



Working Papers

www.cesifo.org/wp

Awareness Reduces Racial Bias

Devin G. Pope
Joseph Price
Justin Wolfers

CESIFO WORKING PAPER NO. 4675
CATEGORY 13: BEHAVIOURAL ECONOMICS
FEBRUARY 2014

An electronic version of the paper may be downloaded

- *from the SSRN website:* www.SSRN.com
- *from the RePEc website:* www.RePEc.org
- *from the CESifo website:* www.CESifo-group.org/wp

Awareness Reduces Racial Bias

Abstract

Can raising awareness of racial bias subsequently reduce that bias? We address this question by exploiting the widespread media attention highlighting racial bias among professional basketball referees that occurred in May 2007 following the release of an academic study. Using new data, we confirm that racial bias persisted in the years after the study's original sample, but prior to the media coverage. Subsequent to the media coverage though, the bias completely disappeared. We examine potential mechanisms that may have produced this result and find that the most likely explanation is that upon becoming aware of their biases, individual referees changed their decision-making process. These results suggest that raising awareness of even subtle forms of bias can bring about meaningful change.

JEL-Code: J150, J700, K310, L830.

Keywords: discrimination, implicit bias.

Devin G. Pope
Booth School of Business
University of Chicago
USA – Chicago, IL 63607
devin.pope@chicagobooth.edu

Joseph Price
Department of Economics
Brigham Young University
USA - Provo, UT 84602
joseph_price@byu.edu

Justin Wolfers
University of Michigan
The Brookings Institution
USA - 48104 Ann Arbor MI
jwolfers@umich.edu

A recent stream of research has documented the existence of in-group racial biases in the employment, criminal, judicial and educational settings (1-5). While social and legal changes have eliminated many institutionalized forms of racial discrimination, the same policy tools may have less leverage against the implicit racial stereotypes that underpin own-race biases.

An example of the continued impact of racial bias on decision making is evident in recent research analyzing the behavior of National Basketball Association (NBA) referees (6). Using NBA data from 1991 to 2002, Price and Wolfers found that personal fouls are more likely to be called against basketball players when they are officiated by an opposite-race refereeing crew than when officiated by an own-race refereeing crew. The own-race bias displayed by NBA referees was large enough to have an appreciable impact on game outcomes and is consistent with a broader literature documenting in-group biases.

In this study, we exploit a natural experiment that occurred in May 2007 when the results of the Price and Wolfers study received wide-spread media attention and examine whether this increased awareness of own-race bias among NBA referees subsequently reduced that bias. The media attention associated with the release of this study included front-page coverage in the New York Times and many other newspapers, extensive coverage on the major news networks, ESPN, talk radio and in the sports media including comments from star players ranging from LeBron James, Kobe Bryant and Charles Barkley, to then NBA Commissioner David Stern. We consider the greater awareness of racial bias that resulted from this to be a quasi-experimental treatment.

Using new data, we replicate the original findings for a sample period after the original study but before the media coverage (2003-2006). Easing the concerns associated with publication bias, we find continued own-race bias during this period that is strikingly similar to

that found in the original 1991-2002 sample. Thus bias continued through the pre-treatment period.

When we conduct the same tests for own-race bias in the period immediately following the media coverage (the post-treatment period, 2007-2010), we find none exists. Moreover, this isn't an artifact of a smaller sample yielding less precise estimates, as the change in bias is itself statistically significant. We argue that this dramatic decrease in bias is a causal result of the awareness associated with the treatment—the *release* and subsequent publicity surrounding the original academic study in 2007. While our inference is only based on a simple pre- versus post-contrast, the magnitude and timing of the effect is notable.

We explore the mechanism for this effect and find no evidence that it is the result of institutional changes made by the NBA. Instead, we argue that the most likely explanation for our finding is that the decisions made by individual referees can be impacted by simply making them aware of their own racial bias. The implication is that racial bias is not a fixed characteristic of an individual decision maker, but rather, it is malleable. In turn, this points to a role that empirical studies of bias may play in changing behavior, and more broadly, it suggests that related policy interventions can play a productive role in reducing racial bias.

Reducing Racial Bias

There has been considerable research about the different influences that can help to reduce racial bias. Important factors that have been shown to reduce racial bias include improved monitoring of the accuracy of individual decision making (9), closer physical proximity to individuals of the other group (10, 11), exposure to multi-cultural education (12, 13), and exposure to situations that contradict the particular bias (14). Additional studies provide examples of how incentives, pressure, and transparency can change racial biases (15-17).

Formulating effective policy solutions requires understanding the type of discrimination that is taking place. For example, in the NBA setting that we study, both referees and players have had a great deal of close physical proximity to members of the opposite race. Thus, increased exposure to members of the other group is likely to have a small impact on eliminating the racial bias found. The decisions that are made by referees in this setting are typically split-second, high-pressure evaluations. This setting, therefore, suggests that implicit discrimination may be important (3).

Implicit discrimination suggests that people have certain mental associations between a group (such as African Americans) and a given attribute. Laboratory research which uses the Implicit Association Test (18, 19) to quantify bias has shown that awareness of subtle biases and willingness to attribute them to internal forces are critical for learning to control them (20, 21), that awareness of racial bias can potentially be channeled into ways to decrease prejudice (22), that contextual variations can impact implicit evaluation (23), and that being motivated to control one's own bias can moderate automatic attitudes (24). Our paper contributes to this literature by testing whether awareness—brought about by a much-publicized empirical study—can reduce implicit bias in an interesting field setting.

Results

Main effects. In Table 1, we start by reproducing the results from the Price and Wolfers (6) study in the first column using the data from the 1991-2002 NBA seasons. In particular, we report the differential number of fouls per 48 minutes played called against a player when facing an out-group refereeing crew, relative to an in-group crew that is composed entirely of referees of the player's own racial group. As in that earlier paper, we find a differential of 0.182 fouls per 48 minutes, and we call this an out-group bias (while acknowledging that it could instead be in-

group favoritism). Relative to the sample mean of 4.46 fouls per 48 minutes, this bias represents about a 4% change in fouls called.

Table 1. Out-group racial bias among NBA referees in three samples

Sample:	Pre-treatment period		Post-treatment period	Change in Bias (out-of-sample)
	Original Study 1991-2002	Out-of-sample 2003-2006	Out-of-sample 2007-2010	From 2003-06 to 2007-10
Out-group bias (Extra fouls per 48 minutes when refereed by an out- group crew, relative to an in-group crew)	0.182*** (0.062)	0.230** (0.093)	-0.0002 (0.089)	0.231** (0.092)
Change in out-group bias, subsequent to “treatment.”				-0.230* (0.130)
N	259,014	93,626	94,682	188,308
Sample mean	4.46	4.43	4.17	4.30

Notes: The years in each column refer to the year in which season started. Each regression includes player and referee fixed effects and controls for home team and whether the player is a starter. Each observation is weighted by the number of minutes the player was in the game. Standard errors in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

In the second column of Table 1, we conduct a simple out-of-sample test in which we estimate the same regression using data from the 2003-2006 seasons. Importantly, this new sample post-dates the sample in the original study but pre-dates the publicity given to that study. In this new sample we find a similar bias as the earlier study; the magnitude of the bias is comparable and is not statistically significantly different from the original sample. This comparison highlights the fact that there was not a downward trend in racial bias prior to May 2007 and casts doubt on concerns that the original findings were a mere statistical aberration, because such an aberration might have been expected to disappear in a new sample. In the third

column, we report the results using data from the 2007-2010 seasons, which include the four seasons that immediately followed the media reporting about racial bias in the NBA. Thus, we consider this sample to be “treated” by greater publicity. In this treated sample, we find that racial bias completely disappeared (the out-group bias falls to a tiny -0.0002 with a standard error of 0.089).

In the final column, we compare our two new samples, estimating the change in out-group bias subsequent to the treatment of greater awareness of racial bias. As in column 2, we find that prior to the treatment, there was a bias of 0.231 more fouls per 48 minutes played when refereed by an entirely out-group refereeing crew relative to an in-group crew. However, during the four years after the treatment, the size of this bias is completely offset—it’s effectively zero—and we find that the difference in racial bias before and after the treatment is statistically significant at the 10% level. In a further (unreported) analysis, we estimated the extent of out-group racial bias in each year and then using these annual estimates, performed a Chow test for a series break, finding statistically significant evidence of a shift in the level of bias in 2007.

In Table 1, all of the results are presented in terms of differences between fouls called when refereed by (racial) out-group versus in-group refereeing crews. Table 2 shows how the raw levels of fouls called varies for each group, in our pre- and post-treatment samples, presenting the (playing-time) weighted average number of fouls called per 48 minutes for black and white players by the time period and the racial make-up of the referee crew, and we make no regression adjustments. As such, we don’t make anything of the raw differences in fouls called against white versus black players, as they tend to reflect the different roles they play on the court. Of greater interest is how this black-white difference varies with the composition of the refereeing crew. Reading down the first two columns shows that the more white referees involved in a game, the fewer fouls are called against white players, with a much smaller decline

for black players. That is, the white-black foul differential declines as the number of white referees increases, which suggests either a negative out-group bias, or a positive in-group bias. Columns 3 and 4 show no such differences for the post-treatment 2007-2010 sample. The last two columns of Table 2 indicate that following the media attention, there was a decrease in fouls being called on both white and black players and for each of the different racial combinations of referee crews. These columns indicate that the elimination of own-race bias was primarily due to a reduction in the correlation between referee race and fouls called on white players and appears to also stem from a particularly large change in the fouls called by crews that involve zero white referees (although these account for only a small proportion of our sample, and so this finding shouldn't be overstated).

Table 2. How foul-calling varies by player race and racial make-up of referee crew

	Pre-treatment			Post-treatment			%Change	
	2003-2006			2007-2010			Post-treatment, relative to pre-treatment	
Referees	White players	Black players	Difference	White players	Black players	Difference	White players	Black players
0 White	5.13	4.4	0.73	4.53	4.05	0.48	-11.63%	-7.91%
1 White	4.92	4.35	0.57	4.51	4.09	0.42	-8.30%	-6.04%
2 White	4.78	4.31	0.47	4.49	4.04	0.45	-6.18%	-6.26%
3 White	4.67	4.23	0.44	4.44	3.97	0.47	-4.89%	-6.13%

Notes: This table provides the average number of fouls called per 48 minutes for Black and White players. Each observation is weighted by the number of minutes the player was in the game. This information is provided separately for the pre-treatment 2003-2006 period, the post-treatment 2007-2010 period, and by the racial make-up of the referee crew (0-3 White referees). The last two columns provide the percentage differences between the 2007-2010 and the 2003-2006 periods.

Mechanism. Our findings suggest that public awareness can reduce the degree of racial bias. While this overall conclusion is important independent of the mechanism, understanding

the channel through which this reduction in racial bias is achieved can be informative about how these results may generalize to other settings.


We begin by looking for evidence that the NBA or referee union took explicit steps to remove racial bias. To do so, we re-analyzed our data at the refereeing level, using the same procedure as above to estimate bias, but doing so separately for games involving each referee. This yields a referee-specific measure of bias. First, we test whether referees that showed a large amount of own-race bias were less likely to continue officiating after the media coverage in 2007. There were 66 referees who officiated at least 100 games between 2003 and 2006. Of these 66 referees, 55 continued and officiated at least 100 games between 2007 and 2010. Of the eleven referees that stopped or significantly reduced the amount of officiating they did starting in the 2007 season, eight were white and three were black. Our estimates of the extent of bias from 2003-06 among these referees who subsequently left the NBA yielded no indication that they were any more favorable to their own race than the average referee of their race during that data period (in fact, the white and black referees that stopped officiating were very slightly less own-race biased on average than their same-race counterparts).

Similarly, we test whether the new referees that started officiating games in 2007 or later were in some way less biased on average. Of the twelve new referees that officiated at least 100 games between 2007 and 2010 (but did not do so in 2003-2006), eight were white and four were black. There is no evidence that these new referees were any less own-race biased than the other referees over this time period (once again, the results go very slightly in the other direction - the new referees if anything showed slightly more own-race bias relative to their colleagues).

Second, we test whether the NBA systematically changed the racial makeup of each crew. The original Price and Wolfers results indicate that most of the racial bias was occurring when all of the referees were of the same race, with little change when moving from one black

referee to two black referees. Thus one easy way for the NBA to reduce the amount of racial bias would have been to increase the fraction of games officiated by mixed-race crews. This policy recommendation is similar to those that followed some high profile police shootings or beatings in which all of the officers involved were white, with the thought being that if at least one of the officers had been black, the incident may not have occurred (25).

Figure 1. Fraction of mixed crews by season



Notes: Each point in the graph indicates the fraction of crews that involved at least one referee that was black and one referee that was not black. The 95% confidence intervals are shaded.

In Figure 1, we document the fraction of games that were officiated by mixed-race crews for each of the seasons between 1991-2010. Comparing changes right around the timing of the media reporting about racial bias, there was an increase from 2005 to 2006 and again from 2006 to 2007, but since 2007 there has been a small but steady decrease in the fraction of mixed-race crews, with an average fraction around 73%. This suggests that the change in bias during this period did not operate through a change in the fraction of mixed-race crews.

We also attempted to verify whether the NBA had instituted a new policy that may have led to a reduction in racial discrimination. A phone conversation with NBA league administrators who oversee the NBA's officiating department suggests that that NBA did not take any specific action to eliminate referee discrimination. Specifically, the administrators to whom we spoke denied that the NBA spoke with the referees about the Price and Wolfers (6) study. They also indicated that the study did not lead to a change in referee incentives or a change in the way they train their referees.

While it is difficult to completely rule out the possibility that the NBA somehow influenced the referees in our study, the evidence presented in this study suggests that the most likely mechanism through which the change in bias occurred is that the media reporting increased the awareness among referees about their own implicit racial bias and that this awareness led to a reduction in such bias. Equally, a purely observational study has obvious limits in assessing the mechanisms at play.

Discussion

In this paper, we examine a real-world setting in which the individuals have large incentives to make correct decisions but were still exhibiting significant amounts of racial bias. Our results suggest that public awareness of racial bias was enough to bring about meaningful change. These results confirm that racial bias is not a fixed characteristic of individual decision making but can be reduced by the efforts of third party observers, particularly those equipped with large and detailed data about past decisions. An open question is whether a similar impact would occur if evidence of racial bias was privately shared with the individual decision makers as opposed to having it publicly disclosed. Our results might encourage organizations to conduct

their own racial bias audits as one of several tools available to reduce racial bias in individual decision making.

Additional Methods

We use box score data from all regular-season games during the 1991-2010 seasons (where the year refers to the year the season starts). We use data from the original Price and Wolfers' study for the 1991-2002 seasons and combine it with data we collected for the more recent years. The box score is a basic summary of the most important statistics for the game including information about individual player statistics (points, minutes, fouls, etc.), the final score of the game, and the names of the three officials. From the boxscore, we extract a dataset describing the performance of each player, in each game. We supplement this with data on the race of the player and referee, using photos to code each individual as either black or not black.

Our main results use the same regression specification as Price and Wolfers in which the main coefficient of interest is an interaction between whether or not the player is black and the fraction of the referees that are white (controlling directly for the race of each). In the absence of any racial bias, this coefficient would be close to zero meaning that the number of fouls that black players receive (relative to white players) does not vary based on the racial composition of the referees. A positive coefficient provides evidence of own-group racial bias in which players receive more fouls when more of the referees assigned to the game are of the other race.

Each regression includes player and referee fixed effects and controls for being on the home team or being one of the starters. Our main dependent variable is a foul rate in which we divide the number of fouls that a player receives by the number of minutes played and multiply by 48 (the number of minutes in a regulation NBA game). We weight each observation by the number of minutes played by each player so that our analysis gives less weight to players who spend very few minutes on the court during the game.

It is important to note that our test is not about differences in fouls received by black players, nor is it about differences in fouls called by white referees. Both of these differences are

controlled for in our player and referee fixed effects. It is possible that black players may systematically receive more fouls or that white referees might call more fouls, but our test is about whether the difference in fouls called between black and white players changes as the racial composition of the refereeing crew differs.

References

1. Anwar, Shamena and Hanming Fang (2004). "An Alternative Test of Racial Prejudice in Motor Vehicle Searches: Theory and Evidence." Yale working paper.
2. Donohue, John and Steven Levitt (2001). "The Impact of Race on Policing and Arrests." *Journal of Law and Economics*, vol. 44, 367-394.
3. Bertrand, Marianne; Dolly Chugh, and Sendhil Mullainathan (2005). "Implicit Discrimination." *American Economic Review*, vol. 95 , no. 2, 94-98.
4. Giuliano, Laura, David Levine, and Jonathan Leonard (2005). "Do Race, Gender, and Age Differences Affect Manager-Employee Relations? An Analysis of Quits, Dismissals, and Promotions at a Large Retail Firm" Working paper.
5. Stoll, Michael, Steven Raphael, and Harry Holzer (2004). "Black Job Applicants and the Hiring Officer's Race." *Industrial and Labor Relations Review*, vol. 57, no. 2, 267-87.
6. Price, J. & Wolfers, J. (2010). Racial Discrimination Among NBA Referees. *The Quarterly Journal of Economics* 125 (4): 1859-1887.
7. Bertrand, M., & Mullainathan, S. (2003). *Are Emily and Greg more employable than Lakisha and Jamal? A field experiment on labor market discrimination* (No. w9873). National Bureau of Economic Research.
8. Dedman, Bill and Latour, Francie (2003). "Traffic Citations Reveal Disparity in Police Searches." *The Boston Globe*, January 6, 2003, pp.A1.
9. Parsons, Christopher A.; Sulaeman, Johan; Yates, Michael C.; Hamermesh, Daniel S. (2011). "Strike Three: Discrimination, Incentives, and Evaluation." *The American Economic Review* Volume 101, Number 4, June 2011 , pp. 1410-1435(26)
10. Marmaros, D., & Sacerdote, B. (2006). How do friendships form?. *The Quarterly Journal of Economics*, 121(1), 79-119.
11. Dovidio, J. F., Eller, A., & Hewstone, M. (2011). Improving intergroup relations through direct, extended and other forms of indirect contact. *Group Processes & Intergroup Relations*, 14(2), 147-160.
12. Richeson, J. A. & Nussbaum, R.J. (2004). The impact of multiculturalism versus color-blindness on racial bias *Journal of Experimental Social Psychology* 40 417-423.
13. Springer, L., Palmer, B., Tenenzini, P. T., Pascarella, E. T. and Nora, A. (1996) Attitudes toward Campus Diversity: Participation in a Racial or Cultural Awareness Workshop *The Review of Higher Education* 20, 53-68.

14. Plant, E. A., Peruche, B. M., & Butz, D. A. (2005). Eliminating automatic racial bias: Making race non-diagnostic for responses to criminal suspects. *Journal of Experimental Social Psychology*, 41(2), 141-156.
15. Gneezy, U., List, J., and Price, M. (2012) Toward an understanding of why people discriminate: Evidence from a series of natural field experiments. *NBER Working Paper #17855*.
16. Zitzewitz, E. (2012) Does transparency reduce favoritism and corruption? Evidence from the reform of figure skating judging. *NBER Working Paper #17732*, *Journal of Sports Economics*, forthcoming.
17. Garicano, L., Palacios-Huerta, I., and Prendergast, C. (2005) Favoritism under social pressure. *Review of Economics and Statistics*, 87, 208-216.
18. Greenwald, Anthony; Debbie McGhee, and Jordan Schwartz (1998). "Measuring Individual Differences in Implicit Cognition: The Implicit Association Test." *Journal of Personality and Social Psychology*, vol. 74, no. 6, 1464-80.
19. Greenwald, Anthony; and Mahzarin Banaji (1995). "Implicit Social Cognition: Attitudes, self-esteem, and stereotypes." *Psychological Review*, vol. 102, no.1, 4-27.
20. Monteith, M. J., Voils, C. I., & Ashburn-Nardo, L. (2001). Taking a look underground: Detecting, interpreting, and reacting to implicit racial biases. *Social Cognition*, 19(4), 395-417.
21. Bargh, J. A. (1999). The cognitive monster: The case against the controllability of automatic stereotype effects.
22. Devine, P. G., Monteith, M. J., Zuwerink, J. R., & Elliot, A. J. (1991). Prejudice with and without compunction. *Journal of Personality and Social Psychology*, 60(6), 817.
23. Mitchell, J. P., Nosek, B. A., & Banaji, M. R. (2003). Contextual variations in implicit evaluation. *Journal of Experimental Psychology: General*, 132(3), 455.
24. Devine, P. G., Plant, E. A., Amodio, D. M., Harmon-Jones, E., & Vance, S. L. (2002). The regulation of explicit and implicit race bias: the role of motivations to respond without prejudice. *Journal of personality and social psychology*, 82(5), 835.
25. Gladwell, M. (2007). *Blink: The Power of Thinking Without Thinking*. Back Bay Books.