, A. L., & Bou Zeineddine, F. (2013). When inequality fails: Power, group dominance, and societal change. *Journal of Social and Political Psychology*, 1(1), 132–160. <https://doi.org/10.5964/jspp.v1i1.97>
- Provine, D. M., & Doty, R. L. (2011). The criminalization of immigrants as a racial project. *Journal of Contemporary Criminal Justice*, 27(3), 261–277. <https://doi.org/10.1177/1043986211412559>
- Reed Jr, A. (2009). The limits of anti-racism. *Left Business Observer*, 121, 6.
- Reicher, S., Hopkins, N., & Condor, S. (1997). Stereotype construction as a strategy of influence. In R. Spears, P. J. Oakes, N. Ellemers, & A. S. Haslam (Eds.), *The Social Psychology of Stereotyping and Group Life* (pp. 94–118). Blackwell Publishers. [https://doi.org/10.1016/0045-7825\(81\)90049-9](https://doi.org/10.1016/0045-7825(81)90049-9)
- Richeson, J. A., & Sommers, S. R. (2016). Toward a social psychology of race and race relations for the twenty-first century. *Annual Review of Psychology*, 67, 439–463. <https://doi.org/10.1146/annurev-psych-010213-115115>
- Rivera, C. (2014). The brown threat: Post-9/11 confluences of Latina/os and Middle Eastern Muslims in the US American imagination. *Latino Studies*, 12(1), 44–64. <https://doi.org/10.1057/lst.2014.6>

- Roth, W. D. (2016). The multiple dimensions of race. *Ethnic and Racial Studies*, 39(8), 1310–1338.
<https://doi.org/10.1080/01419870.2016.1140793>
- Rudat, A., & Buder, J. (2015). Making retweeting social: The influence of content and context information on sharing news in Twitter. *Computers in Human Behavior*, 46, 75–84.
<https://doi.org/10.1016/j.chb.2015.01.005>
- Salter, P. S., Adams, G., & Perez, M. J. (2017). Racism in the structure of everyday worlds: A cultural-psychological perspective. *Current Directions in Psychological Science*, 096372141772423–096372141772423. <https://doi.org/10.1177/0963721417724239>
- Santa Ana, O. (1997). Empirical analysis of anti-immigrant metaphor in political discourse. *University of Pennsylvania Working Papers in Linguistics*, 4(1), 317–330.
- Santa Ana, O. (1999). 'Like an animal I was treated': Anti-immigrant metaphor in US public discourse. *Discourse & Society*, 10(2), 191–224. <https://doi.org/10.1177/0957926599010002004>
- Schaffner, B. F. (2018). *Follow the racist? The consequences of expressions of elite prejudice for mass rhetoric*.
- Sears, D. O., & Henry, P. J. (2003). The origins of symbolic racism. *Journal of Personality and Social Psychology*, 85(2), 259–275. <https://doi.org/10.1037/0022-3514.85.2.259>
- Sears, D. O., & Henry, P. J. (2005). Over thirty years later: A contemporary look at symbolic racism. In *Advances in experimental social psychology* (Vol. 37, pp. 95–150). Elsevier Academic Press.
<http://linkinghub.elsevier.com/retrieve/pii/S006526010537002X>
- Sechrist, G. B., & Stangor, C. (2001). Perceived consensus influences intergroup behavior and stereotype accessibility. *Journal of Personality & Social Psychology*, 80(4), 645–654.
<https://doi.org/10.1037//0022-3514.80.4.645>
- Sellers, R. M., & Shelton, J. N. (2003). The role of racial identity in perceived racial discrimination. *Journal of Personality and Social Psychology*, 84(5), 1079–1092. <https://doi.org/10.1037/0022-3514.84.5.1079>

- Sen, M., & Wasow, O. (2016). Race as a bundle of sticks: Designs that estimate effects of seemingly immutable characteristics. *Annual Review of Political Science*, 19(1), 499–522.
<https://doi.org/10.1146/annurev-polisci-032015-010015>
- Shavelson, R. J., Webb, N. M., & Rowley, G. L. (1989). Generalizability theory. *American Psychologist*, 6, 922–932.
- Shoukri, M. M. (2011). *Measures of Interobserver Agreement and Reliability* (2nd ed., Vol. 20030998). CRC Press.
- Sidanius, J., & Pratto, F. (1999). *Social dominance: An intergroup theory of social hierarchy and oppression* (1st ed.). Cambridge University Press.
- Sides, J., Tesler, M., & Vavreck, L. (2019). *Identity Crisis: The 2016 Presidential Campaign and the Battle for the Meaning of America*.
<https://www.degruyter.com/doi/book/10.23943/9781400888306>
- Sniderman, P. M. ., Piazza, T., Tetlock, P. E. ., & Kendrick, A. (1991). The new racism. *American Journal of Political Science*, 35(2), 423–447. <https://doi.org/10.2307/2111369>
- Sommers, S. R., & Norton, M. I. (2006). Lay theories about white racists: What constitutes racism (and what doesn't). *Group Processes and Intergroup Relations*, 9(1), 117–138.
<https://doi.org/10.1177/1368430206059881>
- Stangor, C., Sechrist, G. B., & Jost, J. T. (2001). Changing racial beliefs by providing consensus information. *Personality and Social Psychology Bulletin*, 27(4), 486–496.
<https://doi.org/10.1177/0146167201274009>
- Stern, C., & Ondish, P. (2018). Political aspects of shared reality. *Current Opinion in Psychology*, 23, 11–14. <https://doi.org/10.1016/j.copsyc.2017.11.004>
- Stern, C., West, T. V., Jost, J. T., & Rule, N. O. (2014). “Ditto heads”: Do conservatives perceive greater consensus within their ranks than liberals? *Personality and Social Psychology Bulletin*, 40(9), 1162–1177. <https://doi.org/10.1177/0146167214537834>

- Tankard, M. E., & Paluck, E. L. (2016). Norm perception as a vehicle for social change. *Social Issues and Policy Review*, 10(1), 181–211. <https://doi.org/10.1111/sipr.12022>
- Tobin, A., Varner, M., & Angwin, J. (2017). *Facebook's uneven enforcement of hate speech rules allows vile posts to stay up*. ProPublica. <https://www.propublica.org/article/facebook-enforcement-hate-speech-rules-mistakes>
- Valentino, N. A., Neuner, F. G., & Vandenbroek, L. M. (2018). The Changing Norms of Racial Political Rhetoric and the End of Racial Priming. *The Journal of Politics*, 80(3), 757–771. <https://doi.org/10.1086/694845>
- van Dijk, T. A. (1993). *Elite discourse and racism*. Sage Publications.
- Vargas, N. (2017). Ideological whitening: Does skintone shape colour-blind ideology adherence for Latina/os? *Ethnic and Racial Studies*, 0(0), 1–19. <https://doi.org/10.1080/01419870.2017.1394478>
- Wacquant, L. (1997). For an analytic of racial domination. *Political Power and Social Theory*, 11, 221–234.
- Warner, J. A. (2005). The social construction of the criminal alien in immigration law, enforcement practice and statistical enumeration: Consequences for immigrant stereotyping. *Journal of Social and Ecological Boundaries*, 1.2, 56–80.
- Westfall, J., Kenny, D. A., & Judd, C. M. (2014). Statistical power and optimal design in experiments in which samples of participants respond to samples of stimuli. *Journal of Experimental Psychology: General*, 143(5), 2020–2045. <https://doi.org/10.1037/xge0000014>
- Williams, A., Oliver, C., Aumer, K., & Meyers, C. (2016). Racial microaggressions and perceptions of Internet memes. *Computers in Human Behavior*, 63, 424–432. <https://doi.org/10.1016/j.chb.2016.05.067>
- Wimmer, A. (2015). Race-centrism: A critique and a research agenda. *Ethnic and Racial Studies*, 38(13), 2186–2205. <https://doi.org/10.1080/01419870.2015.1058510>

- Xie, S. Y., Flake, J. K., & Hehman, E. (2018). Perceiver and target characteristics contribute to impression formation differently across race and gender. *Journal of Personality and Social Psychology*. <https://doi.org/10.1037/pspi0000160>
- Yarkoni, T. (2019). The generalizability crisis. *PsyArxiv*. <https://psyarxiv.com/jqw35>
- Zuberi, T. (2000). Deracializing social statistics: Problems in the quantification of race. *The ANNALS of the American Academy of Political and Social Science*, 568(1), 172–185.
<https://doi.org/10.1177/000271620056800113>
- Zuckerman, M. (1990). Some dubious premises in research and theory on racial differences. *American Psychologist*, 45(12), 1297–1303.

Quantifying shared and idiosyncratic judgments of racism in social discourse: Supplementary Material

Table of Contents:

- Appendix
 - **Appendix A:** List of deviations from pre-registration
 - **Appendix B:** Additional information on variance components and interpretation
- Supplementary
 - **Supplementary Table S1.** Demographic characteristics of Study 1 and Study 2 by ethno-racial group.
 - **Supplementary Table S2.** Tweet stimuli.
 - **Supplementary Table S3.** Demographic characteristics of Study 3, unpopular and popular condition, by ethno-racial group.
 - **Supplementary Figure S1.** Correlation matrices of participant factors in Study 1.
 - **Supplementary Figure S2.** Range of the tweet stimuli's racism across studies.
 - **Supplementary Figure S3.** Variance components for each study.
 - **Supplementary Figure S4.** Correlation matrices of participant characteristics in Study 2.
 - **Supplementary Figure S5.** VPC differences across studies and conditions.
 - **Supplementary Figure S6.** Correlation matrices of participant factors in Study 3, unpopular condition.
 - **Supplementary Figure S7.** Correlation matrices of participant factors in Study 3, popular condition.

Appendix

Appendix A. Deviations from pre-registration

- We pre-registered a *minimum* of 100 participants from each ethno-racial group for each study (N = 300 total each study).
 - While our final samples were close to our pre-registration goal, some of our samples of ethno-racial groups fell slightly below 100 (99 Black participants in Study 2, 98 Latinx participants in Study 3, 97 White participants in Study 3). These deviations were due to the test-retest exclusion criteria (see next point).
- We pre-registered that we would keep participants whose test-retest ICC was below 0 and run sensitivity analyses.
 - It was learned after the pre-registration that participants with a negative test-retest reliability contribute no meaningful variance to the data. These kind of data are uninterpretable in terms of what the ratings represent (e.g., inattention, speeding through task, or not wanting to provide honest responses to a sensitive topic) and importantly amount to no contributions to any of the critical variance components except the residuals, and thus we decided it is better measurement-wise to exclude these data. This exclusion is common in judgment studies.
- We pre-registered that we would run analyses that would pair up groups to identify whether there were coalitions of shared judgments across racial groups (e.g., Black-Latinx, White-Black, Latinx-White).
 - These estimated merely ended up being average effects of their constitutive components, so we do not report them. For example, if the perceiver variance of Black participants was 10 and White participants 2, the perceiver variance when combining Black and White participants would be 6, which provides no new or emergent insights. However, they were left in the analysis scripts on the OSF archive for interested readers.
- Did not test manipulating the identity of the tweeter nor providing interpersonal vs. institutional definitions of racism before the judgment task.
 - We lacked resources and decided to save this idea for a future project.
- Replication (Study 2) is not listed in pre-registration
 - This was an error. Study 1 and 2 were supposed to ask about grammatical correctness and racism of the tweets, respectively, but through accident in the coding of the experiment, we ended up asking about racism twice, providing the opportunity to report a direct replication.

Appendix B. Additional information on variance components and interpretation

In random-effects models, the VPC is equal to the intraclass correlation and represents the proportion of variance explained by a particular component. These indices are useful as they isolate where the relatively largest variance is found: in the participants, stimuli, or their interaction. Specifically, if the stimulus VPC is the largest, one can infer both that participants shared evaluations and that the stimulus range varied. A smaller stimulus VPC *could* indicate more idiosyncratic evaluations suggesting that general perceiver differences or their ranking differences are the underlying mechanism, respectively. However, it may also suggest that the stimuli were more homogenous in their range (i.e., everyone rated all the stimuli the same), which creates ambiguity in interpretation as this would actually reflect agreement. Directly checking the cause of a small stimulus variance is important because it can result from disagreement *or* agreement. In the case of disagreement about the mean of each stimulus, each stimulus is rated across the full scale which compresses between-stimulus averages towards the middle of the scale. This kind of pattern will be reflected in larger perceiver and/or perceiver x stimulus variance components. In the case of agreement, each stimulus is rated within a small range of the scale thus their averages will be similar. This kind of pattern would also show small perceiver or interaction components since there would be very little variance in the judgments overall.

We also computed block-related variance. The block component captures drift in the average stimulus rating across the three repetitions. For example, perceivers may habituate to the stimuli after the first block, shifting their responses on the scale. The block x perceiver component represents the possibility that some perceivers may show more drift in their average across the blocks than others. The block x stimulus component also captures whether average responses to certain stimuli shift more than others across the three blocks. However, given that the blocks were experienced in succession, block-related variance components were small (see **Supplementary Figure S3**).

Supplementary

Supplementary Table S1. Demographic characteristics of Study 1 and Study 2 by ethno-racial group.

	Study 1			Study 2		
	Black (N=103)	Latino (N=103)	White (N=100)	Black (N=99)	Latino (N=101)	White (N=101)
Age						
<i>Mean (SD)</i>	0.121 (15.3)	-5.00 (14.9)	7.08 (15.8)	-0.294 (14.5)	-4.92 (14.5)	6.98 (15.4)
<i>Median [Min, Max]</i>	-0.597 [- 24.6, 37.4]	-7.60 [- 24.6, 37.4]	9.90 [- 22.6, 39.4]	-0.597 [- 24.6, 34.4]	-7.60 [- 23.6, 27.4]	7.40 [- 21.6, 38.4]
Gender						
Don't Say	2 (1.9%)	1 (1.0%)	0 (0%)	0 (0%)	2 (2.0%)	0 (0%)
Female	24 (23.3%)	37 (35.9%)	23 (23.0%)	27 (27.3%)	30 (29.7%)	31 (30.7%)
Male	77 (74.8%)	64 (62.1%)	77 (77.0%)	72 (72.7%)	69 (68.3%)	70 (69.3%)
Non-binary	0 (0%)	1 (1.0%)	0 (0%)	0 (0%)	0 (0%)	0 (0%)
Twitter Use						
Daily	32 (31.1%)	24 (23.3%)	8 (8.0%)	26 (26.3%)	33 (32.7%)	25 (24.8%)
Less Often	11 (10.7%)	14 (13.6%)	12 (12.0%)	10 (10.1%)	13 (12.9%)	13 (12.9%)
Never	42 (40.8%)	46 (44.7%)	64 (64.0%)	41 (41.4%)	31 (30.7%)	45 (44.6%)
Weekly	7 (6.8%)	7 (6.8%)	7 (7.0%)	7 (7.1%)	11 (10.9%)	6 (5.9%)
X a Week	9 (8.7%)	11 (10.7%)	9 (9.0%)	15 (15.2%)	11 (10.9%)	12 (11.9%)
Missing	2 (1.9%)	1 (1.0%)	0 (0%)	0 (0%)	2 (2.0%)	0 (0%)
Political Orientation						
<i>Mean (SD)</i>	3.48 (1.49)	3.83 (1.58)	4.15 (1.61)	3.46 (1.34)	3.49 (1.57)	4.09 (1.71)
<i>Median [Min, Max]</i>	4 [1, 7]	4 [1, 7]	4 [1, 7]	4 [1, 6]	3.5 [1, 7]	4 [1, 7]
Response : 1	11 (10.7%)	10 (9.7%)	8 (8.0%)	13 (13.1%)	15 (14.9%)	10 (9.9%)
2	24 (23.3%)	15 (14.6%)	11 (11.0%)	18 (18.2%)	19 (18.8%)	13 (12.9%)
3	12 (11.7%)	14 (13.6%)	12 (12.0%)	10 (10.1%)	17 (16.8%)	12 (11.9%)

	Study 1			Study 2		
	Black (N=103)	Latino (N=103)	White (N=100)	Black (N=99)	Latino (N=101)	White (N=101)
4	38 (36.9%)	37 (35.9%)	36 (36.0%)	44 (44.4%)	28 (27.7%)	28 (27.7%)
5	9 (8.7%)	10 (9.7%)	8 (8.0%)	8 (8.1%)	8 (7.9%)	17 (16.8%)
6	5 (4.9%)	13 (12.6%)	19 (19.0%)	6 (6.1%)	13 (12.9%)	13 (12.9%)
7	4 (3.9%)	4 (3.9%)	6 (6.0%)	0 (0%)	1 (1.0%)	8 (7.9%)
Symbolic Racism						
<i>Mean (SD)</i>	15.2 (4.66)	17.7 (4.52)	19.4 (4.48)	15.5 (3.90)	17.0 (4.78)	19.1 (5.10)
<i>Median [Min, Max]</i>	15 [9, 27]	18 [8, 27]	20 [9, 29]	16 [8, 24]	17 [8, 29]	19 [8, 29]
Score range: [8,13.8]	42 (40.8%)	18 (17.5%)	12 (12.0%)	32 (32.3%)	26 (25.7%)	19 (18.8%)
[13.8,19.5]	42 (40.8%)	44 (42.7%)	33 (33.0%)	49 (49.5%)	43 (42.6%)	33 (32.7%)
[19.5,25.2]	17 (16.5%)	37 (35.9%)	50 (50.0%)	18 (18.2%)	28 (27.7%)	37 (36.6%)
[25.2,31]	2 (1.9%)	4 (3.9%)	5 (5.0%)	0 (0%)	4 (4.0%)	12 (11.9%)
Blatant Dehumanization - Outgroup						
<i>Mean (SD)</i>	12.0 (24.2)	5.01 (15.5)	9.13 (19.3)	1.85 (24.7)	7.67 (14.0)	7.06 (19.1)
<i>Median [Min, Max]</i>	3 [-50.7, 100]	.33 [-49, 74.3]	0 [-47, 100]	0.33 [-100, 99.0]	1.67 [-41, 50.7]	0 [-68.3, 66.3]
Score range: [-100,-50]	1 (1.0%)	0 (0%)	0 (0%)	3 (3.0%)	0 (0%)	2 (2.0%)
[-50,0]	20 (19.4%)	20 (19.4%)	21 (21.0%)	26 (26.3%)	17 (16.8%)	20 (19.8%)
[0,50]	75 (72.8%)	82 (79.6%)	74 (74.0%)	67 (67.7%)	83 (82.2%)	75 (74.3%)
[50,100]	7 (6.8%)	1 (1.0%)	5 (5.0%)	3 (3.0%)	1 (1.0%)	4 (4.0%)
Blatant Dehumanization - White						
<i>Mean (SD)</i>	14.5 (33.9)	3.82 (22.4)	-	-2.04 (35.7)	10.9 (29.8)	-
<i>Median [Min, Max]</i>	2 [-64, 100]	0 [-47, 100]	-	0 [-100, 100]	0 [-74, 100]	-
Score range: [-100,-50]	3 (2.9%)	0 (0%)	-	8 (8.1%)	1 (1.0%)	-

	Study 1			Study 2		
	Black (N=103)	Latino (N=103)	White (N=100)	Black (N=99)	Latino (N=101)	White (N=101)
[-50,0]	19 (18.4%)	30 (29.1%)	-	25 (25.3%)	24 (23.8%)	-
[0,50]	64 (62.1%)	70 (68.0%)	-	60 (60.6%)	64 (63.4%)	-
[50,100]	17 (16.5%)	3 (2.9%)	-	6 (6.1%)	12 (11.9%)	-
Blatant Dehumanization - Latinx						
<i>Mean (SD)</i>	8.46 (23.8)	-	7.07 (19.8)	1.87 (23.1)	-	5.22 (17.9)
<i>Median [Min, Max]</i>	0 [-55, 100]	-	0 [-63, 100]	0 [-100, 99]	-	0 [-78, 83]
Score range: [-100,-50]	1 (1.0%)	-	1 (1.0%)	2 (2.0%)	-	2 (2.0%)
[-50,0)	20 (19.4%)	-	25 (25.0%)	25 (25.3%)	-	19 (18.8%)
[0,50)	75 (72.8%)	-	68 (68.0%)	69 (69.7%)	-	77 (76.2%)
[50,100]	7 (6.8%)	-	6 (6.0%)	3 (3.0%)	-	3 (3.0%)
Blatant Dehumanization - Black						
<i>Mean (SD)</i>	-	3.49 (18.8)	6.79 (18.1)	-	3.40 (12.8)	6.58 (18.6)
<i>Median [Min, Max]</i>	-	0 [-78, 96]	0 [-26, 100]	-	0 [-31, 49]	0 [-66, 66]
Score range: [-100,-50]	-	1 (1.0%)	0 (0%)	-	0 (0%)	1 (1.0%)
[-50,0]	-	17 (16.5%)	22 (22.0%)	-	16 (15.8%)	16 (15.8%)
[0,50]	-	82 (79.6%)	75 (75.0%)	-	85 (84.2%)	80 (79.2%)
[50,100]	-	3 (2.9%)	3 (3.0%)	-	0 (0%)	4 (4.0%)
Blatant Dehumanization - Muslim						
<i>Mean (SD)</i>	13.0 (26.1)	7.70 (21.6)	13.5 (25.7)	5.73 (25.7)	8.72 (17.0)	9.38 (26.3)
<i>Median [Min, Max]</i>	2 [-33, 100]	0 [-51, 86]	0 [-52, 100]	0 [-100, 100]	0 [-19, 74.0]	0 [-78, 100]
Score range: [-100,-50]	0 (0%)	1 (1.0%)	1 (1.0%)	3 (3.0%)	0 (0%)	2 (2.0%)
[-50,0]	22 (21.4%)	16 (15.5%)	15 (15.0%)	20 (20.2%)	16 (15.8%)	22 (21.8%)

	Study 1			Study 2		
	Black (N=103)	Latino (N=103)	White (N=100)	Black (N=99)	Latino (N=101)	White (N=101)
[0,50]	71 (68.9%)	80 (77.7%)	72 (72.0%)	71 (71.7%)	80 (79.2%)	68 (67.3%)
[50,100]	10 (9.7%)	6 (5.8%)	12 (12.0%)	5 (5.1%)	5 (5.0%)	9 (8.9%)
Social Orientation - Egalitarian						
<i>Mean (SD)</i>	2.39 (1.24)	2.64 (1.25)	3.12 (1.21)	2.52 (1.19)	2.56 (1.34)	2.91 (1.19)
<i>Median [Min, Max]</i>	2.25 [1., 5.5]	2.5 [1, 4.75]	3.25 [1, 6]	2.5 [1, 5.5]	2.5 [1, 5.5]	3 [1, 5.25]
Response Option: 1	47 (45.6%)	36 (35.0%)	19 (19.0%)	36 (36.4%)	41 (40.6%)	21 (20.8%)
2	14 (13.6%)	19 (18.4%)	20 (20.0%)	25 (25.3%)	18 (17.8%)	29 (28.7%)
3	20 (19.4%)	23 (22.3%)	26 (26.0%)	14 (14.1%)	16 (15.8%)	19 (18.8%)
4	21 (20.4%)	25 (24.3%)	31 (31.0%)	22 (22.2%)	22 (21.8%)	29 (28.7%)
5	1 (1.0%)	0 (0%)	2 (2.0%)	2 (2.0%)	4 (4.0%)	3 (3.0%)
6	0 (0%)	0 (0%)	2 (2.0%)	0 (0%)	0 (0%)	0 (0%)
Social Orientation – Dominance						
<i>Mean (SD)</i>	2.27 (1.3)	2.51 (1.26)	2.76 (1.14)	2.44 (1.19)	2.44 (1.31)	2.60 (1.17)
<i>Median [Min, Max]</i>	2 [1, 5]	2.25 [1, 5.25]	3 [1, 4.75]	2.5 [1, 4.75]	2.25 [1, 5.75]	2.5 [1, 6]
Response Option: 1	51 (49.5%)	42 (40.8%)	32 (32.0%)	39 (39.4%)	44 (43.6%)	37 (36.6%)
2	17 (16.5%)	19 (18.4%)	16 (16.0%)	20 (20.2%)	18 (17.8%)	18 (17.8%)
3	12 (11.7%)	18 (17.5%)	27 (27.0%)	21 (21.2%)	18 (17.8%)	28 (27.7%)
4	22 (21.4%)	23 (22.3%)	25 (25.0%)	19 (19.2%)	19 (18.8%)	16 (15.8%)
5	1 (1.0%)	1 (1.0%)	0 (0%)	0 (0%)	2 (2.0%)	1 (1.0%)
6	0 (0%)	0 (0%)	0 (0%)	0 (0%)	0 (0%)	1 (1.0%)

Note: Higher symbolic racism scores indicate higher endorsement of symbolic racism. Negative blatant dehumanization scores indicate ingroup dehumanization compared to the target outgroup and positive scores indicate outgroup target dehumanization compared to the ingroup. Higher social dominance orientation scores indicate higher endorsement of group hierarchies. There were a few notable variables in Study 1 with differing

ethno-racial patterns. Black participants were on average lower in symbolic racism than White participants. Also, Black and Latinx participants showed higher scores on liberal ideology and lower dehumanization scores than did White participants. In Study 2, there were more minoritized participants with lower and more Whites with higher symbolic racism scores. However, this sample was more matched in terms of the racial composition of political orientation. Black and Latinx participants had lower social dominance orientation (egalitarian) scores than Whites.

Supplementary Table S2: Tweet stimuli.

Tweet #	Tweet	Aggressiveness	Target	Hashtag
1	America First fuck everyone else. Let those savages fend for themselves. My taxes protects them? Fuckem.	Aggressive	All	#AmericaFirst
2	Immigration without Assimilation is INVASION!!!	Aggressive	All	#AmericaFirst
3	American grandmother's can look forward to these sadistic attacks if we do not have a Muslim Ban	Aggressive	Muslim	#AmericaFirst
4	Do you get it now libtards? Ban Sharia Law. Totally incompatible with western values, norms, beliefs. No Amnesty!	Aggressive	Muslim	#AmericaFirst
5	And to me, it's not even America First. it's simply protecting us from #terrorists who wish to kill us and our children	Aggressive	Muslim	#AmericaFirst
6	Sharia is diametrically opposed to our Constitution! They will never assimilate- they view infidels as unclean	Aggressive	Muslim	#AmericaFirst
7	The last thing we need is more muslim refugees, just look at all the damage they have caused in Europe	Aggressive	Muslim	#AmericaFirst
8	#freeloaders Please eliminate #AnchorBabies. They now invade state government.	Aggressive	Undocumented	#AmericaFirst
9	I WILL NOT PRESS 2 FOR SPANISH...I LIVE IN AMERICA DAMNIT, I SPEAK ENGLISH!	Aggressive	Undocumented	#AmericaFirst
10	"Fk #Mexico! Bring YOUR people home. We DON'T want ILLEGALS in America!"	Aggressive	Undocumented	#AmericaFirst
11	Round up the #ILLEGALS and toss them back over the border	Aggressive	Undocumented	#AmericaFirst
12	American lives everywhere are being saved by deporting ILLEGAL CRIMINAL ALIENS. Drug dealers, gang members and sex offenders ALL OUT!	Aggressive	Undocumented	#AmericaFirst
13	It's a safe bet that anyone who uses "undocumented worker" instead of "illegal alien" is lying to you.	Aggressive	Undocumented	#AmericaFirst
14	Kind of scary when #Americans think more about other countries than their own.	Political	All	#AmericaFirst
15	"We will defend our FREEDOM, our JOBS, and our ECONOMIC INDEPENDENCE. It's going to be America First!"	Political	All	#AmericaFirst

RUNNING HEAD: Quantifying shared judgments of racism

16	#whyiresist Because our own citizens should come before illegals and terrorist savages.	Political	All	#AmericaFirst
17	FACT 90% Of Recent Middle Eastern Refugees On Food Stamps 70% On Cash Welfare WHILE 16.5 mil American children go hungry	Political	Muslim	#AmericaFirst
18	So well said. Let focus on Americans. Let's help the inner cities. Let's help Detroit before we help Syria	Political	Muslim	#AmericaFirst
19	This isn't an illegal alien playground. America is a nation of law and order.	Political	Undocumented	#AmericaFirst
20	As a legal immigrant & proud US citizen contributing in the workforce, my tax \$ ought to be used for America First	Political	Undocumented	#AmericaFirst
21	undocumented is just PC for Illegal. Illegal means AGAINST THE LAW. Felons are first to be deported.	Political	Undocumented	#AmericaFirst
22	Things you could fund instead of education, healthcare, etc for illegals:	Political	Undocumented	#AmericaFirst
	Bridges Highways VA Healthcare			
23	It is not cruelty to Build The Wall. When we put america first,and citizens of any country are always protected first from any threat.	Political	Undocumented	#AmericaFirst
24	I came here legally. No excuse for illegal immigrants. Assimilate and support this country, or don't come here at all.	Political	Undocumented	#AmericaFirst
25	Deport illegals first, wait a year, you have enough money to build a wall 11.3 times across the border	Political	Undocumented	#AmericaFirst
26	We are worried about a Muslim Ban when illegal aliens have raped and killed far more Americans then those that died	Aggressive	Both	#MuslimBan
27	Leave behind your ancient cultures. Or admittance is rejected.	Aggressive	Muslim	#MuslimBan
28	It may not have been a muslim ban but it should be. It's the only global threat and ideology that wants to target all innocent lives.	Aggressive	Muslim	#MuslimBan

RUNNING HEAD: Quantifying shared judgments of racism

29	Why are we openly welcoming, housing, feeding, and funding the very people who call for our destruction?	Aggressive	Muslim	#MuslimBan
30	No such thing as a #moderate they ALL live by the same evil book! Islam needs to 2 abolished as a religion!	Aggressive	Muslim	#MuslimBan
31	If any one gets hurt or killed by people coming in from these unvetted countries.. let it be on the 9th circuit's heads	Aggressive	Muslim	#MuslimBan
32	We have enough pedophiles and rapists in America, we don't need anymore.	Aggressive	Muslim	#MuslimBan
33	ATTN: refugees: Canada just opened its borders to you. Know go blow then stupid ppl up.	Aggressive	Muslim	#MuslimBan
34	@muslimgirl try spending some time in the Middle East and see how liberated you feel.	Aggressive	Muslim	#MuslimBan
35	Western civilization, greatest humanity has ever known, ancestors died to protect it, yet left wants to let in those who hate us!	Aggressive	Muslim	#MuslimBan
36	islam is an ideology of violence	Aggressive	Muslim	#MuslimBan
37	No White Genocide? What about this poor lady? Murdered by Muslim animals while trying to help them.	Aggressive	Muslim	#MuslimBan
38	Can't wait till U.S. doubles down on travel ban. Make sure NO Muslims can come here next time around. I support a Muslim Ban	Aggressive	Muslim	#MuslimBan
39	We're very close to having Pakis banned from America. Love it!	Aggressive	Muslim	#MuslimBan
40	An untrained dog will always pee on your rug, harsh love is required before it may enter once again.	Aggressive	Muslim	#MuslimBan
41	State after State will either fall under Sharia Law or it will defend itself.	Aggressive	Muslim	#MuslimBan
42	That's reality not Hate they don't assimilate & I would envoke a TOTAL ban of Muslim period.....	Aggressive	Muslim	#MuslimBan
43	These terror attacks at bowling green and Atlanta are exactly why we need border security.	Political	All	#MuslimBan
44	God sent Nehemiah to #Jerusalem to build one around the city to keep the bad guys out.	Political	All	#MuslimBan

45	Biblically we know that there is a legitimate place for walls, for physical barriers designed for protection.	Political	All	#MuslimBan
46	American citizen safety comes 1st. I lock my doors, don't you? Means I'm a snowflake, so be it.	Political	Muslim	#MuslimBan
47	The only time we should allow in immigrants or refugees is when we have a jobs surplus, or when the person is highly qualified.	Political	Muslim	#MuslimBan
48	We should only let in people who are brilliant and driven- even if they're a Muslim from Syria.	Political	Muslim	#MuslimBan
49	There is no Muslim Ban. Only a temporary halt to immigrants from 7 failed nations with no govt control	Political	Muslim	#MuslimBan
50	They fill like they hit the lottery because they know our tax dollars will support them once here. #NoMoreFreebies	Political	Muslim	#MuslimBan

Note: Tweet content is modified from the original (e.g., removing hashtags), but grammar and spelling was kept the same.

Supplementary Table S2. Demographic characteristics of Study 3, popular and unpopular condition, by ethno-racial group.

	Study 3 – Popular			Study 3 - Unpopular		
	Black (N=103)	Latino (N=98)	White (N=102)	Black (N=102)	Latino (N=101)	White (N=95)
Age						
<i>Mean (SD)</i>	-0.257 (15.3)	-9.25 (12.6)	4.68 (16.5)	1.02 (14.9)	-8.62 (14.0)	4.64 (16.6)
<i>Median [Min, Max]</i>	-3.60 [- 23.6, 31.4]	-12.6 [- 24.6, 28.4]	4.40 [- 24.6, 46.4]	0.903 [- 24.6, 30.4]	-11.6 [- 24.6, 39.4]	5.40 [- 24.6, 47.4]
Gender						
Don't Say	1 (1.0%)	3 (3.0%)	0 (0%)	1 (1.0%)	2 (2.0%)	0 (0%)
Female	22 (21.4%)	26 (26.5%)	26 (25.5%)	27 (26.5%)	26 (25.7%)	21 (22.1%)
Male	79 (76.7%)	69 (70.4%)	76 (74.5%)	73 (71.6%)	73 (72.3%)	72 (75.8%)
Non-binary	0 (0%)	0 (0%)	0 (0%)	1 (1.0%)	0 (0%)	2 (2.1%)
Other	1 (1.0%)	0 (0%)	0 (0%)	0 (0%)	0 (0%)	0 (0%)
Twitter Use						
Daily	24 (23.3%)	31 (31.6%)	17 (16.7%)	28 (27.5%)	24 (23.8%)	18 (18.9%)
Less Often	16 (15.5%)	9 (9.2%)	13 (12.7%)	15 (14.7%)	11 (10.9%)	14 (14.7%)
Never	40 (38.8%)	40 (40.8%)	54 (52.9%)	46 (45.1%)	42 (41.6%)	51 (53.7%)
Weekly	7 (6.8%)	5 (5.1%)	7 (6.9%)	3 (2.9%)	6 (5.9%)	4 (4.2%)
X a Week	15 (14.6%)	11 (11.2%)	11 (10.8%)	9 (8.8%)	16 (15.8%)	8 (8.4%)
Missing	1 (1.0%)	2 (2.0%)	0 (0%)	1 (1.0%)	2 (2.0%)	0 (0%)
Political Orientation						
<i>Mean (SD)</i>	3.33 (1.44)	3.91 (1.65)	3.88 (1.56)	3.50 (1.48)	3.37 (1.44)	3.97 (1.54)
<i>Median [Min, Max]</i>	3.00 [1, 7]	4.00 [1, 7]	4.00 [1, 7]	3.50 [1, 7]	3.50 [1, 7]	4.00 [1, 7]
Response option: 1	12 (11.7%)	10 (10.2%)	10 (9.8%)	9 (8.8%)	11 (10.9%)	8 (8.4%)
2	28 (27.2%)	13 (13.3%)	12 (11.8%)	25 (24.5%)	26 (25.7%)	14 (14.7%)
3	15 (14.6%)	15 (15.3%)	18 (17.6%)	18 (17.6%)	15 (14.9%)	6 (6.3%)

	Study 3 – Popular			Study 3 - Unpopular		
	Black (N=103)	Latino (N=98)	White (N=102)	Black (N=102)	Latino (N=101)	White (N=95)
4	29 (28.2%)	33 (33.7%)	33 (32.4%)	32 (31.4%)	34 (33.7%)	41 (43.2%)
5	12 (11.7%)	8 (8.2%)	12 (11.8%)	8 (7.8%)	5 (5.0%)	10 (10.5%)
6	6 (5.8%)	12 (12.2%)	14 (13.7%)	5 (4.9%)	8 (7.9%)	11 (11.6%)
7	1 (1.0%)	7 (7.1%)	3 (2.9%)	5 (4.9%)	2 (2.0%)	5 (5.3%)
Symbolic Racism						
<i>Mean (SD)</i>	14.6 (4.21)	17.0 (4.59)	19.4 (4.96)	14.9 (3.75)	16.3 (4.70)	18.7 (5.13)
<i>Median [Min, Max]</i>	14 [8, 24]	17. [8, 30]	19 [10, 31]	14.5 [8, 25]	17 [8, 30]	20 [8, 30]
Score range: [8,13.8]	49 (47.6%)	24 (24.5%)	14 (13.7%)	41 (40.2%)	33 (32.7%)	17 (17.9%)
[13.8,19.5]	40 (38.8%)	49 (50.0%)	40 (39.2%)	50 (49.0%)	40 (39.6%)	27 (28.4%)
[19.5,25.2]	14 (13.6%)	21 (21.4%)	36 (35.3%)	11 (10.8%)	25 (24.8%)	44 (46.3%)
[25.2,31]	0 (0%)	4 (4.1%)	12 (11.8%)	0 (0%)	3 (3.0%)	7 (7.4%)
Blatant Dehumanization - Outgroup						
<i>Mean (SD)</i>	9.96 (25.7)	7.57 (15.5)	8.84 (20.5)	7.68 (20.9)	5.79 (15.6)	7.94 (21.0)
<i>Median [Min, Max]</i>	1 [-46.7, 100]	1.67 [-32.7, 64.7]	0 [-54.3, 100]	0.333 [-54.7, 78]	0.333 [-47.3, 60.7]	0 [-63.7, 100]
Score range: [-100,-50]	0 (0%)	0 (0%)	1 (1.0%)	1 (1.0%)	0 (0%)	1 (1.1%)
[-50,0]	28 (27.2%)	18 (18.4%)	23 (22.5%)	22 (21.6%)	16 (15.8%)	25 (26.3%)
[0,50]	66 (64.1%)	76 (77.6%)	75 (73.5%)	75 (73.5%)	83 (82.2%)	64 (67.4%)
[50,100]	9 (8.7%)	4 (4.1%)	3 (2.9%)	4 (3.9%)	2 (2.0%)	5 (5.3%)
Blatant Dehumanization - White						
<i>Mean (SD)</i>	10.9 (38.9)	6.93 (25.3)	-	3.28 (32.5)	5.28 (31.2)	-
<i>Median [Min, Max]</i>	1 [-100, 100]	0 [-65, 94]	-	0 [-98, 100]	0 [-100, 100]	-
Score range: [-100,-50]	5 (4.9%)	2 (2.0%)	-	5 (4.9%)	4 (4.0%)	-

	Study 3 – Popular			Study 3 - Unpopular		
	Black (N=103)	Latino (N=98)	White (N=102)	Black (N=102)	Latino (N=101)	White (N=95)
[-50,0]	21 (20.4%)	23 (23.5%)	-	20 (19.6%)	21 (20.8%)	-
[0,50]	64 (62.1%)	66 (67.3%)	-	69 (67.6%)	66 (65.3%)	-
[50,100]	13 (12.6%)	7 (7.1%)	-	8 (7.8%)	10 (9.9%)	-
Blatant Dehumanization						
- Latinx						
<i>Mean (SD)</i>	8.66 (24.9)	-	6.96 (19.0)	9.15 (25.1)	-	5.55 (18.4)
<i>Median [Min, Max]</i>	0 [-59, 100]	-	0 [-49, 100]	0 [-73, 100]	-	0 [-53, 100]
Score range: [-100,-50]	2 (1.9%)	-	0 (0%)	1 (1.0%)	-	1 (1.1%)
[-50,0)	24 (23.3%)	-	20 (19.6%)	18 (17.6%)	-	22 (23.2%)
[0,50)	68 (66.0%)	-	77 (75.5%)	77 (75.5%)	-	69 (72.6%)
[50,100]	9 (8.7%)	-	5 (4.9%)	6 (5.9%)	-	3 (3.2%)
Blatant Dehumanization						
- Black						
<i>Mean (SD)</i>	-	6.93 (16.3)	6.19 (22.9)	-	4.43 (13.3)	5.01 (21.8)
<i>Median [Min, Max]</i>	-	0 [-25, 90]	0 [-100, 100]	-	0 [-29, 71]	0 [-88, 100]
Score range: [-100,-50]	-	0 (0%)	2 (2.0%)	-	0 (0%)	1 (1.1%)
[-50,0]	-	19 (19.4%)	14 (13.7%)	-	15 (14.9%)	21 (22.1%)
[0,50]	-	77 (78.6%)	80 (78.4%)	-	85 (84.2%)	70 (73.7%)
[50,100]	-	2 (2.0%)	6 (5.9%)	-	1 (1.0%)	3 (3.2%)
Blatant Dehumanization						
- Muslim						
<i>Mean (SD)</i>	10.3 (27.2)	8.84 (21.8)	13.4 (28.6)	10.6 (21.7)	7.66 (21.0)	13.2 (29.0)
<i>Median [Min, Max]</i>	0 [-81, 100]	0 [-73, 94]	0 [-62, 100]	0 [-42, 82.0]	0 [-47, 96]	0 [-50, 100]
Score range: [-100,-50]	2 (1.9%)	1 (1.0%)	1 (1.0%)	0 (0%)	0 (0%)	0 (0%)
[-50,0]	21 (20.4%)	19 (19.4%)	21 (20.6%)	20 (19.6%)	15 (14.9%)	21 (22.1%)

	Study 3 – Popular			Study 3 - Unpopular		
	Black (N=103)	Latino (N=98)	White (N=102)	Black (N=102)	Latino (N=101)	White (N=95)
[0,50]	73 (70.9%)	74 (75.5%)	68 (66.7%)	72 (70.6%)	81 (80.2%)	63 (66.3%)
[50,100]	7 (6.8%)	4 (4.1%)	12 (11.8%)	10 (9.8%)	5 (5.0%)	11 (11.6%)
Social Dominance Orientation - Egalitarian						
<i>Mean (SD)</i>	2.38 (1.26)	2.70 (1.24)	2.99 (1.23)	2.52 (1.26)	2.29 (1.19)	2.85 (1.28)
<i>Median [Min, Max]</i>	2.25 [1, 5.25]	3 [1, 5.25]	3.25 [1, 5.75]	2.5 [1, 5]	2 [1, 5]	3 [1, 6]
Response option: 1	45 (43.7%)	34 (34.7%)	23 (22.5%)	42 (41.2%)	49 (48.5%)	28 (29.5%)
2	22 (21.4%)	15 (15.3%)	19 (18.6%)	15 (14.7%)	17 (16.8%)	15 (15.8%)
3	15 (14.6%)	23 (23.5%)	23 (22.5%)	21 (20.6%)	21 (20.8%)	25 (26.3%)
4	20 (19.4%)	25 (25.5%)	34 (33.3%)	23 (22.5%)	13 (12.9%)	25 (26.3%)
5	1 (1.0%)	1 (1.0%)	3 (2.9%)	1 (1.0%)	1 (1.0%)	1 (1.1%)
6	0 (0%)	0 (0%)	0 (0%)	0 (0%)	0 (0%)	1 (1.1%)
Social Dominance Orientation – Dominance						
<i>Mean (SD)</i>	2.36 (1.23)	2.71 (1.32)	2.77 (1.22)	2.53 (1.32)	2.27 (1.15)	2.60 (1.27)
<i>Median [Min, Max]</i>	2.25 [1, 4.75]	2.5 [1, 5.5]	2.75 [1, 5.50]	2.5 [1, 5.25]	2 [1, 4.5]	2.5 [1, 5.25]
Response option: 1	45 (43.7%)	35 (35.7%)	30 (29.4%)	43 (42.2%)	48 (47.5%)	38 (40.0%)
2	23 (22.3%)	19 (19.4%)	21 (20.6%)	15 (14.7%)	22 (21.8%)	12 (12.6%)
3	12 (11.7%)	14 (14.3%)	23 (22.5%)	18 (17.6%)	18 (17.8%)	23 (24.2%)
4	23 (22.3%)	28 (28.6%)	27 (26.5%)	25 (24.5%)	13 (12.9%)	21 (22.1%)
5	0 (0%)	2 (2.0%)	1 (1.0%)	1 (1.0%)	0 (0%)	1 (1.1%)

Note: Higher symbolic racism scores indicate higher endorsement of symbolic racism. Negative blatant dehumanization scores indicate ingroup dehumanization compared to the target outgroup and positive scores indicate outgroup dehumanization compared to the ingroup. Higher social dominance orientation scores indicate higher endorsement of group hierarchies. As in Study 1 and 2, more minoritized participants were liberal and scored lower on symbolic racism than Whites in both conditions.

Supplementary Figure S1. Correlation matrices of participant factors in Study 1. Cooler colors (blue) represent positive correlations and warmer colors represent negative correlations. Hue saturation represents the magnitude of the correlation. The plot quadrants represent the sample in which the correlations were calculated. “NA”s are correlations that could not be calculated because those factors were not available in the specific sample. BD = blatant dehumanization, SDO = social dominance orientation.

Supplementary Figure S2. Range of the tweet stimuli’s racism across studies. Each dot represents a tweet presented to participants. The y-axis is the standard deviation of all ratings given to each tweet while the x-axis the average of all ratings given to each tweet. The blue line is a loess smooth line and the grey shaded area is the standard error of the line. The dashed vertical grey lines represent the 25%, 50%, and 75% quantiles of the rating average. When examining the averages, the stimulus-choosing procedure led to a stimulus range of about 4 to 7 on a racist scale of 1 (not at all) to 9 (extremely). However, examining the standard deviation provides a more nuanced picture. The tweets given the most “neutral” average ratings (around a 5) also showed the largest standard deviations. In other words, these tweets were not truly “neutral”, they were given opposing ratings across the full scale. Even Tweet 39, which stands out as the most racist tweet, has an SD of about 2.6, which means it was also given ratings across the full scale. These plots reveal a more complicated story whereby the “true racist range” of the stimuli cannot easily be inferred from the means for the means are arrived at through major disagreement. If the means of each tweet were less controversial (i.e., very small SDs), then we might have a way to properly discuss the “range” of the stimulus set, but given these results, we can only conclude that our stimulus set evoked major disagreement in perceivers (as concluded by the VCAs in the main paper). Examination of rating histograms for all tweets confirmed that while some tweets showed a skewed distribution, participants used the full scale for every tweet and thus our findings are not explained by ceiling or floor effects. Given the aggressiveness of many of the tweets (Supplementary Table S2), these insights beg the question whether including innocuous or very egregious tweets would have constructed a stimulus set where the range of means reach the ends of the scale, or if disagreement is to always be expected.

Supplementary Figure S3. Variance components for each study. The y axis represents the untransformed amount of variance held by each component. In the legend, P = perceiver, S = stimulus, and B = block. The B x P component is the only block related component with visible variance, this suggests to some small extent, some perceivers changed their overall rating average across blocks.

Supplementary Figure S4. Correlation matrices of participant characteristics in Study 2. Cooler colors (blue) represent positive correlations and warmer colors represent negative correlations. Hue saturation represents the magnitude of the correlation. The plot quadrants represent the sample in which the correlations were calculated. “NA”s are correlations that could not be calculated because those factors were not available in the specific sample. BD = blatant dehumanization, SDO = social dominance orientation.

Supplementary Figure S5. VPC differences across studies and conditions. A) VPC differences between Study 1 and Study 2. B) VPC differences from the unpopular condition (Study 3) and Study 1 & 2 combined. C) VPC differences between popular condition (Study 3) and Study 1 & 2 combined. D) VPC differences between popular and unpopular conditions (Study 3). Error bars represent 95% confidence intervals of bootstrapped differences.

Overall, the pattern of VPC differences in B, C, D shows that including norm information (high and low popularity) seems to slightly decrease disagreement from perceiver variance (i.e., mean differences) but increase disagreement from interaction variance (i.e., stimulus rankings) for Black participants. The opposite pattern occurred for White participants. These results were unexpected and ambiguous to interpret, we do not make further inferences.

Supplementary Figure S6. Correlation matrices of participant factors in Study 3, unpopular condition. Cooler colors (blue) represent positive correlations and warmer colors represent negative correlations. Hue saturation represents the magnitude of the correlation. The plot quadrants represent the sample in which the correlations were calculated. “NA”s are correlations that could not be calculated because those factors were not available in the specific sample. BD = blatant dehumanization, SDO = social dominance orientation.

Supplementary Figure S7. Correlation matrices of participant factors in Study 3, popular condition. Cooler colors (blue) represent positive correlations and warmer colors represent negative correlations. Hue saturation represents the magnitude of the correlation. The plot quadrants represent the sample in which the correlations were calculated. “NA”s are correlations that could not be calculated because those factors were not available in the specific sample. BD = blatant dehumanization, SDO = social dominance orientation.