

# Racial Gaps in the Responses to Hurricane Katrina: An Experimental Study\*

Melissa Harris-Lacewell<sup>†</sup>    Kosuke Imai<sup>‡</sup>    Teppei Yamamoto<sup>§</sup>

May 28, 2007  
(approximately 7,665 words)

## Abstract

In the media reports and public discourses, Hurricane Katrina was portrayed as a racialized disaster that disproportionately affected blacks. Indeed, one in three victims of Katrina was African American, and in the most damaged areas of New Orleans, this proportion was as high as three in four. In this paper, using an original survey, we show that in the aftermath of Katrina blacks and whites held strikingly polarized views about everything from why victims did not evacuate to the appropriateness of the government response. We also conduct a randomized experiment to explore the effect of media framing on this racial gap. The findings suggest that even subtle image manipulations can make blacks more sympathetic towards their own racial group. Similar ingroup bias is found for gender as well as race, with women showing more sympathy when they see an image portraying a female victim with children. Our statistical analysis demonstrates how to avoid usual parametric assumptions about nonresponse when analyzing survey and randomized experiments. Finally, we discuss the implications of our findings for the political psychology of race.

---

\*We thank Cathy Cohen, Michael Dawson, Marty Gilens, Erzo Luttmer, Tali Mendelberg, and Tom Romer for their advice and suggestions. We gratefully acknowledge the funding from the Center for the Study of Race, Politics, and Culture and the Office of the Provost at the University of Chicago as well as the Mamdouha S. Bobst Center for Peace and Justice at Princeton University. Imai also thanks the National Science Foundation (SES-0550873) and Princeton University Committee on Research in the Humanities and Social Sciences for financial support. This study is a part of the Race and Hurricane Katrina research project whose co-principal investigators are Michael Dawson and Cathy Cohen along with Harris-Lacewell.

<sup>†</sup>Associate Professor, Department of Politics and Center for African American Studies, Princeton University, Princeton NJ 08544. Email: lacewell@Princeton.Edu

<sup>‡</sup>Assistant Professor, Department of Politics, Princeton University, Princeton NJ 08544. Phone: 609-258-6610, Fax: 973-556-1929, Email: kimai@princeton.edu, URL: <http://imai.princeton.edu>

<sup>§</sup>Ph.D. student, Department of Politics, Princeton University, Princeton NJ 08544. Email: tyamamot@princeton.edu

# 1 Katrina as a Racialized Disaster

In the final week of August 2005, Hurricane Katrina caused devastating damage to the Gulf Coast region of the United States, costing at least 1,836 lives and 81 billion dollars in total. The American public was highly critical of President Bush’s handling of the disaster, and his approval rate slipped by 4 percentage points while the disapproval rating increased by the same amount (The Pew Research Center for the People and the Press, 2005). Other government organizations, including local administrations and the Federal Emergency Management Agency, also became the targets of public criticism (Lewis, n.d.). In the aftermath of the hurricane, media reports and public discourses portrayed Katrina as a racialized disaster that disproportionately affected African Americans (Harris-Lacewell, 2007). Television news programs repeatedly showed images of predominantly black evacuees, and black community leaders denounced slow government relief actions. Indeed, approximately one in three victims was African American, who represents only about one eighth of the national population (Sherman and Shapiro, 2005). This racial disparity was even more striking in the damaged areas of the City of New Orleans, where as many as 75 percent of the residents were blacks (Logan, n.d.).

Despite the importance of the event for racial politics, Hurricane Katrina has received little attention from political scientists.<sup>1</sup> In this paper, we report the results of an original survey and a randomized experiment which we conducted in November 2005 in order to explore the racial gaps in responses to Katrina. Using a nationally representative sample of black and white respondents, we find that blacks and whites held strikingly polarized views about many important issues that were reported in the media at that time. In particular, by a surprisingly wide margin,

---

<sup>1</sup>Exceptions include a *Washington Post* article by Iyengar and Morin (2006), which we describe later in the paper, as well as a 2006 special issue of *Du Bois Review*, “Katrina: Unmasking Race, Poverty, and Politics in the 21st Century,” to which several political scientists contribute.

blacks expressed more sympathetic attitudes toward the victims, less favorable opinions on the government response, and greater consciousness about racial inequality. Also, we find that within each racial group, women show more overall sympathy with victims compared to men, and citizens with lower income tend to be more supportive of unlimited government assistance.

Where did these racial gaps come from? Political psychology literature suggests that the media coverage might have played an important role by framing the disaster in a highly racialized context (e.g., Nelson *et al.*, 1997a,b; Kellstedt, 2000). Indeed, the media played a key role in the racial politics of Hurricane Katrina and even became a target of public criticism, on the ground that they used racially biased language when describing black victims (Dawson, 2006; Harris-Lacewell, 2007). To explore this possibility, we conducted a randomized experiment within our survey where respondents viewed the image of either black or white victims, taken directly from a news source, before answering several survey questions about Katrina.

Our analysis shows that this subtle image manipulation had different effects on blacks and whites, making black respondents more sympathetic towards their own racial group. For example, black respondents who saw black victims became significantly more favorable to unconditional government assistance than those who saw white victims. In addition, we examine whether these effects also differ depending on other demographic factors, such as gender, education and income. By splitting each racial group into subgroups based on these variables, we find that women become more sympathetic to victims when they see an image featuring a mother holding children. In sum, our findings are consistent with the theory of group identity since blacks and females show more sympathy towards the victims of their own group (e.g., Tajfel and Turner, 1986; Gurin *et al.*, 1980; Kinder and Winter, 2001).

Finally, in analyzing our survey data and randomized experiment, we demonstrate how to deal with survey nonresponse by avoiding usual parametric assumptions about nonresponse mechanism.

In many political science survey randomized experiments, nonresponse is an important complication that needs to be properly addressed (Horiuchi, Imai, and Taniguchi, 2007). We apply the method of bounds to identify what we can learn from the observed data alone without imposing untestable assumptions. Such an approach is deliberately conservative, but it allows for researchers to establish credible conclusions when strong prior information is not available (Manski, 1995). The method we use in this paper is made available as a part of freely available statistical software so that other researchers can apply it to their own research (Imai, 2007).

The rest of the paper is organized as follows. In Section 2, we describe the design of the survey, and report the striking opinion gaps between blacks and whites. We also briefly discuss the related findings of other surveys. In Section 3, we explain the experimental component of our survey and present the findings while comparing our experiment with other Katrina-related experiments. Finally, in Section 4, we conclude by discussing the implications of our findings for the political psychology of race.

## **2 The Katrina Survey**

### **2.1 Survey Data and Questions**

Our survey was administered between October 28 and November 17, 2005, approximately two months after Hurricane Katrina landed on the Gulf Coast. The survey was part of Racial Attitudes and the Katrina Disaster Study conducted by the Center for the Study of Race, Politics and Culture at the University of Chicago. The data were collected by Knowledge Networks, which uses a Random Digit Dialing sampling method to select a nationally representative panel of households and produces appropriate post-stratification weights based on various demographic characteristics. We deliberately oversampled African American respondents so that the sample size of each racial group is sufficiently large and the resulting inference is reliable. The initial sample size is 1,252,

and after excluding respondents who are neither white nor black, we have our final sample of 1,190 respondents, of which 703 are whites and 487 blacks.

All respondents are asked a battery of questions regarding their political attitude and issue opinions, including questions about racial problems and Katrina itself. In this paper, we focus on four of the questions that are designed to measure respondents' feelings and opinions about the disaster victims, government relief actions, and racial inequalities in the American society. The first question measures the preference of respondents regarding government spending on the Katrina victims by asking them to choose one from the following pair of statements:

- The federal government should spend whatever is necessary to rebuild the city and to restore refugees to their homes.
- Although this is a great tragedy the federal government must not commit too many funds to rebuilding until we know how we will pay for it.

Next, the second question tries to measure positive versus negative causal attributions of blame for the victims' situation by asking them to choose from the following two statements:

- Most residents of New Orleans who were trapped by the flood wanted to leave the city but simply did not have the resources to evacuate.
- Most residents of New Orleans who were trapped by the flood were too stubborn to evacuate.

Finally, we ask the following two questions, in order to measure whether respondents understand the disaster in a broader context of racial inequality, and the extent to which respondents evaluate government relief actions in racial terms. Our general question about racial inequality reads:

In your view did this disaster show that racial inequality remains a major problem in this country, or don't you think it this was a particularly important lesson of the disaster?

- Showed that racial inequality remains a major problem.
- Not a particularly important lesson of the disaster.

We ask the respondents to evaluate the appropriateness of the government reactions by framing the question in terms of the belief that was shared by some people at that time:

Most of the people stranded in New Orleans following the hurricane were African American. Do you think the government's response to the situation would have been faster if most of the victims had been white, or don't you think it would have made any difference?

- Yes, would have been faster
- No, wouldn't have made any difference

To ensure the comparability of our findings, the exact wording of the last two questions comes from a survey conducted by the Pew Research Center (The Pew Research Center for the People and the Press, 2005) on September 6 and 7, 2005. The advantage of our survey is that the number of black respondents (487) is more than twice as large as that in the Pew survey (211), permitting us to estimate the black opinions more efficiently. Moreover, as the Pew survey was conducted only a week after the landfall of Katrina, the respondents might have answered questions on the basis of imprecise information in the initial news reports. On the other hand, our survey was carried out after the period of two months, so that citizens were likely to share roughly the same factual information about the hurricane.

## **2.2 Striking Attitudinal Gaps between Blacks and Whites**

### **2.2.1 Methodology**

Before presenting the results, we briefly describe our statistical method which we use to address the problem of survey nonresponse. In particular, we apply the method of bounds to estimate the range of the values our quantities of interest can take, without making any assumption about the nonresponse mechanism (Horowitz and Manski, 1998). In our survey, the quantity of interest is the population proportion of blacks (or whites) who would affirmatively answer each of the four questions. If every respondent in our sample were to answer the questions, we could estimate this quantity by computing the weighted sample proportion of respondents saying "Yes" for each racial group using the survey weights. However, this is impossible in our case because some respondents

give no answer, and such nonresponses are unlikely to have occurred completely at random. Note that throughout the remainder of the paper, for the first two questions described in the previous section, we use “Yes” (“No”) to represent respondents’ answer that is (not) sympathetic towards the victims of Katrina.

The method of bounds deals with this nonresponse problem by considering two extreme situations about the missing data, and taking the corresponding values of the quantity of interest under two scenarios as its upper and lower bounds (see Duncan and Davis, 1953, for an application of this method to ecological inference). The formal presentation of the method is given in Appendix A.1, and here we illustrate the intuition behind the method with a simplified numerical example. Out of 487 black respondents in our sample, 370 answered “Yes” to the first question about the government funds whereas 21 did not provide any answer. The upper bound of the population proportion of blacks who would answer “Yes” results when we suppose that all the nonresponses would have been “Yes”. Hence, assuming, for the sake of simplicity, that all respondents have the same survey weights, we estimate the upper bound by  $(370 + 21)/487 \approx 0.80$ . Likewise, the lower bound is estimated by assuming all the nonresponses to be “No,” i.e.,  $370/487 \approx 0.76$ . Our estimated interval of the population proportion of blacks answering “Yes” is  $[0.76, 0.80]$ . Thus, the method of bounds identifies the extent to which one can learn about the quantity of interest from the data alone without making any assumption about nonresponse mechanism.

In the calculation of our final estimates, we adjust the estimated bounds constructed in this way by using the post-stratification survey weights so that we properly estimate the population proportion rather than the sample proportion. The weights are computed by Knowledge Networks on the basis of demographic and geographic information and provided for each observation. Also, since the resulting estimated bounds have sampling variability, we quantify the uncertainty by constructing confidence intervals (see Appendix A.2 for details). Since the results of our survey

show overwhelming evidence of racial gaps, our statistical analysis does not fundamentally alter the basic conclusion we are about to present. Nevertheless, the formal properties discussed in this section and in the appendices will play a more critical role in Section 3, where we generalize this methodology to analyze our experiment.

### 2.2.2 Findings

Figure 1 presents the estimated bounds and their 95 percent confidence intervals for the population proportion of having an affirmative opinion about each of the questions, reported separately for black and white respondents. The figure reveals striking racial gaps in the responses to Katrina: in every question, blacks are more likely to say “Yes” than whites by a surprisingly wide margin. On the basis of our estimation, when asked if the government should spend whatever necessary for victims, 80.5 to 83.3 percent of blacks say “Yes” with the 95% confidence interval of [78.5, 85.3], while only 32.9 to 34.8 percent of whites do so with the confidence interval of [30.9, 36.8]. Moreover, as many as 87.2 to 88.9 percent of blacks agree with the statement that most of the residents trapped by the flood actually wanted to leave but could not; only 53.0 to 56.4 percent of whites agree with the same statement (the 95% confidence intervals are [85.2, 90.8] and [50.9, 58.5], respectively). The gap between blacks and whites is equally large in the third question (88.1 to 89.9 percent versus 37.1 to 40.9 percent, with the 95% confidence intervals of [86.3, 91.6] and [35.0, 43.1], respectively), where respondents are asked if the disaster showed that racial inequality remains a major problem. Finally, the gap was most striking on the appropriateness of the government response. To the question asking whether the government response would have been faster if the victims had been white, the estimated proportion of answering “Yes” is 78.4 to 84.0 among blacks with the 95% confidence interval of [75.7, 86.3]. In contrast, the proportion is only 19.2 to 21.1 percent among whites, with the confidence interval ranging from 17.5 to 22.9 percent.

We also divided the black and white respondents into two subgroups on the basis of gender,




Figure 1: Racial Gaps in the Katrina Attitudes Survey. The figure shows the estimated bounds (thick lines with horizontal bars) and their 95% confidence intervals (thin lines) for the probability of answering "Yes" to each of the questions. The estimates are given separately for black and white respondents. The solid lines in the center are the estimates for all the respondents (labeled A), and the colored lines on the sides are for women and men (labeled F and M, respectively), in each racial group.

income and education to examine if responses differ by these demographic characteristics within each racial group. The results reveal interesting differences across these demographic subgroups. Among white respondents, females show greater sympathy with victims than males in the first three questions, while we find virtually no gender-based difference in the last “government response” question. On the other hand, there is no systematic difference between black women and men, with the former subgroup being slightly more sympathetic in the first question but less in the third question. For example, in the first question, 34.7 to 36.1 percent of white females support unconditional government spending with the 95% confidence interval of [31.8, 39.2], while 30.9 to 33.3 percent of white males do so with the confidence interval of [28.1, 36.1]. This finding is consistent with the previous evidence that women are more likely to support government social spending compared to men (e.g. Gilens, 1988). The gender gap is larger in the second and third questions, with 56.2 to 60.1 and 40.0 to 44.4 percent of white women answering “Yes” (the confidence intervals ranging [53.2, 63.1] and [37.0, 47.4]), while only 49.6 to 52.4 and 34.1 to 37.1 percent of white men answer affirmatively (the confidence intervals are [46.4, 55.5] and [31.2, 40.1]) to the same pair of questions.

Next, since people with lower economic status are generally thought to be more supportive of government spending, we also expect in this survey that those who earn less will be more likely to support unconditional government assistance in the first question. As expected, people with high income, defined as those who earn \$40,000 or above per year, are less sympathetic toward victims in that question, especially if they are white. While only 28.6 to 30.7 percent of high-income whites support unconditional spending with the confidence interval of [26.0, 33.2], no less than 39.9 to 41.9 percent of low-income whites do so with the confidence interval of [36.3, 45.5]. The evidence for blacks is weaker but still in the expected direction: 76.6 to 81.1 percent of high-income blacks support government spending while 82.1 to 83.3 percent of low-income blacks do so,

with the confidence intervals ranging [73.2, 84.4] and [79.4, 86.0], respectively. Also, whites with lower income levels are more likely to say “Yes” in the other three questions, but the gaps are generally smaller than in the first question. Meanwhile, the difference based on income levels is less clear among blacks, with slightly *more* high-income blacks answering affirmatively to the third and fourth questions. Finally, we find little systematic variation in responses based on educational levels, although there is some indication that citizens with higher education are more likely to say “Yes” in the last three questions regardless of race.

In short, blacks have more sympathetic attitudes toward the victims, less favorable opinions on the government response, and greater consciousness about racial inequality. Also, female and poor whites are more sympathetic toward victims than male and rich whites, respectively, but this pattern does not hold among blacks. Although it is a well-known fact that blacks and whites have different views on a wide array of policy issues (Schuman *et al.*, 1997; Kinder and Winter, 2001), the sizes of these racial gaps we found here appear to be large. Finally, our findings are all the more surprising given that racial responses following the events of September 11, 2001 showed that gaps between whites and blacks narrowed rather than widened in response to national disaster (Harris-Lacewell, 2007).

### **2.3 Comparison with Other Surveys**

Our findings generally agree with the results of other surveys. For example, Huddy and Feldman (2006) interview 507 whites and 267 blacks from late September until late November of 2005 and find similar racial divides in the respondents’ attitudes towards Katrina victims and their support for the government efforts. The authors find that 70 percent of black respondents felt sympathetic towards those left behind in New Orleans, whereas only 42 percent of whites felt the same way. Moreover, 65 percent of blacks thought that the victims deserved a great deal of assistance from government and charities, compared to 50 percent of whites. Also, the Pew survey finds racial



Figure 2: Experimental Image Manipulations. The two images, one showing black mother with children (left) and the other showing a white family (right), are used as image manipulations. The former was taken from the cover of the September 12 issue of *Newsweek*, and the latter appeared inside the same issue.

gaps for the two questions we share, though the racial gaps are less pronounced than our findings (The Pew Research Center for the People and the Press, 2005). For example, in the “government response” question, the Pew survey finds a 49 percent gap between blacks and whites while we find a gap of 57 to 65 percent. This is particularly interesting given that our survey was conducted approximately two months after the Pew’s: one might expect our survey respondents to have more similar understanding of what had actually happened. Some of these differences might have also resulted for methodological reasons. While the other studies assume nonresponses to be missing at random, our approach makes no assumption about the missing data mechanism. Our results are also based on a larger sample of black respondents than the other surveys, yielding more precise estimates about the opinions of African Americans. It is also possible, however, that we found the larger racial gaps in part because of racial differences in the responses to the image stimuli utilized for our randomized experiment, which we will discuss in the next section.

## 3 An Experiment

### 3.1 Design

Having shown striking opinion gaps between black and white respondents, we now turn to the question of where such wide racial gaps come from. We conducted a randomized experiment within our survey to explore the possibility that media framing affects the respondents' views about political and racial issues related to Hurricane Katrina. Before answering those four survey questions described in the previous section, each respondent was randomly assigned to one of the two experimental conditions: respondents in the “black mother” condition saw an image of a black mother holding two children (the left panel of Figure 2), while those assigned to the “white family” condition were shown an image of a white family in a rescue boat (the right panel of Figure 2).<sup>2</sup> Since we were primarily interested in the impact of the media's actual framing of the Katrina victims, we used images that appeared in *Newsweek* issued on September 12, 2005. The “black mother” image was used as the cover, whereas the “white family” image appeared inside the issue. After the randomization of the treatment, 251 black respondents were assigned to the “black mother” condition while 236 of them were shown the “white family” image. Similarly, among white respondents, 345 viewed the “black mother” image whereas 358 of them were shown the other image.

The design of our experiment is in part motivated by the political psychology literature, which has shown that racial framing in media coverage can have a significant impact on citizens' opinions about racial issues (e.g., Nelson *et al.*, 1997a,b; Gilens, 1999; Kellstedt, 2000; Gilliam and Iyengar, 2000). We acknowledge that some proportion of the racial gaps might be attributable to a more

---

<sup>2</sup>Our experiment also contained a random text manipulation, where half of the respondents read a caption that referred to Katrina victims as “Americans” while other respondents read the same caption with the word “refugees.” The implications of this experimental manipulation are discussed elsewhere (Harris-Lacewell, 2007).

fundamental “racial divide” that is known to exist on many policy issues (Kinder and Winter, 2001). However, we still suspect that the media played an important role in the case of Katrina, since given its nature as a large-scale national disaster we would expect less pronounced racial differences. In fact, Gross and Kohenak (n.d.) offer evidence that the coverage of Katrina was markedly different between black press and mainstream newspapers, in that the former more often referred to race and poverty when describing the victims, while it rarely attributed the situation of those left in New Orleans to their individual responsibility. Also, Hopkins (2007) shows that media coverage actually made the residents of local communities that hosted Katrina evacuees less supportive of the poor and African Americans. Another motivation is the theory of group identity in the social psychology literature which implies that people tend to have favorable attitudes toward their own social group while they subconsciously discriminate against others from different socio-economic backgrounds (e.g., Tajfel and Turner, 1986; Gurin *et al.*, 1980; Kinder and Winter, 2001). Although our experimental treatment consists of subtle image manipulations, the literature suggests that implicit visual cues are often sufficient to elicit respondents’ latent racial attitudes (e.g., Mendelberg, 2001; Valentino *et al.*, 2002). Therefore, we expect respondents to be more sympathetic toward victims and more critical of government response when they see the picture of victims of their own race.

In addition, we explore the possibility that the causal effects of our experimental manipulation might also depend on gender, income, and education. Since the “black mother” image portrays a woman who looks needier than the family featured in the “white family” image, the manipulation can also convey implicit cues about gender and economic status as well as race. Therefore, as the theory of group identity predicts, we expect that women and citizens in lower income strata exhibit more sympathetic attitude toward victims when they see the “black mother” image. Our expectation about gender is also reinforced by Gilligan (1982)’s classical thesis that women tend to

acquire the “ethics of care” toward the vulnerable through the process of cognitive development. Finally, in a recent study on the effects of racial cues, Huber and Lapinski (2006) finds that low-educated whites are more susceptible to visual cues than high-educated whites. Accordingly, we hypothesize that the treatment effect is stronger among whites with low levels of education.

## **3.2 Causal Effects of the Image Manipulation**

### **3.2.1 Methodology**

The goal of our experiment is to explore the possibility that media framing might have affected people’s reactions towards Hurricane Katrina and those effects might have been different across racial groups as well as demographic subgroups. The methodological challenge, however, is to estimate the causal effects of interest when the outcome variable is not observed for a subset of observations due to survey nonresponse. To address this problem, we extend the method of bounds described in Section 2 to the analysis of randomized experiments. Our quantity of interest is the average treatment effect (ATE) between the “black mother” and “white family” treatments for each racial group. Thus, the ATE represents the population proportion of blacks (or whites) who would change their answers from “No” to “Yes” if they saw the black mother image instead of the white family image. However, for the same reason as described in Section 2, we cannot identify the ATE without making untestable assumptions about the nonresponse mechanism.

Hence, we take a conservative approach to estimate its bounds by considering two extreme situations about the missing data, and thereby avoid such assumptions. The details of our method are given in the appendix, and here as done in Section 2.2, we explain the intuition behind the method with the following simple numerical example. Out of 251 black respondents who were in the black mother condition, 202 answered “Yes,” 41 answered “No” and 8 refused to answer to our first question, whereas 168 said “Yes,” 55 said “No” and 13 refused among the 236 blacks assigned to the white family condition. The upper bound corresponds to the situation where all the nonresponses

in the black mother condition were “Yes” and all the nonresponses in the white family condition were “No.” Thus, we can estimate the upper bound of the ATE by  $(202+8)/251 - 168/236 \approx 0.125$ . Conversely, the estimated lower bound results when all the nonresponses were “No” in the black mother condition and “Yes” in the white family condition, i.e.,  $202/251 - (168 + 13)/236 \approx 0.038$ . Therefore, our estimated bound for the ATE is given by  $[0.038, 0.125]$ . Note that for the sake of simplicity, this calculation ignores survey weights as well as the existence of another treatment. In the actual calculation, we properly adjust our estimates of the bounds by incorporating survey weights and the possible effect of the other treatment variable (see Appendix A.1 for details). Finally, we derive confidence intervals by using the same method as in Section 2 (see Appendix A.2).

### 3.2.2 Findings

Figure 3 presents the estimated bounds of the average treatment effect and their 95 percent confidence intervals, calculated separately for blacks and whites. Section 3.1 outlines our expectation that respondents become more sympathetic when they see the image of victims of their own racial group. This hypothesis is consistent with the results for our first question, which asks respondents whether they support unconditional government assistance to Katrina victims. Showing the black mother image instead of the white family image to black respondents increases the probability of supporting such governmental assistance by 8.0 to 13.8 percentage points with the 95 percent confidence interval of  $[4.1, 18.2]$ , despite the fact that these blacks already have high levels of support on average regardless of their treatment status. The evidence for whites is weaker, but seeing the black mother image makes them less likely to affirmatively answer the question by 1.7 to 5.6 percentage points with the 95 percent confidence interval of  $[-9.7, 2.4]$ . It is noteworthy that this finding also supports the argument made by Gilens (1999) that white Americans tend to oppose welfare spending because they believe the recipients are predominantly black and “undeserved.”

In contrast, the evidence is mixed for our second question about victims’ intentions to leave the

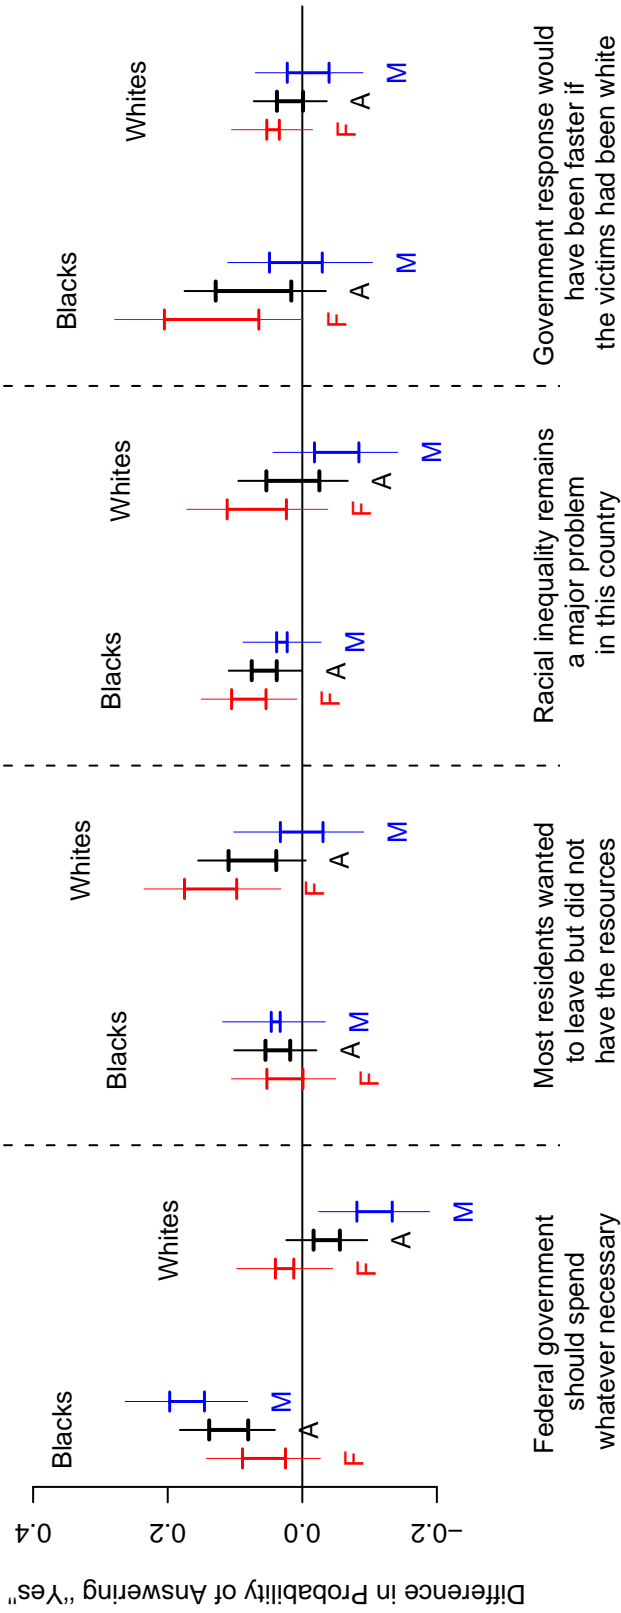


Figure 3: The Average Treatment Effect of the Black Mother Image versus the White Family Image. The figure shows the estimated bounds (thick lines with horizontal bars) and their 95% confidence intervals (thin lines) for the probability of answering “Yes” to each of the questions. The estimates are given separately for black and white respondents. The solid lines in the center are the estimates for all the respondents (labeled A), and the colored lines on the sides are for women and men (labeled F and M, respectively), in each racial group.

city. As expected, the probability of agreeing with the statement that most victims wanted to leave but could not increases by 1.8 to 5.5 percentage points for blacks when they see the black mother image (the 95 percent confidence interval ranging from  $-2.1$  to  $10.1$  percentage points). However, contrary to our expectation, we find that even for whites, seeing the black mother image somewhat increases the probability by 3.9 to 11.0 percentage points with the 95 percent confidence interval of  $[-0.05, 15.5]$ . We conjecture that this finding has little to do with race and arises from the fact that the black mother and children do look more badly in need than the white family sitting in a boat.

The last two questions do not provide as definitive answers as the first two questions, but the results are still informative. In the “racial inequality” question, the average effect of showing the black mother image on the probability of recognizing racial inequality as a major problem is estimated to range from 3.8 to 7.5 percentage points among blacks and from  $-2.5$  to 5.3 percentage points among whites, with the confidence intervals of 0.03 to 10.9 and  $-6.8$  to 9.5, respectively. In the “government response” question, the image manipulation has little effect on whites, while showing the black mother image to black respondents increases the probability of accepting the racial bias in government response by 1.6 to 12.9 percentage points though the associated confidence interval is wide,  $[-3.5, 17.5]$ . Thus, overall, our findings are somewhat mixed but generally consistent with our expectation about racial differences.

Finally, we examine the degree to which the causal effects of the image manipulation depend on gender, education and income within each racial group. In short, our findings are consistent with our expectations described earlier for gender and income, but not for education. First, in every question, white women become more sympathetic toward victims when shown the “black mother” image compared to white men. For instance, white men become less likely to support unconditional government assistance by 8.1 to 13.4 percentage points when they see the “black mother” image

with the 95% confidence interval of  $[-18.9, -2.4]$ , but this effect is reversed among white women, who become somewhat *more* likely to support assistance by 1.3 to 4.0 percentage points when shown the “black mother” image though with the wide confidence interval of  $[-4.5, 9.7]$ . The same pattern is found among whites in the other questions, with almost equally large gender gaps in the second and third and a smaller difference in the fourth. In fact, this finding reveals one reason why we consistently found weaker treatment effects on whites than blacks: the effects of racial framing were partially offset by those of gender framing and thus appeared to be lessened at the aggregate level. Meanwhile, evidence for blacks is more mixed, and our expectations are validated only in the last two of the four questions.

The evidence about income is also more or less consistent with our hypothesis that people in lower income strata become more sympathetic when shown the “black mother” image. For each racial group, this hypothesis is supported in three out of the four questions. For instance, the low-income blacks become more likely to attribute blame to victims’ financial necessity when shown the “black mother” image by 2.8 to 6.9 percentage points (the confidence interval ranges  $[-2.3, 12.4]$ ), but showing this image to high-income blacks decreased the probability by 3.5 to 4.6 percentage points, with the confidence interval of  $[-15.5, 4.4]$ . The same pattern is observed among whites for this question, with the “black mother” image decreasing the probability of making such causal attribution by 4.4 to 11.1 and 14.6 to 18.6 percentage points among high-income and low-income whites, respectively (the confidence intervals are  $[-1.3, 17.0]$  and  $[7.7, 26.6]$ ).

Finally, we find no systematic differences across education groups for white respondents, contrary to our expectation from Huber and Lapinski (2006). Interestingly, however, blacks with higher levels of education are found to be *more* responsive to the image treatment than those with lower education in all but the fourth question, although the differences are generally small.

### 3.3 Comparison with Other Experimental Studies

While there exists a large body of experimental literature on racial politics, we were able to find only two experimental studies conducted in the context of Hurricane Katrina. Both studies share our interest in the racial effects of media framing on the citizens' attitudes toward the Katrina victims, and they both use image manipulations to measure this effect. Iyengar and Morin (2006) examine a sample of approximately 2,300 people to see how racial cues conveyed in news coverage affected viewers' response to Katrina. Unlike our experiment, their sample was predominantly white (86 percent), and over-represented Democrats (only 12 percent identified as Republican). The respondents also tended to be liberal (86 percent were critical of Bush) and have higher levels of education (84 percent completed at least a bachelor's degree). In their experiment, participants were shown a Katrina-related news story that features a victim varying in ethnicity and skin color, and then asked how much money hurricane victims should be awarded as government assistance. Participants were also asked how long victims should receive such assistance. The results turned out to be somewhat mixed: while the duration of assistance varied according to the manipulated victim's ethnicity, with those who saw an African-American awarding a shorter period of assistance, the amount of monthly support did not differ by the race of the hypothetical victim. Although these findings are based on a non-representative sample, they are consistent with the results of our experiment: as described in Section 3.2, we find that whites are in general less sensitive to image manipulations than blacks.

While Iyengar and Morin (2006) focus on the effects of victims' ethnicity per se, the experimental design of Fong and Luttmer (2007) is similar to ours in that they examine the interactions between victims' and respondents' racial backgrounds. Using 1,101 respondents drawn from the Knowledge Networks panel, of which 247 were African American, the authors manipulated photos so that respondents would see victims of different racial compositions (In addition, they varied

the audio information that accompanied the photos). While our experiment uses the images that actually appeared in the media, Fong and Luttmer (2007) took a different approach and carefully manipulated the images in order to extract the pure effect of race. After viewing a randomly selected audiovisual presentation, respondents were asked to allocate 100 dollars between themselves and Katrina victims. Surprisingly, despite their sophisticated image manipulation, the authors found little evidence of the effect of race: on average, the respondents' willingness to give more to the victims was not affected either by their race or the race of the victims featured in the image. However, when divided into subgroups based on the strength of ethnic identity, the black and white respondents did respond to the image manipulation differently, with those with stronger ingroup identification giving more when seeing the victims of their own racial group. Thus, while they fail to find the effect of race on average, their substantive findings about ingroup bias are consistent with ours.

In short, the existing experimental studies on Hurricane Katrina do not give definitive answers to the question of how media framing affected citizens' attitudes towards victims. The contribution of our experiment is to answer this question with a similar experimental setup, but with a larger sample of blacks. As we showed in Section 3.2, the results suggest that racial contents in the media did play an important role in the case of Katrina. Thus, as largely consistent with the previous findings, the effects of media framing appear to be different depending on the race of both victims and respondents. However, the exact conditions under which such media effects occur remain to be investigated.

## 4 Concluding Remarks

To summarize, using an original survey, we found striking racial gaps in the responses to Hurricane Katrina. Our analysis of a nationally representative sample of blacks and whites revealed that blacks held more sympathetic attitudes toward the victims, less favorable opinions on the govern-

ment relief actions, and greater consciousness about racial inequality. Moreover, our randomized experiment showed that the subtle image manipulations about race could have different effects on black and white respondents even in a highly racialized context. For example, black respondents who saw the black mother image became more sympathetic toward victims than those who saw the white family image. These findings are based on the method of bounds, which makes no parametric assumptions and thus allows us to draw more credible statistical inference.

Our findings have implications for political psychology and the study of racial politics. In particular, they are consistent with the theory of social identity, which maintains that people subconsciously favor their own social group while forming less favorable attitude toward the groups to which they do not belong (Tajfel and Turner, 1986). Furthermore, our result that blacks are more sensitive to the image manipulation confirms the earlier finding that blacks tend to have stronger in-group identification than other racial groups (Gurin *et al.*, 1980; Kinder and Winter, 2001). Finally, our experiment illuminates the importance of media framing by showing that the images, which actually appeared in the media, can cause changes in respondents' opinions about racial politics and that these effects are different between blacks and whites. This suggests that the striking racial gaps we found earlier might have been caused in part by the actual media coverage, which did contain many implicit and explicit cues about the victims' racial background.

## A Technical Appendix

In this appendix, we formally describe the method we use to analyze the survey and the experiment.

### A.1 The Method of Bounds

Let  $Y_i$  represent the binary outcome variable which is equal to 1 if the answer to the survey question of interest is “Yes” and is equal to 0 if “No.” Furthermore, let  $R_i$  represent the response indicator which is equal to 1 if  $Y_i$  is observed for respondent  $i$  and is equal to 0 if it is missing. Then, our quantity of interest is the probability of answering “Yes” conditional on a respondent’s race, i.e.,  $\Pr(Y_i = 1 \mid X_i = x)$  where  $X_i$  represents the respondent’s race. Now, note the following simple equality:

$$\begin{aligned} \Pr(Y_i = 1 \mid X_i = x) &= \Pr(Y_i = 1 \mid R_i = 1, X_i = x) \Pr(R_i = 1 \mid X_i = x) \\ &\quad + \Pr(Y_i = 1 \mid R_i = 0, X_i = x) \{1 - \Pr(R_i = 1 \mid X_i = x)\}. \end{aligned} \quad (1)$$

In the right hand side of this equation,  $\Pr(Y_i = 1 \mid R_i = 0, X_i = x)$  is not identifiable from the observed data and its value can range from 0 to 1. The observed data do not provide any information about what value is more likely within this range. Thus, the following bounds can be derived by substituting the extreme values for  $\Pr(Y_i = 1 \mid R_i = 0, X_i = x)$  in equation 1 (Horowitz and Manski, 1998):

$$\begin{aligned} &\Pr(Y_i = 1 \mid R_i = 1, X_i = x) \Pr(R_i = 1 \mid X_i = x) \\ &\leq \Pr(Y_i = 1 \mid X_i = x) \leq 1 - \{1 - \Pr(Y_i = 1 \mid R_i = 1, X_i = x)\} \Pr(R_i = 1 \mid X_i = x). \end{aligned} \quad (2)$$

The estimation of the bounds is straightforward. In fact, they can be non-parametrically and consistently estimated by using appropriate sample averages to estimate the conditional probabilities. In our analysis, we use survey weights to compute weighted sample averages to consistently estimate the bounds in equation 2.

Next, we extend the method of bounds to the analysis of our experiment. To describe our method, we begin by defining the binary treatment variable  $T_i \in \{0, 1\}$  where  $T_i = 1$  if the  $i$ th respondent is assigned to the black mother image and  $T_i = 0$  if he/she is assigned to the white family image. Following the statistical literature of causal inference (e.g., Holland, 1986), we define the potential outcomes  $Y_i(t)$  which would be observed if respondent  $i$  receives the treatment value of  $t$ . Since all of our outcome variables of interest are binary, the observed outcome is denoted by  $Y_i \equiv Y_i(T_i) \in \{0, 1\}$ .

Our quantity of interest is the average treatment effect (ATE) between the two different treatments for each racial group. We denote this quantity  $\tau_x$  and define it as:

$$\tau_x \equiv E[Y_i(1) - Y_i(0) | X_i = x] = \Pr(Y_i = 1 | T_i = 1, X_i = x) - \Pr(Y_i = 1 | T_i = 0, X_i = x), \quad (3)$$

where the equality follows from the randomization of the treatment. To non-parametrically control for the randomized text manipulation, in our analysis we further decompose  $\tau_x$  as  $\frac{1}{2} \sum_{s \in \{0,1\}} [\Pr(Y_i = 1 | T_i = 1, S_i = s, X_i = x) - \Pr(Y_i = 1 | T_i = 0, S_i = s, X_i = x)]$  where  $S_i$  is the indicator variable representing whether respondent  $i$  is assigned to “American” rather than “Refugee.” For the sake of notational simplicity (but without loss of generality), however, we explain our method without the reference to this aspect of the experiment.

Equation 3 implies that the identification of the ATE, i.e.,  $\tau_x$ , requires the identification of the conditional probability, i.e.,  $\Pr(Y_i = 1 | T_i = t, X_i = x)$ . However, in the presence of missing data, we can only identify the bounds of this quantity. Thus, using the result in equation 2, the following bounds on  $\tau_x$  can be derived (Horowitz and Manski, 2000; Imai, 2006):

$$P_1(x)Q_1(x) - P_0(x)Q_0(x) + Q_0(x) - 1 \leq \tau_x \leq P_1(x)Q_1(x) - P_0(x)Q_0(x) - Q_1(x) + 1. \quad (4)$$

where  $P_t(x) = \Pr(Y_i = 1 | T_i = t, R_i = 1, X_i = x)$  and  $Q_t(x) = \Pr(R_i = 1 | T_i = t, X_i = x)$  for  $t = 0, 1$ . For example, the lower bound is obtained when non-respondents in the group with  $T_i = t$

have the opinion of “No” with probability 1, while non-respondents in the group with  $T_i = s$  have the opinion of “Yes” with probability 1. The width of the bounds is equal to  $2 - Q_t(x) - Q_s(x)$ , which confirms our intuition that the bounds are shorter if the response rates are higher. As before, the bounds in equation 4 can be non-parametrically and consistently estimated by using the appropriate sample weighted averages (with survey weights) to estimate  $P_t(x)$  and  $Q_t(x)$  for  $t = 0, 1$ .

## A.2 Confidence Intervals of the Estimated Bounds

While the estimation of the bounds is straightforward, these estimates themselves have sampling variability in finite samples and hence we compute the confidence interval to quantify the estimation uncertainty. One straightforward way to construct the  $(1 - \alpha)$  level confidence interval for the estimated bounds is to compute the confidence interval separately for the estimated upper and lower bounds based on their (asymptotic) sampling distributions and take the upper confidence limit of the estimated upper bound and the lower confidence limit of the estimated lower bound. Such confidence intervals are known as the Bonferroni confidence intervals in the statistics literature.

However, these Bonferroni confidence intervals are necessarily conservative and hence too wide. To see this formally, use a general notation and let  $\theta_L$  and  $\theta_U$  be the true upper and lower bounds of the parameter of interest,  $\theta$ . If we further denote their  $(1 - \alpha)$  level Bonferroni confidence interval by  $\tilde{\theta}_{L_\alpha}$  and  $\tilde{\theta}_{U_\alpha}$ , then for any  $\alpha \in (0, 1)$  the Bonferroni inequality implies:  $\Pr([\theta_L, \theta_U] \subset [\tilde{\theta}_{L_\alpha}, \tilde{\theta}_{U_\alpha}]) \leq \Pr(\theta_L \geq \tilde{\theta}_{L_\alpha}) + \Pr(\theta_U \leq \tilde{\theta}_{U_\alpha}) - 1 = 1 - \alpha$ , where  $\Pr(\theta_L \geq \tilde{\theta}_{L_\alpha}) = \Pr(\theta_U \leq \tilde{\theta}_{U_\alpha}) = 1 - \alpha/2$  by construction. Thus, the inequality shows that the coverage probability of the Bonferroni confidence intervals is greater than or equal to their nominal coverage probability.

To overcome this problem, we follow and Imai and Soneji (2007) and apply the bootstrap method developed by Beran (1988). The resulting confidence interval has attractive properties: if we use  $[\hat{\theta}_{L_\alpha}, \hat{\theta}_{U_\alpha}]$  as this confidence interval, then (1) its (asymptotic) coverage probability is exact,

i.e.,  $\Pr([\theta_L, \theta_U] \subset [\hat{\theta}_{L\alpha}, \hat{\theta}_{U\alpha}]) = 1 - \alpha$ , and (2) the interval is balanced thereby treating the upper and lower bounds equally, i.e.,  $\Pr(\theta_L \geq \hat{\theta}_{L\alpha}) = \Pr(\theta_U \leq \hat{\theta}_{U\alpha})$  holds asymptotically. Thus, like the estimation of the bounds, the construction of the confidence interval is nonparametric and does not require any distributional assumption.

To apply this method, let  $\theta_L^*$  and  $\theta_U^*$  represent bootstrap replicates of the estimated lower and upper bounds. We then compute  $\tilde{c}_{L\alpha} = \tilde{F}_L^{-1}[\tilde{F}^{-1}(1 - \alpha)]$  and  $\tilde{c}_{U\alpha} = \tilde{F}_U^{-1}[\tilde{F}^{-1}(1 - \alpha)]$ , where  $\tilde{F}_L$  and  $\tilde{F}_U$  are the empirical distribution functions for  $\theta_L^* - \hat{\theta}_L$  and  $\hat{\theta}_U - \theta_U^*$ , and  $\tilde{F}$  is the empirical distribution function of  $\max\{\tilde{F}_L(\theta_L^* - \hat{\theta}_L), \tilde{F}_U(\hat{\theta}_U - \theta_U^*)\}$ . Finally, the confidence interval is given by  $[\hat{\theta}^L - \tilde{c}_{L\alpha}, \hat{\theta}^U + \tilde{c}_{U\alpha}]$ .

## References

- Beran, R. (1988). Balanced simultaneous confidence sets. *Journal of the American Statistical Association* **83**, 403, 679–686.
- Dawson, M. C. (2006). After the deluge: Publics and publicity in Katrina’s wake. *Du Bois Review* **3**, 1, 239–249.
- Duncan, O. D. and Davis, B. (1953). An alternative to ecological correlation. *American Sociological Review* **18**, 665–666.
- Fong, C. M. and Luttmer, E. F. (2007). What determines giving to Hurricane Katrina victims? experimental evidence on income, race and fairness. Unpublished manuscript.
- Gilens, M. (1988). Gender and support for Reagan: A comprehensive model of presidential approval. *American Journal of Political Science* **32**, 1, 19–49.
- Gilens, M. (1999). *Why Americans Hate Welfare: Race, Media, and the Politics of Antipoverty Policy*. University of Chicago Press, Chicago, IL.
- Gilliam, F. D. and Iyengar, S. (2000). Prime suspects: The influence of local television news on the viewing public. *American Journal of Political Science* **44**, 3, 560–573.
- Gilligan, C. (1982). *In a Different Voice: Psychological Theory and Women’s Development*. Harvard University Press.
- Gross, K. and Kohenak, M. (n.d.). Race, poverty and causal attributions: Media framing of the aftermath of Hurricane Katrina. Unpublished manuscript.
- Gurin, P., Miller, A. H., and Gurin, G. (1980). Stratum identification and consciousness. *Social Psychology Quarterly* **43**, 1, 30–47.

- Harris-Lacewell, M. (2007). Do you know what it means... mapping emotion in the aftermath of Katrina. *Souls* **9**, 1.
- Holland, P. W. (1986). Statistics and causal inference (with discussion). *Journal of the American Statistical Association* **81**, 945–960.
- Hopkins, D. J. (2007). Politicized change: How local reactions to the post-Katrina migration were shaped by the media. Unpublished manuscript.
- Horiuchi, Y., Imai, K., and Taniguchi, N. (2007). Designing and analyzing randomized experiments: Application to a Japanese election survey experiment. *American Journal of Political Science* **51**, 3, 669–687.
- Horowitz, J. L. and Manski, C. F. (1998). Censoring of outcomes and regressors due to survey non-response: Identification and estimation using weights and imputations. *Journal of Econometrics* **84**, 37–58.
- Horowitz, J. L. and Manski, C. F. (2000). Nonparametric analysis of randomized experiments with missing covariate and outcome data. *Journal of the American Statistical Association* **95**, 449, 77–84.
- Huber, G. A. and Lapinski, J. S. (2006). The “Race Card” revisited: Assessing racial priming in policy contests. *American Journal of Political Science* **50**, 2, 421–440.
- Huddy, L. and Feldman, S. (2006). Worlds apart: Blacks and whites react to Hurricane Katrina. *Du Bois Review* **3**, 1, 97–114.
- Imai, K. (2006). Statistical analysis of randomized experiments with nonignorable missing binary outcomes. Tech. rep., Department of Politics, Princeton University.

- Imai, K. (2007). experiment: R package for designing and analyzing randomized experiments. available at The Comprehensive R Archive Network (CRAN). <http://cran.r-project.org>.
- Imai, K. and Soneji, S. (2007). On the estimation of disability-free life expectancy: Sullivan's method and its extension. *Journal of the American Statistical Association* Forthcoming.
- Iyengar, S. and Morin, R. (2006). Natural disasters in black and white: How racial cues influenced public response to Hurricane Katrina. *Washington Post* .
- Kellstedt, P. M. (2000). Media framing and the dynamics of racial policy preferences. *American Journal of Political Science* **44**, 2, 239–255.
- Kinder, D. R. and Winter, N. (2001). Exploring the racial divide: Blacks, whites, and opinion on national policy. *American Journal of Political Science* **45**, 2, 439–453.
- Lewis, D. E. (n.d.). FEMA's politicization and the road to Katrina. Unpublished manuscript.
- Logan, J. R. (n.d.). The impact of Katrina: Race and class in storm-damaged neighborhoods. *Spacial Structures in the Social Sciences, Brown University* .
- Manski, C. F. (1995). *Identification Problems in the Social Sciences*. Harvard University Press.
- Mendelberg, T. (2001). *The Race Card: Campaign Strategy, Implicit Messages, and the Norm of Equality*. Princeton University Press, Princeton, NJ.
- Nelson, T. E., Clawson, R. A., and Oxley, Z. M. (1997a). Media framing of a civil liberties conflict and its effect on tolerance. *American Political Science Review* **91**, 3, 567–583.
- Nelson, T. E., Oxley, Z. M., and Clawson, R. A. (1997b). Toward a psychology of framing effects. *Political Behavior* **19**, 3, 221–246.

Schuman, H., Steeh, C., Bobo, L., and Krysan, M. (1997). *Racial Attitudes in America: Trends and Interpretations*. Harvard University Press, Cambridge, MA.

Sherman, A. and Shapiro, I. (2005). Essential facts about the victims of Hurricane Katrina. *Center on Budget and Policy Priorities* .

Tajfel, H. and Turner, J. C. (1986). The social identity theory of inter-group behavior. In S. Worchel and L. W. Austin, eds., *Psychology of Intergroup Relations*. Nelson-Hall, Chicago.

The Pew Research Center for the People and the Press (2005). Huge racial divide over Katrina and its consequences: Two-in-three critical of Bush's relief efforts. Tech. rep., The Pew Research Center for the People and the Press.

Valentino, N. A., Hutchings, V. L., and White, I. K. (2002). Cues that matter: How political ads prime racial attitudes during campaigns. *American Political Science Review* **96**, 1, 75–90.