

Understanding Explicit and Implicit Attitudes: A Comparison of Racial Group and Candidate  
Preferences in the 2008 Election

Shanto Iyengar, Stanford University  
([siyengar@stanford.edu](mailto:siyengar@stanford.edu))

Kyu Hahn, Yonsei University  
([khahn@yonsei.ac.kr](mailto:khahn@yonsei.ac.kr))

Christopher Dial and Mahzarin R. Banaji, Harvard University  
([cdial@wjh.harvard.edu](mailto:cdial@wjh.harvard.edu))    ([mahzarin\\_banaji@harvard.edu](mailto:mahzarin_banaji@harvard.edu))

## Abstract

Using data from a national sample, we show that a measure of implicit racial bias -- the race IAT -- reveals significantly higher levels of anti-black bias than standard survey measures of racial prejudice and that there is only weak correspondence between implicit and explicit measures, thus replicating in this sample previous results from drop-in, web-based samples. In the same sample, we show that a candidate IAT measuring implicit preference for McCain or Obama yields strong explicit-implicit correspondence. Third, we investigate the antecedents of implicit-explicit attitude consistency and find that individuals who face stronger conformity pressures are especially prone to under-report their level of race prejudice. Finally, we report an analysis of the overlap between racial attitudes and candidate evaluations. Although one particular racial attitude -- racial resentment -- proved a robust predictor of both explicit and implicit candidate evaluations, attitudes toward the individual candidates proved more influential than attitudes toward racial groups.

The measurement of Americans' racial attitudes has become especially challenging in the post-civil rights era. On the one hand, there are few traces of overt bigotry. The percentage of white Americans who use stereotypic and derogatory terms such as "lazy" or "unintelligent" to describe African-Americans, for instance, has declined sharply since the 1960s (Gaertner and Dovidio 2005; Virtanen and Huddy 1998; Taylor, Sheatsley, and Greeley 1978) and in 2004, white Americans evaluated black Americans just as favorably as their own group. On the other hand, when racial attitudes are recorded using more indirect questions, there is considerable evidence of persisting anti-black and more general anti-minority group biases in American public opinion (Schuman et al. 1997; Sears and Henry 2005; Kuklinski et al. 1997).

To some extent, the sharp decline in self-reported racial prejudice may represent an artifact of survey research rather than meaningful attitude change. In the social (and sometimes interpersonal) setting of an opinion survey, whites may be motivated to conform to widely-shared egalitarian norms and respond in a manner that suggests the absence of racial bias (see McConahay, Hardee, and Batts 1981). When survey questions are framed so as to disguise the racial cues, however, the results typically indicate that "blatantly prejudiced attitudes still pervade the white population" (Kuklinski et al. 1997, p. 403; also see Crosby et al. 1980). Thus, when people do not recognize that they are violating the norm of racial equality, they feel free to express preferences and stereotyped judgments that are hostile to minorities.

Evidence of lingering racial bias in Americans' policy preferences raises further doubts about the decline of prejudice (see Fording 2003; Quillian 2006). In the case of crime, support for punitive policies such as the death penalty increases significantly when whites learn that the criminal perpetrator is non-white rather than white (Gilliam and Iyengar 2000; Hurwitz and Peffley 2007; Eberhardt et al. 2004). Race bias also characterizes employment decisions; job

applicants with European-sounding first names are preferred (by 50 percent) over applicants with identical resumes, but African American-sounding names (Bertrand and Mullainathan 2004). In short, Americans say they are free of racial bias, but their attitudes and behaviors frequently indicate otherwise.

In order to better detect lingering racial animus, researchers have advocated shifting the definition of prejudice away from explicit racial animus in favor of more indirect and diffuse measures of “symbolic racism” or “racial resentment.” In this revisionist view, prejudice in the modern era is some blend of racial animus and mainstream cultural values that is best captured by focusing on beliefs about minorities’ adherence to the American way (Kinder and Sears, 1981; Kinder and Sanders, 1996; Feldman and Huddy, 2005). Although survey indicators of symbolic racism or racial resentment are known to predict a variety of race-related policy preferences e.g. affirmative action (see Sears and Henry 2005), they have been challenged on the grounds that their content has little to do with race per se (see Sniderman and Piazza 1993; Carmines and Sniderman 1997).

### Implicit Versus Explicit Racial Attitudes

Over the past 25 years, psychologists have arrived at the very same place via a different path. Experiments on the most fundamental aspects of the human mind, such as the ability to perceive (e.g., vision) and remember (memory) have shown not only that the human brain can operate outside conscious awareness, but also that such unintended thought and feeling may even be the dominant mode of operation (Bargh 1999). Evidence from behavior and direct measures of the brain suggest it may be useful to think about two separate systems that have evolved to support the unconscious and conscious aspects of thought. Greenwald and Banaji (1995) offered that the analysis of attitudes, stereotypes, and self-concept could gain from an analysis of

relatively more automatic versus reflective forms of operation and labeled the new system of interest as one that tapped implicit social cognition as distinct from explicit social cognition.

Contemporary psychologists have been less interested in the idea that people may deliberately misrepresent their attitudes and beliefs and have largely assumed that even if that were not the case, the conscious aspect of preferences and beliefs are likely to be a thin sliver of the mind's overall work. In other words, psychologists now believe that the mind's architecture precludes introspective access for the most part and have sought to develop measures of preferences and beliefs (see Banaji and Heiphetz 2010, for a review) that have an existence independent of consciously stated ones. The assumption is that although explicit attitudes do in fact reflect genuine conscious preferences (which, in the case of race, have indeed changed over the course of the past 100 years), they shed no light on less conscious and therefore inaccessible preferences that may nevertheless influence behavior. In the area of race, there is now an extensive literature on implicit attitudes, their relationship to explicit attitudes, and their prediction of behaviors (see Wittenbrink, Judd and Park 1997; Dovidio et al. 2002; McConnell and Liebold 2001). A recent meta-analysis of research using a particular measure of implicit bias, the Implicit Association Test (IAT) showed that implicit measures are better at predicting behavior and incrementally so over explicit measures in the discrimination context (Greenwald et al. 2009).

In general, research on implicit social cognition is marked by a strong effort to develop methods that bypass the standard posing of questions altogether and relies instead on rapid responses to concepts (such as Black and White) and attributes (such as good and bad). Based on the idea that that which has come to be automatically associated will be responded to faster and with fewer errors, these measures focus on the error rates and time taken to respond to

pairings of say {White+good and Black+bad} and the opposite concept+attribute pairs such as {Black+good and White+bad} to generate an indirect measure of racial preference as well as other aspects of social cognition such as stereotypes and identity. There are several such methods, of which the Implicit Association Test (IAT; Greenwald, McGhee, and Schwarz, 1998) and evaluative priming are the most common (see Banaji and Heiphetz 2010; Petty, Fazio, and Brinol 2007).

Just as survey research using newer questions led to the discovery that old-fashioned and modern versions of racial attitudes may be distinct psychological constructs, research on implicit social cognition has shown an even sharper divide between the attitudes towards race expressed on survey questions and those revealed on more automatic measures of implicit bias involving response latency.

#### Overview

Conceptually, we are interested in mapping the distribution of implicit and explicit versions of racial and political candidate attitudes. More than a million implicit association tests have been collected at [implicit.harvard.edu](http://implicit.harvard.edu), but these data are based entirely on self-selected participants. The first test we will provide is to compare data from our representative national sample with these non-random samples. This in itself is an important contribution because there is no evidence as yet that the data generated from large web samples are generalizable. Because data about levels of bias, implicit or explicit, play an important role in policy decisions as well as in shaping the public's understanding of the impact of racial attitudes on significant aspects of life from education and health care to employment, it is especially important to know whether the results reported on group race bias by Nosek, Banaji, and Greenwald (2002) hold up when superior methods of sampling are undertaken.

Second, we introduce two types of race comparisons, one involving attitudes toward the social group Black vs. White (the race IAT) and a second test involving a comparison between two candidates, one of whom is Black and the other White (the candidate IAT). This particular pair of tests has not been administered to the same individuals before and it allows us to observe in this more representative sample, the relationship between group-level attitudes and those toward well-known political candidates who belong to the group.

At the most basic level, these two tests provide the opportunity to evaluate a fundamental question: to what extent does an attitude toward a social group (e.g., black, white) teach us about attitudes toward individual members of the group (Obama, McCain). On the one hand there are many studies showing that one's attitude toward a category predicts attitude toward an instance of that category: loving oceans more than forests should predict a preference for the Aruba coast instead of a Costa Rican rainforest; a strong preference for White over Black Americans should predict a preference for McCain over Obama. On the other hand, when categories are complex, the generic attitude toward the category may only weakly predict attitudes toward a particular instance of the category. One may have a strong preference for White Americans over Black Americans, but may choose to vote for Obama over McCain, because these candidates also vary in many other features such as age, party affiliation, and policy positions, differences that may lead to a break between group attitude and individual attitude. Fiske and Neuberg (1990) in their continuum model of social perception extending from categorical perception to individuated perception laid the foundation for accommodating both group-based perceptions of people versus the piecemeal perception of them as individuals.

In short, the within-subject administration of the two IATs can provide evidence concerning the nature of group versus individual attitudes and the complex pattern of implicit-

explicit relationships for group attitudes (e.g., black vs. white Americans) versus individual attitudes (e.g., Obama vs. McCain). Insofar as the candidate tests involved (a) two well-known and highly scrutinized individuals (Obama and McCain), and (b) the data were collected close enough to the election that most voters' minds were likely made up, we have optimal conditions for observing consistency between explicit and implicit attitudes. Specifically, given the degree of involvement and deliberation over the 2008 election, we expect that explicit and implicit candidate attitudes should be less divergent from each other than implicit and explicit racial group attitudes. We use confirmatory factor analyses to provide evidence of the magnitude of separation between conscious and less conscious preferences when they concern racial groups versus political candidates from these groups.

Following the analysis of attitude consistency across implicit and explicit measures, we turn to identifying a particular source of inconsistency, namely, the tendency of individuals to under-report racial bias in explicit attitudes. We identify respondents especially prone to under-report racial bias, i.e. individuals who report lower levels of explicit bias than their own implicit bias reveals. In effect, we identify individuals with inconsistent explicit and implicit attitudes. Finally, we assess the level of overlap between racial attitudes, both implicit and explicit, and candidate preference. We expect, given the level of attention and deliberation accorded the 2008 election, to find that implicit racial group attitudes (black/white) will not necessarily predict candidate attitudes.

## Indicators

### Implicit Racial Preference

The IAT (Greenwald, McGhee, and Schwartz 1998) is a computer-based task that requires participants to rapidly sort items into categories. Based on the time it takes to sort these



items and the errors made in sorting, the IAT measures the strength of association between any category (say animals vs. plants, Hispanics vs. Africans) and attributes (good vs. bad, strong vs. weak). Most IATs contain four distinct categories consisting of a pair of targets (e.g., African American and European American) and a pair of attributes (e.g., good and bad). These category labels are displayed on either the left or right side of the screen while words or pictures representing those categories appear one by one in the center of the screen.

Participants sort each item as it appears into its corresponding category using only two computer keys: 'E' for items representing category A (say African American) on the left, 'I' for items representing category B (say white American) on the right. The same occurs for classifying attributes "good" and "bad" using the same keys, with the critical blocks of trials merging the two: for half the trials, African American and good share a response key while white American and bad share a different key; for the other half of the critical trials African American and bad share a response key while white American and good share a different key. For a demonstration, readers can visit <http://implicit.harvard.edu> and sample one of 14 tests at the demonstration website or many more at the research website.

In the case of the race IAT, the target categories African American and European American are represented by images of black and white faces (available at <http://www.projectimplicit.net/research.php>), while the attribute categories good and bad are represented by words conveying positive and negative concepts (e.g., wonderful, joy, laughter and terrible, hurt, failure). Implicit race attitudes are assessed by subtracting the response times during blocks with hypothesized compatible pairings (e.g., African American paired with bad & European American paired with good) from the response times during blocks with hypothesized

incompatible pairings (e.g., African American paired with good & European American paired with bad).

For the race IAT used in this study, positive values represent faster sorting when African American is paired with bad and European American is paired with good (compared to the inverse); negative values represent faster sorting when African American is paired with good and European American is paired with bad (compared to the inverse). In short, positive IAT scores represent a race preference for whites. An effect size, or “IAT score,” ranging from -2 to 2 is calculated for each participant based on this difference (Full details on scoring an IAT are presented in the Appendix; see Greenwald et al, 2003 for a detailed description for computing the D score, a measure of effect size related to Cohen’s d.).

Since it was developed in the 1990s, the race IAT has been used in dozens of papers as a measure of implicit race bias and in studies of intergroup variation in race attitudes (for a review see, Nosek et al. 2002; for critical commentary on the IAT and responses, see Blanton & Jaccard 2006; Greenwald, Nosek, and Sriram 2006).

### Explicit Racial Preference

We relied on two widely utilized survey indices of explicit racial attitudes -- overt racism and racial resentment. The former is based on a set of four trait ratings that respondents apply to African-Americans and whites.<sup>1</sup> The latter is based on a set of four agree-disagree items that tap

---

<sup>1</sup> The first item in the set was worded as follows: “We’re interested in your opinions about different groups in our society. Using the scale shown below, where a score of 1 would mean that you think most of the people in the group tend to be “hard working,” while a score of 7 would mean that most of the people are “lazy,” where would you place African-Americans.” This was followed by scales with end points of “violent” and “peaceful,” “self-reliant” and “prefer to be on welfare,” and “interact with people of different backgrounds” and “stick to themselves.”

We converted each item to a 0-1 metric, summed the four responses aimed at each group and divided by four. The final indicator was the difference between the ratings of whites and blacks. The Alpha values for the African-American and White indices were .77 and .67 respectively.

beliefs about minorities, individualist cultural values, and support for racial equality.<sup>2</sup> In addition to the indices of overt racism and racial resentment, we also compare respondents' thermometer ratings (on a 1-10 scale) of self-reported warm or cold feelings towards African-Americans and European-Americans.

### Implicit Candidate Preference

Since the development of the race IAT, the methodology has been extended to several other attitude domains including gender, skin color, body weight, nationality, sexual orientation, disability and politics. The candidate IAT is based on the same procedures and measurement as the race IAT. However, the target labels European American and African American are instead represented by targets labeled John McCain and Barack Obama. Multiple images of each candidate constituted the stimuli for the Obama and McCain categories and were matched along obvious dimensions such as clarity, pose, facial expression and background. To make interpreting the relationship between group and candidate IATs intuitive, positive candidate IAT scores represent faster sorting of Barack Obama paired with bad and John McCain paired with good (compared to the inverse); negative values represent the opposite, i.e., a relatively more positive implicit attitude toward Obama over McCain. Scores for this IAT are interpreted as an implicit measure of candidate preference; the higher the candidate IAT D score, the stronger the preference for McCain over Obama.

---

<sup>2</sup> The items, taken from Kinder and Sanders (1996) were as follows. (1) "Over the past few years, blacks have got less than they deserve." (2) "The Irish, Italians, Jews, Vietnamese and other minorities overcame prejudice and worked their way up. Blacks should do the same without any special favors." (3) "It's really a matter of some people not trying hard enough; if blacks would only try harder they could be just as well off as whites." (4) "Generations of slavery and discrimination have created conditions that make it difficult for blacks to work their way out of the lower class." Respondents answered each item along a four-point scale that ranged from "strongly agree" to "strongly disagree." Items 2 and 3 were reflected, the items were converted to a 0-1 metric and an index score was computed as the average of the six items. Coefficient Alpha was .89.

## Explicit Candidate Preference

The pre-election survey included an extensive set of questions measuring respondents' preference for John McCain and Barack Obama. Respondents indicated their feelings (warm or cold) towards each candidate on a 100-point thermometer scale. They also indicated whether a set of positive and negative emotions described their feelings about Obama and McCain.<sup>3</sup>

## The Sample

Our study utilizes a matched online sample of 1100 registered voters recruited from the Polimetrix National Panel. Polimetrix (PMX) maintains a large online panel of American adults (N in excess of one million) who agree to participate in surveys in exchange for accumulating credit points applicable towards acquiring various consumer products (e.g. an Ipod). PMX has developed a matching-based methodology for sampling from their pools of opt-in respondents (details of the sampling methodology are available at [www.polimetrix.com](http://www.polimetrix.com).) First, PMX constructs a sampling frame from the American Community Study with additional data from the Current Population Survey voter supplement and the Pew Religious Life study.<sup>4</sup> From this frame, PMX draws a stratified random sample (the target sample) of people similar in size to the desired sample from their opt-in panel. Next, PMX searches their opt-in online panel for respondents who most closely match the individuals in the target sample on the variables of race, gender, age, education, and imputed party identification. On average, 2-3 matches are drawn for

---

<sup>3</sup> “Now we would like to know something about the feelings you have toward the candidates for President. For each of the two major candidates running for President, please indicate whether something the candidate has done has made you have certain feelings like anger or pride. Has Barack Obama – because of the kind of person he is, or because of something he has done, ever made you feel: angry, hopeful, afraid, proud, happy, sad, and disgusted.” For each candidate, we computed indices of positive and negative affect. (Cronbach’s Alpha ranged from .73 to .85.) We then created a measure of net affect for each candidate (positive affect-negative affect). Finally, we took the difference of these two net indices.

<sup>4</sup> The 2006 American Community Survey (ACS), conducted by the U.S. Bureau of the Census, is based on a probability sample of size 1,194,354 with a response rate of 93.1 percent.

every person in the target sample all of whom are invited to complete the study. From this set of completed interviews, PMX draws the final matched-sample taking the panelists who most closely match the target sample counterparts. The end result is a sample of opt-in respondents with equivalent characteristics as the target sample on the matched characteristics listed above; under most conditions, the matched sample will converge with a true random sample (see Rivers 2005).<sup>5</sup>

The panelists for this study were recruited to participate in a survey of election-related attitudes. PMX fielded the online survey during the second week in October. On completing the survey, respondents were directed to the Project Implicit website where they were given a “warm-up” IAT designed to acclimatize them to the reaction time protocol followed by the race and candidate IATs. Finally, the IAT data were merged with the survey data.

### Analysis

The data analysis proceeds in several stages. First, we compare the distributions of the implicit indicators of racial preference and candidate preference in this national sample with previously reported findings based on opt-in samples. Second, focusing on the national sample, we compare the distribution of the explicit and implicit indicators of racial and candidate attitudes. Our third objective is to examine whether implicit-explicit attitude inconsistency can

---

<sup>5</sup> The fact that PMX matches according to a set of demographic characteristics does not imply that their samples are unbiased. All sampling modes are characterized by different forms of bias and opt-in Internet panels are no exception. Systematic comparisons of PMX matched samples with RDD (telephone) samples and face-to-face interviews indicate trivial differences between the telephone and online modes, but substantial divergences from the face-to-face mode (Hill, Vavreck, and Zaller 2007; Malhotra & Krosnick 2007). In general, the online samples appear biased in the direction of politically attentive voters. For instance, in comparison with National Election Study respondents (interviewed face-to-face), PMX respondents were more likely by eight percentage points to correctly identify the Vice-President of the US. Because attentiveness is likely to be associated with recognition of cultural norms, it is possible that the level of under-reporting of racial bias may be somewhat higher in online samples in comparison with RDD samples.

be attributed in part to a systematic underreporting of racial bias in surveys. Finally, we measure the degree of overlap between racial attitudes on the one hand, and evaluations of an African-American candidate on the other.

#### Implicit Attitudes: Comparing National and Opt-In Samples

This study provides the first administration of the IAT with a representative national sample making it possible to speak to the robustness of the opt-in data by observing whether the data from this culled sample converges or diverges from it. We begin by comparing the distribution of the race and candidate IATs in our national sample with the corresponding distribution in the pooled, drop-in samples collected at [www.implicit.harvard.edu](http://www.implicit.harvard.edu). As shown in Figure 1, the level of implicit racial bias is remarkably consistent across the opt-in and representative samples. The overwhelming majority of respondents – 79 percent of the opt-in sample and 81 percent of the national respondents revealed an implicit preference for whites.<sup>6</sup>

(Figure 1 here)

The consistency of the two distributions of the race IAT in the two samples is further demonstrated by comparisons within racial groups. In both samples, white and Hispanic respondents indicated a stronger preference for white rather than black Americans. For black American participants, the IAT is distributed more evenly with the negative mean indicating a slight preference for blacks over whites. In the [implicit.harvard.edu](http://implicit.harvard.edu) database, 47 percent of blacks showed a pro-white preference, in the PMX sample, it is 45 percent.

The striking correspondence between the opt-in and national samples suggests that implicit racial preference is driven more by racial (and ethnic) affiliation and less by attributes

---

<sup>6</sup> The figure shows “violin plots” -- a combination of standard box plots with a smoothed histogram.

such as education or age, both of which are associated with willingness to take online surveys.<sup>7</sup> In fact, using hierarchical regression (see Table 1), we find that most of the variance in the race IAT is explained by the race of the respondent. Respondents' performance on the race IAT is only weakly correlated with level of education, age, gender, political party identification, or support for egalitarian values.<sup>8</sup> In other words, implicit racial preference primarily reflects the individual's group membership and little else.

(Table 1 here)

Next we turn to the candidate IAT. As shown in Figure 2, the level of implicit preference for Obama differed across the opt-in and national samples. The mean of -.12 in the opt-in sample indicates a clear preference for Obama over McCain, while the mean of .05 shows that the national sample is more evenly divided with a slight preference for McCain.

The considerable variation in implicit candidate preference across the two samples is attributable to the over-representation of Democrats among opt-in participants. (Democrats account for nearly two-thirds of the Project Implicit participant pool.) When we compare the mean IAT score within partisan groups, however, the results prove generally consistent: McCain is favored by over 80 percent of the Republicans in both samples, while Democrats show an equally strong preference for Obama. In other words, when the opt-in sample is brought into line with the national sample on the percentage representation from both parties, the correspondence in candidate preference is again comparable.

---

<sup>7</sup> In the Project Implicit database, the median age of study participants is 26. In the national sample it is 49. As might be expected the education profile of the two groups is also at odds; in the PMX sample, 30% are college graduates; in the Project Implicit database, however, college graduates account for more than 60 percent of the participant pool.

<sup>8</sup> We used two agree-disagree questions to measure egalitarianism. (1) Our society should do whatever is necessary to ensure that everyone has an equal opportunity to succeed. (2) This country would be better off if we worried less about how equal people are. The correlation between the two was .44. The egalitarianism score is based on the average response, scaled from 0-1.

(Figure 2 here)

Party affiliation and egalitarianism are the strongest predictors of implicit candidate preference (see Table 2). Once the effects of these political predispositions are accounted for, respondents' race contributes very little additional explanatory leverage. In short, implicit attitudes towards individual candidates are driven by political considerations, while implicit attitudes concerning racial groups are driven by individuals' racial identity.

(Table 2 here)

### Consistency of Explicit and Implicit Attitudes

We turn next to examining the level of consistency across implicit and explicit attitudes within the race and candidate evaluation domains, presenting the percentage of the national sample favoring whites and Obama (see Table 3). Where appropriate, we compute Cohen's *d* as an approximate measure of effect size.<sup>9</sup> We also present the simple correlations (*r*) between the implicit and explicit indicators.

(Table 3 here)

There is an unmistakable pattern to the data -- implicit and explicit preferences diverge in the arena of race, but converge in the case of well known candidates for elective office. The significantly lower level of preference for whites (and the corresponding smaller values of Cohen's *d*) associated with the explicit indicators suggests a considerable mismatch of explicit with implicit racial attitudes. As generally documented in previous studies based on less

---

<sup>9</sup> Cohen's *d* requires comparability of stimuli across the implicit and explicit domains (Cohen 1982). In the case of race, we have full comparability between the IAT, the survey measure of overt racism, and the race thermometers. In all these cases, the responses indicate positive or negative affect for blacks/whites. The index of racial resentment, however, mixes items about race with items about political values. Accordingly, it is not possible to calculate any measure of effect size attributable to race per se. Strictly speaking, comparing effect size across indicators assumes equivalent midpoints (and endpoints). The items we compare here have very different metrics; the *d* values are thus presented as rough approximations of effect size.



representative samples (Nosek et al., 2002), explicit indicators significantly understate the level of race bias in American society. The estimate of racial preference based on the feeling thermometers, for instance, is 41 points lower than the estimate based on the IAT; while 81 percent of the sample has a preference for whites on the IAT, only 40 percent show a similar preference on the feeling thermometers. Although the mean level of race attitudes diverges when comparing implicit and explicit attitudes, the average correlation of the three explicit measures with the race IAT is .25, suggesting that those who rank high in explicit anti-black attitudes are also those who rank high in implicit anti-black attitudes.

Explicit and implicit evaluations of the presidential candidates, on the other hand, prove generally consistent on both comparisons of mean levels of preference and implicit-explicit correlation. Cohen's *d* shows a relatively modest and uniform effect size associated with the race of the candidate and the average spread in support for Obama between the three explicit measures and the IAT is less than five points. Nonetheless, there is some evidence of a "Bradley effect" -- higher levels of explicit than implicit support for Obama. Thus Obama "loses" the election on the basis of the candidate IAT (where McCain obtains 54 percent of the "vote"). The overall correspondence of implicit and explicit evaluations is clearly high -- the average correlation between the implicit indicator and the survey measures is .67, significantly higher than the same correlation for black and white social groups of .25.<sup>10</sup>

### Factor Analysis

The varying level of implicit-explicit attitude convergence across the race and candidate domains raises the basic question of construct validity. Are explicit and implicit attitudes indicators of the same underlying concept (generic racial bias or candidate preference), or do

---

<sup>10</sup> Greenwald et al. (2009a) report a slightly higher level of convergence between the candidate IAT and survey indicators of candidate preference.

they instead represent distinct concepts? Confirmatory factor analysis (Klein 1994) provides an appropriate method for comparing the fit of a measurement model that combines indicators of explicit and implicit preference with models that treat implicit and explicit attitudes as separate concepts. Our baseline model subsumes implicit and explicit attitudes and posits three generic attitudes -- overt racism, racial resentment, and candidate preference.

The race IAT is considered a measure of implicit racism and the candidate IAT an indicator of implicit candidate preference. Given our results concerning the divergence between the race IAT and the survey measures of racial attitudes, we first compare the baseline model with a model that introduces implicit racial preference as a separate factor. Next, we differentiate between explicit and implicit candidate preference by adding the candidate IAT as a separate factor.

Our baseline measurement model consists exclusively of explicit attitudes -- overt racism, racial resentment, and candidate preference. Overt racism and racial resentment are known to tap distinct ingredients of prejudice (see Sears and Henry, 2005). We force the race IAT to be part of the overt racism factor and the candidate IAT to load on the candidate preference factor. We tested the fit of this three-factor model (Model 1 in Figure 3) against the four-factor model (Model 2 in Figure 3) that separates the race IAT from the survey measures of overt racism and the five-factor model (Model 3 in Figure 3) that further distinguishes between explicit and implicit candidate preference.<sup>11</sup>

(Figure 3 - Table 4 here)

---

<sup>11</sup> CFA requires at least two operational indicators of any latent variable. We therefore computed the race IAT score separately for the even and odd blocks (for a similar approach, see Nosek and Smyth 2007). These within-block IAT scores may be treated as “spilt-halves” and are highly correlated. In the case of the race IAT, the correlation between the two blocks is .72; for the candidate IAT, the correlation is .81.

As shown in Table 4, the addition of the race IAT to the baseline model produced a significant improvement in model fit according to the Chi-Square/degrees of freedom, CFI, FMIN and ECVI criteria (for similar results, see Cunningham, Preacher, and Banaji 2001; Nosek and Smyth 2007). Moreover, the improvement in fit caused by the addition of implicit race bias generally surpassed the further improvement associated with the introduction of the candidate IAT as a separate factor.<sup>12</sup> The loadings of the candidate and race IAT on their respective explicit factors are also revealing. While both candidate IAT scores have an average loading of .70 on the generic candidate preference factor, the corresponding average loading for the race IAT on the overt racism factor is around .35. (The full set of factor loadings is available from the authors.) In short, although both IATs represent separate implicit attitudes, the degree of separation between the implicit and explicit attitudes is greater in the area of race; the candidate IAT is not as distinct an implicit attitude as the race IAT, a result suggested by the zero-order correlations and confirmed by the present analysis.

#### The Underreporting of Racial Bias

To this point, we have shown that the consistency of implicit and explicit attitudes is lower for race attitudes than for candidate attitudes. One possible explanation for this result, which we pursue here, is that survey respondents recognize contemporary societal norms and respond in a manner consistent with these norms. They are disinclined to rate minorities negatively (or whites favorably) and, when given a choice between a black and white candidate,

---

<sup>12</sup> The deviation of the RMSEA from this general pattern may be attributed to the sensitivity of this statistic to the degrees of freedom in any given model (see Savalei and Bentler 2006). The degrees of freedom associated with the three models ranges between 160 and 167. A more appropriate RMSEA test is one that is invariant across degrees of freedom. We carried out such a test by comparing two different four-factor models in which we either added the race IAT or candidate IAT to the baseline model. In this comparison, the improvement in the RMSEA associated with the addition of the race IAT proved larger than the comparable improvement associated with the addition of the candidate IAT.

are likely to underreport their support for the latter.<sup>13</sup> In both domains, therefore, although especially in the area of racial attitudes, we expect a systematic tendency to underreport explicit pro-white preferences.

Our methodology for assessing individual-level underreporting is based on a comparison of rankings. Because the implicit and explicit measures are based on different scoring procedures and metrics we first group respondents into ten quantiles based on their attitude scores. Our measure of underreporting is the ratio of the individual respondent's quantile rank on any given pair of implicit-explicit ranks. Since there are ten quantiles, the implicit-explicit rank ratio can range from .1 to 10. A ratio of one would indicate perfect consistency in the two sets of rankings while a ratio of 10 would indicate the extreme pattern of downward (underreporting) bias in the explicit measures, i.e. respondents' implicit rankings exceeding their explicit rankings.<sup>14</sup>

We present the distribution of the four relevant rank ratios in Figure 4. There are two clear patterns. First, noticeably higher mean ratios obtain for the pairings of implicit and explicit racial attitudes. In all four comparisons, the difference in the mean ratios between the two attitude domains proved statistically significant. (The relevant t-statistics ranged between 6.356 and 8.638.) Second, the rank ratios for the racial attitudes show significantly more asymmetry – there are considerably more respondents who score higher on implicit than explicit bias.<sup>15</sup> Both patterns suggest that respondents either explicitly mask their explicit attitudes when answering

---

<sup>13</sup>The so-called “Bradley effect” suggests the operation of such masking mechanisms in election polling. Recent research suggests that the overreporting of support for black candidates has waned in the past decade (see Hopkins 2009).

<sup>14</sup> Conversely, a ratio of .1 would indicate the extreme value of the opposite pattern of explicit rankings > implicit rankings.

<sup>15</sup> Using a sign test, the level of asymmetry is significantly higher in both comparisons involving racial attitudes.

questions about race or have genuine conscious attitudes that are more pro-black and are unaware of their less conscious anti-black attitudes.

(Figure 4 here)

Last, we turn to identifying the individual-level predictors of implicit-explicit consistency. Based on work by Nosek and others (Nosek 2005; Hofmann et al. 2005), we expect underreporting of explicit racial bias to be especially pronounced among respondents for whom questions of race pose self-presentation conflicts. For instance, respondents who are more likely to recognize and endorse egalitarian norms and who affiliate with a party that has nominated a minority candidate are likely to feel greater pressure to report an absence of bias or have acquired a conscious attitude that is genuinely positive. Thus, we predict higher levels of explicit-implicit attitude inconsistency among whites, especially those who are Democrats and more educated, and especially in the arena of race attitudes. Table 5 presents the results of a regression analysis of the four rank ratios in relation to race, education and party identification.<sup>16</sup> At the bottom of the table we present the results of Wald tests comparing the magnitude of the effect of each predictor across the race and candidate domains.<sup>17</sup>

As anticipated, more educated respondents show a stronger tendency to underreport race bias in both attitude domains, but the impact of education is strengthened in the case of racial attitudes. Thus, the more educated are especially likely to underreport their explicit race bias. A similar pattern holds for party identification -- Democrats exhibit more disparity between their

---

<sup>16</sup> Positive regression coefficients indicate increased underreporting. The predictor variables were scored as follows: -3 (Strong Democrat) to 3 (Strong Republican); 0 (African-American), 1 (Whites, Hispanics, Asians); 1 (less than high school), 2 (high school graduate), 3 (some college), 4 (college graduate), 5 (graduate work).

<sup>17</sup> In order to compute the Wald test statistic, we first estimated a set of four seemingly unrelated regressions (SUR) with one of the racial attitude and candidate evaluation ratios as the dependent variables. In each of these regressions, we then applied the Wald test to compare the coefficient estimates for education, race and party identification across attitude domains.

implicit and explicit attitudes, but the Wald tests indicate that the effects of partisanship are magnified for racial attitudes. The finding that Republicans' survey responses are more commensurate with their IAT scores suggests that their racial attitudes and candidate evaluations are relatively "principled," an interpretation offered by several scholars of racial attitudes (e.g. Sniderman et al. 1991; for an opposing view, see Sidanius et al., 1996). In effect, Republicans are less motivated to mask their survey responses because the survey questions implicate not only their group attitudes, but also their conservative ideology.

(Table 5 here)

Racial differences in the level of attitude consistency were not as clear as anticipated.<sup>18</sup> In each of the attitude domains, only one of the two coefficients associated with race proved significant, indicating higher levels of inconsistency among whites. But, when compared with the results for education and partisanship, the effects of race on attitude consistency proved relatively uniform across attitude domains. Unlike the more educated, whites did not feel greater pressure to underreport race bias; instead, they were equally likely to underreport racial prejudice and support for McCain.

#### Racial and Candidate Preference: The Question of Overlap

The multiple comparisons between implicit and explicit measures of race and candidate preference show divergence in the case of race and convergence in the latter case despite the presence of an African-American candidate. We surmise that the enhanced consistency of candidate evaluation reflects differences in both normative pressures and the information context. In the case of race, generally accepted egalitarian norms motivate some respondents to

---

<sup>18</sup> The relatively weak effects of race may be attributable in part to the small number of African American respondents in our sample.

underreport their explicit preference for whites. These same norms are not only less applicable to evaluations of presidential candidates, but they are also trumped by any number of attitude cues that most voters have internalized since childhood, most notably, their sense of party identification and a whole constellation of election-related considerations derived from an accepted and reinforced partisanship. In the context of the 2008 presidential campaign, for instance, the war in Iraq and the state of the American economy dominated the content of everyday news coverage and interpersonal discussions for several months (see Holbrook 2009). In effect, when the attitude targets are Obama and McCain, individuals have access to highly salient partisan affiliations and related attitudes that structure both implicit and explicit candidate evaluations and override any possible effects of the candidates' race.

Our final analysis pits racial attitudes against the standard predictors of presidential vote choice including party identification, assessments of the state of the national economy, policy preference concerning Iraq, and support for egalitarian values. We ran the analysis using both an explicit (the difference in the candidate feeling thermometers) and implicit (the candidate IAT) indicator of candidate preference.<sup>19</sup> The results appear in Table 6.

(Table 6 here)

As expected, in the context of a campaign waged over highly salient issues having little to do with race, the effects of implicit racial attitudes were limited to implicit candidate preference. Explicit racial attitudes, however, were at the forefront of voters' candidate preferences – both explicit and implicit. At the level of explicit candidate preference, while the race IAT proved irrelevant, both measures of explicit racial bias exerted strong effects on the thermometer ratings: Obama was favored by those scoring lower on the racial resentment and

---

<sup>19</sup> The results are no different using either the net affect index or self-reported vote choice.

overt racism indices. In fact, resentment – the combination of racial animus and support for mainstream values -- proved to be the dominant predictor of the thermometer ratings exceeding even the effects of partisanship and respondents’ position on the Iraq War (for similar evidence on the importance of racial resentment in 2008, see Tesler and Sears in press; Jackman and Vavreck 2010). In the case of implicit candidate preference, racial resentment proved just as influential a predictor as the race IAT; the more resentful expressed higher levels of implicit preference for McCain. While racial resentment was a pivotal cue for both implicit and explicit candidate evaluations, the effects of party affiliation, overt racism, issue positions, and beliefs about the economy either dissipated or disappeared altogether when moving from the explicit to implicit level of candidate preference.

There are two interpretations of the pattern of results in Table 6. First, the presence of an African-American candidate elevated the importance of explicit racial attitudes despite the presence of “distractions” in the form of an economic crisis and ongoing military conflicts. In this sense, the presence of Obama racialized the 2008 election (Tesler and Sears in press). The alternative view, however, is that even at the level of explicit racial attitudes, the overlap between group preference and candidate preference is far from complete. Over 60 percent of our sample expressed a preference for whites on the measures of racial resentment and overt racism. Yet this degree of racial bias provided an insufficient impetus to the candidacy of McCain. In this sense, attitudes toward the individual candidates took precedence over attitudes toward racial groups.

### Conclusion

Some 80 percent of Americans harbor implicit bias against blacks. Yet their implicit racial attitude did not spillover to influence preference for a black candidate. One explanation



for the low correlation between the race and candidate IATs is that individuating information about Barack Obama and John McCain proved sufficient for voters to disassociate evaluations of the candidates from their racial group preferences (see Fiske and Neuberg 2001). Alternatively, the availability of a strong anchor (party identification) in the area of candidate evaluation may have served to suppress affective spillover between group and candidate preference.

The substantial discrepancy in the level of race bias elicited by implicit and explicit measures confirm that survey responses underestimate actual levels of bias, sometimes by a considerable margin. Our results are likely to provide a lower bound on the level of underreporting since the online survey platform provides relative anonymity; telephone or in-person interviews would no doubt reveal higher levels of inconsistency. Scholars interested in mapping the role of race in contemporary public opinion would be well advised to utilize both explicit and implicit indicators of race bias. The recent development of an abbreviated, brief race IAT, which can be administered in less than five minutes, means that the inclusion of an implicit measure is both relatively inexpensive and imposes insignificant opportunity costs in the form of displaced survey questions.

One caveat is in order. The logic underlying our comparative analysis -- i.e., that explicit measures of race are suspect since their divergence from the corresponding implicit measures is greater than the divergence observed for candidate preference -- can be challenged on the grounds that the closer correspondence between the candidate measures may in this case be the product of the specific context of our study. These measures were taken during the closing stages of a historic, closely contested, and polarizing presidential campaign. At the time of the study, explicit attitudes towards Barack Obama and John McCain were well developed and based on considerable cognitive investment in the ongoing campaign. The presence of strong explicit

candidate preferences may have facilitated respondents' performance on the candidate IAT, making for an artificially high correspondence between the explicit and implicit measures. With less prominent candidates as the "targets," it is possible that the level of consistency between implicit and explicit measures would have been attenuated.<sup>20</sup>

In closing, we note several further implications of this study for ongoing research on race and politics. First, the discrepancy between the race IAT and the standard survey indicators of race preference suggests that political science research results based on survey data will understate the effects of racial attitudes. The literature on "racial priming" (Mendelberg 2001; Valentino et al. 2002) suggests that exposure to subtle racial cues during campaigns (rather than blatant racist appeals) has the effect of making explicit indicators of racial preference stronger predictors of vote choice or candidate evaluation. However, since the explicit indicators are contaminated with systematic measurement error, estimates of the racial priming effect are likely understated. We might expect stronger priming effects if the measure of racial preference is implicit rather than explicit. Second, in comparison with explicit racial attitudes, we would expect the IAT to more strongly structure policy preferences on "racialized" issues. (The comparison is inapplicable to racial resentment since that measure is more policy relevant by definition.) Positions on the death penalty, drug abuse, welfare reform, affirmative action, and immigration are likely to show greater covariation with the IAT than the measure of overt racism.

---

<sup>20</sup> In other research, we know that implicit candidate attitudes do in fact predict intentions and behavior. Devos (in press) took implicit measures of Obama and McCain's association with "American" and showed that it predicted likelihood of voting; those who implicitly associated Obama more with "foreign" and less with "American" relative to McCain were also less likely to vote for him.

Finally, we hope that this comparative analysis of survey-based and latency based measures will stimulate greater inter-disciplinary conversations about the measurement of racial and political attitudes. Political scientists have for too long relied exclusively on survey instruments, while psychologists have proved reluctant to extend their laboratory procedures to large-scale, more representative subject pools. With the advent of low-cost, Internet-based surveys and the increased availability of software for implementing measurement of implicit attitudes, we look forward to more work that combines these disparate traditions.

## References

- Banaji, Mahzarin R., Larisa Heiphetz. 2010. "Attitudes." In *Handbook of Social Psychology*, eds. Susan T. Fiske, Daniel T. Gilbert, Gardner Lindzey. New York: John Wiley & Sons, 348-88.
- Bargh, John A. 1999. "The Cognitive Monster: The Case Against the Controllability of Automatic Stereotype Effects." In *Dual Process Theories in Social Psychology*, eds. Shelly Chaiken, Yaacov Trope. NY: Guilford Press, 361-82.
- Bertrand, Marianne, Sendhil Mullainathan. 2004. "Are Emily and Greg More Employable Than Lakisha and Jamal? A Field Experiment on Labor Market Discrimination." *American Economic Review* 94(4): 991-1013.
- Blanton, Hart, James Jaccard. 2006. "Arbitrary Metrics in Psychology." *American Psychologist* 61(1): 27-41.
- Cohen, Jacob. 1982. "A Power Primer." *Psychological Bulletin* 112(1): 155-59.
- Crosby, Faye, Stephanie Bromley, Leonard Saxe. 1980. "Recent Unobtrusive Studies of Black and White Discrimination and Prejudice: A Literature Review." *Psychological Bulletin* 87(3): 546-63.
- Cunningham, William A., Kristopher J. Preacher, Mahzarin R. Banaji. 2001. "Implicit Attitude Measures: Consistency, Stability, and Convergent Validity." *Psychological Science* 12(2): 163-70.
- Devos, Thierry. In press. "The Role of Race in American Politics: Lessons Learned from the 2008 Presidential Election." In *The Obamas and a (Post) Racial America*, eds. Gregory S. Parks, Matthew W. Hughey. Oxford: Oxford University Press.

- Dixon, Wilfred J., Alexander M. Mood. 1946. "The Statistical Sign Test." *Journal of the American Statistical Association* 41(236): 557-66.
- Eberhardt, Jennifer L., Phillip A. Goff, Valerie J. Purdie, Paul G. Davies. 2004. "Seeing Black: Race, Crime, and Visual Processing." *Journal of Personality and Social Psychology* 87(6): 876-93.
- Feldman, Stanley, Leonie Huddy. 2005. "Racial Resentment and White Opposition to Race-Conscious Programs: Principles or Prejudice?" *American Journal of Political Science* 49(1): 168-83.
- Fiske, Susan T., Steven L. Neuberg. 1990. "A Continuum of Impression Formation, from Category-Based to Individuating Processes: Influences of Information and Motivation on Attention and Interpretation." In *Advances in Experimental Social Psychology*, Vol. 23, ed. Mark P. Zanna. San Diego: Academic Press, 1-74.
- Fording, Richard C. 2003. "'Laboratories of Democracy' or Symbolic Politics? The Racial Origins of Welfare Reform." In *Race and the Politics of Welfare Reform*, eds. Sanford Schram, Joe Soss, Richard C. Fording. Ann Arbor: University of Michigan Press, 77-100.
- Gaertner, Samuel L., John F. Dovidio. 2005. "Understanding and addressing Contemporary Racism: From Aversive Racism to the Common In-Group Identity Model." *Journal of Social Issues* 61(3): 615-39.
- Gilliam, Franklin D., 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., Mahzarin R. Banaji. 1995. "Implicit Social Cognition: Attitudes, Self-Esteem, and Stereotypes." *Psychological Review* 102(1): 4-27.

- Greenwald, Anthony G., Brian A. Nosek, Mahzarin R. Banaji. 2003. "Understanding and Using the Implicit Association Test: I. An Improved Scoring Algorithm." *Journal of Personality and Social Psychology* 85(2): 197-216.
- Greenwald, Anthony G., Brian A. Nosek, N. Sriram. 2006. "Consequential Validity of the Implicit Association Test: Comment on Blanton and Jaccard (2006)." *American Psychologist* 61(1): 56-61.
- Greenwald, Anthony G., T. Andrew Poehlman, Eric L. Uhlmann, E., Mahzarin R. Banaji. 2009. "Understanding and Using the Implicit Association Test: III. Meta-Analysis of Predictive Validity." *Journal of Personality and Social Psychology* 97(1): 17-41.
- Hill, Seth J., James Lo, Lynn Vavreck, L., John Zaller. 2007. "The Opt-In Internet Panel: Survey Mode, Sampling Methodology and the Implications for Political Research." Unpublished Paper, Department of Political Science, UCLA.
- Hofmann, Wilhelm., Tobias Gschwendner, Brian A. Nosek, Manfred Schmitt. 2005. "What Moderates Implicit-Explicit Consistency?" *European Review of Social Psychology* 16: 335-390.
- Holbrook, Thomas M. 2009. "Economic Considerations and the 2008 Presidential Election." *PS: Political Science and Politics* 42: 473-78.
- Hopkins, Daniel J. 2009. "No More Wilder Effect, Never a Whitman Effect: When and Why Polls Mislead About Black and Female Candidates." *The Journal of Politics* 71(3): 769-81.
- Jackman, Simon, Lynn Vavreck. 2010. "Primary Politics: Race, Gender, and Age in the 2008 Democratic Primary." *Journal of Elections, Public Opinion and Parties*, 20(2): 153-86.

- Kinder, Donald R., David O. Sears. 1981. "Prejudice and Politics: Symbolic Racism Versus Racial Threats to the Good Life." *Journal of Personality and Social Psychology* 40(3): 414-31.
- Kinder, Donald R., Lynn M. Sanders. 1996. *Divided by Color: Racial Politics and Democratic Ideals*. Chicago: The University of Chicago Press.
- Kuklinski, James H., Michael D. Cobb, Martin Gilens. 1997. "Racial Attitudes and the 'New South.'" *The Journal of Politics* 59: 323-49.
- Kuklinski, James H., Paul M. Sniderman, Kathleen Knight, Thomas Piazza, Philip E. Tetlock, Gordon R. Lawrence, Barbara Mellers. 1997. "Racial Prejudice and Attitudes Toward Affirmative Action." *American Journal of Political Science* 41(2): 402-19.
- Malhotra, Neil, Jon A. Krosnick. 2007. "The Effect of Survey Mode and Sampling on Inferences about Political Attitudes and Behavior: Comparing the 2000 and 2004 ANES to Internet Surveys with Nonprobability Samples." *Political Analysis* 15: 286-323.
- McConahay, John B., Betty B. Hardee, Valerie Batts. 1981. "Has Racism Declined in America? It Depends Upon Who is Asking and What is Asked." *Journal of Conflict Resolution* 25(4): 563-79.
- McConnell, Allen R., Jill M. Leibold. 2001. "Relations among the Implicit Association Test, Discriminatory Behavior, and Explicit Measure of Racial Attitudes." *Journal of Experimental Social Psychology* 37(5): 435-42.
- Mendelberg, Tali. 2001. *The Race Card: Campaign Strategy, Implicit Messages, and the Norm of Equality*. Princeton: Princeton University Press.


- Nock, Matthew K., Mahzarin R. Banaji. 2007. "Prediction of Suicide Ideation and Attempts Among Adolescents Using a Brief Performance-Based Test." *Journal of Consulting and Clinical Psychology* 75(5): 707–15.
- Nosek, Brian A. 2005. "Implicit-Explicit Relations." *Current Directions in Psychological Science* 16(2): 65-69.
- Nosek, Brian A., Mahzarin R. Banaji, Anthony G. Greenwald. 2002. Harvesting Intergroup Attitudes and Stereotypes from a Demonstration Website. *Group Dynamics* 6(1): 101-115.
- Nosek, Brian A., Frederick L. Smyth, Jeffrey J. Hansen, Thierry Devos, Nicole M. Lindner, Kate A. Ranganath, Colin T. Smith, Kristina R. Olson, Dolly Chugh, Anthony G. Greenwald, Mahzarin R. Banaji. 2007. "Pervasiveness and Correlates of Implicit Attitudes and Stereotypes." *European Review of Social Psychology* 18: 36-88.
- Nosek, Brian A., Anthony G. Greenwald, Mahzarin R. Banaji. 2007. "The Implicit Association Test at Age 7: A Methodological and Conceptual Review." In *Automatic Processes in Social Thinking and Behavior*, ed. John A. Bargh. New York: Psychology Press, 265-92.
- Nosek, Brian A., Frederick L. Smyth. 2007. "A Multitrait-Multimethod Validation of the Implicit Association Test: Implicit and Explicit Attitudes are Related but Distinct Constructs." *Experimental Psychology* 54: 14-29.
- Peffley, Mark, Jon Hurwitz. 2007. "Persuasion and Resistance: Race and the Death Penalty in America." *American Journal of Political Science* 51(4): 996-1012.
- Petty, Richard E., Russell Fazio, Pablo Brinol. 2007. *Attitudes: Insights from the New Implicit Measures*. New York: Psychology Press.



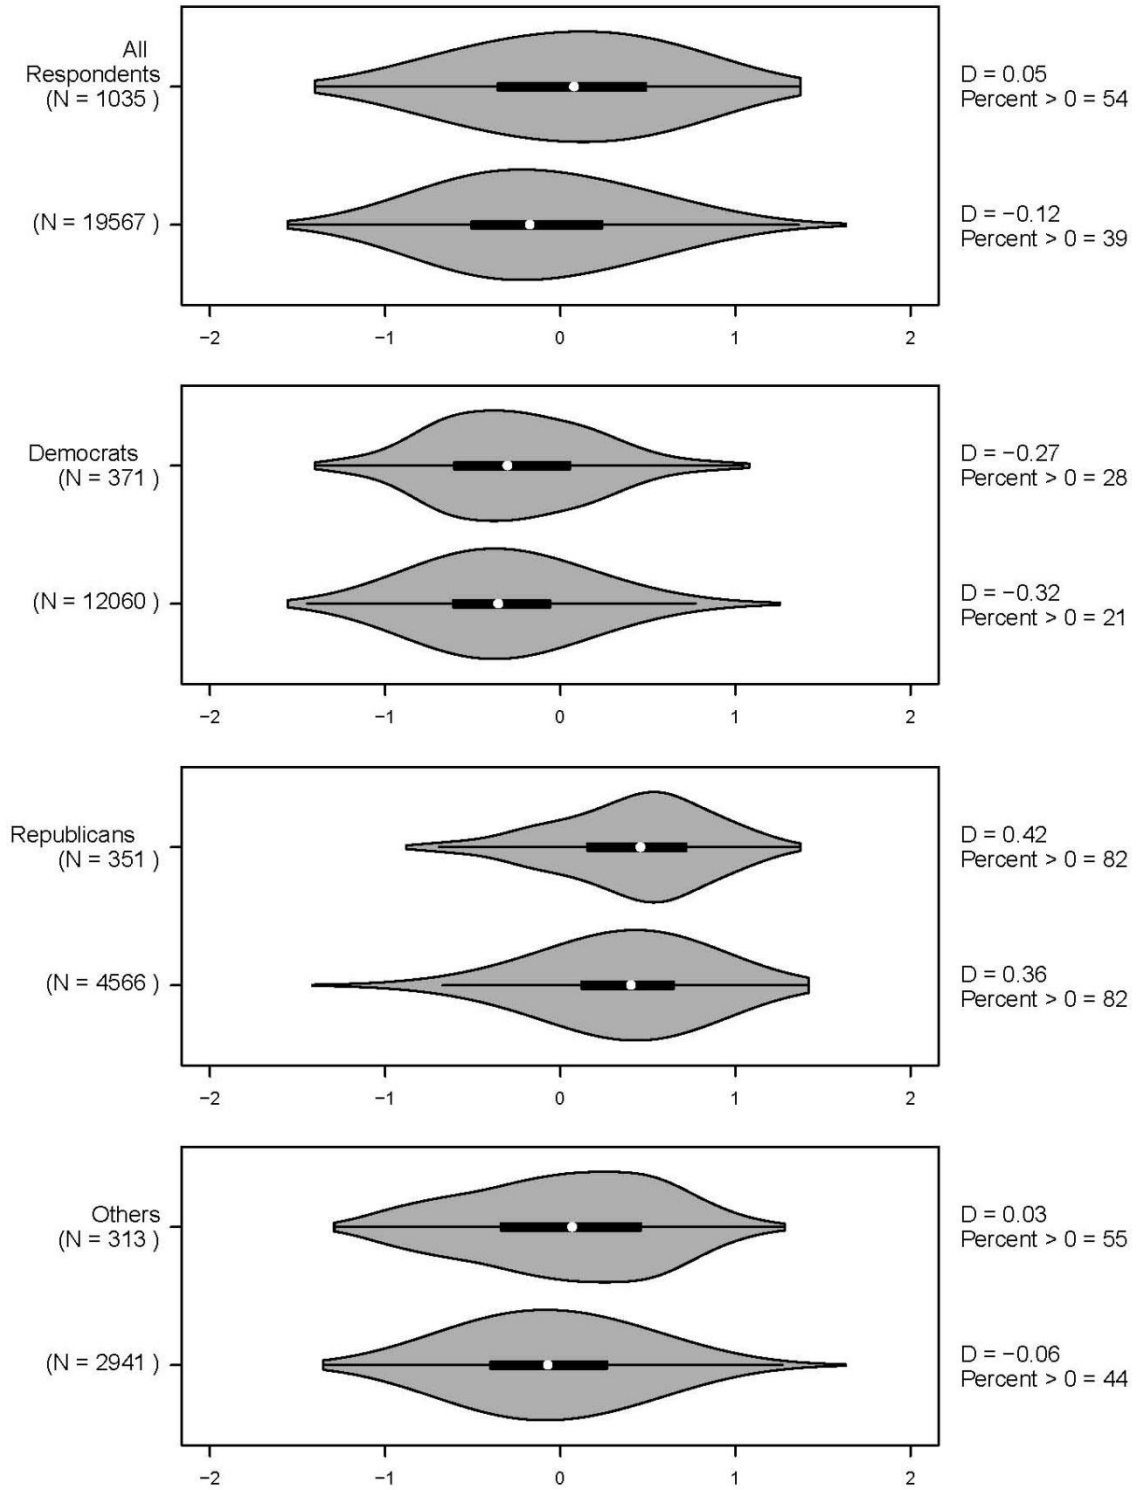
- Quillian, Lincoln. 2006. "New Approaches to Understanding Racial Prejudice and Discrimination." *Annual Review of Sociology* 32: 299-328.
- Rivers, Douglas. 2005. "Sample Matching: Representative Sampling from Internet Panels." Unpublished paper, Department of Political Science, Stanford University.
- Savalei, Victoria, Peter M. Bentler. 2006. "Structural Equation Modeling." In *The Handbook of Marketing Research: Uses, Misuses, and Future Advances*, eds. Rajiv Grover, Marco Vriens. Thousand Oaks, CA: Sage Publications.
- Schuman, Howard, Charlotte Steeh, Lawrence D. Bobo, Maria Krysan. 1998. *Racial Attitudes in America: Trends and Interpretations, Revised Edition*. Cambridge, MA: Harvard University Press.
- Sears, David O., P. J. Henry. 2005. "Over Thirty Years Later: A Contemporary Look at Symbolic Racism". In *Advances in Experimental Social Psychology*, Vol. 37 ed. Mark P. Zanna. San Diego: Elsevier Academic Press.
- Sidanius, Jim, Felicia Pratto, Lawrence Bobo. 1996. "Racism, Conservatism, Affirmative Action, and Intellectual Sophistication: A Matter of Principled Conservatism or Group Dominance?" *Journal of Personality and Social Psychology* 70(3): 476–90.
- Sniderman, Paul M., Edward G. Carmines. 1997. *Reaching Beyond Race*. Cambridge, MA: Harvard University Press.
- Sniderman, Paul M. Philip E. Tetlock. 1986. "Symbolic Racism: Problems of Political Motive Attribution." *Journal of Social Issues* 42(2): 129–50.
- Sniderman, Paul M., Thomas Piazza, Philip E. Tetlock, Ann Kendrick. 1991. "The New Racism." *American Journal of Political Science* 35(2): 423–47.

- Taylor, D. Garth, Paul B. Sheatsley, Andrew M. Greeley. 1978. "Attitudes Toward Racial Segregation." *Scientific American* 238(6): 42-49.
- Tesler, Michael, David O. Sears. 2010. *Obama's Race: The 2008 Election and the Dream of a Post-Racial America* (Chicago Studies in American Politics). Chicago: University of Chicago Press.
- Valentino, Nicholas A., Vincent L. Hutchings, Ismail K. White. 2002. "Cues that Matter: How Political Ads Prime Racial Attitudes During Campaigns." *American Political Science Review* 96(1): 75-90.
- Virtanen, Simo V., Leonie Huddy. 1998. "Old-Fashioned Racism and New Forms of Racial Prejudice." *The Journal of Politics* 60(2): 311-32.
- Wittenbrink, Bernd., Charles M. Judd, Bernadette Park. 1997. "Evidence for Racial Prejudice at the Implicit Level and Its Relationship with Questionnaire Measures." *Journal of Personality and Social Psychology* 72(2): 262-74.

**Figure 1: Comparing the Distribution of the Race IAT in the National and Opt-In Samples**



**Figure 2: Comparing the Distribution of the Candidate IAT in the National and Opt-In Samples**



**Table 1: Hierarchical Regression of Race IAT**

Block	$\Delta$ d.f.	df	F-statistic	p-value	R <sup>2</sup>	$\Delta R^2$
Race	2	925	52.27	.000	.102	
Political Predispositions	2	923	4.82	.008	.111	.009
Education, age	2	921	3.24	.039	.117	.006

Block 1 Dummy variables for African-American, White and Hispanic respondents

Block 2 Party identification and egalitarianism

Block 3. Education and age

**Table 2: Hierarchical Regression of Candidate IAT**

Block	$\Delta$ d.f.	df	F-statistic	p-value	R2	$\Delta$ R2
Political Predispositions	2	935	276.5	.000	.372	
Race	2	933	4.85	.008	.378	.007
Education, Age	2	931	2.11	.122	.381	.003

Block 1 Party identification and egalitarianism

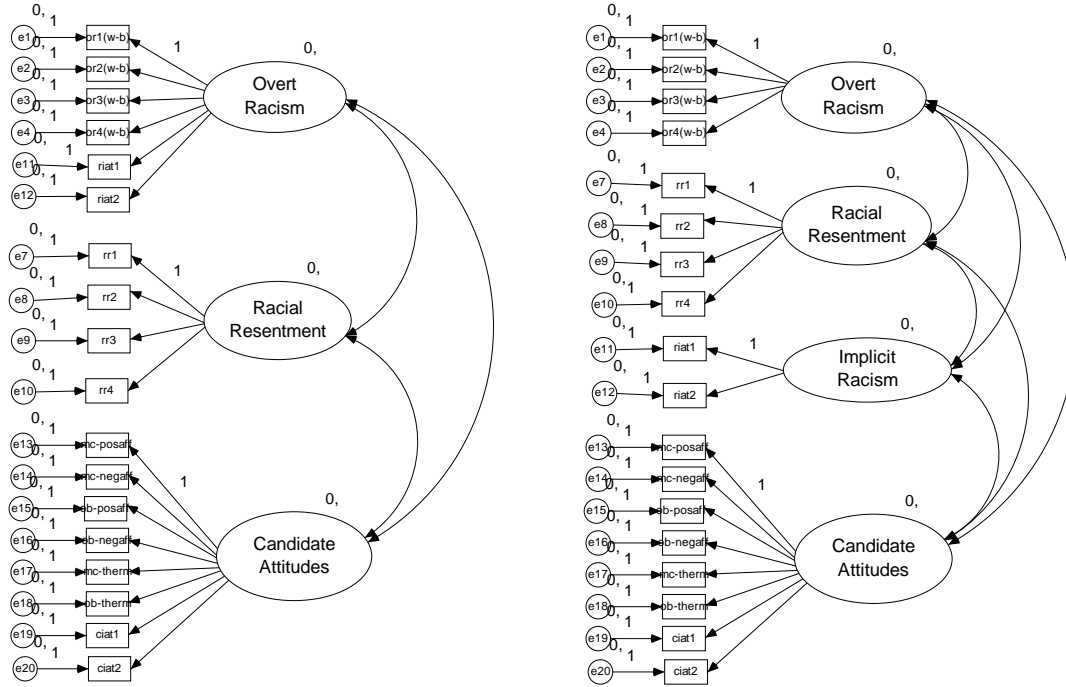
Block 2 Dummy variables for African-American, White and Hispanic respondents

Block 3 Education and age

**Table 3: Comparing the Distribution of Implicit and Explicit Attitudes**

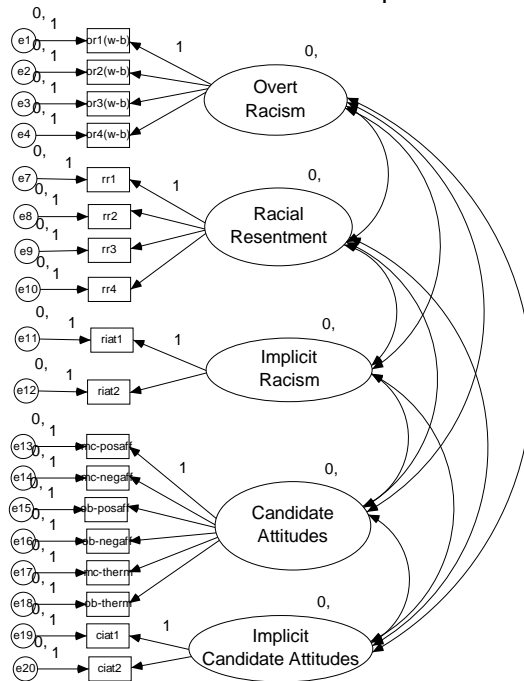
Race Attitudes	% Pro-white	Cohen's d	Imp-Exp Correlation
Race IAT	80.9	.969	
Overt Racism	62.5	.802	.27
Racial Resentment	61.8	--	.23
Race Thermometer	40.4	.435	.25
Candidate Preference	% Pro-Obama	Cohen's d	Imp-Exp Correlation
Candidate IAT	45.4	.094	
Vote Intention	47.5		.67
Candidate Thermometer	52.9	.11	.67
Candidate Affect	50.1	.265	.68

**Figure 3: Confirmatory Factor Analysis Models**



**Model 1: Explicit Attitudes**

**Model 2: Explicit Attitudes and Implicit Racism**



**Model 3: Explicit Attitudes, Implicit Racism, and Implicit Candidate Preference**




**Table 4: CFA Goodness of Fit Statistics\***

	Adding Race IAT to Explicit Attitudes	Adding Cand IAT to Explicit Attitudes and Race IAT
$\Delta\chi^2/\Delta df$	200.939	129.838
CFI	-.047	-.041
NCP	599.816	515.351
FMIN	.548	.473
RMSEA	.017	.017
ECVI	.543	.466

\*Table entries are differences in the value of each statistic between the three- and four-, and between the four- and five-factor models respectively. 200.939, for example, is the difference in the Chi-Square between Models 1 and 2 divided by the difference in the d.f. The statistics showing the fit of each of the three models are available from the authors.

**Figure 4: Consistency of Implicit and Explicit Attitudes**



**Table 5: Effects of Respondent Race, Education and Party Identification On Implicit-Explicit Attitude Consistency**

	Race IAT: Racism	Race IAT: Resentment	Cand IAT: Thermometers	Cand IAT: Affect
Race	-.006 (.199)	-.507** <sup>a</sup> (.177)	-.045 (.111)	-.246* (.108)
Education	.115* <sup>b</sup> (.054)	.223** <sup>c</sup> (.048)	-.021 (.030)	.042 (.029)
Party Id	-.152** <sup>d</sup> (.024)	-.264** <sup>e</sup> (.021)	-.083** (.019)	-.084** (.018)
N	998	1011	915	1023
Adj. R <sup>2</sup>	.046	.163	.071	.082

Table entries are unstandardized regression coefficients with standard errors in Parentheses

<sup>a</sup> Significantly different from the Cand IAT/Therm coefficient estimate at  $p < .01$

<sup>b</sup> Significantly different from the Cand IAT/Therm coefficient estimate at  $p < .05$

<sup>c</sup> Significantly different from both candidate evaluation coefficients at  $p < .01$

<sup>d</sup> Significantly different from both candidate evaluation coefficients at  $p < .05$

<sup>e</sup> Significantly different from both candidate evaluation coefficients at  $p < .01$

\*  $p < .05$ ; \*\* $p < .01$

**Table 6: Effects of Racial Attitudes on Candidate Evaluations**

	McCain-Obama Thermometer	Candidate IAT
Constant	-47.309** (9.795)	-.234* (.115)
Racial Resentment	83.634** (10.562)	.658** (.121)
Overt Racism	-49.887** (12.636)	-.355* (.146)
Race IAT	.341 (3.175)	.163** (.037)
Party Identification	13.103** (.750)	.060** (.009)
Egalitarianism	-12.292 (9.556)	-.068 (.109)
National Economy	11.321** (2.904)	.052 (.034)
Iraq Pullout	-13.984** (1.511)	-.087** (.018)
White Respondent	1.197 (3.750)	-.002 (.043)
Black Respondent	5.326 (5.610)	.009 (.066)
Adj R <sup>2</sup>	.753	.493
N	795	870

Entries are unstandardized regression coefficients with standard errors in parentheses; \* $p < .05$ ; \*\* $p < .01$

## Appendix

The IAT uses response latencies to calculate a “D score” that indicates a participant’s relative association strength for two pairs of concepts (e.g. African American + good, European American + bad) relative to their inverse pairings (e.g. African American + bad, European American + good). Depending on the difference between a participant’s speed at responding to these pairings, the D score can range from -2.0 to 2.0, where zero indicates identical response times for each pairing, and greater or lesser values indicate a relatively stronger (i.e. faster) association for one pairing relative to the other. Negative or positive values indicate the direction of the association (e.g. pro-White or pro-Black).

Standard IATs consist of “blocks” of presented stimuli. The first block introduces the target stimuli (e.g. African American & European American); block two introduces the attribute stimuli (e.g. good & bad); block three presents a pairing of both targets and attributes (e.g. African American + good, European American + bad); block four re-presents the attribute stimuli, though now on the opposite side of the screen (e.g. if good was presented on the right during block two, it now appears on the left); and the final block presents the inverse of the pairing presented in block three. Only blocks three and five (the pairings) are used in the scoring. Blocks one, two, and four are provided for familiarization and do not influence the D score.

The procedure for calculating D scores as follows:

1. Delete trial latencies greater than 10,000 milliseconds
2. Compute a single standard deviation for trials in blocks three and five
3. Compute the mean latency for trials in block three and again for trials in block five
4. Compute the difference between the block three and block five mean latencies\*
5.  $D = \text{the difference from step four divided by the standard deviation from step two}$

\*Note: To minimize the influence of order effects, well-designed IATs counterbalance the presentation of pairings across participants. For example, if even numbered participants first see African American paired with good and European American paired with bad, odd numbered participants first see the inverse pairing. When counterbalancing is used, to ensure that the negative and positive values of D scores remain consistent across participants, the difference computed in step four should correspond to the stimuli presentation order (e.g. for odd participants,  $\text{Mean}_{\text{Block 5}} - \text{Mean}_{\text{Block 3}}$ ; for even participants,  $\text{Mean}_{\text{Block 3}} - \text{Mean}_{\text{Block 5}}$ ). For further details on designing, analyzing, and reporting IAT research, see Greenwald, Nosek, and Banaji (2003), and Lane, Banaji, Nosek, and Greenwald (2007)