

Zingales + Hansen JAP

p3 SS8:

NS 99-101

ICC white opp: -0.02 (w)

ICC black opp: 0.29 (R)

IE w/ ATB = $.12$

IE w/ MBS = $.12$

IE w/ EMS = $.07$

ECC: white applicant

ATB: -0.15 (w)

MBS: -0.22 (w)


EMS: 0.03 (R)

ECC: black Applicant

ATB: 0.19 (R)

MBS: $.18$ (R)

EMS: $.02$ (w)



Behavior:

IM: race attitude (AT), pg 556

EM: MRS, pg 556
ATB, pg 556
EMS, pg 556

BEH: Hiring recommendations, pg 556

- (10) (A) race, pg 553
- (11) ~~single (1)~~ average (2), pg ~~556~~ 557
- (12) attitude (1), pg 556
- (13) attitude (MRS/ATB), belief (EMS), pg 556
- (14) not (0), pg 556
- (15) judge (4), pg 556
- (16) "3" - combined means & errors, pg 556
- (17) words, pg 556
- (18) 1, pg 556
- (19) 3, pg 557-558
- (20) 1
- (21) 2
- (22) after (2)
- (23) before (1)
- (24) same (0)
- (25) different (1)

(26)	7	pg 556
(27)	7	pg 556
(28)	9	pg 556
(29)	3.5 3.5	" "
(30)	2.5	" "
(31)	opp: 2.5	pg 556
(32)	racism (1)	
(33)	dist (2)	

☒ **DONE** nm

RESEARCH REPORTS

Employment Discrimination: The Role of Implicit Attitudes, Motivation, and a Climate for Racial Bias

Jonathan C. Ziegert and Paul J. Hanges
University of Maryland

This study is an attempt to replicate and extend research on employment discrimination by A. P. Brief and colleagues (A. P. Brief, J. Dietz, R. R. Cohen, S. D. Pugh, & J. B. Vaslow, 2000). More specifically, the authors attempted (a) to constructively replicate the prior finding that an explicit measure of modern racism would interact with a corporate climate for racial bias to predict discrimination in a hiring context and (b) to extend this finding through the measurement of implicit racist attitudes and motivation to control prejudice. Although the authors were unable to replicate the earlier interaction, they did illustrate that implicit racist attitudes interacted with a climate for racial bias to predict discrimination. Further, results partially illustrate that motivation to control prejudice moderates the relationship between explicit and implicit attitudes. Taken together, the findings illustrate the differences between implicit and explicit racial attitudes in predicting discriminatory behavior.

Keywords: discrimination, implicit attitudes, racism, prejudice, IAT

In recent years, racist attitudes have evolved from being blatant and hostile in nature to being more subtle and ambivalent (Brief, Dietz, Cohen, Pugh, & Vaslow, 2000). Indeed, whereas traditional self-report measures have indicated that there has been a decline in racist attitudes, discrimination continues in employment decisions (Maass, Castelli, & Arcuri, 2000). This discrepancy and shift in the nature of racist attitudes has prompted social scientists to design new measures that are consistent with the more modern expression of racism (McConahay, 1986; McConahay, Hardee, & Batts, 1981). These scales attempt to get around self-presentation bias and identify individuals with negative racial attitudes by using questions in which the prejudiced response could be attributed to "racially neutral ideology" (Fazio, Jackson, Dunton, & Williams, 1995).

Recently, some researchers have moved away from these self-report measures to physiological, response latency, or priming measures to assess an individual's level of racist attitudes (Fazio & Olson, 2003). These more implicit measures are believed to be less susceptible to self-presentation biases and thus are more successful at assessing prejudices. Although research has documented that these implicit measures correlate with other attitudes and predict microlevel behavior, there is currently little evidence indicating

that such implicit attitudes are useful for predicting more macrolevel behavior, such as discriminatory hiring decisions.

The present study was designed to be a constructive replication of the Brief et al. (2000) study, which showed that modern racists act on their prejudices in particular social climates. Lykken (1968) defined a constructive replication as a study that not only tests the validity of prior findings but also tests new hypotheses. The present study extends the Brief et al. (2000) study in several ways. First, in addition to assessing modern racism, we included a measure of more traditional racist beliefs (i.e., old-fashioned or hostile racism) to ascertain whether these more modern racism assessments were really needed to identify racist individuals. Second, in addition to using these two self-report (i.e., explicit) measures of racism, we included an implicit racial attitudes measure. Third, we included a measure of motivation to control prejudice to test whether it is indeed a self-presentation bias that accounts for different results obtained by explicit and implicit measures. Finally, we used a more sensitive measure of racial discrimination to test our hypotheses by comparing differences in the ratings of Blacks and Whites with a more sophisticated statistical technique: hierarchical linear modeling (HLM; Bryk & Raudenbush, 1992; Kreft & De Leeuw, 1998). Taken together, these additions allowed us to test new hypotheses that provide a greater understanding of employment discrimination. In summary, the present study merged the literatures on racist attitudes and their measurement, self-presentation bias, and organizational social norms and climates to understand the differential utility of implicit and explicit measures in predicting individuals' behavior in a selection context.

Employment Discrimination and Racist Attitudes

Many studies have looked at the relationship between race considerations and employment discrimination (Roberson &

Jonathan C. Ziegert and Paul J. Hanges, Department of Psychology, University of Maryland.

We thank Brent Smith and Susan Taylor for their comments on drafts of this article. An earlier version of this article was presented at the 17th Annual Conference of the Society for Industrial and Organizational Psychology, April 2002, Toronto, Ontario, Canada.

Correspondence concerning this article should be addressed to Jonathan C. Ziegert, Department of Psychology, University of Maryland, College Park, MD 20742. E-mail: jziegert@psyc.umd.edu

Method

Participants

Participants were 103 undergraduates from a large mid-Atlantic public university who received course credit. Ages ranged from 18 to 30 with a mean of 18.8 ($SD = 1.5$) years. The majority of the participants were women (61.2%; men = 38.8%), and because the purpose of the experiment was to assess bias against Blacks, all participants were non-Black with Whites as the majority (White = 81%; Asian = 7%; Latino/Hispanic = 6%; Arab = 2%; and Other = 4%). There were no significant differences among these groups for any of the study variables.

Measures

Explicit racial attitudes. Racism was measured explicitly with two self-report scales: the Attitudes Toward Blacks Scale (ATB; Brigham, 1993) and the MRS. The ATB is a 20-item measure of old-fashioned racist attitudes on a 7-point Likert-type scale. For example, participants are asked if they would dislike living near Black people and if they feel that Black and White people are equal. The reliability of participants' scores on this scale was acceptable with a coefficient alpha of .89. In contrast to the ATB, the MRS assesses modern racism. The MRS contains seven items measured on a 7-point Likert-type scale. An example item is "Discrimination against Blacks is no longer a problem in the United States." The modern racism items were embedded with 14 other items assessing attitudes toward other issues (e.g., homosexuality and abortion) to limit potential reactivity effects. The reliability of the scores of this scale was acceptable with a coefficient alpha of .81.

Implicit racial attitudes. Implicit racism was assessed with the IAT. The IAT measures racist attitudes by recording the speed and accuracy with which participants can categorize words. Participants first sorted words into a descriptive category relevant to race (i.e., names that Greenwald et al. [1998] had determined to be perceived as prototypical of Whites or Blacks) and then sorted the next group of words into an evaluative category (i.e., pleasant vs. unpleasant words). In particular, names appeared on the computer screen one at a time, and participants had to categorize these names as being either "Black" or "White" using the left or right key on the computer keyboard. In the next trial, words appeared on the computer screen one at a time and participants had to categorize them as being either "pleasant" or "unpleasant." These categorization tasks were then combined and participants had to categorize 40 words as either belonging to the "White or pleasant" category or the "Black or unpleasant" category. The speed and accuracy of the categorization process were recorded. After this combined task trial, the race categories were switched to allow for practice and were then combined with the evaluative category. That is, participants determined whether a word belonged to the "White or unpleasant" category versus the "Black or pleasant" category for another 40 words.¹ An example of this sequence of trials for the IAT can be seen in Figure 1.

We focused on both the speed (response latency) as well as accuracy (percentage of errors) in the word classifications as our assessment of implicit attitudes. We examined the degree to which participants were slower in their responding and made more errors (initially placed the word in the wrong category) when the "White and unpleasant" categories were paired than when "White and pleasant" were paired.² In particular, latencies and percentage of errors for the "White and pleasant" trial were subtracted from the "White and unpleasant" trial, which results in positive latencies and percentage of errors being indicative of implicit racist attitudes (individuals can associate "White and pleasant" faster and make less errors than they can with "Black and pleasant"). We followed the procedure of Greenwald et al. (1998) and recoded responses below 300 ms and above 3,000 ms, and we eliminated cases with an average error rate of over 20% (indicative of random responding), which resulted in only a small number of cases being deleted (less than 5% of the sample; after deletion of these cases, the final sample was 103 participants).

Average response latencies ranged from -368.03 ms to 739.15 ms with a mean of 233.52 ms ($SD = 168.36$). This positive value indicates that the average implicit racial attitude for our sample was somewhat negatively biased against Blacks (e.g., faster classification for the "White and pleasant" than for the "White and unpleasant"). The error rate percentage ranged from -10% to 15% with an average of 1.5% ($SD = 4.63$), which also indicates that the average implicit attitude was somewhat negatively biased against Blacks.

We subjected these two measures to an exploratory principal components factor analysis, and by examining the eigenvalues we found evidence that only one factor emerged. Both scales had factor loadings of .79 with this single factor. Using these factor loadings and the reliability formula provided by Bollen (1989), we estimated the reliability of each scale to be .63. As these two measures are on different metrics, we created a composite implicit racist attitude score by using z scores to standardize the response latency and the error rate percentage and then by adding these two z scores together. Using the linear composite reliability formula provided by Nunnally and Bernstein (1994), we estimated the reliability of the composite implicit measure to be .70.

Motivation to control prejudice. Motivation to control prejudice was assessed through the use of Plant and Devine's (1998) External subscale on the Motivation to Respond Without Prejudice measure. This External subscale measures motivation to hide racial prejudices in order to conform to societal norms and appear nonprejudiced to others. The five items for this measure were responded to on a 9-point Likert-type scale. An example item is "I try to hide any negative thoughts about Black people in order to avoid negative reactions from others." The reliability of participants' scores on this scale was acceptable with a coefficient alpha of .77. Validity evidence and freedom from social desirability concerns for this scale have been demonstrated on the basis of small correlations with several social desirability scales (Plant & Devine, 1998).

Experimental task and racial discrimination measure. Participants completed the in-basket exercise used by Brief et al. (2000). It contained many tasks typically encountered by managers (e.g., determining compensation for a newly hired employee). Embedded in this in-basket was a "hiring recommendation" task. Participants were provided with the dossiers of eight job applicants and were instructed to evaluate them. The dossiers provided information about each applicant's education, prior work experience, race, gender, and hobbies. These dossiers were written such that six of the eight applicants had outstanding qualifications. Participants rated each applicant to the degree to which they were an exceptional referral on a 5-point Likert-type scale ranging from 1 (should not have been referred) to 5 (excellent referral). Prior work has shown that there are no differences among these six candidates when race information is removed (Brief, Buttram, Elliott, Reizenstein, & McCline, 1995). The race of these applicants was randomly assigned with half (three) of the qualified applicants as Black and the other half (three) as White. In addition to race, the sex of the applicants was also randomly assigned so that one each of the qualified Black and White candidates was a woman (i.e., there were two Black men, two White men, one Black woman, and one White woman).

¹ This order was counterbalanced such that half of the participants first completed the pairing of "White or pleasant" versus "Black or unpleasant," whereas the other half of the participants first completed "White or unpleasant" versus "Black or pleasant." Further, the order of the specific words that appeared was randomized. Finally, consistent with Greenwald et al. (1998), there was one practice trial that took place before each of the combined categorization tasks in order to familiarize participants with the task. Data were not recorded for these practice trials.

² Alternatively, because of the nature of the measure, this could have been equally phrased regarding the degree to which participants were slower in their responding and made more errors when "Black and pleasant" categories are paired than when "Black and unpleasant" are paired.

Trial 1: Initial Target Categorization		
WHITE	Sample Stimuli	BLACK
✓	HEATHER	
✓	ANDREW	
	ALONZO	✓
✓	EMILY	
	LATONYA	✓
	TYREE	✓
✓	HARRY	
	TAWANDA	✓

Trial 2: Initial Attribute Categorization		
pleasant	Sample Stimuli	unpleasant
	sickness	✓
✓	freedom	
	death	✓
✓	miracle	
✓	happy	
✓	paradise	
	poison	✓
	disaster	✓

Trial 3: Initial Combined Categorization		
WHITE or pleasant	Sample Stimuli	BLACK or unpleasant
	ALONZO	✓
✓	paradise	
	disaster	✓
✓	HEATHER	
✓	miracle	
	LATONYA	✓
✓	ANDREW	
	sickness	✓

Trial 4: Reversed Target Categorization		
BLACK	Sample Stimuli	WHITE
✓	LATONYA	
	ANDREW	✓
	HARRY	✓
✓	ALONZO	
✓	TYREE	
	EMILY	✓
	HEATHER	✓
✓	TAWANDA	

Trial 5: Reversed Combined Categorization		
BLACK or pleasant	Sample Stimuli	WHITE or unpleasant
	sickness	✓
✓	miracle	
✓	ALONZO	
✓	LATONYA	
	disaster	✓
	ANDREW	✓
✓	paradise	
	HEATHER	✓

Figure 1. Sample illustration of the Implicit Association Test (IAT). Participants completed a series of five trials in which the target concepts and attributes were introduced in the first two trials. The targets and attributes were combined in the third trial and reversed in the fifth trial after initially reversing the target categories in the fourth trial. The correct categorization of the sample stimuli that appears one at a time is illustrated with a check mark. Implicit racist attitudes exist if a participant takes longer and makes more errors when "Black or pleasant" (also "White or unpleasant") are paired than when "Black or unpleasant" (also "White or pleasant") are paired.

In their second study, Brief and colleagues (2000) used the average rating of the three qualified Black applicants as their measure of racial discrimination. Our measure was obtained by conducting within-individual regression analyses through the use of a dummy-coded race variable (i.e., 0 = White applicant; 1 = Black applicant) to predict a participant's ratings of the six applicants. More specifically, a hierarchical linear model (HLM) was conducted with the dummy-coded race variable used as a Level 1 (i.e., within-individual) predictor. The slopes for this dummy-coded variable were used as our dependent measure of racial discrimination. A negative slope was indicative of bias against the Black applicants (i.e., the mean for the Black applicants was lower than the mean for the White applicants), a zero slope was indicative of no racial bias (i.e., equal mean ratings of Black and White applicants), and a positive slope was indicative of a pro-Black

applicant bias (i.e., the mean for the Black applicants was higher than the mean for the White applicants). The use of the within-individual slope as our assessment of racial discrimination produced a sensitive measure because it accounts for any disparate ratings between Black and White applicants.

Procedure

At the beginning of the semester, introductory psychology students participated in a mass testing session in which they were asked to complete a variety of measures. Included were the two explicit racist attitude measures (ATB and MRS) as well as the Motivation to Respond Without Prejudice Scale. These measures were spaced equally among a number of

19-25
#1

Table 1
Means, Standard Deviations, and Intercorrelations Among Predictor Variables

Variable	M	SD	1	2	3	4	5	6	7
1. Attitude Toward Blacks (ATB) Scale	2.64	0.88	—						
2. Modern Racism Scale (MRS)	2.79	1.05	.73**	—					
3. Implicit z-score ^a composite	0.00	1.59	.12	.12	—				
4. Motivation to control prejudice	4.98	1.62	.06	-.17	-.07	—			
5. Mean rating of White applicants ^c	4.30	0.53	-.15	-.22*	-.02	-.03	—		
6. Mean rating of Black applicants ^c	3.73	0.74	-.19	-.18	-.29**	.02	.29**	—	
7. Climate condition ^b	0.48	0.50	.10	.12	.12	.02	.23*	-.25*	—

Note. As a result of missing data, *n* ranges from 99–101. Higher scores equal greater amounts of the construct.

^a The implicit measure of racism is a standardized z-score composite, thus the mean is 0. This sample indicated implicit racist attitudes on the basis of the unstandardized scores. ^b For the climate condition, 0 = climate for equality and 1 = climate for racial bias. ^c Although our measure of discrimination was the slope obtained by conducting within-individual regression analyses using a dummy-coded race variable of the ratings of the six applicants, we have provided the mean ratings of the White and Black applicants for illustrative purposes.

* $p < .05$. ** $p < .01$.

other scales in the mass testing session to minimize any response confounds from potentially completing the measures sequentially. Approximately 1 month later, participants were recruited from the mass testing sample for the laboratory portion of the study.

Upon arrival, participants were asked to play the role of a manager and complete the in-basket exercise. Participants were randomly assigned to either the climate for equality or the climate for racial bias condition. This manipulation took place through a memo from the president of the company. In both conditions, the president instructed individuals to take into account applicants' education and experience in making their evaluations of each candidate. However, for participants in the climate for racial bias condition, the president's memo also stated the following:

Given that the vast majority of our workforce is White, it is essential we put a White person in the VP position. I do not want to jeopardize the fine relationship we have with our people in the units. Betty (the outgoing vice president) worked long and hard to get those folks to trust us, and I do not want her replacement to have to overcome any personal barriers.

Participants in the climate for equality condition did not receive this statement. After completion of the in-basket and the embedded hiring recommendation task, participants completed a manipulation-check item, the implicit measure of racism, and a demographic questionnaire.

Manipulation check. In order to determine whether participants were cognizant of the instructions in each condition, participants were asked to recall the hiring preferences of the president. The item asked whether the president preferred to hire applicants that were White, Black, Latino, or of no stated preference.

Statistical analyses. We tested the first three hypotheses with a random slope HLM analysis and the fourth with moderated regression. The hypotheses using HLM were tested through the use of between-person (i.e., Level 2) variables to predict the magnitude of the within-person slope measure of racial discrimination.

Results

Manipulation Check and Preliminary Evidence of Racial Discrimination

Analysis of the manipulation check indicated the hiring preferences of the president were correctly recalled as 95.15% of the participants identified the proper preference for their respective condition; this indicates that the climate for racial bias versus climate for equality manipulation worked. We first performed an

HLM analysis to assess the extent to which ratings were biased. With the within-person Black–White slope (i.e., the Level 1 predictor was the dummy-coded race variable) as our measure of racial bias and no Level 2 predictors, results indicated an overall bias against Blacks; that is, the Black applicants were rated lower than were the White applicants (i.e., $R^2_{\text{within}} = .26$, $b_{y,x} = -.58$, $t[96] = -7.25$, $p < .05$).³ We also examined the variance of this slope and found that it was significantly different from zero (i.e., $\sigma^2_{\text{slope}} = 0.27$, $\chi^2(96, N = 99) = 172.92$, $p < .05$), which indicates that some of the participants exhibited more bias than did others.⁴ As Black applicants were rated lower and there were meaningful individual differences in this degree of bias, we proceeded to test our hypotheses in order to explain these slope differences.

Test of Hypotheses

Table 1 presents the means, standard deviations, and intercorrelations among the study variables. H1 predicted that participants in the climate for racial bias condition would exhibit greater discrimination than would participants in the climate for equality condition. To test H1, we ran a randomized slope HLM model in which we predicted the Level 1 Black–White slope through the use of a Level 2 dummy-coded corporate climate predictor (i.e., 0 = climate for equality; 1 = climate for racial bias). The Level 2

³ As indicated earlier, the in-basket task contained six exceptional candidates: two Black men, one Black woman, two White men, one White woman. Given that the applicants differed on race and gender, it is possible that the gender of the applicant interacted with applicant race to influence participant's ratings. To test whether race and gender interacted, we performed an HLM analysis in which the race, gender, and the interaction of these two variables were used to predict participant ratings. The analyses revealed that the Race \times Gender interaction was not significant ($R^2_{\text{within}} = .00$, $b_{y,x} = -.11$, $t[96] = -1.00$, *ns*).

⁴ The HLM analyses indicated that there was a significant applicant gender effect in that male applicants were rated significantly higher than were female applicants ($R^2_{\text{within}} = .05$, $b_{y,x} = -.28$, $t[96] = -4.87$, $p < .05$). Unlike the race of the applicant, however, this gender effect did not significantly differ across participants ($\sigma^2_{\text{slope}} = 0.02$, $\chi^2[96, N = 99] = 86.72$, *ns*).