

Racial Priming: Issues in Research Design and Interpretation

Tali Mendelberg

In a critique of Huber and Lapinski (in this volume) I argued that their 2006 study failed to find evidence of racial priming and that this failure stands out in the recent accumulation of studies that do find racial priming. I argued further that this failure to replicate is the result of deficiencies in Huber and Lapinski's research. Huber and Lapinski (in this volume) respond by claiming that they did find evidence of racial priming among a subgroup, that their research is sound, that my research is flawed, and that the relevant literature does not comment on the differences between implicit and explicit messages. I show that 1) Huber and Lapinski's results demonstrate that their study produced null findings, 2) these null findings are caused by flaws in their study, 3) my research withstands their criticism, and 4) the relevant literature is in fact relevant and highlights the extent to which their null results are anomalous. There are, however, several points of agreement: 1) racial predispositions shape policy views, 2) these predispositions can be primed by cues and messages, and 3) these predispositions are primed by implicit racial messages. What remains at issue is the impact of explicit racial messages.

In their reply to my critique, Professors Huber and Lapinski argue that their study does replicate the literature on race cues and white Americans' opinion. They further argue that this successful replication helps to establish the validity of their study. They then argue that no research other than my own finds a greater power of implicit over explicit cues. They question the validity of the claim that explicit cues are weaker than implicit ones. In reply I show that Huber and Lapinski did not in fact replicate the findings of the race cue literature, that this failure to replicate is likely due to flaws in their research design, and that research on this question provides evidence for the greater power of implicit appeals.

Claim I: Huber and Lapinski Did Not Replicate the Racial Priming of 17 Other Studies

Huber and Lapinski argue that they did find racial priming, replicating previous scholarship, but that the priming they found is much more specific, located only among the poorly educated. However, the priming I reported and, more importantly, reported in the literature at large, exists in well-educated samples. For example, in the largest study, half of the 2,331 respondents had graduated from college and thus were well above the education cut-off used by Huber and Lapinski.¹ Huber and Lapinski thus fail to find any racial priming precisely among the subsample well-represented in the studies that do find racial priming. If Huber and Lapinski are correct that the effect lies entirely among the poorly educated, then high-education samples would find strong effects

only if the effect in their poorly educated subsample was far larger than the effect in Huber and Lapinski's poorly educated subsample. So even if Huber and Lapinski are correct that the effect obtains only among the poorly educated, the size of their effect would have to be smaller by orders of magnitude than the size in other studies, which represents a failure to replicate. Thus Huber and Lapinski's results are a failure to replicate; the results stand apart from 17 other studies in failing to find racial priming in the sample as a whole and in the education subgroup well represented in other studies.

Moreover, contrary to Huber and Lapinski's claim that they found priming among the poorly educated, the results displayed in their figure 5 show that they found no statistically significant racial priming among any subgroup (despite a large number of cases).² No line deviates above the confidence interval in 8 of the 8 tests shown, telling us that no racial cue had a statistically discernable effect on any of the four issue areas for either educational level. In addition, the same article notes that while both implicit and explicit conditions raise the impact of racial predispositions relative to the control, "*neither effect is statistically significant.*"³ Neither low-nor high-education respondents are moved by either implicit or explicit cues in the Huber and Lapinski study.

In their response Huber and Lapinski introduce a new table A and conclude based on it that implicit appeals prime racial predispositions among the poorly educated relative to the control.⁴ They conclude that their study did in fact find racial priming and therefore does replicate the literature and thus is not flawed. However,

this table directly contradicts the results in their original article.

I tried to replicate this table A from the data and codes supplied by Huber and Lapinski for the 2006 article, but the replication failed. Instead, I obtained the following incremental effects for racial predisposition in the implicit condition (using the effect in the control condition as the baseline): .99 (SE .48), $p = .04$ for decrease spending; .50 (SE .47), $p = .29$ for strengthen work requirements; .19 (SE .95) $p = .85$ for government help blacks; 1.00 (SE .55), $p = .07$ for affirmative action.⁵ These results show that in only one of the four analyses is there a statistically significant increase in the impact of racial predispositions; this effect obtains only for “decrease spending,” where the baseline effect has an unusual negative sign that would not be expected by any model. A reanalysis thus reinforces the conclusion that their study failed to replicate the literature: Huber and Lapinski’s data yield null results in every case except one analysis where the control represents an outlier baseline.⁶

In sum, the Huber and Lapinski study did not replicate the basic finding that race cues or messages racialize the opinions of white American. It failed to find evidence that a racial cue both increases the effect of racial views and that this effect differs statistically from 0. The failure to replicate not only holds for the sample as a whole; it extends to the low-education subgroup that Huber and Lapinski target.

Claim II: Huber and Lapinski’s Study’s Serious Deficiencies Caused the Failure to Replicate

I suggested in my critique that failures of experimental design and procedure can account for the discrepancy between Huber and Lapinski’s study and the literature. The results of their manipulation check show that these messages are nearly indistinguishable (2006). This can explain why Huber and Lapinski failed to find a greater impact of implicit than explicit messages.

The manipulation check results provided in the reply by Huber and Lapinski (table B) merely strengthen the conclusion that the experimental manipulation failed. Huber and Lapinski focus on the large percent increase in negative perception of the message from the implicit to the explicit condition. However, the large percentage increase is misleading because the baseline numbers are extremely low. The perceived differences between the implicit and explicit messages are in the single digits on a percentage point scale. So the impressive proportionate increase of 94.8 percent, for example, is actually meaningless because it represents an increase from 5 percent to 11 percent. Only a few people perceived either the implicit or explicit message negatively, so there is a deceptively

large proportionate increase but no meaningful difference between the messages.

Contrast this single-digit difference to differences reported in other studies. For example, Gilliam and Iyengar conducted a manipulation check to verify that subjects perceived the racial manipulation as intended.⁷ They exposed some subjects to a news story about a black criminal and others to a news story about a white criminal (among other conditions). Subjects were clearly able to distinguish between the two treatments. Of subjects who were exposed to the black suspect, 70 percent perceived a black suspect, while only 10 percent of subjects exposed to a white suspect perceived him to be black. This difference of 60 percentage points dwarfs the largest difference reported by Huber and Lapinski, which is 9 percentage points.

Even the small differences Huber and Lapinski obtain in perception of the experimental treatments are difficult to interpret because the manipulation check is only tenuously related to the treatment. Asking if issue ads are bad for democracy taps neither violations of the norm of equality nor the racial content of the message. Huber and Lapinski reply that they did not ask about the racial nature of the ad because they did not want this question to contaminate other variables of interest by introducing racial considerations artificially. But contamination can be avoided easily by asking about the racial content of the message at the end of the experiment. Regardless, the study failed to conduct an appropriate manipulation check

Huber and Lapinski further argue that any individual differences in interpretation of “issue ads” and “democracy” are irrelevant in an experiment. However, the problem is not the random differences of interpretation by respondents but rather the possibility that many people interpret “issue ads” and “democracy” similarly across conditions. This common nonracial interpretation may well dominate any racial differences in the treatment and consequently flatten the response across the conditions. And flattening we observe in the nine-point and other very small differences in perception.

In sum, the small differences that Huber and Lapinski obtain in perception of the treated messages are further clouded by the difficulty of interpreting the probe, whose nonracial content (“issue ads are bad for democracy”) may swamp the effect of racial differences across messages.

In my critique I also raised questions about rates of non-reception in Huber and Lapinski’s study. Huber and Lapinski referred readers seeking details to an earlier study conducted by Professor Lapinski. In that study a large number of respondents did not receive their ad. In their reply Huber and Lapinski provide additional data from a measure that asks respondents to indicate if they were able to download the video. Only a small percentage of respondents (8 percent and 11 percent in the two studies) indicated an inability to view the video. However,

Huber and Lapinski do not report the percentages of people who did not respond to this question. Consequently we do not know what percent did view the ad, only the percentage who did not view the ad. In addition, Huber and Lapinski report that they included in their statistical analyses even those respondents who indicated that they had not viewed the ad. Their reason is that non-viewing may contain systematic measurement error. However, including cases that may contain systematic error is not a compelling solution. At a minimum, we need to know how the results change by this estimation choice.

Finally, I criticized Huber and Lapinski's procedure of asking racial predisposition questions before the treatment in all conditions. This likely primes everyone before they get the treatment, and renders the entire sample one big treatment group, washing out any effect of the racial messages. Huber and Lapinski respond by arguing that they did find racial priming among the low-education subgroup. I have already responded to that assertion above and shown it to be incorrect. Huber and Lapinski further respond by arguing that such universal priming by measurement is irrelevant. Subjects exposed to the explicit message would, they argue, decrease their reliance on their racial predispositions despite having been primed by the questions on racial predispositions. However, it is possible that the initial priming by questions is not fully neutralized by subsequent exposure to the explicit message. It is also possible that the explicit message does not decrease the power of racial predispositions to levels below the control but rather to levels comparable to the control. The implicit-explicit model does not predict that explicit messages result in effects *lower* than the control. Explicit messages are expected to fail as racial primes, not necessarily to entirely obliterate the effects of racial attitudes. Thus, it is consistent with Huber and Lapinski's results to argue as I do that priming via question order *both* pre-empted the increased effect of racial predispositions in the implicit condition *and* masked the suppression effect of the explicit message.

In sum, Huber and Lapinski's results show no difference between conditions even in the lower-education group they target. This is likely due to one or more of the flaws I identified in their study: a nine percentage-point difference on the manipulation check, testimony to the ineffective experimental treatment; the introduction of measurement error by treating respondents who did not receive the assigned message as if they did and not showing whether this estimation choice affects the results; and priming racial predispositions for all respondents, washing out whatever distinctive effects might have been obtained in each condition. This study thus represents primarily a failure to replicate the literature. Consequently, the study does not afford a test of any of the propositions of the implicit-explicit model.

Claim III. Existing Research Supports the Claim That Racial Priming Happens Implicitly

In my critique I pointed out that the existing literature not only finds evidence of racial priming, but that the priming is implicit. I reviewed 17 studies that find racial effects in response to implicit cues. Huber and Lapinski respond by eliminating nearly all the relevant studies. But there is no good reason to dismiss the studies I reviewed in my critique. I agree with Huber and Lapinski that studies directly contrasting explicit and implicit versions of the same message are the optimal research strategy because they are diagnostic in distinguishing implicit from explicit messages. I also agree with Huber and Lapinski that these studies are rare. But this does not mean that the remaining studies are irrelevant. The remaining studies are relevant in providing evidence that racial cues work, that they work in a racial fashion, that they work implicitly, and that they do not work very well when they are explicit. Huber and Lapinski's study has been able to replicate none of these findings.

In my critique (this volume) I provided many examples of studies that document the existence of racial priming and the evidence that racial priming happens in response to implicit racial cues. At issue now is whether the existing literature also provides good evidence that explicit messages work less than implicit messages.

Huber and Lapinski incorrectly dismiss two relevant studies that contrast explicit and implicit cues. White finds that a more coded reference to "inner city" elicits a bigger racial effect on opinion about food stamps than the direct reference to African Americans.⁸ Huber and Lapinski omit this study from their list in table C and dismiss its importance in a footnote because of its small sample. However, a small sample does not by itself constitute grounds for elimination. This sample yields statistically significant effects. In addition, this sample is part of a set of independently drawn small or medium samples, including mine, which provide the consistent result that explicit cues and messages fail. Huber and Lapinski thus incorrectly dismiss this article.⁹

Huber and Lapinski also incorrectly dismiss the study by Terkildsen in which an implicit candidate cue elicits more racial priming than an identical explicit cue.¹⁰ Huber and Lapinski are correct to argue that skin tone does not neatly map onto the explicit-implicit distinction. However, the process by which skin tone affects racial priming fits neatly with the IE model: self-monitors—self-censors—repress the impact of their racial predispositions in evaluating the candidate, but only if he is dark-skinned. In other words, this study shows that skin tone functions along the explicit-implicit dimension: dark tone elicits prejudice but suppresses its use, while light tone elicits prejudice and allows its expression. This article does not suffer from any of the criticisms levied by Huber and Lapinski

against other studies. Thus, Huber and Lapinski incorrectly dismiss this article as well.

In sum, the literature does provide useful examples beyond my work or Huber and Lapinski's of the greater power of implicit versus explicit cues and messages.

Claim IV. Criticisms of My Research

Huber and Lapinski criticize my research and argue that flaws in that research cast doubt on the implicit-explicit model. Their criticism focuses on the validity of the inference that implicit messages work better than explicit messages. Most relevant in Huber and Lapinski's discussion are two of my studies, the 1998 Horton study and the welfare experiment.¹¹ I take up each in turn.

The 1988 Horton study

Huber and Lapinski criticize my investigation of racial priming in the 1988 presidential campaign on the grounds that it is correlational rather than experimental.¹² In this study I capitalized on the fact that for roughly three weeks in October, voters were exposed to a large volume of implicitly racial campaign messages, while in the subsequent period the racial messages were conveyed explicitly. Huber and Lapinski are concerned that campaign events correlated with either the implicit or explicit phases of the campaign are doing the real work behind the apparent racial priming of implicit messages and the apparent decrease of racial effects in the explicit period. First, Dukakis appeared ineffectual on crime during the implicit phase of the campaign, and this could produce a decrease in his popularity that has nothing to do with the implicit message. Second, the decrease of racial effects I attribute to the explicit phase could instead be due to other co-occurring campaign developments. These include the possibility that the Dukakis campaign regrouped and may have more effectively provided its own positive message about Dukakis's effectiveness on crime and his credentials as an old-fashioned Democrat, and criticisms of the Bush campaign's use of Horton as a dirty tactic. All these present an alternative to the racial message. However, none of these covarying events can explain the pattern of activation and depression of racial resentment.

First, Dukakis appearing ineffectual on crime and then appearing effectual cannot explain why people would change their reliance on racial resentment. Racial resentment would have to stand in for something else, something related to effectiveness on crime. That would be the concern over crime, yet my analysis showed that this concern did not rise and fall in impact the way racial resentment did. Racial resentment and worry about crime move in opposite directions across the periods, one responding to racial messages and the other to crime messages. Huber and Lapinski argue that the crime analysis suffers by relying on a measure of crime views collected after the cam-

paign, but the same holds for racial resentment, so the post-election measurement cannot explain the pattern. Huber and Lapinski also worry that fear of crime is correlated with racial resentment and consequently that its effects are masked by those of racial resentment. But as reported in the original study, this is not the case, and neither does this correlation change with time period.¹³ The implicit/explicit variable thus does not prime fear of crime but only racial predispositions, suggesting that the variable is measuring implicitness rather than Dukakis's perceived effectiveness on crime.

In addition, Dukakis's invigorated appeal to core Democrats in the explicit phase is accounted for in the analysis by controlling on the respondent's party identification. Being racially resentful thus has effects above and beyond the effects of being a Democrat. The control operates in each period so that the estimated effect of racial resentment is always the effect net of party identification. This explanation also cannot account for the rise in the power of racial resentment from September, when the implicit message was infrequent, to October when it rose in frequency, since Dukakis's rallying of the Democratic faithful did not increase with this change.

In sum, the only explanation that accounts both for the rise and the decline of racial predispositions is the switch from implicit to explicit phases of the campaign. The only event that corresponds to the two time phases and that would prime and then suppress racial predispositions—and only racial predispositions—is the escalation of implicit messages followed by the change to explicit discourse.

What remains to consider is the possibility that the criticism of the implicit message rather than its transformation to an explicit message caused the decrease in racial priming. True, this design cannot tell us whether it is the explicit version of the anti-black message or instead the attack on that message as racist that caused the decrease in racial priming. But for the purpose of the IE model, either outcome counts as support. The theory argues that either an explicit negative message or a challenge to that message that characterizes it as such will work similarly. So this ambiguity in the study design does not undermine the model.

The Welfare Experiment

Huber and Lapinski also criticize my welfare experiment.¹⁴ They argue that this study did not use appropriate controls and had a small and unrepresentative sample that casts doubt on the findings. Both claims are incorrect.

The experiment used two contrast conditions, not just the one to which Huber and Lapinski refer. To be clear, in this study the nonracial control group—and only that group—turned out by the vagaries of random assignment to differ from the remaining conditions on important demographic dimensions, and was therefore unusable.

However, setting aside this control group does not damage my argument. This experiment used another condition that functions as a baseline. This condition duplicated the implicit and explicit racial messages exactly but featured white instead of black targets. Huber and Lapinski themselves write that this qualifies as an appropriate control (“a more persuasive design would compare a Horton-like [racial] appeal with an identical one in which the individual portrayed in the advertisement was white” (in this volume, p. 133 note 20). A message featuring white targets thus qualifies as an appropriate contrast, and the study Huber and Lapinski criticize indeed finds that only the implicit and not the explicit message produces results relative to it.

Huber and Lapinski apply the “no control” criticism inconsistently by exempting their own no-control study from it. Only one of their two studies used a nonracial control condition, yet they include both studies as relevant while excluding others’ studies on these grounds.

Huber and Lapinski’s second criticism of this study—that it is small and unrepresentative—also does not travel far. The sample’s higher education actually makes the obtained results all the more powerful according to Huber and Lapinski’s own logic. In addition, it is simply incorrect to doubt the validity of a result simply because of its small sample. Small samples that consistently produce the same statistically significant result provide more confidence than one large study whose results stand out from the rest.

Most importantly, by this point in Huber and Lapinski’s reply the issue is no longer whether racial priming occurs relative to a control or among a particular slice of the education distribution, but rather whether implicit appeals work better than explicit ones. This is the only point of contention remaining. The only relevant criticism of the experiment, which contrasts identical implicit and explicit messages using random assignment, is that it has a high education level. Specifically, Huber and Lapinski argue that the racial predispositions of the better-educated already chronically affect their opinions since they have higher levels of attitudinal constraint in general, and this leaves less room for priming effects. But this argument predicts a small or no difference between implicit and explicit messages, with priming failing in each. But in fact, I found strong differences between the implicit and explicit messages. So the high educational level of my sample actually works in favor of the IE model’s prediction—I found strong effects despite reason to expect weak ones.

Conclusion

In sum, with only one exception, Huber and Lapinski in fact agree with all of my findings: 1) racial predispositions shape political opinions; 2) racial appeals cause racial effects, including the priming of these racial predispositions; 3)

these racial effects obtain from implicit cues and messages; 4) the racial effects from implicit messages obtain relative to control messages. All this is replicated by many studies and is not in dispute by Huber and Lapinski.

What is in dispute, then? All that remains of the disagreement is the claim that implicit messages work better than explicit messages. Huber and Lapinski argue that there is no good evidence for this claim. I disagree, and discuss evidence from several studies, including my own, showing why.

The implications of the remaining disagreement are important because they speak to the immutability of racism in politics. Huber and Lapinski believe that they failed to replicate evidence that explicit appeals are weak. They infer that Americans dislike these messages but nevertheless respond to them. Overtly racist campaign ads may work as intended, mobilizing racially hostile or fearful white voters, worsening race relations, enhancing racial inequality, and inhibiting democratic progress. At least, so Huber and Lapinski argue.¹⁵

But if so, then why do we not find such ads in American politics? Why the absence of negative characterizations of blacks in any mainstream forum for public communication? Why do politicians react with horror and panic when accused of using explicit messages? Why do they instead use coded references to race?¹⁶ Huber and Lapinski cannot answer these questions, while the IE model does.

Do racial appeals work? Do they work racially? Do they shrivel under the right conditions? The answer to all three questions is affirmative. More research on racial discourse will further refine our understanding of the continuing cleavage of race in politics. We need a better understanding of the conditions under which explicit appeals might indeed work. But for now, the weight of the evidence clearly points to the unique power of implicit communication in an age of egalitarian norms.

Notes

- 1 Gilliam and Iyengar 2000.
- 2 Huber and Lapinski 2006.
- 3 Huber and Lapinski 2006, 436.
- 4 In note 5 of their response Huber and Lapinski write that figure 5 in the original 2006 article is based on the newly introduced table A but do not explain the discrepancy.
- 5 I used the same estimation and control variables reported in 2006 and obtained the same analysis N reported in table A, but the coefficients and t-values differed from table A.
- 6 Table A seems to have mislabeled t-ratios as standard errors, but even if we ignore this error we can see that table A presents largely null results. The literature finds racial priming in the sense that the racial

cue not only raises the effect of racial attitudes above some baseline level but also causes that effect to deviate from 0. In the Huber and Lapinski study, by contrast, in only one of the four analyses—on affirmative action—does the effect of racial predispositions exceed its standard error by a comfortable margin *and* increase substantially from the control condition. In this sense Huber and Lapinski find null results and do not replicate the literature.

- 7 Gilliam and Iyengar 2000.
- 8 White 2007.
- 9 Huber and Lapinski also criticize White's findings because the cue "poor" does not result in the same racial priming as the cue "inner city" and argue that this is failure to find consistent implicit priming and thus grounds to dismiss the study. However, Hurwitz and Peffley 2005 also found, consistent with White, that "inner city" racializes white Americans' views. More importantly, the cue of "poor" may not elicit negative racial views unless the poor are characterized as black and undeserving; Gilens 1999.
- 10 Terkildsen 1993.
- 11 Mendelberg 2001, ch. 6 and 7.
- 12 Mendelberg 2001, ch. 6. The first study Huber and Lapinski take up is Mendelberg 1997. They note that the 1997 study did not include a condition with explicit messages. This is true; that study was designed to test the simple proposition that an implicit message produces more racial consequences than a nonracial control, and it made no claim beyond that.
- 13 Mendelberg 2001, 180.
- 14 Mendelberg 2001, ch. 7.
- 15 Huber and Lapinski mention the example of the 2006 Tennessee Senate contest in which Harold Ford, an African American candidate, was attacked by an ad featuring suggested sexual allegations by a white woman. Huber and Lapinski point to the ostensible (undocumented) failure of the NAACP's condemnation of this ad as evidence that rendering a message explicit does not neutralize it. But there is no evidence for or against the success of the NAACP's challenge. The mere fact that Ford lost does not diagnose; Ford may have lost despite the NAACP's success, or

might have won if the rebuttal had come earlier and played longer and wider. But while the Ford ad is a stylized example and has no force for being such, it does illustrate what is at stake in this debate: whether and how political actors inject race into politics, and how this injection can be neutralized.

- 16 Not only politicians but the media as well tend to avoid verbal mentions of race, according to Gilens, who had trouble finding verbal characterizations of the race of poor people but found plenty of visual images conveying this information (1999, ch. 5, 111–114).

References

- Gilens, Martin. 1999. *Why Americans Hate Welfare*. Chicago: University of Chicago Press.
- Gilliam, Frank D., and Shanto Iyengar. 2000. Prime suspects: The influence of local television news on the viewing public. *American Journal of Political Science* 44 (3): 560–73.
- Greenwald, Anthony G., and Linda Krieger. 2006. Implicit bias: Scientific foundations. Symposium on Behavioral Realism. *California Law Review* 94 (4): 945–67.
- Huber, Gregory A., and John Lapinski. 2006. The "race card" revisited: Assessing racial priming in policy contests. *American Journal of Political Science* 48 (2): 375–401.
- Hurwitz, Jon, and Mark Peffley. 2005. Playing the race card in the post-Willie Horton era. *Public Opinion Quarterly* 69 (1): 99–112.
- Mendelberg, Tali. 1997. Executing Hortons: Racial crime in the 1988 presidential campaign. *Public Opinion Quarterly* 61(1): 134–57.
- . 2001. *The Race Card: Campaign Strategy, Implicit Messages, and the Norm of Equality*. Princeton, NJ: Princeton University Press.
- Terkildsen, Nayda. 1993. When white voters evaluate black candidates: The processing implications of candidate skin color, prejudice and self-monitoring. *American Journal of Political Science* 37: 1032–53.
- White, Ismail. 2007. When race matters and when it doesn't: Racial group differences in response to racial cues. *American Political Science Review* 101 (2): 1–16.