

Diversity and Implicit Race Bias in the United States

James R. Rae

A thesis

submitted in partial fulfillment of the

requirements for the degree of

Master's of Science

University of Washington

2014

Committee:

Kristina R. Olson

Cheryl R. Kaiser

Program Authorized to Offer Degree:

Psychology

©Copyright 2014
James R. Rae

University of Washington

Abstract

Diversity and Implicit Race Bias in the United States

James R. Rae

Chair of the Supervisory Committee:

Kristina R. Olson, Ph.D.

Psychology

The world is becoming more ethnically diverse. Research from psychology and other disciplines, primarily examining explicit attitudes among majority-group members, has suggested that greater diversity may be associated both with more positive *and* more negative intergroup attitudes. In the United States (U.S.), given that states are often at the center of debate concerning diversity-related public policy, the present research examined how diversity is associated with state-level implicit race-based bias among White and Black Americans. Findings suggest that greater diversity across U.S. states and counties was associated with stronger implicit ingroup bias among both White and Black respondents. Furthermore, the present research investigated alternative explanations proposed by previous research, concluding that diversity, rather than factors proposed by alternative explanations, best explained the obtained results. The findings convey the importance of investigating why diversity may not always have the positive impact on intergroup relations that one might expect or hope.

Diversity and Implicit Race Bias in the United States

Virtually all modern societies are becoming more ethnically and racially diverse (Putnam, 2007). For example, the United States (U.S.) Census Bureau has projected that racial/ethnic minority groups will make up a majority of the U.S. population by 2042 (U.S. Census Bureau, 2008). Perhaps as a consequence of increasing global diversity, researchers have become increasingly interested in understanding the consequences and correlates of diversity.

Interestingly, research from numerous academic disciplines has found both negative and positive implications of diversity: While a large body of interdisciplinary research suggests that greater diversity is associated with important benefits such as improved group problem solving (Page, 2007) and enhanced educational outcomes (Chang, 1999), other work has controversially suggested that diversity has significant *downsides*, including decreased trust and increased social isolation (Putnam, 2007).

Likewise in the field psychology, research on diversity—which has largely focused on its relationship with *individual-level* explicit intergroup attitudes (the explicit attitudes that individuals endorse about racial groups)—has reached similarly opposing conclusions: A considerable amount of work suggests that greater diversity may be associated with more positive explicit intergroup attitudes (e.g., Sigelman & Welch, 1993). By contrast, other research suggests just the opposite—that greater diversity may be associated with more negative explicit intergroup attitudes at the individual-level (e.g., Bobo, 1988). The majority of work on intergroup attitudes has focused on individual level attitudes, however, the focus of this project was to investigate how *group-level* intergroup attitudes (the aggregated, mean-level attitudes that a group of individuals within a specific geographical region endorse about racial groups) covary with diversity within in a geographical unit. Specifically, the current work examined the

relationship between diversity and state-level implicit race-based bias among both White and Black Americans.

Diversity and Intergroup Attitudes

Negative Effects of Diversity

In U.S. samples, several studies have suggested that diversity can negatively impact intergroup attitudes measured at the group level. For example, Pettigrew (1959) found that among White respondents sampled from four comparable cities in the North ($N = 180$) and South ($N = 186$), participants residing the South had significantly higher scores on a 12-item self-report measure assessing stereotype beliefs (e.g. “Most Negroes are happy-go-lucky and irresponsible”) and exclusion-discrimination beliefs (e.g. “Most Negroes scarcely deserve the right to vote”). Although the slavery era may have created a unique subculture of especially pronounced racial bias in the South, one alternative explanation for this pattern is that the greater proportion of Black residents in southern cities leads to increased racial bias among White respondents in those cities. Strengthening this hypothesis, Pettigrew (1959) also found that increased diversity significantly predicted interracial bias *within* the South: Southern communities with greater proportions of Black residents (38% – 45% of the total population) had significantly higher levels of prejudice than communities with smaller proportions of Black residents (10% – 18% of the total population).

In a later study using data from the 1972 National Election Study (NES) and both the 1976 and 1977 General Social Surveys (GSS; National Opinion Research Center), Fossett and Kiecolt (1989) investigated how both White respondents’ perceived threat from Blacks and support for racial integration covaried with the proportion of Black residents within a community (metropolitan estimates for those residing in cities or county estimates for nonmetropolitan

residents). The authors found that the proportion of Blacks within a community was negatively correlated with community level support for racial integration and positively associated with White respondent's perceived threat from Black residents. Also, further analyses indicated that perceived threat from Black residents predicted lower levels of community support for racial integration. However, even when controlling for the effect of perceived threat from Black residents, the percentage of Black respondents within a community remained a significant— and negative — predictor of support for racial integration among White respondents. The authors argue that this pattern suggests that perceived threat may only partially explain why larger proportions of Black residents within a community negatively correlate with explicit interracial bias among White respondents and that further research is needed to determine what other variables (e.g. anxiety from interracial interactions) may explain this relationship.

Again using GSS data (collected in 1990), Taylor (1998) investigated how the proportion of Black respondents within a community (e.g. metropolitan region), measured by 1990 U.S. Census estimates, correlated with White respondent's explicit interracial attitudes within three domains: Traditional anti-Black attitudes (e.g. aversion to interracial contact), opposition to race-target intervention (e.g. providing opportunities for Black Americans like special scholarships), and policy beliefs (e.g. beliefs concerning whether Black Americans have structural barriers to success). Results indicated that even after controlling for location in the South and total population size of the locality, the percentage of Black respondents within a community positively predicted explicit prejudice within each of the three attitude domains. Interestingly, the level of segregation within each locality did not moderate the effect of the percentage of Black residents on explicit racial prejudice among White respondents. Further, the authors found that

the effect of diversity on explicit racial attitudes among Whites was still significant after controlling for perceived community level political and economic threat from Black residents.

In summary, the positive interregional and intraregional association between percentage of Black residents and explicit race-based bias among White respondents in numerous domains (Fosset & Kiecolt, 1989; Pettigrew, 1959; Taylor 1998) provides support for the controversial hypothesis that increased diversity may actually be harmful to intergroup relations. Interestingly, the negative effect of diversity persists even after controlling for effects of overall perceived threat more generally (Fossett & Kiecolt, 1989) and both economic and political threat (Taylor, 1998) from Black residents. These findings suggest that many different factors may mediate the effect of the proportion of Black residents on prejudice among White residents. Lastly, given that segregation level did not moderate the effect of the proportion of Black residents on explicit attitudes among White respondents, it seems possible that ingroup favoritism among Whites may not be directly activated by living in a racially integrated neighborhood and instead may be a psychological byproduct of living within a larger— and more diverse— geographical unit. Indeed, if direct interactions between White and Black neighbors contributed to explicit racial attitudes among White respondents, we would expect to see higher levels of segregation to attenuate prejudice and lower levels of segregation to exacerbate prejudice (or vice versa).

It should be noted that although the scope of the studies described above have broadly been restricted to intergroup attitudes (e.g. support for racial integration or stereotype beliefs), other research using non-attitudinal dependent variables have also reached the similar conclusion that diversity correlates with suboptimal outcomes: disparities in income, education, and occupation level (Wilcox & Roof, 1978), occurrences of hate crimes (Green, Wong, & Strolovitch, 1996), and certainty of punishment by law enforcement (Liska, Chamlin, & Reed,

1985) increase as a function of racial/ethnic diversity within a geographical unit. However, it is not possible to observe racial disparities in income or hate crimes in the absence of diversity.

Positive Effects of Diversity

In contrast the research carried out in the U.S. suggesting that diversity may in fact harm intergroup relations, a smaller body of research considering primarily European samples has suggested that increased racial/ethnic diversity within a locality actually lower levels of explicit prejudice among majority group members. For example, using two different data sets, Wagner and colleagues (2003) investigated whether proportions of non-German residents could explain differences in levels of prejudice towards non-Germans in East and West Germany. In sample of 2,893 adults (ALLBUS data set), the authors found that opposed to living in West Germany (more non-Germans), residing in East Germany (fewer non-Germans) decreased the likelihood of intergroup friendship with non-Germans ($r = -.14$). Additionally, more intergroup friendships with non-Germans in turn reduced levels of prejudice towards non-Germans ($r = -.23$). In a second sample of 3,560 adolescents and young adults ranging from 15 years to 24 years of age (Shell Youth Study), the authors again found that living in West Germany provided more opportunity for intergroup friendships with non-Germans, which in turn directly reduced prejudice ($r = -.41$). Interestingly, residence in East Germany was associated with higher levels of prejudice towards non-Germans even after controlling for the effect of all contact related variables (e.g. intergroup friendship; $r = .22$).

In another sample comprised of 2,619 German adults with data from the Group Focused Enmity (GFE) survey from 2002, a related study indicated that as the proportion of non-German residents within German districts (approximately 180,000 people) increased, negative explicit evaluations of non-Germans among majority-group members decreased (Wagner et al., 2006).

Further, this negative relationship between district level diversity and prejudice remained significant even after controlling for other potential variables that may explain changes in community level prejudice, including migration levels both in and out of the district and total district size. Similar to the findings of Wagner and colleagues (2003), the authors argue that ethnic diversity increases the opportunity for intergroup friendships, which ultimately results in reduced community level prejudice.

In summary, it appears diversity can at least indirectly provide positive benefits for intergroup relations by providing an *opportunity* for individuals to interact with outgroup members. Although studies demonstrating the positive impact of diversity on group level attitudes has solely consisted of German samples, this research demonstrates that at both the country (East and West Germany; Wagner et al., 2003) and district level (Wagner et al., 2006), larger proportions of non-Germans is associated with reduced levels of explicit prejudice among majority group members. However, given that intergroup friendships only partially explain the positive effects of diversity on majority group members explicit attitudes (Wagner et al., 2003)—similar to how perceived threat only partially explains the negative effects of diversity on racial attitudes in the U.S. (Fossett & Kiecolt, 1989; Taylor, 1998)—it appears that community level diversity, even while having a positive effect, has a complex and multifaceted influence on explicit racial/ethnic attitudes.

Implicit versus Explicit Attitudes

One crucial feature of all of the research reviewed above is that all studies investigating the relationship between diversity and intergroup attitudes have measured *explicit* intergroup attitudes, typically assessed by asking people to self-report their attitudes. Therefore, in addition to having produced inconsistent findings, existing research on the relationship between diversity

and intergroup attitudes is limited by reliance on attitude measures that suffer from social desirability concerns—especially problematic when considering sensitive topics such as race (Crowne & Marlowe, 1960; McConnell & Leibold, 2001). In part because of this concern, within social psychology, and especially in laboratory research, researchers have largely moved from studies of explicit measures of intergroup attitudes to using *implicit* measures of intergroup attitudes (Fazio & Olson, 2003). Implicit measures have allowed researchers to investigate intergroup attitudes with fewer concerns about impression management and conscious control. Further, research has revealed that explicit and implicit intergroup biases are associated with different dimensions of intergroup behavior (Dovidio, Kawakami, & Gaertner, 2002) and separate facets of experience with outgroups (Prestwich, Kenworthy, Wilson, & Kwan-Tat, 2008). Specifically, quality of interactions with outgroup members predicts scores on explicit measures, while quantity of intergroup interactions appears to better predict scores on implicit measures. Accordingly, to fully understand the association between diversity and intergroup attitudes, it is critical to also investigate implicit bias.

Implicit Association Test

The most commonly used implicit attitude measure is the Implicit Association Test (IAT; Greenwald, McGhee, & Schwartz, 1998). The IAT is a response latency measure designed to measure the association between a target concept and an attribute concept. For example, the IAT adapted to studying race (race IAT) typically uses photographs of *White* and *Black* Americans as targets and *good* and *bad* words as concepts of an evaluative attribute. The basic premise of the IAT is that concepts that have come to be closely associated should be paired more rapidly than concepts that are not as closely associated.

The typical procedure of the IAT contains seven blocks (although shorter versions have been formulated: see Sriram & Greenwald, 2009), split between two phases, and administered via a computer. In phase one, during the first block of the race IAT, respondents are instructed to rapidly categorize photographs of each target category (Black and White Americans) as they appear on the computer screen with two computer keys on a computer keyboard. Participants use one computer key in response to all photographs of White Americans and the other computer key in response to all photographs of Black Americans. In the second block, respondents use the same computer keys from the first block to then rapidly categorize words with either positive or negative connotations from the evaluative attribute (e.g. *hate/love*). In the third block, the tasks from the first two blocks are combined and participants practice using the same two computer keys to classify pairings of the target and attribute. So, for example, one computer key may be used in response to either photos of Black Americans or good words and the other computer key may be used in response to photos of White Americans or bad words. Finally, in the fourth block, participants repeat the same categorization procedure from the third block.

In the second phase of the race IAT, the alternative pairing of target and attribute (opposite of the pairings in the third and fourth blocks) are tested. Similar to the first and second blocks described above, the fifth block is a practice block where participants practice with the reversed pairings of computer key to target and attribute categories. For example, the computer key formerly used to respond to photos of White Americans may now be used to respond to photos of Black Americans (or vice versa). Again similar to the procedures described above for the fourth block, in the sixth block, participants practice using two computer keys to categorize the opposite target and attribute pairings. So, for example, participants would now use two computer keys to categorize photos of Black Americans or bad words and photos of White

Americans or good words. In the seventh block, participants repeat the same procedure from the sixth block.

The basic assumption of the IAT is that for any participant, the stronger the association between a paired target and attribute, the faster and more accurately they should be able to categorize the stimuli using the same computer key. For example, a participant with a stronger association between White Americans as bad and Black Americans as good should be able to respond more quickly and more accurately when the same computer key is used to respond to both Black American photos and good words in comparison to when Black American photos are paired with bad words. To directly index the association between the two possible pairings of target and attribute, the average response times between the presentation of stimuli and a response to the categorization task in blocks three and four (e.g., Black Americans with good words and White Americans with bad words) are compared to the response times of the paired in blocks six and seven (e.g., Black Americans with bad words and White Americans with good words). To encourage accurate categorization, participants are only allowed to proceed to the next trial following a correct response and the order presentation of the target and attribute pairings is counterbalanced to avoid biasing the results of the test. After removing careless responders with very slow (any trials greater than 10,000 milliseconds) and very quick response latencies (10% of trials less than 300 milliseconds), standardizing the ratio of the differences in response latencies between both blocks three and four and blocks six and seven, the IAT yields an effect size measure called D (related to Cohen's d) that can take on values ranging from -2 to +2 (Greenwald, Nosek, & Banaji, 2003).

One key advantage of using the IAT is that the measure has been found to have greater predictive validity relative to parallel explicit measures in the intergroup domain (Greenwald,

Poehlman, Uhlmann, & Banaji, 2009; cf. Oswald, Mitchell, Blanton, Jaccard, & Tetlock, 2013).

Therefore, the IAT not only provides an index of socially sensitive attitudes that participants may not be willing, or able, to report on explicit measure, but also provides researchers with more accurate information on how participants may behave in real world interracial interactions.

Unit of Analysis

In the current work, the relationship between racial diversity and implicit bias was investigated at the level of U.S. states because states are often at the center of debate about diversity-related politics and public policy (e.g., regarding immigration; Archibold, 2010). Thus, understanding how diversity relates to group-level attitudes may be especially helpful in contextualizing the environment within which diversity-related policy and politics occur. Furthermore, previous research has suggested that important structural and historical characteristics (e.g., diversity or segregation) may be influential in shaping intergroup attitudes within a spatial unit (e.g., U.S. South; Wagner et al., 2006). Estimating an aggregated implicit bias score for each state, thus averaging across individual sources of variability on the IAT that were not of direct interest to the present research question and that have been thoroughly documented in previous research (Nosek et al., 2007), provided an opportunity to empirically test the hypothesis that state-level diversity impacts the social climate within each state.

Accordingly, adapting the strategy used in prior work on diversity and explicit intergroup attitudes (Christ et al., in press), the current study investigated whether the racial diversity within a U.S. state is associated with the average amount of implicit race-based bias among the residents of that state. Prior work has demonstrated that aggregated IAT scores are useful measures of group-level attitudes that relate to important social phenomena: For instance, Nosek and colleagues (2009) found that nation-level scores on a gender-science IAT predicted national

differences in gender disparities in math and science performance (see also Marini et al., 2013, for a similar approach in the domain of weight bias). Moreover, few studies have investigated the relationship between diversity and minority-group members' attitudes (Shelton, 2000). Thus, the current study further extended prior work by examining the impact of diversity on implicit bias among both Black and White American respondents. Using data aggregated from nearly 1.1 million race IATs taken on the Project Implicit website, this study conducted state-level analyses of the implicit race bias of White and Black Americans. The scope of the present investigation was restricted to White and Black participants respondents because they are the focal groups for the race IAT on the Project Implicit website.

Methods

Measures

Implicit Race Bias. The dependent measure in the analyses were D scores (Greenwald et al., 2003) on the race IAT collected via a virtual laboratory operated through Project Implicit (<https://implicit.harvard.edu/implicit/>) between the years of 2002 and 2012 and recently made available for public use (Xu, Nosek, & Greenwald, in press). All IAT scores were coded to represent ingroup preference. Accordingly, for White respondents, D scores greater than 0 indicated a bias favoring Whites over Blacks (i.e., ingroup bias), D scores at 0 indicated no bias, and D scores below 0 indicated a bias favoring Blacks over Whites (i.e., outgroup bias). Conversely, for Black respondents, D scores greater than 0 indicated a bias favoring Blacks over Whites, D scores at 0 indicated no bias, and D scores below 0 indicated a bias favoring Whites over Blacks.

Because geographic location was of primary interest in the present study, the original data set—which contains over 3.6 million observations—was restricted to observations from

participants indicating they lived/have lived in the U.S. on any of the following variables: "country", "country reside", or "county citizenship," which reduced the sample size to 2,031,503 observations. The sample was further reduced to those respondents that reported their state of residence and fully completed all seven blocks of the IAT testing procedure ($N = 1,837,594$). The sample was further restricted based on the recommendations of the improved IAT scoring algorithm (see Greenwald et al., 2003 and Nosek et al., 2009), such that data from participants who made errors on over 30% of trials and had more than 10% of trials with reaction times less than 300 milliseconds were excluded. No respondents in the public use data set had response latencies greater than 10,000 milliseconds. Given that individual trial IAT performance is not included in the public use data, individual trial exclusion recommendations were not used. These screening procedures reduced the sample size by approximately 2%. In comparison, these screening procedures for a similar analysis with IAT data excluded approximately 9% of the sample (Nosek et al., 2009).

Given research suggesting that the impact of diversity on intergroup attitudes (Tropp & Pettigrew, 2005) and performance on the race IAT (Nosek et al., 2009) differ based on respondents' race, one question of interest in the present study was whether racial diversity within a state differentially predicted implicit bias among White and Black respondents. Although this research question could be investigated in numerous ways, for ease of reporting and given little concerns for losing statistical power because of the large sample size contributing to the state-based analyses, separate analyses were performed on the subsamples of White ($n = 923,001$) and Black ($n = 164,996$) respondents. That is, each analysis investigating the relationship between diversity and implicit race-based bias was performed separately for data collected from Black and White respondents.

Diversity within States. Estimates used to calculate the racial diversity within each U.S. state were obtained via the American Community Survey (ACS; <http://www.census.gov/acs/www/>), which is conducted yearly by the U.S. Census Bureau. To maximize the accuracy of state-level estimates and to best correspond with the time period during which most of the IAT scores were collected, the ACS 5-year averages of data collected between 2008 and 2012 were used to estimate the number of White and Black residents in each state. Because the race IAT measures relative bias favoring Whites over Blacks (or vice versa), state-level diversity was indexed by creating a ratio of the estimated number of Black to White residents in each state (with higher numbers indicating more racial diversity). Given that this ratio was quite small in states with very few Black residents, for all analyses these values were natural-log-transformed to correct for the negative skew. However, for ease of interpretation in descriptive statistics and figures, the diversity index refers to the untransformed number of Black residents for every 100 White residents within a state. Diversity varied greatly across states, with the number of Black residents per 100 White residents in a state ranging from 0.47 to 62.3 (*Mdn* = 9.28).

State-Level Control Variables. Again using ACS 5-year estimates, information about state median income, degree of economic inequality, average education level, and percentage of U.S. citizens within each state were used as covariates in subsequent analyses (Putnam, 2007). The median state household income in 2012 inflation-adjusted dollars was used to measure state income level. The state Gini coefficient, representing state-level economic inequality and ranging between 0 (complete equality) and 1 (complete inequality), was obtained directly from the ACS 5-year estimates. State education level was operationalized as the percentage of state residents holding a bachelor's degree or higher. The percentage of citizens was measured by

using ACS estimates of the percentage of state residents holding U.S. citizenship.

State level statistical analyses were also adjusted for state political orientation, operationalized as the percentage of voters endorsing the Democratic candidate 2012 U.S. presidential election (“Election 2012: Results,” 2012). Although the relative size of party-registered voters within a state was considered as a measure of political orientation, Pew Research Center polling (“Section 9: Trends in Party Affiliation,” 2012) has suggested that nearly a quarter of voters are independents and have no party affiliation. Therefore, the operationalization of political orientation as the percentage of Democratic voters in a presidential election was chosen because the election generally corresponds to the time period in which most of the IAT data were collected, captures the opinions of independent voters, and voting turnout tends to be greater for Presidential rather than mid-term elections.

Analytic Strategy

Adapting the approach used by Nosek and colleagues (2009) to measure implicit bias at the level of nations, state-level implicit race-based bias was estimated by separately averaging IAT scores within each state for White and Black respondents. Then, state-level racial diversity estimates (as described above) were used to predicted state-level mean implicit bias among White and Black respondents. Critically, in response to recent concerns about the use of covariates in psychological research (e.g., Simmons, Nelson, & Simonsohn, 2011), the effect of racial diversity—the primary predictor of interest—is reported without including any adjustment variables in the model. Then, changes in the predictive strength of diversity after sequentially including covariates in the model are reported. Importantly, the relationship between diversity and implicit bias is consistent whether all, some, or none of these control variables are included in the model.

Results

Diversity and Implicit Race Bias at the State Level

Mean IAT *D* scores of White and Black respondents in each state were used as the dependent measure. Recall, for both Black and White respondents, *D* scores greater than 0 indicate ingroup preference while scores less than 0 indicate outgroup preference. On the national level (i.e., collapsing across states), the mean IAT *D* score among Black respondents was 0.05 ($SD = 0.44$); among White respondents, it was 0.46 ($SD = 0.40$). This pattern is consistent with results obtained in U.S. laboratory studies in which White participants typically express stronger implicit ingroup bias than do Black participants. For the sample consisting of Black respondents, the number of observations per state ranged from 48 to 12,734, with aggregated IAT *D* scores ranging from -0.11 to 0.12 ($SD = 0.04$). For the sample consisting of White respondents, the number of observations per state ranged from 1,271 to 76,077, with aggregated IAT *D* scores ranging from 0.35 to 0.46 ($SD = 0.03$). For descriptive statistics for all study variables, including state-level population estimates used to calculate the measure of diversity used in the current study, see Table 1.

To validate the state-level IAT scores, aggregated IAT estimates were correlated with the estimated number of racially charged Google search queries in each state, an alternative method of circumventing the shortcomings of self-report measures in the assessment of group-level racial attitudes (Stephens-Davidowitz, 2012). The results indicated that this alternative measure positively correlated with state-level IAT scores for both White respondents, $r(48) = .78, p < .001$, and Black respondents, $r(48) = .56, p < .001$. Thus, states with relatively higher frequencies of racially charged Google search queries also tended to have higher group levels of implicit ingroup bias among both White and Black respondents. The strong positive correlation between

state-level IAT scores and racially charged Google search queries, while using quite different methodologies, validates the use of aggregated IAT scores in the present study.

White and Black respondents' IAT *D* scores were positively correlated at the state level, $r(48) = .64, p < .001$ (see Figure 1). Thus, states with relatively higher levels of implicit ingroup bias among White respondents also tended to have higher levels of implicit ingroup bias among Black respondents. More central to testing the relationship between diversity and implicit bias, the ratio of Black to White residents within a state (with higher numbers indicating more diversity) was positively correlated, with similar magnitudes, with implicit bias among both White respondents, $r(48) = .84, p < .001$ (see Figure 2), and Black respondents, $r(48) = .82, p < .001$ (see Figure 3). Thus, higher levels of diversity were associated with stronger implicit ingroup bias among both groups.

White Respondents. First, a simple linear regression model predicting IAT *D* scores in the White sample from the ratio of Black to White residents within states explained 71% of the variance, $F(1, 48) = 119.40, p < .001$. As the ratio of Black to White residents increased (indicating increasing diversity), state-level implicit ingroup bias among White respondents increased, $\beta = .84, p < .001$ (see Model 1 in Table 2). Though redundant with the previously presented correlation analysis, this result provides a reference for subsequent models and indicates a strong effect.

Then, the control variables described above were added to a second multiple regression model (see Model 2 in Table 2). None of the control variables were significantly related to IAT *D* scores among White respondents and the proportion of variance explained increased only slightly ($R^2 = .76$). Thus, even after adjusting for other state-level variables that may influence societal-level implicit bias, the ratio of Black to White residents within a state remained a

positive and significant predictor of White respondents' IAT *D* scores at the state level, $\beta = .87$, $p < .001$.

Black Respondents. A simple linear regression model predicting IAT *D* scores in the Black sample from the ratio of Black to White residents within states explained 67% of the variance, $F(1, 48) = 96.20$, $p < .001$. The results indicated that as the ratio of Black to White residents increased (indicating increasing diversity), state-level implicit ingroup bias among Black respondents increased, $\beta = .82$, $p < .001$ (see Model 1 in Table 3).

In a second multiple regression model containing state-level control variables, the results were similar to those observed in the White sample; none of the control variables were significantly related to aggregated IAT *D* scores among Black respondents and only provided a slight increase to the proportion of variance explained (see Model 2 in Table 3). Thus, again even after adjusting for other state-level variables that may influence societal-level implicit bias, the ratio of Black to White residents within states remained a positive and significant predictor of Black respondents' IAT *D* scores at the state level, $\beta = .83$, $p < .001$.

Ruling Out Alternative Explanations

Southern versus Northern States. One possible explanation for the results reported above is that systematic variation from a specific region of the country is driving the observed pattern. For example, previous research has suggested higher levels of prejudice in southern samples (Pettigrew, 1959) and lower levels of support for racial integration and higher levels of perceived threat for Black individuals by Whites in the U.S. South (Fossett & Kiecolt, 1989). Indeed, observing that most studies conducted in the U.S. that have revealed a negative relationship between diversity and interracial relations have surveyed intergroup attitudes in the

South, Wagner et al. (2006) hypothesized that unique features of the South may explain the failure of diversity to attenuate prejudice.

Broadly, the results from the current work do not support this hypothesis: the findings of the primary analysis indicate that greater diversity is positively associated with stronger implicit ingroup bias among both White and Black respondents across *all* U.S. states. However, the possibility that lingering historical biases particular to, or more pronounced in, Southern states could explain the pattern observed was tested. Specifically, a binary variable indicating whether a state was a member of the Confederate States of America was added to Model 2. This variable was not associated with aggregated IAT *D* scores among White respondents, $\beta = .33, p = .223$, or Black respondents, $\beta = -.35, p = .239$, and its inclusion did not diminish the predictive strength of diversity (see Model 3 in Tables 2 and 3).

Overall Racial/Ethnic Diversity. Given that the race IAT available in the public-use dataset (Xu et al., in press) measures relative implicit bias favoring Whites over Blacks (or vice versa), the diversity measure employed in the current study as the main predictor of interest was the ratio of these two racial groups. However, one potential explanation for the obtained results is that the overall racial/ethnic diversity of each state, rather than the ratio of Black to White residents specifically, may predict implicit bias. For example, states with relatively larger proportions of Black residents may also contain relatively larger proportions of other ethnic minority groups. In this case, perhaps it is the additive role of minority groups within a state that increases a sense of intergroup threat among Whites, thus triggering stronger implicit ingroup bias. This possibility was investigated by using the ratio of all racial/ethnic minority residents (e.g., Black, Asian, Hispanic, and multiracial residents) to White residents within each state as an index of overall diversity (see Model 4 in Tables 2 and 3) to predict implicit attitudes among

White and Black respondents. Again, the ACS 5-year averages of data collected between 2008 and 2012 were used to estimate the number of White residents, the number of non-White residents, and the total state population. A square-root transformation was used on these values because of the moderate positive skew of the distribution. Whether this index predicted aggregated IAT *D* scores as well as the ratio of Black to White residents was investigated.

For White respondents, the ratio of racial/ethnic minority residents to White residents was positively associated with higher state-level implicit ingroup bias, $\beta = .36, p < .01$. In this model, former Confederate states also had higher state-level implicit bias among White respondents, $\beta = 1.00, p < .001$. However, the predictive strength of overall racial/ethnic diversity was weaker than that of the ratio of Black to White residents (which was $\beta = .84, p < .001$ in Model 1; $\beta = .87, p < .001$ in Model 2). Similar to the White sample, higher levels of overall racial/ethnic diversity predicted stronger aggregated implicit ingroup bias among Black respondents, $\beta = .55, p < .01$, but overall diversity was not as strong of a predictor as was the ratio of Black to White residents (which was $\beta = .82, p < .001$ in Model 1; $\beta = .83, p < .001$ in Model 2). Moreover, for both Black and White respondents, using overall diversity as a predictor in Model 4 explained less variance in state-level IAT scores than did the ratio of Black to White residents in Model 3 (see Tables 2 and 3). Thus, while similarly predicting higher levels of implicit ingroup bias among both White and Black respondents, overall state-level diversity had less predictive utility than did the more specific ratio of Black to White residents within each state. Additionally, testing the effectiveness of overall diversity as a predictor of implicit bias at the county level (see below) yielded the same pattern as the state-level analyses. Therefore, these results help validate the use of the ratio of Black to White residents as the primary index of diversity in the present research.

Economic Threat. Another possible explanation for the results of the primary analyses is that some other type of intergroup threat that is not related to group size is driving the results (Taylor, 1998). For example, starting in roughly 2008, the U.S. economy experienced a substantial recession. In the current dataset, over 60% of both Black (61.4%) and White respondents (61.3%) completed the race IAT between 2009 and 2012. Therefore, if states more severely impacted by the recession also contain relatively higher levels of diversity, perhaps intergroup threat stemming from difficult economic conditions (Quillian, 1995) rather than group size may explain the obtained results? This possibility was tested by investigating whether state-level implicit ingroup bias was stronger for either Blacks or Whites in states with more difficult economic conditions, as measured by 2009 unemployment rates (Bureau of Labor Statistics, 2014); the results of which are included in Model 3 (containing the aforementioned state-level predictors). For Whites, results indicated that unemployment rate was not a significant predictor of state-level IAT scores, $\beta = -.08$, $p = .074$, and the sign of the coefficient was also contrary to what would be expected if harsher economic conditions resulted in greater ingroup bias. For Blacks, state unemployment rate was also not a statistically significant predictor of state-level IAT scores, $\beta = .05$, $p = .438$. In sum, these analyses indicate that more difficult state-level economic conditions did *not* result in stronger implicit biases.

Alternative Analytic Methods. Another possible explanation for the results of the primary analyses is that the pattern observed are a byproduct of the analytic method employed, and that different analytic method would yield different results. For example, although differentially weighting state estimates to account for variation in sample size was considered, this plan was abandoned given weighting schemes based on sample size would be confounded with the primary predictor: the ratio of Black to White residents. For example, as the ratio of

Black to White residents decreases, with equal sampling there should be fewer IAT scores for Black respondents. Thus, there should be a positive correlation between the number of IAT scores from Black respondents in a state and the number of Black residents in a state. This was the case, $r(48) = .54, p < .001$. Therefore, weighting state estimates according to the number of observations from Black respondents would necessarily be assigning less weight to states as the index of diversity decreases.

Crucially, however, results were essentially the same regardless of whether such weighting was employed. To illustrate, the red lines in Figure 4 displays the ordinary least squares (OLS) or non-weighted regression lines. The black line is a weighted least squares (WLS) regression line where each data point is weighted by the number of IAT observations from the state. For the White sample, there is very little difference between the OLS and WLS estimates, although the difference is more pronounced among states lower in diversity. For the Black sample (Figure 5), the WLS line is somewhat flatter than the OLS line and tends to somewhat overestimate implicit bias in states with low diversity.

Although it was decided not to use a weighting scheme based on sample size, the pattern of findings is substantively the same regardless of whether estimates are based on WLS or OLS. Using WLS, for the White sample, the ratio of Black to White residents predicted aggregated IAT scores in Model 1, $\beta = .93, p < .001$, and remained significant when controlling for covariates in Model 2, $\beta = .91, p < .001$ (other predictors *ns*). In Model 3, the ratio of Black to White residents remained significant, $\beta = .84, p < .001$ (other predictors *ns*). In Model 4, overall state-level diversity was a significant predictor, $\beta = .46, p = .023$, but was weaker than the ratio of Black to White residents in Model 3. Membership in the Confederacy was a significant

predictor in Model 4, $\beta = .72$, $p = .010$, with implicit bias among White residents from former Confederate (vs. non-Confederate) states being stronger.

For the Black sample, the ratio of Black to White residents predicted aggregated IAT scores in Model 1, $\beta = .48$, $p < .001$. This relationship remained significant when controlling for covariates in Model 2, $\beta = .40$, $p = .001$, while the only other significant predictor was state education level, $\beta = -.27$, $p = .023$. In Model 3, the ratio of Black to White residents was the strongest predictor, $\beta = .60$, $p < .001$, with state education level, $\beta = -.29$, $p = .012$, and membership in the Confederacy, $\beta = -.44$, $p = .019$, also being significant predictors. In Model 4, overall state-level diversity was a significant predictor, $\beta = .54$, $p < .001$, but was weaker than the ratio of Black to White residents in Model 3.

Level of Analysis. Another potential limitation to the results reported above is the focus on relatively large geographical regions—U.S. states—which could appear more or less diverse simply because of their inclusion of urban centers with high concentrations of Black residents and rural areas with high concentrations of White residents. Indeed, prior work examining explicit intergroup attitudes has focused on smaller geographical regions with an average population of approximately 180,000 residents (Wagner et al., 2006). In an exploratory analysis, the level of analysis was reduced from U.S. states to U.S. counties (which on average contain about 100,000 residents and represent the smallest geographical unit available in the public-use IAT dataset; Xu et al., in press) and utilized the same model-testing procedure as described above. Unfortunately, the public-use dataset contains IAT scores for Black respondents from only about two-thirds of the roughly 3,100 U.S. counties or county equivalents. Further, only counties with at least 10 observations from White and Black respondents were considered, which reduced the final sample to 929 counties that were mostly located in states that were high in

racial diversity. Diversity varied greatly across counties with the number of Black residents per 100 White residents in a county ranging from 0.40 to 432.8 ($Mdn = 14.37$), including 57 counties where Black residents outnumber White residents.

Although these counties do not represent a sample that can be generalized to the rest of the U.S., the county-level findings mirrored our findings at the state level: For the White sample, the ratio of Black to White residents (with higher numbers indicating greater diversity) in a county positively predicted county-level IAT D scores, $\beta = .36, p < .001$, indicating that greater diversity in a county was associated with stronger implicit ingroup bias among White respondents. Including Model 2 control variables at the county level did not diminish this relationship, $\beta = .40, p < .001$. For the Black sample, the ratio of Black to White residents in a county positively predicted county-level IAT D scores, $\beta = .41, p < .001$, indicating that greater within-county diversity was associated with stronger implicit ingroup bias among Black respondents. After including the Model 2 control variables at the county level, the ratio of Black to White residents remained the strongest predictor of IAT scores among Black respondents, $\beta = .36, p < .001$. Finally, as in the state-level analysis, IAT scores among White and Black respondents were positively correlated at the county level, $r(927) = .18, p < .001$, indicating that counties with higher levels of implicit ingroup bias among White respondents also tended to have higher implicit ingroup bias among Black respondents.

Segregation. Lastly, one could argue that the ratio of Black to White individuals within a geographical unit, whether at the state or county level, is a suboptimal index of diversity given that regions that appear racially diverse on indices based on proportions of residents from different racial groups may in reality be divided into homogeneous subregions that are anything but racially diverse (Holloway, Wright, & Ellis, 2012). Therefore, it was investigated whether an

index of integration between White and Black residents in a spatial unit produced a different pattern relative to using the proportions of each group as the primary predictor. Among the many segregation indices developed in sociology and demography, one that is especially relevant to the present work is the *interaction index*, which can be interpreted as the chances of a Black resident encountering a White resident in a given spatial unit (or vice versa; Massey & Denton, 1988). Interaction index scores range from 0 to 1, with scores near 0 indicating that Black and White residents in an area are unlikely to meet (i.e., high segregation and low diversity) while scores near 1 indicate that encounters between individuals from different racial groups are likely (i.e., low segregation and high diversity). Given that segregation is typically measured at the level of relatively small geographical units (e.g., metropolitan areas or the U.S. Census tract level; e.g., Iceland & Scopilliti, 2008) and because the state-level and county-level analyses were congruent, predictive utility of the interaction index between White and Black residents at the county level using data from the 2010 U.S. Census was tested (<http://www.census.gov/2010census/>).

Results from a simple linear regression model indicated that higher county-level interaction index scores for White residents (indicating a greater probability that a White resident will meet a Black resident within a specific county) predicted stronger county-level ingroup bias among White respondents, $\beta = .38, p < .001$. This relationship was not diminished when control variables were included in Model 2. By contrast, higher county-level interaction index scores for Black residents (indicating a greater probability that a Black resident will meet a White resident within a specific county) predicted *weaker* county-level ingroup bias among Black respondents, $\beta = -.23, p < .001$. This relationship was not diminished by the inclusion of county-level control variables in Model 2, $\beta = -.25, p < .001$.

Discussion

The current study was the first to investigate the relationship between racial diversity and implicit race bias across U.S. states and counties and found that greater diversity predicted stronger implicit ingroup-favoring bias among both White and Black respondents. Thus, states like Mississippi and South Carolina that have a higher proportion of Black residents tended to have stronger ingroup bias among both White and Black respondents, relative to states like Montana and Vermont that have few Black residents. Furthermore, variables such as state-level income, education, and economic inequality had no significant relationship with state-level implicit bias, nor did they diminish the relationship between state-level racial diversity and implicit bias. Even the overall racial/ethnic diversity of a state, while a predictor of stronger ingroup bias among both Black and White respondents, did not reach the explanatory power provided by simply accounting for the ratio of Black to White residents in each state. In addition, several alternative explanations that may influence group level implicit attitudes, such as historical effects of state membership in the Confederacy, were ruled out.

The finding that greater diversity—measured as either group proportions or estimated exposure to Black residents—predicted stronger implicit ingroup bias among White respondents may appear to contradict the contact hypothesis, according to which positive contact between members of different racial groups reduces race-based prejudice, especially when it occurs under optimal conditions (Allport, 1954; Pettigrew & Tropp, 2006). However, the findings for Black respondents may actually be interpreted as supporting the contact hypothesis: Greater diversity as measured by group proportions, which represents *less* opportunity for contact with White residents, predicted stronger implicit ingroup bias among Black respondents. Furthermore, higher county-level scores on the interaction index, indicating a greater chance for a Black

resident to encounter a White resident, predicted lower implicit ingroup bias among Black respondents. However, aggregated IAT scores among Black respondents ranged from slight pro-Black (ingroup) bias to slight pro-White (outgroup) bias. Whether reduced ingroup favoritism is in fact beneficial, if ultimately resulting in outgroup-favoring biases, is questionable.

Additionally, while diversity may seem an appropriate predictor in the context of the contact hypothesis, examining the effects of diversity and intergroup contact on intergroup relations represent distinct, albeit related, research questions. For example, although greater diversity provides more opportunities for contact (Pettigrew, 1997), it does not guarantee that people will choose to actually engage in positive intergroup contact (Shelton & Richeson, 2005). Conversely, residing in a racially diverse state, county, or neighborhood is not necessary for individuals to engage in positive interracial contact and cross-group friendship. Moreover, greater diversity may also increase *negative* contact experiences, which may cancel out and even counteract the beneficial effects of positive contact (Barlow et al., 2012). Thus, greater diversity alone is not sufficient to guarantee contact, let alone positive contact, between majority- and minority-group members.

Indeed, the diversity measures used do not index the optimal conditions of positive intergroup contact that Allport (1954) suggested would maximally reduce prejudice. However, recent work has argued that researchers should abandon the search for utopian conditions that maximally reduce prejudice and instead explore how mundane cross-group interactions can improve intergroup relations (Dixon, Durrheim, & Tredoux, 2005). The diversity measures employed in this study, especially the interaction index (which represents the chances that minority- and majority-group members will meet), may provide useful indices of intergroup contact in everyday life—occurring while riding buses, buying groceries, or going bowling

(Putnam, 2000). Taken together with the large body of evidence demonstrating that positive intergroup contact, exemplified by cross-group friendship, effectively reduces prejudice (Christ et al., in press; Pettigrew & Tropp, 2006) and a recent national poll indicating that many Americans (e.g., 40% of Whites) have no cross-race friends (Dunsmuir, 2013), the results of this study underscores the need to understand when diversity can encourage people to interact positively across group boundaries.

At the same time, understanding how diversity is associated with intergroup bias, without considering the additional role of intergroup contact, is also critical—especially given that political discourse related to diversity (e.g., immigration, affirmative action) is often contentious and that ever-increasing diversity is inevitable. As previously mentioned, projected U.S. population estimates predict that racial/ethnic minority groups will comprise a majority of the U.S. population by 2042 (U.S. Census Bureau, 2008). As such, the findings of this study fit particularly well within an emerging area of research examining the consequences of such diversity-related population shifts. For example, when presented with information about the shifting racial composition of the U.S., White participants expressed more race-based bias on both implicit and explicit measures (Craig & Richeson, in press). Taken together with the results from the current study, this finding suggests the unfortunate possibility that increases in actual or perceived diversity, contrary to what many would hope or believe, may serve to trigger intergroup hostility.

The results of the current study are consistent with a body of research suggesting that diversity is associated with more negative explicit intergroup attitudes because it increases perceived threat and competition (e.g., Taylor, 1998). However, explicit attitudes best predict deliberative (e.g., verbal) behavior whereas implicit bias best predicts spontaneous behavior

(e.g., non-verbal friendliness; Dovidio et al., 2002). Given that diversity impacts the likelihood of ordinary, everyday interactions between White and Black individuals, this investigation into the relationship between diversity and group-level implicit bias makes an important contribution by characterizing how bias may be elicited in such everyday interactions. For example, while it may be possible to avoid conveying explicit bias (e.g., racially charged language) while sitting next to an outgroup member on a bus, it is substantially more difficult to avoid showing non-verbal cues (e.g. anxiety or discomfort) of subtler bias that may nonetheless prove harmful in future interracial interactions.

Although the results of this study contribute to a more comprehensive understanding of the relationship between diversity and intergroup bias, there are several important limitations that must be considered. First, the data used in this study are correlational and causal relationships cannot be inferred. Therefore, it is technically possible that implicit attitudes *cause* specific levels of diversity within states; however, it appears to be more likely that diversity impacts implicit bias. It is also technically possible that a third variable, one that was not investigated in the current study, influences both. Second, participants self-selected into our sample by completing the race IAT via the Project Implicit website; reasons for participation may include a particular interest in the topic or fulfillment of an educational requirement (Nosek, Banaji, & Greenwald, 2002). The sample used in this study may not be representative of the U.S. more broadly in factors such as education level. One additional limitation or at least qualification of these results is that the relative variation in state level IAT D scores for White respondents (Max $D = .456$, Min $D = .346$) and Black respondents (Max. $D = .115$, Min. $D = -.125$) were modest. Additionally, given that there are no U.S. states in which Blacks represent the majority, we cannot definitively extend our findings to geographic locations in which a lower-status racial

group is a numerical majority, as in much of Africa or the Caribbean. Indeed, such locations represent an interesting test case for future work. Lastly, it is paramount to recall that IAT scores in the present work were aggregated by state, causing individual-level variation to be averaged out of the analyses. Therefore, it is crucial to not fall victim to the ecological fallacy by interpreting our findings as explaining variance among *individuals* and instead to recognize that our findings relate to the proportion of variance explained across *states*. A productive direction for future research would be to explain individual-level variation in implicit bias by examining how individual differences interact with societal-level variables (such as diversity).

In conclusion, aligning with findings from political science (Putnam, 2007) and sociology (Quillen, 1995), we found that greater diversity is associated with stronger implicit ingroup bias among both White and Black Americans. Although we attempted to isolate the relationship between racial diversity and implicit race bias (e.g., by ruling out alternative predictions and including control variables), it remains unclear *why* greater diversity is associated with stronger implicit race bias. Our results thus underscore the importance of further investigating why diversity—a feature that we strive to increase in our schools, workplaces, and communities—may not ultimately have the beneficial impact we hope or expect it to have.

References

- Allport, G. W. (1954). *The nature of prejudice*. Reading, MA: Addison-Wesley.
- Archibold, R. C. (2010, April 23). Arizona enacts stringent law on immigration. *New York Times*. Retrieved from http://www.nytimes.com/2010/04/24/us/politics/24immig.html?_r=0.
- Barlow, F. K., Paolini, S., Pedersen, A., Hornsey, M. J., Radke, H. R., Harwood, J., Rubin, M., & Sibley, C. G. (2012). The contact caveat: Negative contact predicts increased prejudice more than positive contact predicts reduced prejudice. *Personality and Social Psychology Bulletin*, 38, 1629-1643. doi:10.1177/0146167212457953.
- Bobo, L. (1988). Group conflict, prejudice, and the paradox of contemporary racial attitudes. In P. A. Katz & D. A. Taylor (Eds.), *Eliminating racism: Profiles in controversy* (pp. 85-114). New York, NY: Plenum Press.
- Chang, M. J. (1999). Does racial diversity matter?: The educational impact of a racially diverse undergraduate population. *Journal of College Student Development*, 40, 377-395.
- Christ, O., Schmid, K., Lolliot, S., Swart, H., Stolle, D., Tausch, N., ... & Hewstone, M. (in press). Contextual effect of positive intergroup contact on outgroup prejudice. *Proceedings of the National Academy of Sciences of the United States of America*.
- Craig, M. A., & Richeson, J. A. (in press). More diverse yet less tolerant? How the increasingly-diverse racial landscape affects White Americans' racial attitudes. *Personality and Social Psychology Bulletin*.
- Dixon, J., Durrheim, K., & Tredoux, C. (2005). Beyond the optimal contact strategy: A reality check for the contact hypothesis. *American Psychologist*, 60, 697-711. doi:10.1037/0003-066X.60.7.697

- Dovidio, J. F., Kawakami, K., & Gaertner, S. L. (2002). Implicit and explicit prejudice and interracial interaction. *Journal of Personality and Social Psychology, 82*, 62-68. doi:10.1037/0022-3514.82.1.62
- Dunsmuir, L. (2013, August 8). *Many Americans have no friends of another race: poll*. Retrieved from <http://www.reuters.com/article/2013/08/08/us-usa-poll-race-idUSBRE97704320130808>
- Election 2012: Results. (2012). Retrieved March 16, 2014, from <http://www.cnn.com/election/2012/results/main>.
- Fazio, R. H., & Olson, M. A. (2003). Implicit measures in social cognition research: Their meaning and use. *Annual Review of Psychology, 54*, 297-327. doi:10.1146/annurev.psych.54.101601.145225
- Fossett, M. A., & Kiecolt, K. J. (1989). The relative size of minority populations and white racial attitudes. *Social Science Quarterly, 70*, 820-835.
- Green, D. P., Wong, J., & Strolovitch, D. (1996). *The Effects of Demographic Change on Hate Crime* (Yale University Institution for the Social and Policy Studies Working Paper No. 96-06). New Haven, CT: Institution for the Social and Policy Studies, Yale University.
- Greenwald, A. G., McGhee, D. E., & Schwartz, J. L. (1998). Measuring individual differences in implicit cognition: The Implicit Association Test. *Journal of Personality and Social Psychology, 74*, 1464-1480. doi:10.1037/0022-3514.74.6.1464
- Greenwald, A. G., Nosek, B. A., & Banaji, M. R. (2003). Understanding and using the implicit association test: I. An improved scoring algorithm. *Journal of Personality and Social Psychology, 85*, 197-216. doi:10.1037/0022-3514.85.2.197

- Greenwald, A. G., Poehlman, T. A., Uhlmann, E. L., & Banaji, M. R. (2009). Understanding and using the Implicit Association Test: III. Meta-analysis of predictive validity. *Journal of Personality and Social Psychology, 97*, 17-41. doi:10.1037/a0015575
- Holloway, S. R., Wright, R., & Ellis, M. (2012). The racially fragmented city? Neighborhood racial segregation and diversity jointly considered. *The Professional Geographer, 64*, 63-82. doi: 10.1080/00330124.2011.585080
- Iceland, J., & Scopilliti, M. (2008). Immigrant residential segregation in US metropolitan areas, 1990–2000. *Demography, 45*, 79-94. doi: 10.1353/dem.2008.0009
- Liska, A. E., Chamlin, M. B., & Reed, M. D. (1985). Testing the economic production and conflict models of crime control. *Social Forces, 64*, 119-138.
- Marini, M., Sriram, N., Schnabel, K., Maliszewski, N., Devos, T., Ekehammar, B., ... & Nosek, B. A. (2013). Overweight people have low levels of implicit weight bias, but overweight nations have high levels of implicit weight bias. *PloS ONE, 8*(12), e83543. doi: 10.1371/journal.pone.0083543
- Massey, D. S., & Denton, N. A. (1988). The dimensions of residential segregation. *Social Forces, 67*, 281-315. doi: 10.1093/sf/67.2.281
- McConnell, A. R., & Leibold, J. M. (2001). Relations among the Implicit Association Test, discriminatory behavior, and explicit measures of racial attitudes. *Journal of Experimental Social Psychology, 37*, 435-442.
- Nosek, B. A., Banaji, M., & Greenwald, A. G. (2002). Harvesting implicit group attitudes and beliefs from a demonstration web site. *Group Dynamics: Theory, Research, and Practice, 6*, 101-115. doi:10.1037/1089-2699.6.1.101

- Nosek, B. A., Smyth, F. L., Hansen, J. J., Devos, T., Lindner, N. M., Ranganath, K. A., ... & Banaji, M. R. (2007). Pervasiveness and correlates of implicit attitudes and stereotypes. *European Review of Social Psychology, 18*, 36-88. doi:10.1080/10463280701489053
- Nosek, B. A., Smyth, F. L., Sriram, N., Lindner, N. M., Devos, T., Ayala, A., ... & Greenwald, A. G. (2009). National differences in gender–science stereotypes predict national sex differences in science and math achievement. *Proceedings of the National Academy of Sciences of the United States of America, 106*(26), 10593-10597. doi:10.1073/pnas.0809921106
- Oswald, F. L., Mitchell, G., Blanton, H., Jaccard, J., & Tetlock, P. E. (2013). Predicting ethnic and racial discrimination: A meta-analysis of IAT criterion studies. *Journal of Personality and Social Psychology, 105*, 171-192. doi:10.1037/a0032734
- Page, S. E. (2007). *The difference: How the power of diversity creates better groups, firms, schools, and societies*. Princeton, NJ: Princeton University Press.
- Pettigrew, T. F. (1959). Regional differences in anti-Negro prejudice. *The Journal of Abnormal and Social Psychology, 59*, 28.
- Pettigrew, T. F. (1997). Generalized intergroup contact effects on prejudice. *Personality and Social Psychology Bulletin, 23*, 173-185. doi:10.1177/0146167297232006
- Pettigrew, T. F., & Tropp, L. R. (2006). A meta-analytic test of intergroup contact theory. *Journal of Personality and Social Psychology, 90*, 751-783. doi:10.1037/0022-3514.90.5.751
- Prestwich, A., Kenworthy, J. B., Wilson, M., & Kwan-Tat, N. (2008). Differential relations between two types of contact and implicit and explicit racial attitudes. *British Journal of Social Psychology, 47*, 575-588. doi:10.1348/014466607X267470

- Putnam, R. D. (2000). *Bowling alone: The collapse and revival of American community*. New York, NY: Touchstone Books/Simon & Schuster.
- Putnam, R. D. (2007). E pluribus unum: Diversity and community in the twenty-first century: The 2006 Johan Skytte Prize Lecture. *Scandinavian Political Studies*, *30*, 137-174. doi: 10.1111/j.1467-9477.2007.00176.x
- Quillian, L. (1995). Prejudice as a response to perceived group threat: Population composition and anti-immigrant and racial prejudice in Europe. *American Sociological Review*, *60*, 586-611. doi:10.2307/2096296
- Section 9: Trends in Party Affiliation. (2012). Retrieved March 16, 2014, from <http://www.people-press.org/2012/06/04/section-9-trends-in-party-affiliation/>.
- Shelton, J. N. (2000). A reconceptualization of how we study issues of racial prejudice. *Personality and Social Psychology Review*, *4*, 374-390. doi:10.1207/S15327957PSPR0404_6
- Shelton, J. N., & Richeson, J. A. (2005). Intergroup contact and pluralistic ignorance. *Journal of Personality and Social Psychology*, *88*, 91-107. doi:10.1037/0022-3514.88.1.91
- Sigelman, L., & Welch, S. (1993). The contact hypothesis revisited: Black-white interaction and positive racial attitudes. *Social Forces*, *71*, 781-795. doi:10.2307/2579895
- Simmons, J. P., Nelson, L. D., & Simonsohn, U. (2011). False-positive psychology: Undisclosed flexibility in data collection and analysis allows presenting anything as significant. *Psychological Science*, *22*, 1359-1366. doi:10.1177/0956797611417632
- Sriram, N., & Greenwald, A. G. (2009). The Brief Implicit Association Test. *Experimental Psychology*, *56*, 283.

- Stephens-Davidowitz, S. I. (2013). The cost of racial animus on a Black presidential candidate: Using Google search data to find what surveys miss. Retrieved from <http://ssrn.com/abstract=2238851>
- Taylor, M. C. (1998). How white attitudes vary with the racial composition of local populations: Numbers count. *American Sociological Review*, *63*, 512-535. doi:10.2307/2657265
- Tropp, L. R., & Pettigrew, T. F. (2005). Relationships between intergroup contact and prejudice among minority and majority status groups. *Psychological Science*, *16*, 951-957. doi:10.1111/j.1467-9280.2005.01643.x
- U.S. Census Bureau. (2008, August 14). An older and more diverse nation by midcentury. Retrieved from <http://www.census.gov/newsroom/releases/archives/population/cb08-123.html>.
- Wagner, U., Christ, O., Pettigrew, T. F., Stellmacher, J., & Wolf, C. (2006). Prejudice and minority proportion: Contact instead of threat effects. *Social Psychology Quarterly*, *69*, 380-390. doi:10.1177/019027250606900406
- Wagner, U., Van Dick, R., Pettigrew, T. F., & Christ, O. (2003). Ethnic prejudice in East and West Germany: The explanatory power of intergroup contact. *Group Processes & Intergroup Relations*, *6*, 22-36. doi: 10.1177/1368430203006001010
- Wilcox, J., & Roof, W. C. (1978). Percent Black and Black-White Status Inequality: Southern versus Nonsouthern Patterns. *Social Science Quarterly*, *59*, 421-34.
- Xu, K., Nosek, B. A., Greenwald, A. G. (in press). Psychology data from the Race Implicit Association Test on the Project Implicit demonstration website. *Journal of Open Psychology Data*