

What Counts as Racism?

American Norms, Partisan Norms, and Personal Beliefs About Racism

Mark Pickup
(Simon Fraser University)

Erik O. Kimbrough
(Chapman University)

Eric William Groenendyk
(University of Memphis)

Antoine Jevon Banks
(University of Maryland)

Abstract: According to conventional wisdom, the 1960s Civil Rights Movement caused a shift in American cultural norms. Although many Americans continued to harbor racial animosity, they came to share an understanding that it was no longer acceptable to express this sentiment openly. Thus, while hidden racial resentment still influenced public opinion and political behavior, shared norms against racism helped to quell its effect and forced openly racist political actors into the fringe of society. Recent evidence suggests these shared norms may be eroding. As political polarization intensifies and scholarly conceptualizations of racism evolve, separate norms may be emerging among Republicans and Democrats around the very notion of what does and does not count as racism. To test this hypothesis, we measure Americans' beliefs across four different conceptualizations of racism—explicit racism, subtle racism, institutional racism, and colorblind racism—as well as their beliefs about what their fellow partisans and their fellow Americans believe counts as racism, providing a monetary incentive for correctly identifying the most common response within the group. Our findings suggest shared American norms are indeed eroding in favor of distinct partisan norms of what counts as racism.

Key words: identity, partisanship, racism, norms

Word count: 8,655

Conflict over how to address the legacy of slavery, segregation, and racial prejudice is nothing new in American politics. Affirmative action, busing, and welfare programs have long been hot-button issues, but recent discourse has shifted away from policies designed to address racism to a more basic and fundamental question: what should count as racism in the first place?

Although the conceptualization of racism has been a rich topic of discussion in academic journals and college classrooms for decades, debate has recently spilled out into mainstream politics, leading to significant confusion over use of the term “racism” within the public sphere and providing ample opportunities for politicians to take advantage. As many on the left express outrage over systemic racism, a firestorm over critical race theory, “wokeness,” and “PC culture” has been ignited on the right. This, it seems, is the new front in the racial politics debate.

If Americans cannot agree on what counts as racism, attempts to enforce norms against racism are likely to be met with confusion and anger. And since anger has been shown to enflame racial policy divisions (Banks and Valentino 2012; Banks 2014), this is hardly a recipe for bringing people together. To have any hope for a productive dialog over matters of race, Americans must possess a shared vocabulary. Thus, we must determine whether Americans agree on what counts as racism, and if not, what is causing the division. While numerous studies show that Americans have sorted into parties based on race (e.g., Carmines and Stimson 1989, Abramowitz 2016), we possess only anecdotal evidence of growing partisan division over the questions of what counts as racism. How divided are partisans on this topic? On what points do they agree and disagree? Is this really a partisan division at all, or is partisanship simply standing in for race? And, finally, is this division merely the product of

sorting, or might parties actually be shaping these beliefs, leading to further polarization? To answer these questions, we examine what Americans believe constitutes racism today, whether shared American norms against racism exist, and the role partisan identities may play in weakening these norms—perhaps replacing them with distinct partisan norms.

Despite the growing prevalence of these issues and their important implications for racial policy discourse, relatively little is known about what Americans believe does and does not constitute racism, what causes division, and if this division can be overcome. This paper seeks to address these important gaps in the literature by introducing three measurement innovations. First, we design a question battery to tap whether or not Americans believe various scholarly conceptualizations of racism indeed constitute racism. These statements were carefully designed and pretested to capture four prominent concepts of racism: explicit racism, subtle racism, institutional racism, and colorblind racism. Second, in addition to measuring Americans' own beliefs about what does and does not constitute racism, we also examine *social norms*—people's beliefs about what others (either fellow Americans or fellow partisans) believe counts as racism. Third, we conduct a norm priming experiment by manipulating the salience of group norms prior to measuring personal beliefs. If group norms are 1) strong and 2) salient as people report their personal beliefs, individuals should feel increased pressure to bring their beliefs into alignment with the expectations of their group. This experimental manipulation allows us to draw causal inferences regarding how American, Republican, and Democratic norms affect beliefs about what counts as racism.

By leveraging these methods, we are able to draw a number of important insights. First, turning to our nationally representative sample of Americans' beliefs, we show that while a

large majority of Americans agree that typical examples of explicit racism count as racism, there is considerably less agreement over whether typical examples of subtle racism, institutional racism, and colorblind racism constitute racism. Perhaps more surprising, the most pronounced divisions over what does and does not count as racism are not observed between races, age groups, or education levels, but between Republicans and Democrats, suggesting distinct norms may be forming within these groups. Second, turning to our norms data, we show that American norms regarding what constitutes racism are significantly less clear than partisan norms on all types of racism except explicit racism. However, when we examine the consistency between perceived norms and personal beliefs, we find that Republicans are much more likely than Democrats to hold beliefs about racism that are consistent with perceived partisan norms. As a result of this partisan asymmetry, norm priming tends to affect Democrats more than Republicans. Whereas most Republicans express beliefs consistent with party norms, regardless of whether those norms are made salient, many Democrats hold norm inconsistent beliefs until party norms are primed, leading to post-treatment consistency levels more on par with Republicans. Moreover, Democrats who believe other Americans hold narrow conceptualization of racism also respond to American norm priming by endorsing broader conceptualizations of racism, thereby exhibiting a backlash effect. As a result, in both cases, we see greater evidence of belief polarization after priming norms, with Republicans endorsing narrower conceptualizations of racism and Democrats endorsing broader conceptualizations.

Conceptualizing Racism

Racism, and scholarship relating to it, each have a long and varied history in the United States. In the years leading up to the Civil War, as the abolition movement grew stronger, slaveholders and sympathizers sought to justify slavery through an ideology of white supremacy (Sears, Hetts, Sidanius, and Bobo 2000). This racism was built on the belief that whites were biologically superior to blacks. Attempts were even made to “prove” white racial superiority through “scientific” analysis.

In the 20th Century, however, things changed. The atrocities committed in Nazi Germany highlighted the danger of ideologies built on prejudice, and beginning with Allport’s (1954) classic work, researchers began to investigate racial prejudice in America (Allport 1954). With Blacks leading the charge (Lee 2000), many Americans pushed back against white supremacist ideology and segregation. And, by the end of the Civil Rights Era, public opinion surveys showed significant change—or at least they appeared to. The sort of overt racial animosity associated with *explicit or “old-fashioned” racism* had declined substantially (Kinder and Sanders 1996; Schuman et al. 1997), but did this mean racism was dissipating?

This apparent reduction in “old-fashioned” racism struck many scholars as difficult to reconcile with widespread opposition, among whites, to government policies aimed at integrating schools and neighborhoods and providing equal opportunity for African Americans. Theoretically, opposition to integration might be explained by self-interest, but self-interest was shown to be a poor predictor of most policy positions (Kinder and Kiewiet 1981). Alternatively, this opposition might be explained by ideological commitment to limited government, but Americans were famous for their ideological “innocence” (Converse 1964;

Kinder 1984). Thus, many scholars suspected racial prejudice remained high, while *social norms* had changed, prohibiting the *expression* of explicitly racist attitudes. As a result, racial attitudes could still play an important role in American politics. They were simply more difficult to measure. And politicians could still take advantage of the white's racial animosity by using code words such as "welfare recipients" and "inner cities" to evoke race without violating norms (Gilens 1999; Mendelberg 2001).

To capture this sort of *subtle racism*, scholars began to devise measures that provided survey respondents with a way to express their frustration, without violating social norms against explicit racism. These measures, referred to in various works as indicators of symbolic racism (Kinder and Sears 1981), modern racism (McConahay 1983), and racial resentment (Kinder and Sanders 1996), blended "anti-black affect and the kind of traditional American moral values embodied in the Protestant Ethic" (Kinder and Sears 1981, p 416.). The result was a series of new racial attitudes measures that correlated strongly with policy opinions—particularly policies dealing with issues of race.

In contrast to both explicit racism and this new subtle racism, each of which are understood to be rooted in the prejudices of individual Americans, other scholars draw attention to racism perpetuated at the social level, through social structures and institutions that provide certain groups with a consistent advantage over others (Carmichael and Hamilton 1967; Bonilla-Silva 2018). From this perspective, individuals need not necessarily be motivated by personal prejudices to participate in systems that promote racism. For example, incarceration rates are much higher among African Americans than among white Americans, due in large part to disparate sentencing practices, differences in policing policies between black and white

neighborhoods, stiffer penalties for possession of crack cocaine compared to powdered cocaine, etc. The individual police officers, prosecutors, and judges who participate in these systems may not harbor racial animosity, and may even be African Americans themselves, but the effect of these institutions is to perpetuate the subjugation of African Americans. We will refer to this conceptualization as *institutional racism*.

The final conceptualization of racism we examine is a related idea that Bonilla-Silva (2010) has termed *colorblind racism*. Colorblind racism occurs when people accept the idea that the best way to deal with matters of race is to ignore race. Advocates of “colorblindness” argue that, if we set race aside and treat everyone equally, our problems will be resolved. However, Bonilla-Silva and others argue that this ideology itself constitutes racism, because it serves to protect the racial status quo, perpetuates racial inequity, and maintains white privilege. The term ideology is important, because colorblind racism is conceptualized, not merely as individual racial animosity or prejudice, but as a broadly shared belief system. And this belief system tends to dominate because it favors those with power. Thus, in contrast to both the explicit and subtle racism described previously, colorblind racism, like institutional racism, is structural (Bonilla-Silva p. 54). According to this point of view, what makes racism so endemic is that it no longer even requires animosity, merely acceptance of the ideology of colorblindness.

How Should We Conceptualize Racism?

While scholars have long debated how racism should be conceptualized, the discussion has recently moved out of academic journals and college classrooms into the center of American politics. In the past, the one thing that scholars seemed to agree on was that norms

against racism were emerging. While some thought this was reducing racism (Hyman and Sheatsley 1954; 1964), others argued it was merely reducing Americans' willingness to express explicitly racist attitudes in surveys (Kinder and Sanders 1996; Bonilla-Silva 2010). But both sides of the debate agreed norms were changing and influencing behavior.

Although often overlooked, this point of agreement is vital. While racial animosity has continued to influence public opinion and political behavior, norms can exert a powerful influence on opinions and behavior (for recent reviews see Miller and Prentice 2016; Legros and Cislighi 2019).

Norms exercise both intrinsic and extrinsic influence on our behavior. Their intrinsic influence arises because people "internalize" norms, evaluating their own behavior in light of (what they imagine to be) the normative expectations of those whose judgments they value, such as friends, family, co-religionists, co-partisans, etc. People feel psychological discomfort when they violate those internalized expectations, and they feel pleasure when they live up to or exceed them. Thus, priming individuals to think about the norms associated with their political identity has been shown to reduce peoples' willingness to "vote against their interests", sacrificing personal benefits to avoid violating group norms (Pickup et al. 2020, 2015 White et al. 2014).

At the same time, norms provide extrinsic motivation because they are typically supported and sustained by patterns of "punishments" and "rewards" meted out by third party observers (Kandori 1992). When we violate someone's normative expectations, they may be slower to return our texts or shun us entirely; they may gossip about our behavior; they may withdraw their support or proactively seek to harm us. Even if they sympathize with us and

understand our reasons, they may feel obligated to join others (who are less sympathetic) in meting out punishment. Similarly, when we behave appropriately in others' eyes, they may seek us out for friendship or reward us with gifts, praise, and introductions to powerful and influential others.

Crucially for politics, norms can shape not only what we do, but also what we believe and are willing to admit that we believe. One need only witness a single public shaming on Twitter and the apologetic contortions of the shamed to see how this works. Explicit social pressures exert a powerful and observable influence on stated opinions; mere belief that such pressures exist has a similar impact. People tend to report opinions that conform more strongly to their group's norms after being primed to think about those norms (Groenendyk et al. 2022).

Individuals may identify with multiple groups and the salience of those identities varies with time and context. When a particular group identity is made salient, the pull of its norms is amplified. If Americans agree that racism is wrong and runs counter to American values, individuals who identify most strongly as Americans will feel pressure to avoid committing acts deemed racist, since the individual would risk, not only being labeled racist, but also unamerican. Fear of violating such norms, scholars surmised, was why subtle racial appeals seemed to influence political preferences more than explicit appeals. Explicit appeals violated social norms and were thus rejected, while more subtle appeals activated racial resentment without violating norms (Mendelberg 2001; Valentino et al. 2002). Thus, office-seekers like former Ku Klux Klan leader, David Duke, were viewed as fringe candidates. Now, however, it seems shared American norms against acts and systems deemed racist may be eroding. Recent studies have found that explicit racial appeals exert effects similar to more subtle racial appeals

(Valentino et al. 2018), and efforts to call out racism appear to have limited influence (Banks and Hicks 2019).

To develop a clearer understanding why this is happening, we treat racism as a social construct and take a norms-focused approach, which makes our work very different from much of the existing literature, especially other quantitative works. Rather than thinking of racism as a theoretical construct that can be measured correctly or incorrectly—a long debated topic within the existing literature—we view racism as an intersubjective social construct. In other words, the meaning of racism does not exist independently from society but instead depends on our shared understanding of the term. If different groups within society understand racism to mean different things, then communication becomes difficult, if not impossible. The same sorts of cultural and linguistic misunderstandings that can lead to miscommunication and offense between individuals from different parts of the world (e.g., expectations surrounding greetings, table manners, etc.) may occur between people who inhabit different social circles within the same geographic area.

This way of thinking about racism leads us to set aside the question of how any particular scholarly conceptualization of racism influences behavior and instead focus on understanding what influences how the American public conceptualizes racism. To do this we designed a question battery in which survey respondents read a variety of scenarios and decided whether or not they constitute an example of racism. Each of the scenarios was designed to capture one of the four prominent conceptualizations of racism described in the previous section—'explicit racism', 'subtle racism', 'institutional racism', and 'colorblind racism'. This

design allows us to examine what Americans, themselves, believe constitutes racism, what does not, and who tends to disagree with whom.

Taking our approach of treating racism as a social construct a step further, we then turn our attention to social norms—people’s beliefs about what others in their group believe constitutes racism. Specifically, we compare norms among Republicans, Democrats, and Americans, allowing us to examine the relationship between group norms and personal beliefs, and how priming those norms effects the personal beliefs people report. While individual prejudices certainly play an important role in motivating behavior, it is also vital to understand the social processes that govern the expression of these feelings. In other words, if what changed from the Jim Crow era to the post-civil rights era was not primarily Americans’ underlying prejudices, but rather the social norms governing their expression (Kinder and Sanders 1996) and the ideology used to justify maintenance of white dominance (Bonilla-Silva 2010), then it seems norms should be at the very center of research on race in politics. Americans’ beliefs about how their opinions and behaviors will be perceived by others (e.g., as condoning racism or not) play an enormous role in the policies they support and the actions they take in politics (e.g., Pickup, Kimbrough, and de Rooij 2021; White and Laird 2020; Anoll 2022; Groenendyk, Kimbrough, and Pickup 2022). If previously shared American norms and understandings regarding racism are giving way to separate partisan norms and understandings, denunciation of racism may cease to be an effective method for motivating social change. Certainly, Americans have always been more willing to condemn racism in theory than in practice, but if we can no longer even agree in theory, the potential for future progress on matters of race would seem to be quite limited.

Personal Beliefs regarding Racism

Research Design: We conducted a survey of a nationally representative sample of Americans over 18 years of age, using NORC's AmeriSpeak panel (conducted January 2022; $N = 2,030$). The survey began with a consent form, and made use of the standard ANES measure of party identification to construct the standard 7-point scale (see Supplementary Information 1). The survey measured respondent's personal beliefs about what constitutes racism, with the following question.¹

Measure of personal beliefs

"The next questions will ask for your personal opinion.

Next, we will describe a series of situations and ask you to tell us whether you consider the situation to be an instance of racism or you do not consider the situation to be an instance of racism. In some cases, it may be difficult to decide, but we would like you to give us your best answer."

1. A white person uses the n-word referring to a black person.
2. A white person objects to a member of their family dating a black person.
3. A white person says that welfare recipients could get ahead if they were willing to work hard.
4. A white person says inner cities are full of criminals.
5. Black Americans are incarcerated at more than five times the rate of white Americans.
6. Black Americans are more than twice as likely to be denied home mortgages compared to white Americans.
7. A company does not consider race when making hiring decisions. As a result, only 2% of their employees are black.
8. A prestigious university does not consider race when making admissions decisions. As a result, black students are underrepresented by 50%.

¹ The NORC study also included a survey experiment in which response options were manipulated, but most of our analyses focus on the 1,019 of the sample who were asked what constituted an example of racism. The other 1,011 respondents were asked what constituted either racial insensitivity or racial injustice. The results of the experiment appear in SI2.

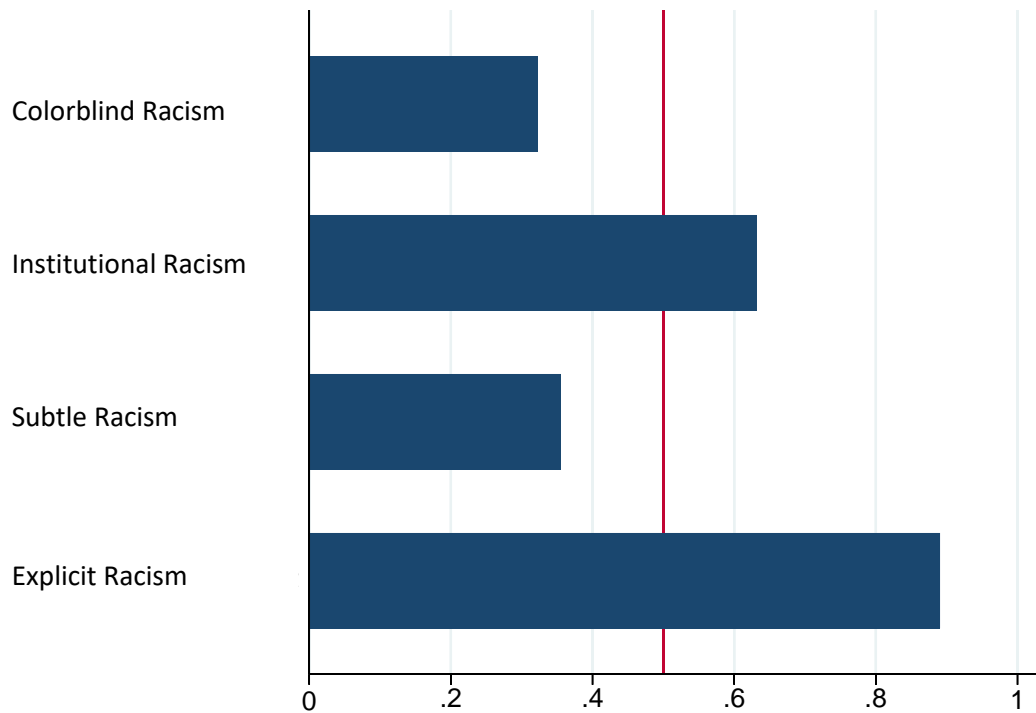
I would consider this situation an instance of racism

I would NOT consider this situation an instance of racism

Results: For each conceptualization of racism (explicit, subtle, institutional, and colorblind) respondents were asked to judge two scenarios. Figure 1 present the percentage of NORC respondents who expressed the belief that a given conceptualization constitutes racism, averaged across the two items.² Across the sample, we see a clear consensus (89%) that the scenarios designed to represent explicit racism constitute instances of racism. However, there is substantially less agreement on the other scenarios: 36% of the sample believes the subtle racism scenarios constitute racism, while the remaining 64% of the sample believes these scenarios should not be counted as racism; 63% of the sample believes the examples of institutional racism constitute racism, while 37% of the sample disagrees; and 32% of the sample believes the colorblind scenarios constitute examples of racism, while 68% believes they do not. As a whole, this first look at the data reveals substantial disagreement amongst Americans about what constitutes all types of racism besides “old fashioned” or explicit racism.

² Cronbach’s alpha values: .531 (explicit), .626 (subtle), .682 (institutional) .749 (colorblind).

Figure 1. Percentage “Agree” the Scenario is an Instance of Racism – NORC Study




N = 1,019

Next, we examine what underlies this disagreement. As discussed in the introduction, there is a long history of racial division in American politics, particularly between Blacks and Whites (e.g., Kinder and Sanders 1996), and, in recent generations, Americans have sorted into parties based on race (e.g., Carmines and Stimson 1989; Abramowitz 2016). Given the centrality of race in American politics, and the nature of the topic under investigation, one might expect this to be the dominant line of cleavage. On the other hand, the partisan division has come to rival racial division in American politics (Iyengar and Westwood 2015), and disagreements between those seeking social justice and those decrying political correctness have moved to the center of political discourse. Thus, Figure 2 compares partisan division to racial division across our four conceptualizations of racism. Independent variables are coded as dummies with the excluded

categories of political Independent (party identification) and any other race (race) serving as the baseline for the respective comparisons. Each point represents the OLS coefficient for a given dummy variable with controls for education and age in addition to party identification and race. Again, the dependent variables are averages across the two scenarios for each conceptualization of racism, scaled to run from 0 to 1.

Results show that both race and partisanship are strong predictors of beliefs about what constitutes racism, but the divide between Republican versus Democratic identifiers tends to be even greater than that between those who identify as Black versus White. Although enormous partisan divisions have become commonplace in American politics, this finding strikes us as particularly arresting. Parties are often defined as coalitions of social groups (Key 1942), and people are thought to develop partisan attachments because their party tends to serve the interests of their group (Campbell et al. 1960). From this point of view, one might only expect to see partisan division on this topic because party identification tends to be aligned with racial identification. Our results suggest something more is happening. Party identification is not simply serving as a proxy for race. When we look within groups, we see that while African Americans are relatively united on what counts as racism (although interesting points of disagreement are visible), Whites are intensely divided. In fact, if anything, White Democrats endorse broader conceptualizations of racism than Black Democrats, and White Republicans endorse far narrower conceptualization of racism than the small number of Black Republicans in our sample (see SI4 for additional information). Finally, one can see that Republicans and Democrats are each differentiated from Independents, suggesting the division is being driven by both sides of the debate.

Figure 2. Racial Versus Partisan Division Over What Counts as Racism – NORC Study



Note: $N = 1,019$. Dots represent OLS regression coefficients and whiskers represent 95% confidence intervals.

Given the lack of agreement among Americans regarding beliefs about what constitutes racism, and the clear partisan divide in those beliefs, we next use data from a YouGov study to examine the relationship between those beliefs and the norms of both partisan and American identity groups.³

Normative Expectations regarding Racism

³ SI4 includes a figure showing the distribution of partisan and American identity group strength.

Research Design: We conducted a survey of a nationally representative sample of Americans over 18 years of age, using YouGov’s panel (June 23 to July 1, 2022; $N = 3,000$). The YouGov survey included a consent form and asked the same partisan identity question as the NORC survey. It also included some attention focusing/testing text (see S11). Participants were then randomly assigned (with equal probability) to the partisan or American identity version of the questionnaire. If they were assigned to the partisan identity branch of the questionnaire, Republican identifiers were asked about the strength of their Republican identity as a social identity, and Democrat identifiers were asked about the strength of their Democrat identity as a social identity. If they are assigned to the American identity version, they were asked about the strength of their American identity as a social identity, with questions adapted from Huddy and Khatib (2007) (see S11).

Respondents to the YouGov so they were asked the same questions about their personal beliefs regarding racism as those in the NORC survey. They were also asked about the normative expectations of their partisan or American identity group (depending on which survey they were assigned).⁴ Prior to asking about normative expectations or personal beliefs, respondents read the following preamble:

“There is a lot of debate these days about what does and does not constitute racism. In the following questions, we ask what you believe and what you think others believe.”

The questions asking about the respondents’ personal beliefs were as follows:

⁴ Respondents who were assigned to the partisan identity survey and did not indicate a partisan identity, were still asked about their own personal beliefs and with 50/50 probability asked about the beliefs of either Democrats or Republicans. For these respondents, we are only using the personal beliefs data in this study.

The questions asking about the respondents' normative expectations of their group (American or partisan) were as follows:

Measures of partisan and American normative expectations

“Next, we will describe a series of situations and ask you to tell us whether other *Republicans/Democrats/Americans* would consider the situation to be an instance of racism or whether other *Republicans/Democrats/Americans* would not consider the situation to be an instance of racism. In some cases, it may be difficult to decide, but we would like you to give us your best answer.

These questions are intended to examine how much *Republicans/Democrats/Americans* know about other *Republicans'/Democrats'/Americans'* opinions. We will select one of the questions and **you receive an additional \$0.50 if you correctly choose the most popular response among *Republicans/Democrats/Americans***. Remember: these questions are not about your own opinions. They are about choosing the same option as other *Republicans/Democrats/Americans*.”

1. A white person uses the n-word referring to a black person.
2. A white person objects to a member of their family dating a black person.
3. A white person says that welfare recipients could get ahead if they were willing to work hard.
4. A white person says inner cities are full of criminals.
5. Black Americans are incarcerated at more than five times the rate of white Americans.
6. Black Americans are more than twice as likely to be denied home mortgages compared to white Americans.
7. A company does not consider race when making hiring decisions. As a result, only 2% of their employees are black.
8. A prestigious university does not consider race when making admissions decisions. As a result, black students are underrepresented by 50%.

Other *Republicans/Democrats/Americans* would consider this situation an instance of racism

Other *Republicans/Democrats/Americans* would NOT consider this situation an instance of racism

These questions measure both the individual normative expectation and, when aggregated, the group norm. The group norm is defined as the majority agreement on normative expectations. The size of that majority reveals the strength of the norm. A 50/50 split on normative expectations (e.g., one half of Republicans believe other Republicans would say that the situation is not racism and the other half of Republicans believe other Republicans would say the situation is racism) would reveal that there is no group norm. The more that normative expectations are in agreement with each other, the stronger is the underlying norm.

Results: Figures 3 and 4 plot the personal beliefs and perceived partisan and American norms of Democrats and Republicans.⁵ From Figure 3, we can see that partisan norms and personal beliefs are in line with each other. The group norm is the normative expectation answered by the majority of group members. The majority of Democrats and Republicans express personal opinions that are consistent with the group norm. In other words, Democrats and Republicans are both, on average, correct about what others within their party believe.

The differences between Democrats and Republicans are striking. A large majority of both groups agree that the ‘explicit racism’ situations are instances of racism. And, on average, both expect other co-partisans to agree that the situations are instances of racism. Among Democrats, this is also true for the ‘subtle racism’, ‘institutional racism’ and ‘colorblind racism’, although the majorities are not as large. Among Republicans, however, a sizable majority

⁵ The order in which the personal beliefs and normative expectations questions were asked was randomized. For the purposes of these figures, the personal beliefs reported are those of respondents asked the personal beliefs question first and the normative expectations reported are those of respondents asked normative expectations first. SI6&7 include figures where the personal beliefs and normative expectations of all respondents are included.

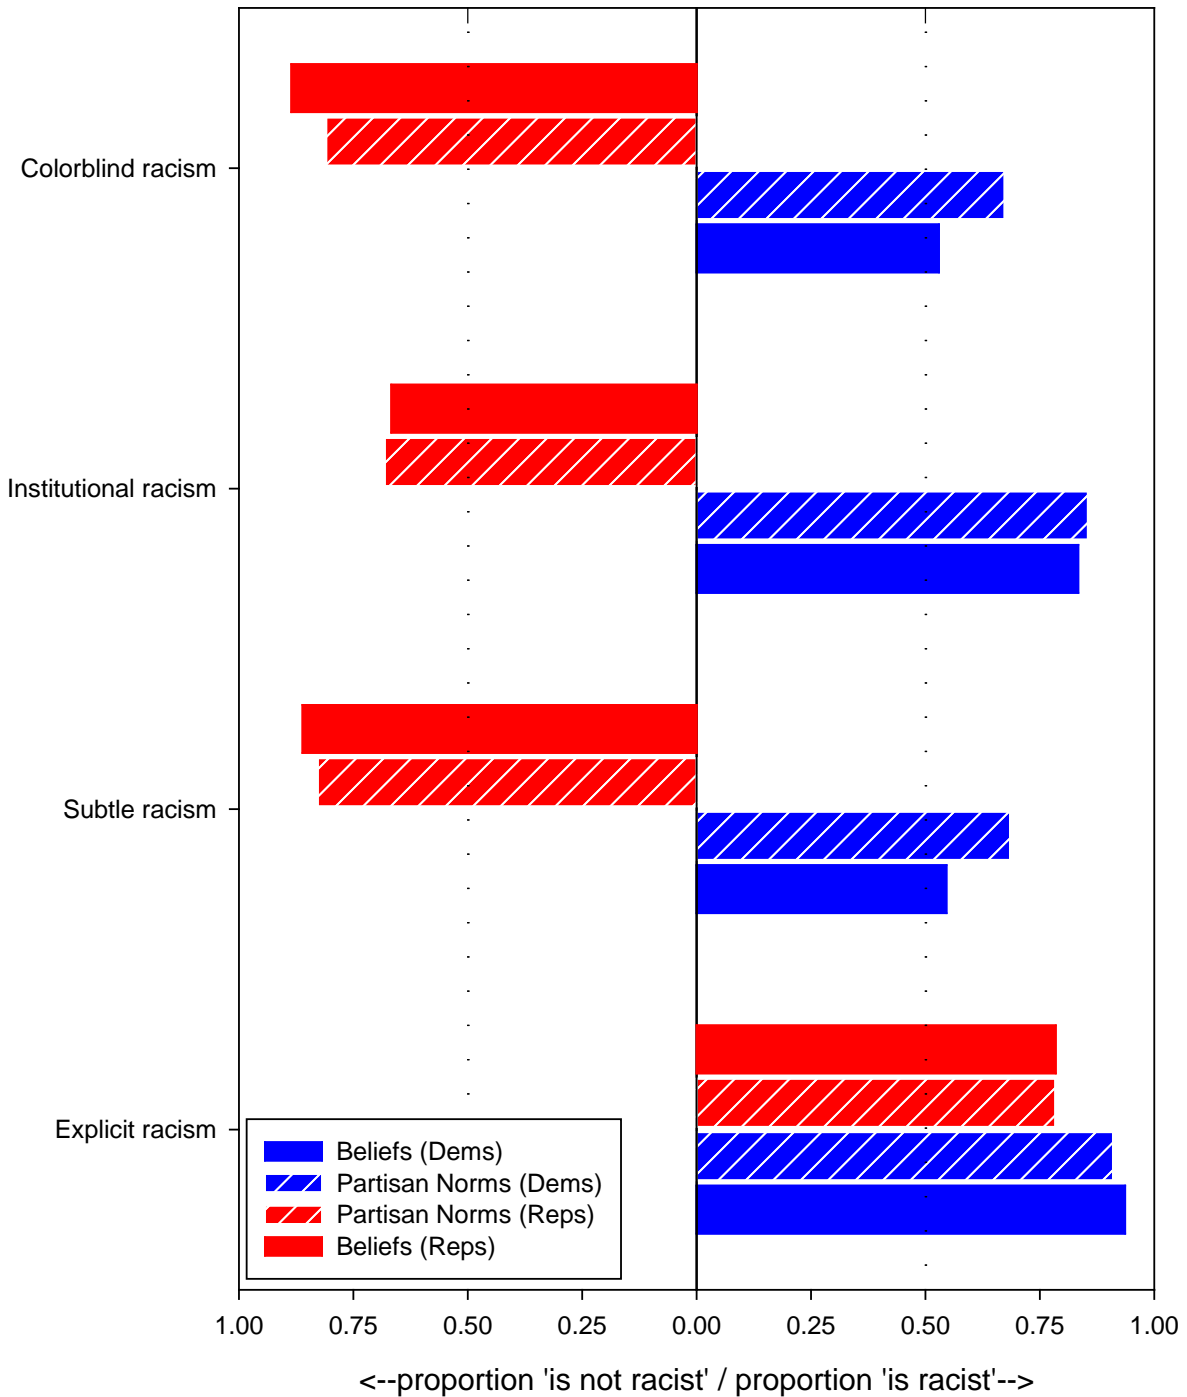
believes these situations are not instances of racism. And, on average, expect other Republicans to agree that these situations are not instances of racism. With the exception of 'explicit racism,' the average Democrat disagrees with the average Republican on what constitutes racism, and each correctly believes their opinion is in line with the majority of their co-partisans. The other distinction between Democrats and Republicans is that there is greater agreement amongst Republicans on normative expectations. In other words, the group norms are stronger.

In contrast, Figure 4 shows that Republicans and Democrats have normative expectations of Americans that do not always align with their own personal beliefs. The figure also includes the personal beliefs of all Americans to demonstrate the accuracy of these normative expectations. Regarding the 'explicit racism' situations, both Democrats and Republicans correctly expect other Americans, on average, to agree with them that these situations are instances of racism. However, a (small) majority of Democrats correctly expect most Americans to disagree with them that the 'subtle racism' and 'colorblind racism' situations are instances of racism. Similarly, a (small) majority of Republicans correctly expect most Americans to disagree with them that the 'institutional racism' situations are not instances of racism. Interestingly, Republicans' and Democrats' expectations about other Americans are in line with each other. They correctly believe that Americans are fairly close to a 50-50 split on whether the 'subtle racism', 'institutional racism' and 'colorblind racism' situations are instances of racism. In other words, American identity norms regarding what counts as racism are perceived as weak by both Democrats and Republicans, except for those situations that

reflect 'explicit racism'. This is in stark contrast to the strength of the partisan norms, which are much stronger, especially among Republicans.


So far, we have focused on aggregate effects, but how do things look when we zoom in on individual citizens? We have already seen that the majority of partisans disagree with American norms on some conceptions of racism. Further, the fact that partisans do not 100% agree that each situation is an instance of racism or not, reveals that some partisans do not have beliefs that match party norms. But how many deviate from group norms, and are these deviations due to lack of awareness or intentional deviation?

Figure 3. Partisan Norms – YouGov Study



Note: solid bars denote the proportion of respondents indicating the situations for each category of racism are racist or are not racist (whichever was greater). N = 335(Dems); 263(Reps). Striped bars denote the proportion of respondents that expect other partisans would indicate the situation is racist or not racist (whichever was greater). N = 338(Dems); 273(Reps).

Figure 4. American Norms – YouGov Study



Note: solid bars denote the proportion of respondents indicating the situations for each category of racism are racist or are not racist (whichever was greater). N = 330(Dems); 254(Reps); 3,000(Americans). Striped bars denote the proportion of respondents that expect other Americans would indicate the situation is racist or not racist (whichever was greater). N = 345(Dems); 275(Reps).

Figure 5 plots the proportion of Republicans and Democrats that express beliefs (regarding racism) that deviate from the partisan norm, and divides these deviations between cases where the individual is aware they are deviating from the partisan norm and cases where the individual is unaware.⁶ Figure 6 does the same for deviations from the American norm. Both figures are based on the results from the YouGov Study. For our purposes, a group norm can be identified as the belief that a majority of respondents within the group holds, or the belief that a majority of respondents expect others within the group to hold. Because Republicans and Democrats within our sample have, on average, accurate opinions about what the majority within their (partisan and national) group believes, the two possible ways of identifying the group norm amount to the same thing. A respondent is then categorized as having a belief that is deviant from the group norms when they express a belief that differs from the majority belief/expectation. They are categorized as being aware of the deviation when they express an opinion that differs from the majority but correctly identifies the majority opinion. In other words, they correctly believe they are expressing a deviant belief. Respondents are categorized as being unaware of the deviation when they express an opinion that differs from the majority but incorrectly identifies the majority opinion. In other words, they incorrectly believe they are expressing the majority belief.

From Figure 5, we see that less than 20% of Republicans express a belief deviant from the partisan norm for all but one of the situations they were asked about. The lone exception is institutional racism, specifically item 2, which described disproportionately high mortgage denial

⁶ Figures 5 and 6 includes only the respondents that were asked about their personal beliefs before normative expectations.

racism for African Americans compared to White Americans. For each of the other scenarios, more than half, or about half, of the time a deviant belief was expressed, the Republican respondent was aware of it. There is greater deviation from the partisan norm amongst Democrats. Approximately 40% of Democrats expressed deviant beliefs on 'colorblind racism' and 'subtle racism' situations. These Democrats did not agree that the scenarios described constituted instances of racism, although the norm among Democrats is to say they are instances of racism. For each situation, more than half, or about half, of the time a deviant belief was expressed, the Democrat was aware of it. Given the substantial number of Democrats that disagree with the group norm but are aware of the normative expectations of the group, there is significant potential for norm priming to affect personal beliefs. This is not the case among Republicans. In other words, since the vast majority of Republicans already hold beliefs that align with their party on the question of what counts as racism, there is little additional room for movement. Among Democrats, however, there is plenty of room. We will return to this shortly.

Turning to American norms (Figure 6), we see that the proportion of Republicans who express beliefs deviant from the American norms is similar to the proportion of Republicans who express beliefs deviant from partisan norms. In other words, there is not a great deal of deviation. The beliefs of Democrats, on the other hand, deviate even more from American norms than from partisan norms. About 60% of Democrats express beliefs that deviate from the American norm for situations that correspond to colorblind and subtle racism. About 80% of Democrats expressed beliefs that deviate from the American norm for situations that correspond to institutional racism. A majority of the time, respondents were aware that they

were expressing beliefs deviant from the American norm. Normally this would indicate a potential for norm priming effects, but the weakness of the American norms suggests such effects might not be realized.

Figure 5. Partisan Norm Deviations – YouGov Study



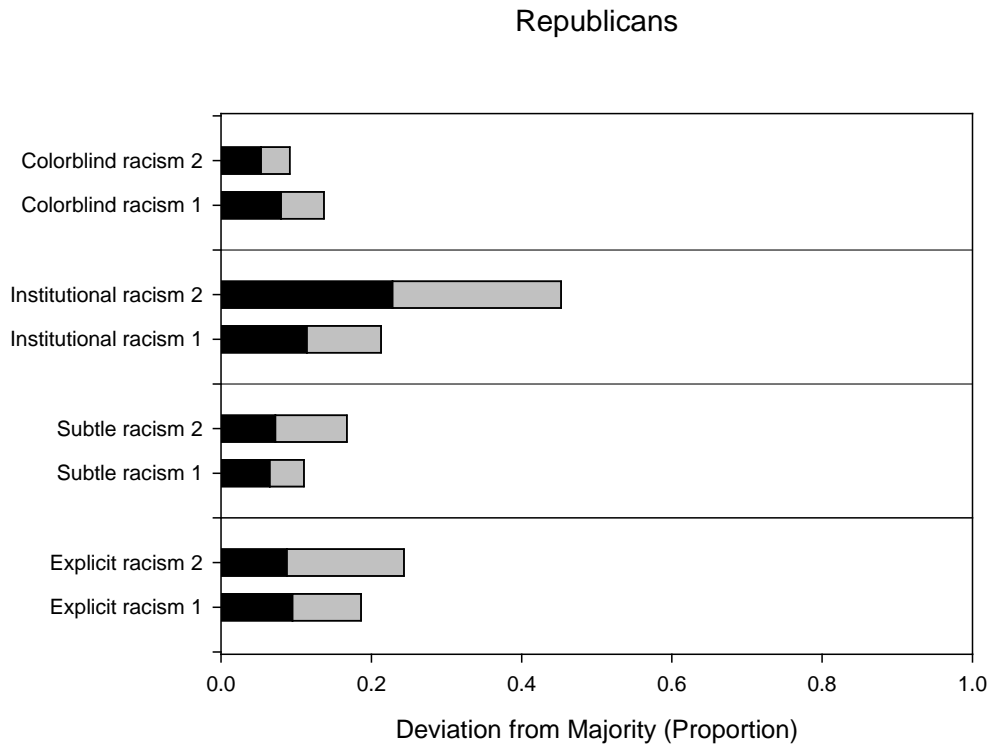
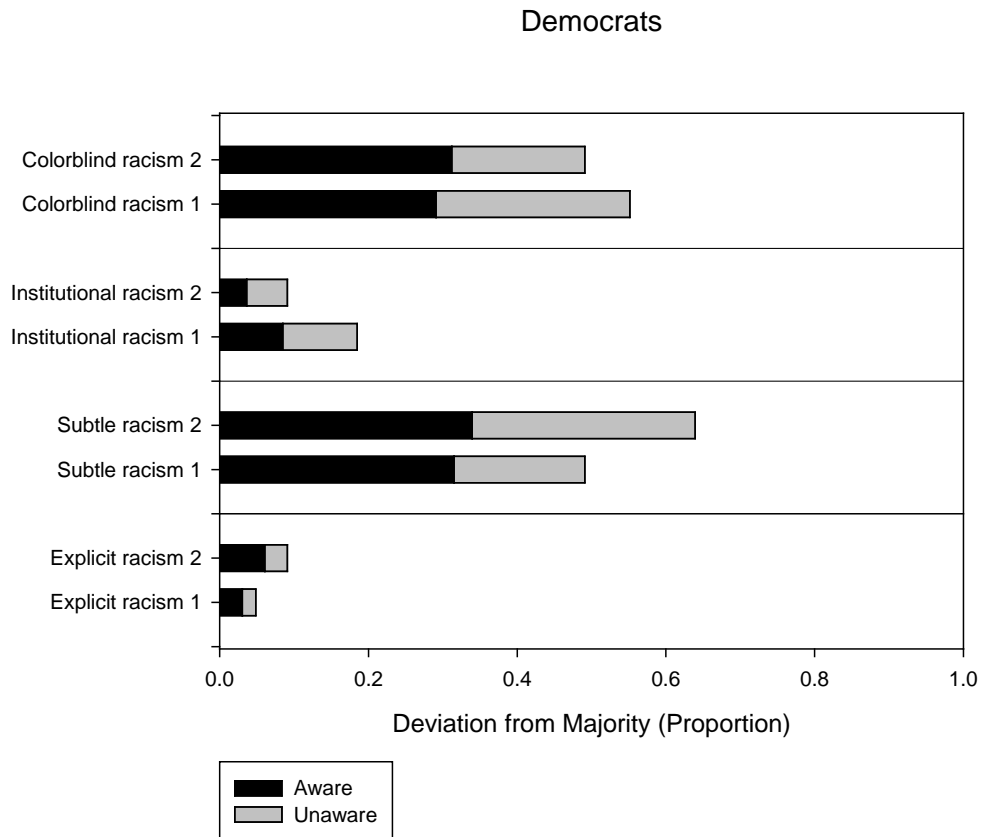
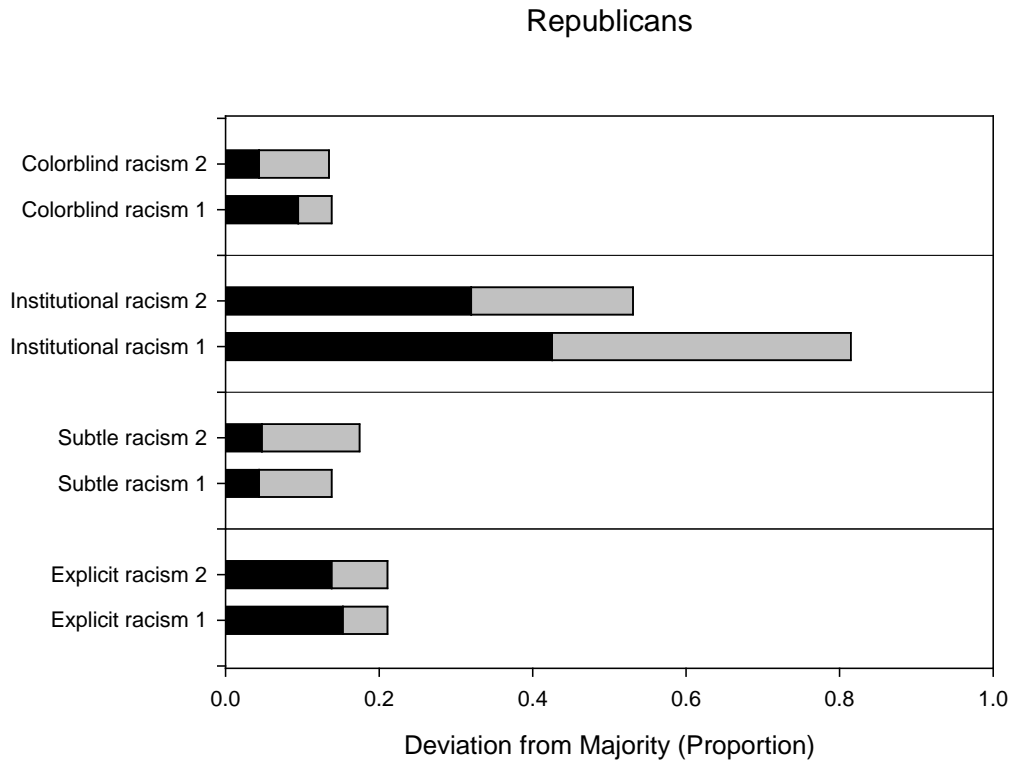


Figure 6. American Norm Deviations – YouGov Study



Finally, we turn to the question of whether partisans' beliefs about what does and does not constitute racism are exogenous or whether they can be altered by priming partisan group norms.

Priming Norms to Change Personal Beliefs

Research Design: These YouGov survey also included a norm priming experiment, which we repeated in three other surveys of Americans who had previously reported a partisan identity. These additional surveys were conducted through the crowdsourcing platform, Prolific. For Prolific Study 1 (October 18 – 19, 2022), we recruited a sample of 1,314 respondents, and for Prolific Study 2 (February 21 – 22, 2022), we recruited 1,458 respondents. For Prolific Study 3 (April 4 – 12, 2022), respondents from Prolific studies one and two were recontacted, with 2,329 respondents completing Prolific Study 3.

Respondents to the YouGov and Prolific studies were all randomized (with equal probability) to be asked: 1) first about their own personal beliefs and then the norms of their identity group; or 2) first about the norms of their identity group and then their own personal beliefs.⁷ Those who were first asked about the norms of their identity group, we refer to as the treatment group, because group norms were primed before respondents were asked about their personal beliefs. Those who were first asked about their own personal beliefs we refer to as the control group. Note that the priming in the treatment group does not occur through the

⁷ Table 1 presents the flow of the YouGov and Prolific survey questionnaires.

provision of any external information. The priming occurs by asking the respondent what they already believe about the group norms, and in doing so they remind themselves of those norms.

Prolific Study 3 was a recontact survey. Those who participated in Prolific studies 1 and 2 were recontacted and again asked about their personal beliefs and their normative expectations. Those that were asked about partisan [American] norms in studies 1 and 2 were again asked about partisan [American] norms in Study 3. Prolific Study 3 participants were randomized to be in the norm prime (treatment) group or the control group. Prolific Study 3 then allows us to estimate the effects of this norm prime conditioning on their prior expectations regarding the norm as measured in Prolific studies 1 and 2.

Results: Figures 7 and 8 plot the effects of priming partisan and American norms on expressed beliefs regarding racism. These results are again from the YouGov Study, but we now supplement these results with those of Prolific studies 1 and 2. As explained in the research design section, respondents were randomly assigned to either the treatment (norm prime) group or to the control group. Those in the treatment group were asked about the norms of their group before being asked their own personal beliefs. They were also given a monetary incentive to identify the norms of the group and not just their own beliefs. In this way, those in the treatment group were primed to think about the norms of their group before reporting their own beliefs. Those in the control group were asked about their personal beliefs before being asked about norms. Therefore, the control group was not primed to think about norms before reporting their personal beliefs. If norms have the potential to influence beliefs, we

might expect those that are primed to think about those norms to express beliefs that are more consistent with norms, compared to those who were not primed to think about norms.

It should be noted that some individuals (a minority) incorrectly perceive the norms of their group. It is possible that these individuals will be influenced by these incorrect norm perceptions, which would move them away from the actual group norm. The treatment effects plotted in figures 7 and 8 are the probability that an individual expresses a belief that the situation constitutes an instance of racism (regardless of whether their perception of the norm is correct or incorrect). We will subsequently examine treatment effects conditioning on perceptions of the norm, using Prolific Study 3.

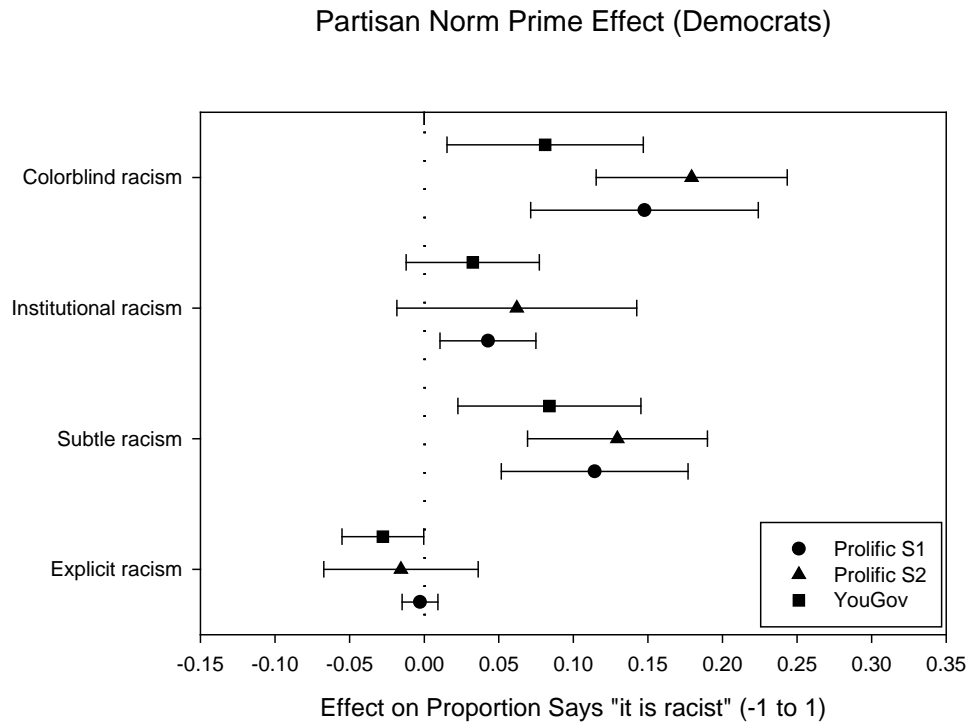

In Figure 7, we see few effects for priming partisan norms on Republicans. In two out of three studies, the norm prime increases the probability that a Republican respondent agrees that a situation corresponding to colorblind racism is indeed racism. These effects are a 6% (YouGov) and 8% (Prolific 2) increase in the probability. However, given that the norm among Republicans is to say that all the scenarios, other than the explicit racism scenarios, are not examples of racism, it will be important to examine this effect conditioning on expectations regarding the norm. Overall though, the weakness of norm priming effects is consistent with the fact that most Republicans express norm consistent opinions, without being primed.

There is more evidence of norm priming effects amongst Democrats. In all three studies, the norm prime increases the probability of a Democrat agreeing that situations which correspond to colorblind and subtle racism are instances of racism. For 'colorblind racism', the effects range from an 8 to 18 percent increase. For 'subtle racism', the effects range from an 8 to 13 percent increase. The norm amongst Democrats for both are that these are instances of

racism. Therefore, the norm prime is moving opinion in the direction of the norm. These norm priming effects are consistent with the size of the group of Democrats that express beliefs that they know are inconsistent with the group norm. Since Republicans hold firm to their narrow beliefs about what counts as racism and Democrats move toward party's norm of conceptualizing racism more broadly, the overall result is increased belief polarization.

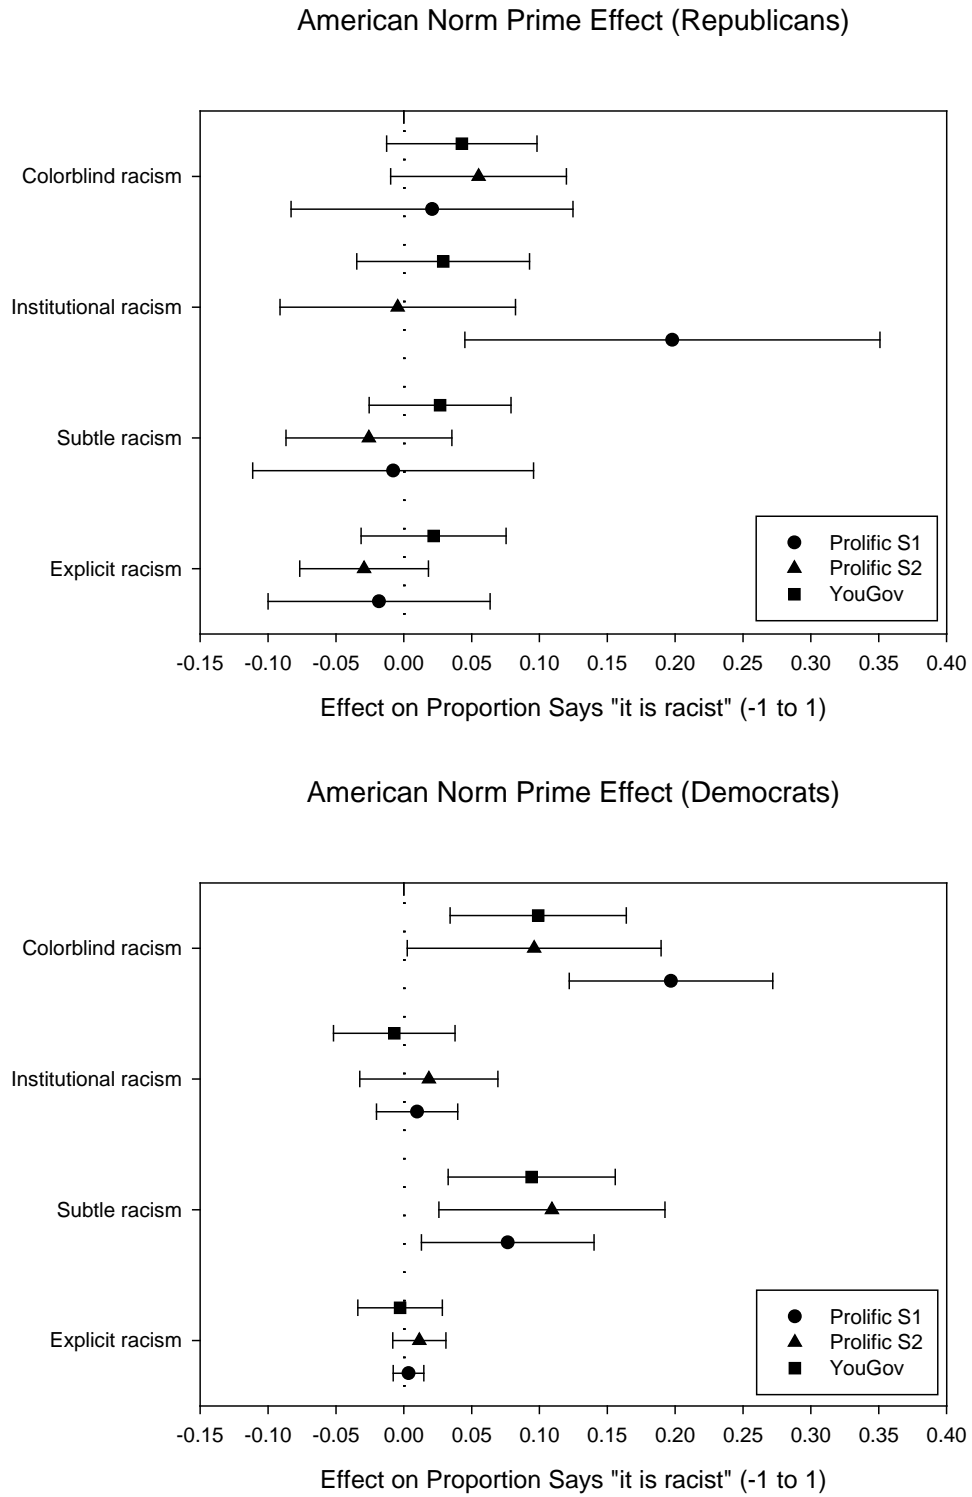
With respect to American norms (Figure 8), there is again little evidence that priming such norms affect Republicans beliefs about what counts as racism. Among Democrats, the norm prime in all three studies increases the probability of a Democrat agreeing that situations that correspond to colorblind and subtle racism are instances of racism. Given that Democrats (correctly) believe that the American norm is to disagree that these are instances of racism, it again will be important to examine this effect conditioning on expectations regarding the norm. For now, it appears that priming Democrats to think about the fact that the average American would not describe the scenarios as examples of racism motivates them to insist that they are examples of racism. Again, the overall result of Republicans being consistent with but unaffected by the American norm to conceptualize racism form narrowly, and Democrats shifting away from the American norms, is further polarization in Democrats' and Republicans' beliefs about what constitutes racism.

Figure 7. Partisan Norm Prime Effects



Note: average effect of priming partisan identity norms on proportion of respondents indicating the situations for each category of racism are racist. $N =$ Reprs, 113(Prolific S1), 396(Prolific S2), 536(YouGov); Dems, 541(Prolific S1), 387(Prolific S2), 673(YouGov).

Figure 8. American Norm Prime Effects



Note: average effect of priming American identity norms on proportion of respondents indicating the situations for each category of racism are racist. $N =$ Reps, 112(Prolific S1), 364(Prolific S2), 529(YouGov); Dems, 548(Prolific S1), 384(Prolific S2), 675(YouGov).

As explained in the research design section, Prolific Study 3 was a recontact survey. This allows us to estimate the effects of norm priming in Prolific Study 3, conditioning on the respondent's prior expectations regarding the norm as measured in Prolific studies 1 and 2. There are two instances in which this was of particular importance: the effect of partisan norm priming on Republicans and the effect of American norm priming on Democrats. Figures 9 and 10 plot these norm priming effects, conditioning on prior normative expectations.

Unfortunately, this study was unable to replicate the Republican priming effect. This time around, we found no evidence that priming partisan norms increased the likelihood of Republicans characterizing the colorblind racism scenarios as examples of racism. However, we were able to replicate the American norm priming effect among Democrats (Figure 9). By conditioning on prior normative expectations, we see that the effects are all among those Democrats who believe the American norm is to characterize the colorblind racism scenarios as examples of *non*-racism. This is true of situations related to colorblind and subtle racism, as we saw when we estimated the priming effects without conditioning on prior normative expectations. This also appears to be the case for our institutional racism scenarios. When Democrats believe the American norm is to conceptualize racism more narrowly (i.e., NOT count these scenarios as examples of racism), forcing them to think about these American norms makes them more likely to report believing the opposition—that these institutional and colorblind racism scenarios DO count racism. In other words, these Democrats are repelled when they believe other Americans conceptualize racism narrowly.

Figure 9. Heterogeneity in Effect of Partisan Prime on Republicans – Prolific Study 3





Figure 10. Heterogeneity in Effect of American Prime on Democrats – Prolific Study 3



Conclusion

What counts as racism and what does not, according to Americans? Do shared American norms exist regarding the concept of racism, or have Republicans and Democrats come to conceptualize racism differently? We find that a clear majority of Americans characterize scenarios designed to represent the concept of ‘explicit racism’ as examples of racism. We find much less agreement when it comes to other conceptualizations of racism. A small majority of Americans also characterize scenarios designed to represent ‘institutional racism’ as examples of racism. On the other hand, only a minority of Americans believe scenarios designed to represent ‘subtle racism’ and ‘colorblind racism’ actually constitute racism.

Across these conceptualizations of racism, party identification constitutes the strongest predictor of division. Apart from ‘explicit racism’, the average Democrat disagrees with the average Republican on what constitutes racism, and, on average, each correctly believes their opinion is in line with that of their co-partisans. The norm among Democrats is to endorse a broad conceptualization of racism, characterizing scenarios designed to represent explicit, subtle, institutional and colorblind racism all count as racism. The majority of Democrats express opinions in line with these norms. The norm among Republicans is to endorse a much narrower conceptualization of racism, characterizing only scenarios designed to represent explicit racism as examples racism. And, like Democrats, the majority of Republicans express opinions in line with their party’s norms.

When it comes to American norms regarding what does and does not constitute racism, Republicans and Democrats tend to agree on what other Americans believe, but they do not necessarily express personal beliefs in line with those norms. Most Democrats disagree with the

American norm that scenarios designed to represent 'subtle racism' and 'colorblind racism' do not constitute examples of racism. Most Republicans disagree with the American norm that scenarios designed to represent 'institutional racism' constitute examples of racism. Republican and Democrat also correctly believe that Americans norms regarding 'subtle racism', 'institutional racism' and 'colorblind racism' are relatively weak. This weakness is due to the lack of agreement amongst Americans about what constitutes racism.

When it comes to partisan norms, a substantial number of Democrats knowingly disagree with their party's norms, in particular norms regarding subtle and colorblind racism. An even larger number of Democrats also knowingly disagree with American norms, in particular norms regarding institutional racism. Republicans are far less likely to report beliefs that are inconsistent with the norms of their party or American norms regarding what does and does not constitute racism.

What happens when survey participants are primed to think about either partisan norms or American norms prior to reporting their own personal beliefs? Among Democrats, we see clear effects. Partisan norm prime increases the probability of a Democrat agreeing that scenarios which correspond to the concepts of colorblind and subtle racism constitute examples of racism. The American norm prime also increases the probability of a Democrat agreeing that scenarios designed to correspond to colorblind and subtle racism constitute examples of racism across all three studies. Among Republicans we see minimal evidence of priming, but this is likely do to the fact that most Republicans already hold beliefs that are consistent with both partisan and American norms. Since Republicans already hold beliefs that are consistent with norms, there is little room for movement.

In both cases, the effect of norm priming is to move Democrats even further away from Republicans regarding beliefs about what constitutes racism. The priming effect of American norms is interesting because it seems to be backfire, pushing Democrats away from normative expectations. We see that when Democrats believe the American norm is to conceptualize racism narrowly and they are primed to think about these American norms, they themselves become more likely to endorse broader conceptualizations of racism.

Overall, there do not appear to be strong shared American norms regarding what does and does not count as racism, the only exception being explicit racism. Partisan norms, on the other hand, are quite strong and partisans tend to report personal preferences consistent with these norms. Among Democrats, this consistency appears to be at least partially attributable to conformity pressure, since Democrats bring their reported beliefs into closer alignment with their party's norms after being primed to think about them. Although it is impossible to determine for sure without more extensive longitudinal data, Republicans appear to have either already adjusted their beliefs about racism to conform to the norms of their party or this is why they identify as Republicans in the first place. This is evident from the fact that they came into the study with personal beliefs that already matched both partisan and American norms quite closely, making priming effects difficult to detect. If distinct norms regarding what does and does not constitute racism are indeed developing within parties, and shared American norms do not appear to exist, the prospect for productive dialog on matters of racial justice does not look promising.

References

- Allport, Gordon W. 1954. *The Nature of Prejudice*. New York, NY: Doubleday Books.
- Anoll, Allison. 2022. *The Obligation Mosaic*. Chicago, IL: University of Chicago Press.
- Banks, Antoine. 2014. *Anger and Racial Politics: The Emotional Foundation of Racial Attitudes in America*. New York, NY: Cambridge University Press.
- Banks, Antoine and Heather Hicks. 2019. "The Effectiveness of Racialized Counterstrategy." *American Journal of Political Science* 63(2): 305-322.
- Banks, Antoine and Nicholas Valentino. 2012. "The Emotional Substrates of White Racial Attitudes." *American Journal of Political Science* 56(2): 286-297.
- Bonilla-Silva, Eduardo. 2018. *Racism without racists: Color-blind racism and the persistence of racial inequality in America* (5th Ed.). New York, NY: Rowman and Littlefield.
- Carmichael, Stokely and Charles V. Hamilton. 1967. *Black power: The politics of liberation*. New York, NY: Random House.
- Converse, Philip. 1964. "The nature of belief systems in mass publics." In *Ideology and discontent*, ed. David E. Apter. New York, NY: Free Press.
- Gilens, Martin. 1999. *Why Americans Hate Welfare: Race, Median, and the Politics of Antipoverty Policy*. Chicago, IL: University of Chicago Press.
- Groenendyk, Eric, Erik Kimbrough and Mark Pickup. 2022. "How Norms Shape the Nature of Belief Systems in Mass Publics." *American Journal of Political Science* (Published Online).
- Huddy, L. and N. Khatib. (2007) "American Patriotism, National Identity, and Political Involvement." *American Journal of Political Science* (51)1:63-77.

- Hyman, Herbert H. and Paul B. Sheatsley. 1956. "Attitudes Toward Desegregation." *Scientific American*, 19 (6): 35-39.
- Hyman, Herbert H. and Paul B. Sheatsley. 1964. "Attitudes Toward Desegregation." *Scientific American*, 195 (6): 16-23.
- Kinder, Donald R. and D. Roderick Kiewiet. 1981. "Sociotropic politics: The American case." *British Journal of Political Science* 11: 129-161.
- Kinder, Donald R. and Lynn M. Sanders. 1996. *Divided by color: Racial politics and democratic ideals*. Chicago, IL: University of Chicago Press.
- Kinder, Donald R. and David O. Sears. 1981. "Prejudice and politics: Symbolic racism versus racial threats to the good life." *Journal of Personality and Social Psychology* 40: 414-431.
- Lee, Taeku. 2000. *Mobilizing Public Opinion: Black Insurgency and Racial Attitudes in the Civil Rights Era*. Chicago, IL: University of Chicago Press.
- Legros, Sophie and Beniamino Cislighi. 2020. "Mapping the Social-Norms Literature: An Overview of Reviews." *Perspective on Psychological Science* 15(1): 62-80.
- McConahay, John. B. 1983. "Modern racism and modern discrimination: The effect of race, racial attitudes, and context on simulated hiring decision." *Personality and Social Psychology Bulletin* 9: 551-558.
- Mendelberg, Tali. 2001. *The race card: Campaign strategy, implicit messages, and the norm of equality*. Princeton University Press.
- Miller, Dale T. and Deborah A. Prentice. 2016. "Changing Norms to Changing Behavior. Annual Review of Psychology 67: 339-361.

Pickup, Mark, Erik Kimbrough and Eline de Rooij. 2021. "Expressive Politics as (Costly) Norm Following." *Political Behavior* (Published online).

Sears, David O., John J. Hetts, Jim Sidanius, and Lawrence Bobo. 2000. "Race in American politics: Framing the debates." In *Racialized Politics: The debate about racism in America*, ed. D. O. Sears, J. Sidanius, and L. Bobo. Chicago, IL: Chicago University Press.

Schuman, Howard, Charlotte Steeh, Lawrence Bobo, and Maria Krysan. 1997. *Racial attitudes in America: Trends and interpretations*. Cambridge, MA: Harvard University Press.

Valentino, Nicholas A., Vincent L. Hutchings, and Ismail K. White. 2002. Cues the Matter: How Political Ads Prime Racial Attitudes During Campaigns. *American Political Science Review* 96(1): 75-90.

Valentino, Nicholas A., Fabian G. Neuner, and L. Matthew Vandenberg. 2018. "The Changing Norms of Racial Political Rhetoric and the End of Racial Priming." *Journal of Politics* 80(3): 757-771.

White, Ismail K. and Chryl N. Laird 2020. *Steadfast Democrats: How Social Forces Shape Black Political Behavior*. New Jersey: Princeton University Press.

Supplementary Information 1: Survey Questions

The YouGov Study and Prolific Study 2 included a survey preamble (see below) and then after age and gender questions, respondents were given an attention check (see below). Respondents in Prolific Study 3 (the recontact study) were only given the survey preamble.

Survey preamble

This is an important survey.

There will be one or more attention checks.

And there is an opportunity to increase your payment depending on your responses.

Attention check

We'd like to know how you feel about local news coverage. Please read this short article. On the next page, we will ask you a few questions about your reactions to this article.

SUSPECT ARRESTED FOR STRING OF BANK THEFTS

New Westminster Police have arrested a suspect who gave their driver's license to a bank teller during a robbery.

According to court documents, Alex Lee is accused of robbing four New Westminster banks between October 3 and November 5, 2020.

During a robbery on November 5, the police say Lee was tricked into giving the teller their driver's license.

According to court documents, Lee approached the counter and presented a demand note for money that said "I have a gun." The teller gave Lee about \$500, which the suspect took.

Documents say Lee then asked the teller for more money. The teller told Lee a driver's license was required to use the machine to get out more cash. Lee reportedly then gave the teller their license to swipe through the machine and then left the bank with about \$1000 in additional cash, but without the ID.

Detectives arrested Lee later that day at the address listed on the ID.

How was Lee identified by police for the allegedly committed crime?

A police officer recognized Lee

From video surveillance

Because Lee left their ID
Lee surrendered to authorities
None of the above

Measure of partisan identity

“Generally speaking, do you usually think of yourself as a Republican, a Democrat, an Independent, or what?”

Republican
Democrat
Independent
Other Party⁸

If ‘Republican’ or ‘Democrat’: “Would you call yourself a strong Republican or not a very strong Republican/Democrat?”

Strong
Not very strong

If ‘Independent’ or ‘Other Party’: “Do you think of yourself as closer to the Republican or Democratic party?”

Closer to the Republican Party
Closer to the Democratic Party
Neither

Measures of partisan and American identities as social identities

“How important is being a Republican/Democrat/American to you?”

0 (Not very important) to 10 (Very important)

How well does the term Republican/Democrat/American describe you?

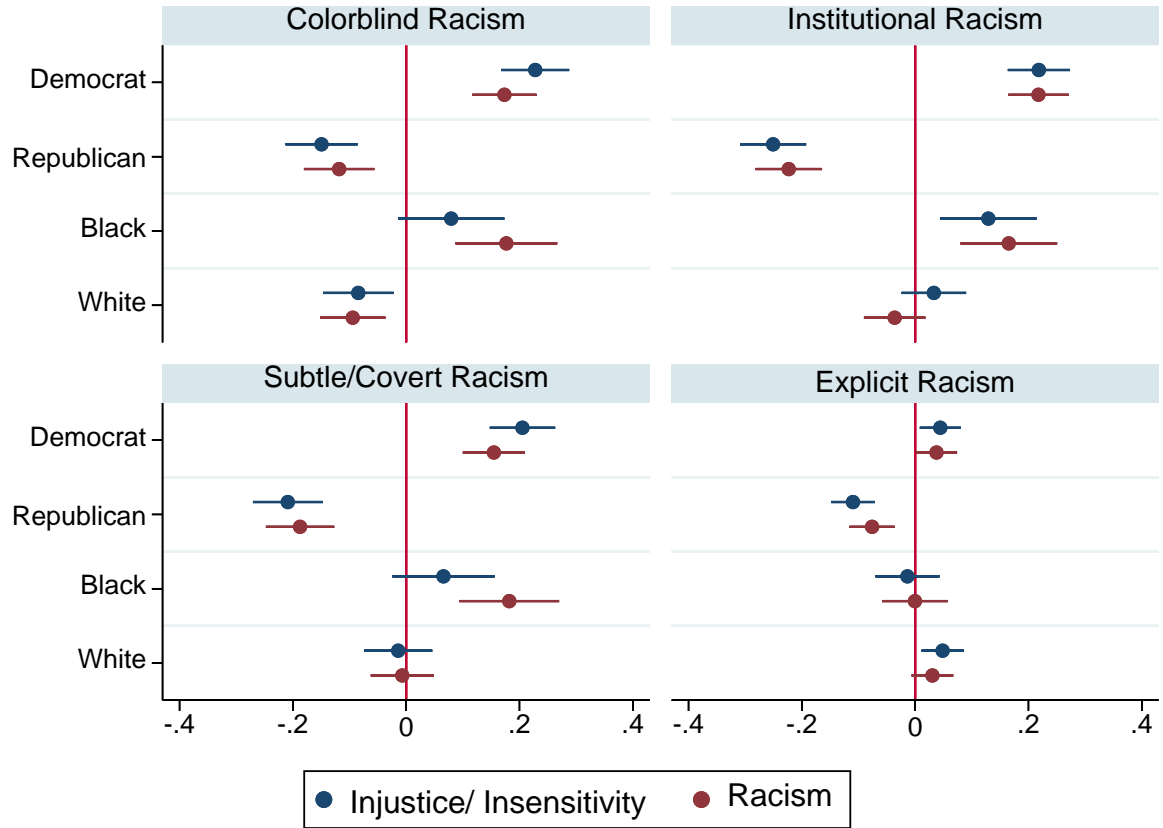
0 (Not very well) to 10 (Very well)

To what extent do you see yourself as typical for someone that is a Republican/Democrat/American?

0 (Not very typical) to 10 (Very typical)

⁸ The YouGov Study included a “Not Sure” response category.

Supplementary Information 2: Wording Experiment Results



Supplementary Information 3: Survey Flow

Table S1. Survey flow

Prolific Study 1		Prolific Study 2/YouGov Study		Prolific Study 3	
Consent					
Attention text					
Partisan identity					
Assignment to American or partisan identity questionnaire				Same assignment as Study 1 or 2	
Personal beliefs	Identity strength	Identity strength	Identity strength	Identity strength	Identity strength
Identity strength	Identity norms	Personal beliefs	Identity norms	Personal beliefs	Identity norms
Identity norms	Personal beliefs	Identity norms	Personal beliefs	Identity norms	Personal beliefs

*Bolded cells identify the identity norm prime; shaded cells identify the outcome measure.

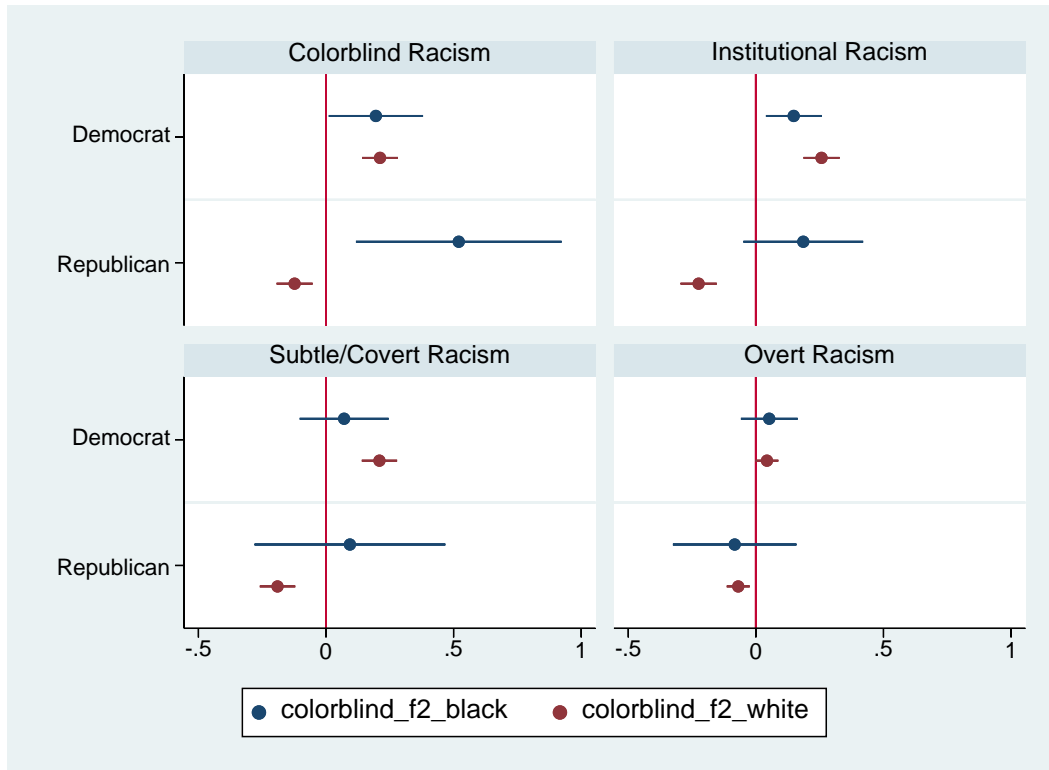
**Respondents in studies one and two were asked about either their American or partisan identity. Respondents in study three (recontact study) were asked about the same identity as they were in study one or two.

The survey questionnaires for all five studies began with a consent form. The YouGov study and Prolific studies 2 and 3 included some attention focusing and/or testing text (see SI1). The YouGov studies made use of the standard ANES measure of party identification to construct the standard 7-point scale. Prolific studies 1 and 2 confirmed each respondent’s party identification using the same measure (see SI1). In all studies the outcome variables of interest are respondents’ expressed personal beliefs (shaded cells in Table S1).


In Prolific Study 1, the treatment is the combination of the strength of identity and identity norms questions. This is illustrated in Table S1. In the other studies, everyone is asked about the strength of identity prior to assignment to treatment and control, and the treatment is only the norms question. In Prolific Study 1, both group identity and norms were primed as part of the treatment, which may have led to bigger effects. In the other studies, only norms

were primed to ensure treatment effects could be attributed specifically to norms, independent of identity priming.

Supplementary Information 4: Race Crossed with Party Identification



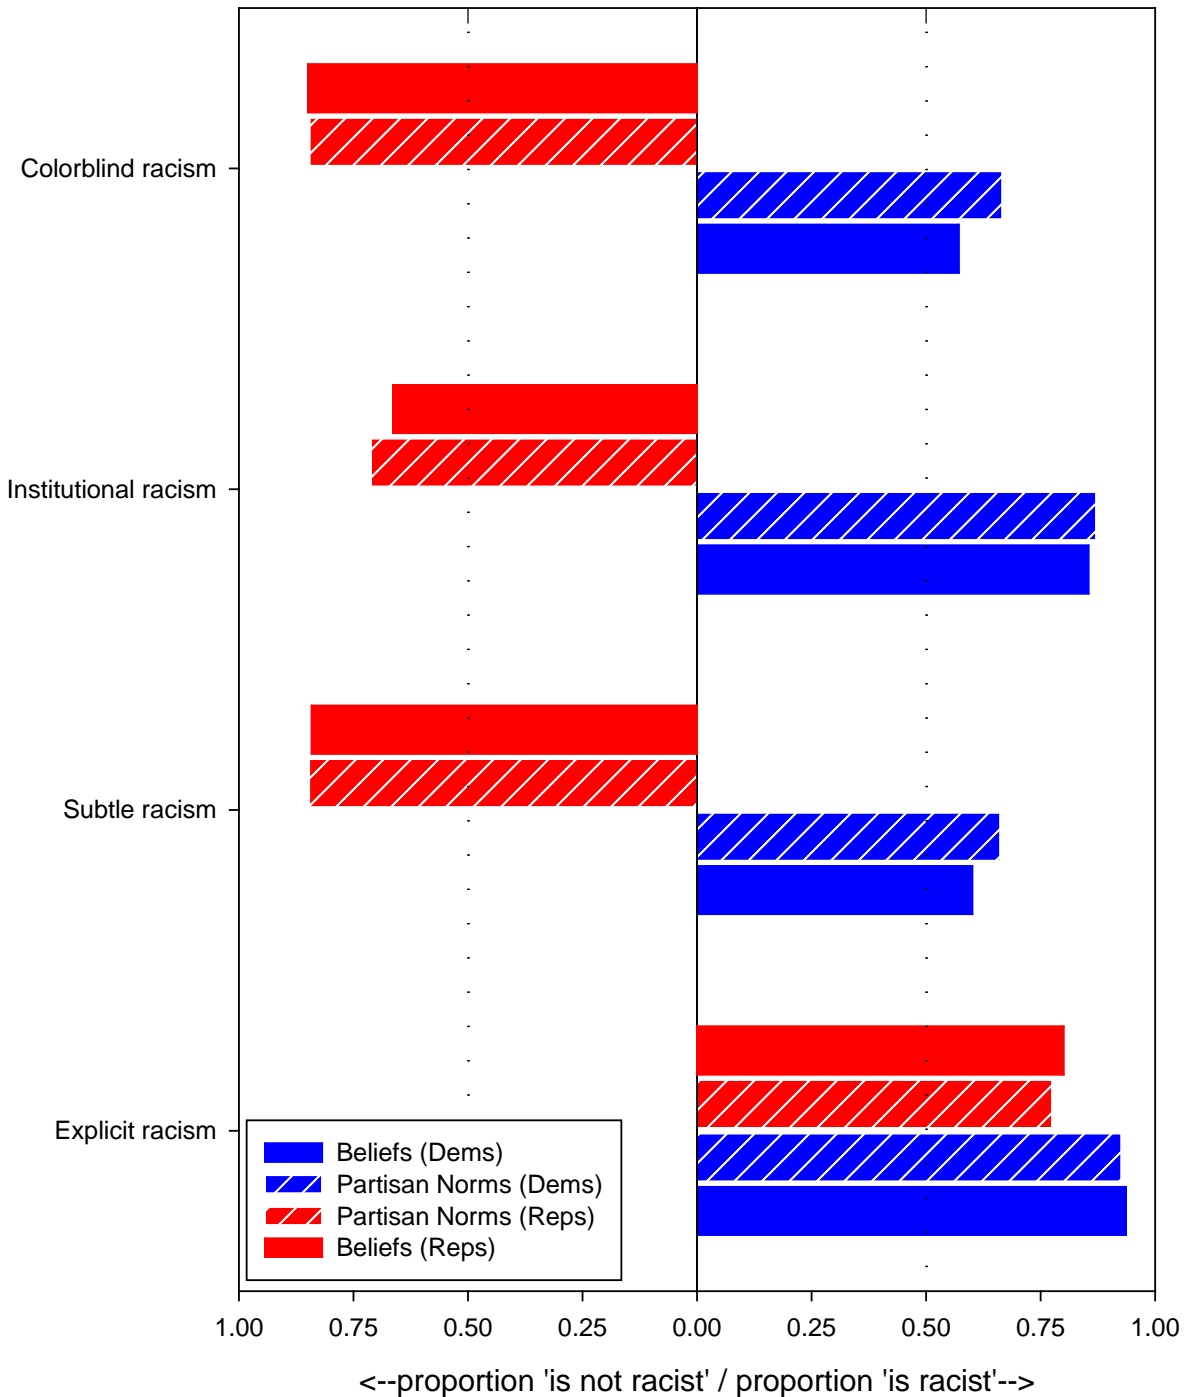
Supplementary Information 5: Strength of Identity as a Social Identity – YouGov Study



Note: average score on 0 to 10 scale for respondents from YouGov on three partisan/American identity as a social identity questions. $N = 673$ (Partisan ID (Dems); 536 (Partisan ID, Reps); 675 (American ID, Dems); 529(American ID, Reps).


The SI5 figure shows the distribution of partisan and American identity strengths using boxplots. These scores come from averaging across the three partisan/American social identity questions described in the SI1 (Cronbach’s alpha: Republican ID 0.91; Democrat ID 0.92; American ID 0.90). The theoretical range is 0 to 10. American identity strength among Republicans stands out as being stronger than American identity strength among Democrats. Republicans’ American identity strength is also stronger than their partisan identity as well as the party identity strength of Democrats.

Supplementary Information 6: Partisan Norms – YouGov Study



Note: solid bars denote the proportion of respondents indicating the situations for each category of racism are racist or are not racist (whichever was greater). N = 1,348(Dems); 1,065(Reps). Striped bars denote the proportion of respondents that expect other partisans would indicate the situation is racist or not racist (whichever was greater). N = 673(Dems); 536(Reps).

Supplementary Information 7: American Norms – YouGov Study



Note: solid bars denote the proportion of respondents indicating the situations for each category of racism are racist or are not racist (whichever was greater). N = 1,348(Dems); 1,065(Reps); 3,000(Americans). Striped bars denote the proportion of respondents that expect other Americans would indicate the situation is racist or not racist (whichever was greater). N = 675(Dems); 529(Reps).

