

Awareness Reduces Racial Bias

Devin Pope
The Booth School, University of Chicago

Joseph Price
Brigham Young University

Justin Wolfers
The Wharton School, University of Pennsylvania

Abstract

Can empirical studies that raise awareness of racial bias impact decision making? We address this question by exploiting the widespread media attention of racial bias among NBA referees that occurred in May 2007 following the release of an academic paper. Using new data, we find that racial bias continued to exist in the NBA prior to the media coverage, but then completely disappeared during the four years thereafter. These results suggest that publicly documenting cases of racial bias can bring about meaningful change.

*This draft is preliminary and incomplete. Any comments are extremely welcome: joe_price@byu.edu. We are grateful for excellent research assistance from Craig Palsson and Bryson Pope.

Although the US and other countries have made dramatic progress on issues of race over the last century, racial bias continues to influence decisions in a number of important settings such as employment, criminal, judicial, and educational settings.¹

An example of the continued impact of racial bias on decision making is recent research analyzing the behavior of National Basketball Association (NBA) referees (Price and Wolfers, 2010). 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.

An important policy issue is how to eliminate racial bias in settings like those mentioned above. While legal changes have eliminated many of the institutionalized forms of racial discrimination, an equally important question relates to how to eliminate the types of implicit racial stereotypes that continue to influence individuals' decisions such as those demonstrated by NBA referees.

In this study, we exploit a natural experiment that occurred in May 2007 when the results of Price and Wolfers (2010) received considerable media attention and ask whether this increased awareness of racial bias changed the level of own-race bias displayed. Using new data, we begin by replicating the original findings for the data 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, if anything, larger than the bias found in the original sample period (1991-2002).

¹ Anwar and Fang 2004, Donohue and Levitt 2001, Bertrand, McAdam, and Mullainathan 2005, Giuliano, Levine, and Leonard 2005, Stoll et al. 2004

We then test for bias in the data period immediately following the media coverage (2007-2010) and find zero evidence of racial bias. We argue that this dramatic decrease in bias is a causal result of the awareness associated with the release of the original academic study. We explore the mechanism for this effect and find no evidence that it is the result of institutional changes made by the NBA. 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.

Our paper provides a unique test for how empirical studies of discrimination can play an important role in bringing about meaningful change. There have been many studies of racial bias both by academics and investigative reporters that have garnered considerable attention (for example, Bertrand and Mullainathan's (2004) work in the labor market and Dedman's investigative reporting of racial bias by police officers in Boston (2003-2004)). It is possible that these high-impact studies affected decisions made by employers and police officers whose awareness of potential racial bias increased. Our analysis is unique in that we directly test this question.

I. Past research

There has been considerable research about the different influences that can help to reduce racial bias. As expected, racial bias can be manipulated through changes in incentives. For example, Parsons et al. (2011) find that racial bias in baseball is sensitive to how easy it is to scrutinize the umpire's decision (e.g. whether a computerized monitoring system is being used to evaluate the accuracy of judgments). Beyond incentives, some of the important factors that have been shown to reduce racial bias are closer physical proximity to individuals of the other group

(Marmaros and Sacerdote 2006, Dovidio, Eller, and Hewstone 2011), exposure to multi-cultural education (Richeson 2004, Springer 1996), and exposure to situations that contradict the particular bias (Plant 2005).

The NBA setting that we study in this paper is unique in many ways. It is a setting in which both referees and players have had a great deal of close physical proximity to members of the opposite race. Furthermore, the decisions that are made by referees are typically split-second, high-pressure evaluations. This setting, therefore, suggests that implicit discrimination (see Bertrand, Chugh, and Mullainathan 2005 for a primer) may be important.

Implicit discrimination suggests that people have certain mental associations between a group (such as African Americans) and a given attribute. Research in this area has benefitted from the introduction of the Implicit Association Test (IAT), which provides a measure of implicit attitudes (Greenwald et al. 1998; Greenwald and Banaji 1995). Importantly for our study, there exists research on how to moderate the amount of implicit attitudes exhibited by participants in an IAT-type setting. These studies - conducted in the laboratory - have noted that awareness of subtle biases and willingness to attribute them to internal forces are critical for learning to control them (Monteith 2001; Bargh 1999), that awareness of racial bias can potentially be channeled into ways to decrease prejudice (Devine, Monteith, Zuwerink, & Elliot 1991), that contextual variations can impact implicit evaluation (Mitchell, Nosek, and Banaji 2003), and that being motivated to control one's own bias can moderate automatic attitudes (Devine, Plant, Amodio, Harmon-Jones, & Vance, 2002). Our paper contributes to this literature by testing for whether awareness - brought about by an empirical study - can reduce implicit bias in an interesting field setting.

II. Empirical Strategy

The NBA provides a unique setting in which to study racial bias. Referees are quasi-randomly assigned to games. These referees must make quick decisions under pressure, which provide an opportunity to show revealed bias. Price and Wolfers (2010) find that players earn up to 4 percent more fouls and score up to 2½ percent fewer points when there are three referees of the opposite race compared to three referees of their own race. This bias is large enough that the probability of a team winning is noticeably affected by the racial mix of the officials assigned to the game.

In May 2007, the New York Times published a front-page article about Price and Wolfers' research on racial bias among referees in the NBA. Following this article, there was a large amount of media coverage in newspapers, radio, and television. The widespread media coverage of the research at a single point in time provides a nice discontinuous change in awareness about the amount of racial bias in the NBA.

We use box score data from all regular-season games during the 1991-2010 seasons. We use data from the original Price and Wolfers' study for the 1991-2003 seasons (where the year refers to the years the season starts) and then use data from espn.com 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. We use player and referee photos to assign a race to each player and referee and code each individual as either black or not black (see Price and Wolfers (2010) for more detail).

Our main results use the same regression specification as Price and Wolfers (2010) 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. In the absence of any racial bias, this coefficient would be close to zero meaning that number of fouls that black players receive (relative to white players) does not vary based on the racial composition of the referees.² 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.

III. Results

Main effects. In Table 1, we start by replicating the results from the Price and Wolfers study in the first column using the data from the 1991-2002 seasons. We find that the difference in fouls called on black and white players shifts by .197 fouls per 48 minutes played as we change from a situation where all three referees are black to one in which none of them are black. Relative to the sample mean of 4.43, this represents about a 4% change in fouls called.

One of the criticisms of the original study was that Price and Wolfers did not use information about who blew the whistle, but rather used information on the total number of fouls called by the referee crew as a group. One solution to this concern is to limit the observations to

² It is important to note that our test is not about difference in fouls received by black players or 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 there are more white referees officiating the game.

just those games in which either all of the referees are black or none of the referees are black, in which case, we would know the race of the person who blew the whistle. This restriction reduces our original sample from 266,984 observations down to 86,170. In this restricted sample we find very similar results and, if anything, an even larger estimate for the amount of racial bias (Panel B of Table 1).

One concern with any type of empirical estimate of a particular pattern is a publication bias in which only significant results are published and that when a statistical test of the pattern is conducted in out-of-sample data, the pattern is no longer present or is greatly diminished. 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. During these years (which were not part of the original Price and Wolfers study) we still find a significant amount of racial bias and we find that the magnitude of the coefficient has actually increased (though the difference in the coefficients between columns 1 and 2 is not statistically significant). This comparison, though, highlights the fact that there was not a downward trend in racial bias prior to May 2007.

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.³ For our full sample of observations we find a coefficient that indicates that the racial bias had completely disappeared (the coefficient is -0.0002 with a standard error of 0.09). The coefficient increases slightly when we restrict the sample to single-race crews (0.047 with a standard error of 0.12) but again is close to zero and much smaller than the estimated bias during the four seasons prior to the media reporting.

³ The media coverage occurred in May of 2007 - after regular season games for the 2006-2007 season had ended. We test for a change in racial bias starting with the 2007-2008 season which began in October of 2007.

In the final column, we pool the data during the four seasons before and the four seasons after the media reporting to directly test the interaction between racial bias and a post-media-coverage dummy. The main coefficient represents the amount of racial bias during the pre-period and the interaction term indicates by how much this estimate changed after the media reporting. As in the previous column, we find that black players received .231 more fouls per 48 minutes played when switching from an all-black to an all-white referee crew prior to the 2007 NBA season. However, during the four years after the media coverage, the size of this coefficient is completely offset towards zero and we find that the difference in racial bias before and after is statistically significant at the 10% level.

It is also possible to analyze the data at the referee level. Using the same methodology as the original Price and Wolfers (2010) paper, Figure 1 illustrates the estimates of the number of extra fouls per 48 minutes that an individual referee calls on black versus white players for the 55 referees that officiated at least 100 games in both the 2003-2006 and 2007-2010 data periods. The referees are sorted by their 2003-2006 estimates (black-filled dots). The black-filled dots illustrate the own-race bias that is estimated in column 2 of Table 1 by the fact that more squares (black referees) are located on the left side of the figure and more circles (white referees) are located on the right. The unfilled points illustrate the bias estimates for the 2007-2010 data periods. While one might have expected some degree of persistence (referees that are more favorable to black players in the 2003-2006 period are more favorable to black players in the 2007-2010 period - a positive correlation is found when comparing the 1991-2002 estimates with the 2003-2006 estimates), we actually find a negative correlation between estimates for 2003-2006 and 2007-2010. The figure provides a nice visual representation of how the behavior changes for each individual referee in the data.

Mechanism. Our findings suggest that public awareness of racial bias brought about by an empirical study can reduce the degree of racial discrimination. 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 took explicit steps to remove racial 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 11 referees that stopped or significantly reduced the amount of officiating they did starting in the 2007 season, 8 were white and 3 were black. The 2003-2006 individual estimates for these referees show 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).

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 (Gladwell 2007).

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.

It is difficult to completely rule out the possibility that the NBA organization somehow influenced the referees in our study. However, the evidence that we have is 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 this awareness reduced the amount of racial bias that was occurring.

IV. Conclusion

There is a long history of studies showing that racial bias impacts how individuals are treated in a variety of settings. An important question for academia and the media is what impact the awareness of these results may have on individual decision makers. In this paper, we examine a real-world setting in which the individuals had 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.

References


- Bonilla-Silva, E., & Forman, T. (2000). "I am not a racist, but.. .": Mapping White college students' racial ideology in the USA. *Discourse and Society*, 11, 50-85.
- Dovidio, J. F., & Gaertner, S. L. (2000). Aversive racism and selection decisions: 1989 and 1999. *Psychological Science*, 11, 315-319.
- Gaertner, S. L., & Dovidio, J. F. (1986). The aversive form of racism. In J. F. Dovidio, & S. L. Gaertner (Eds.), *Prejudice, discrimination, and racism* (pp. 61-86). Orlando, FL: *Academic Press*.
- Gladwell, Malcolm (2007). *Blink: The Power of Thinking Without Thinking*. Back Bay Books.
- McConahay, J. B., Hardee, B. B., & Batts, V. (1981). Has racism declined in America? It depends on who is asking and what is asked. *Journal of Conflict Resolution*, 25, 563-579.
- Monteith, M. J., Devine, P. G., & Zuwerink, J. R. (1993). Self-directed versus other-directed affect as a consequence of prejudice-related discrepancies. *Journal of Personality and Social Psychology*, 64, 198-210.
- 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.
- Plant, E., Peruche, B. and Butz, D. (2004). Eliminating automatic racial bias: Making race non-diagnostic for responses to criminal suspects. *Journal of Experimental Social Psychology* 41 141-156.
- Son Hing, L. S., Li, W., & Zanna, M. P. (2002). Inducing hypocrisy to reduce prejudicial responses among aversive racists. *Journal of Experimental Social Psychology*, 38, 71-78.
- 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.
- Weisbuch M, Pauker K, Ambady N (2009) The subtle transmission of race bias via televised nonverbal...

Table 1. Changes in the amount of racial bias among NBA referees over time.

	Original Study 1991-2002	Pre-period 2003-2006	Post-period 2007-2010	Change in Bias 2003-2010
A. All Games				
Black player * %White refs	0.182*** (.062)	0.230** (0.093)	-0.0002 (0.089)	0.231** (0.092)
Post*Black* %White refs				-0.230* (0.130)
N	259,014	93,626	94,682	188,308
Sample mean	4.46	4.43	4.17	4.30
B. Single race crews				
Black player * %White refs	0.288*** (0.100)	0.253** (0.129)	0.0469 (0.117)	0.254** (0.127)
Post*Black* %White refs				-0.206 (0.174)
N	83,922	21,710	20,471	42,181


Notes: The years in each column refer to the year in which season started. Each regression includes player and refereed 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

Figure 1. Distribution of Racial Bias by Referee Race and Pre/Post Period



Notes: Each vertical column represents an estimate of the number of extra fouls per 48 minutes an individual referee calls on black versus white players in 2003-2006 (black-filled points) and 2007-2010 (no-fill points). Data are presented for all referees that officiated at least 100 games in the 2003-2006 period and the 2007-2010 period. Specifically, we run separate regressions for each referee, regressing the number of fouls earned per 48 minutes for each player-game observation in which the referee participated, against an indicator variable for whether the offending player is black, controlling for year fixed effects and the full set of player, team-game, and player-game controls and career statistics used in Price and Wolfers (2010).

Figure 2. 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.