Management Science

Publication details, including instructions for authors and subscription information:
http://pubsonline.informs.org

Awareness Reduces Racial Bias

Devin G. Pope, Joseph Price, Justin Wolfers

To cite this article:
Devin G. Pope, Joseph Price, Justin Wolfers (2018) Awareness Reduces Racial Bias. Management Science

Published online in Articles in Advance 06 Feb 2018

. https://doi.org/10.1287/mnsc.2017.2901

http://pubsonline.informs.org/journal/mnsc/

MANAGEMENT SCIENCE
Articles in Advance, pp. 1–8
ISSN 0025-1909 (print), ISSN 1526-5501 (online)

Awareness Reduces Racial Bias

Devin G. Pope,[a,b] **Joseph Price,**[b,c] **Justin Wolfers**[b,d]

[a] University of Chicago, Chicago, Illinois 60637; [b] National Bureau of Economic Research, Cambridge, Massachusetts 02138; [c] Brigham Young University, Provo, Utah 84602; [d] University of Michigan, Ann Arbor, Michigan 48109

Contact: devin.pope@chicagobooth.edu, http://orcid.org/0000-0002-6567-9654 (DGP); joseph_price@byu.edu (JP); jwolfers@umich.edu (JW)

Received: February 24, 2015
Revised: August 25, 2016; November 16, 2016; January 6, 2017
Accepted: January 17, 2017
Published Online in Articles in Advance: February 6, 2018

https://doi.org/10.1287/mnsc.2017.2901

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 disappeared. Several potential mechanisms may have produced this result, including voluntary behavior changes by individual referees, adjustments by players to new information, and changes in referee behavior due to institutional pressure. These results suggest a new kind of Hawthorne effect in which greater scrutiny of even subtle forms of bias can bring about meaningful change.

History: Accepted by John List, behavioral economics.

Keywords: behavioral economics • racial discrimination

1. Introduction

A recent stream of research has documented the existence of in-group racial biases in the employment, criminal, judicial, and educational settings (Donohue and Levitt 2001, Bertrand et al. 2005, Giuliano et al. 2011, Stoll et al. 2004). 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 in-group favoritism.

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 (Price and Wolfers 2010). Using NBA data from 1991 to 2003, Price and Wolfers found that white and black players receive relatively fewer personal fouls when more of the referees officiating the game are their own race. The in-group favoritism 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 particularly unusual natural experiment that occurred in May 2007 when the results of the Price and Wolfers study received widespread media attention, and examine whether this increased awareness of in-group favoritism among NBA referees subsequently impacted the observed 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, and talk radio, and in the sports media, including comments from star players LeBron James, Kobe Bryant, and Charles Barkley, and 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 analysis for a sample period after the original study but before the media coverage (2004–2006). Easing the concerns associated with publication bias, we find continued in-group favoritism during this period that is similar to that found in the original 1991–2003 sample. When we conduct the same tests for in-group favoritism in the period immediately following the media coverage (the posttreatment period, 2007–2010), we find that none exists. Moreover, this is not an artifact of a smaller sample yielding less precise estimates, as the change in bias is itself statistically significant.

We argue that this 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. As a robustness check, we demonstrate that other biases (unrelated to race) that NBA referees exhibit (Price et al. 2012)—but which have not been the subject of similar publicity and scrutiny—continued to persist after 2007. This is supporting evidence that the disappearance of racial bias that we find is a direct result of the media exposure to Price and Wolfers (2010) that focused only on racial bias as opposed to general improvements in refereeing since 2007.

We explore several possible mechanisms for this effect, including voluntary changes by referees in their decision making, behavioral adaptation by players, institutional pressure by the NBA on referees to change their decisions, and more dramatic institutional changes (firing of certain referees, changes in how referees are assigned to games, etc.). While the exact mechanism is hard to pin down, this paper illustrates the broader point that empirical studies of bias can play an important role in changing behavior, and more broadly, it suggests that related policy interventions can reduce racial bias.

2. Methods and Empirical Results

Several empirical strategies have been used to test for evidence of discrimination. One of these methods is in-group favoritism: showing that a decision maker favors members of his/her own race over those of other races. Price and Wolfers (2010) provided evidence of in-group favoritism of NBA referees who were effectively randomly assigned to games. The authors found that more personal fouls were called against players when they were officiated by a referee crew of the opposite race relative to an own-race refereeing crew. This empirical strategy effectively controls for possible differences in style, aggressiveness, and position that black and white players might exhibit by looking at the *difference* in fouls called on the same players by referees of the player's own race versus opposite race.

In this paper, we will use the exact same methodology (even the same regression specifications) as used by Price and Wolfers (2010). However, our interest is in extending the earlier analysis by looking at how in-group favoritism of referees has changed over time. Specifically, we are interested in how in-group favoritism changed after the Price and Wolfers (2010) paper was released and garnered substantial media attention. We will argue that a shift in the level of in-group favoritism that occurs right around the media release of the Price and Wolfers (2010) findings is likely to be a causal result of referees and players being made aware of this bias.

In Table 1, we start by reproducing the results from column 1 of table 4 of the Price and Wolfers (2010) study using the data from the 1991–2003 NBA seasons.[1] In particular, we calculate the number of fouls a player is called for per 48 minutes for each game and report a coefficient that provides the change in the difference in fouls called on white and black players when switching from a refereeing crew with three officials of the player's race to a refereeing crew with three officials not of the player's race. As in the earlier paper, we find a differential of 0.182 fouls per 48 minutes, and we call this our measure of 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.

The original Price and Wolfers paper used data from 1991–2003 but was not publicized until May 2007. Appendix Figure A.1 uses Nexis Lexis data to illustrate that the attention that the Price and Wolfers paper garnered occurred almost exclusively in May 2007 (none before and very little after). Therefore, data from the 2004–2006 seasons can be used as an out-of-sample test since it postdates the sample in the original study but predates the publicity given to that study. In the second column of Table 1, we estimate the same regression using data from the 2004–2006 seasons. In this new sample, we find a similar bias as the earlier study; the magnitude of the bias is comparable and is not statistically distinguishable from the original sample. This comparison casts doubt on concerns that the original findings were a mere statistical aberration. In the third column of Table 1, we report the results using data from the 2007–2010 seasons, which include the four

Table 1. Out-Group Racial Bias Among NBA Referees in Three Samples

	Pretreatment		Posttreatment	Change in bias
Sample	Original study 1991–2003	Out-of-sample 2004–2006	Out-of-sample 2007–2010	From 2004–2006 to 2007–2010
Out-group bias (extra fouls per 48 minutes when refereed by an out-group crew, relative to an in-group crew)	0.192*** (0.059)	0.214** (0.107)	−0.0002 (0.089)	0.214** (0.105)
Change in out-group bias, subsequent to "treatment"				−0.214 (0.139)
N	282,175	70,465	94,682	165,147
Sample mean	4.44	4.46	4.17	4.30

Notes. The years in each column refer to the year in which the 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 are in parentheses.

$^{***}p < 0.01$; $^{**}p < 0.05$; $^{*}p < 0.1$.

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 −0.0002 with a standard error of 0.089).

In the final column of Table 1, 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 the second column, 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 is effectively zero—and we find that the difference in racial bias before and after the treatment is statistically significant at the 10% level.

Figure 1 illustrates the extent of out-group racial bias season-by-season, including an extension of the data to the 2014–2015 season.[2] Each dot in the figure represents the point estimate of out-group bias for a given season, with an accompanying 95% confidence interval. Estimating out-group bias separately by year leads to noisy coefficients but allows one to see the temporal pattern over time in out-group bias. We see evidence of out-group bias in nearly every basketball season from 1999 to 2006. After the publicity treatment, however, we find very little evidence of out-group bias in the following four years. Specifically, we find the opposite of out-group bias in 2007, returning (but statistically insignificant) evidence of out-group bias in 2008, and close to zero out-group bias in 2009 and 2010. The noisy time-series evidence is consistent with our finding in Table 1 that out-group bias significantly decreased in the posttreatment period.

Figure 1. (Color online) Out-Group Bias Effect Size by Season

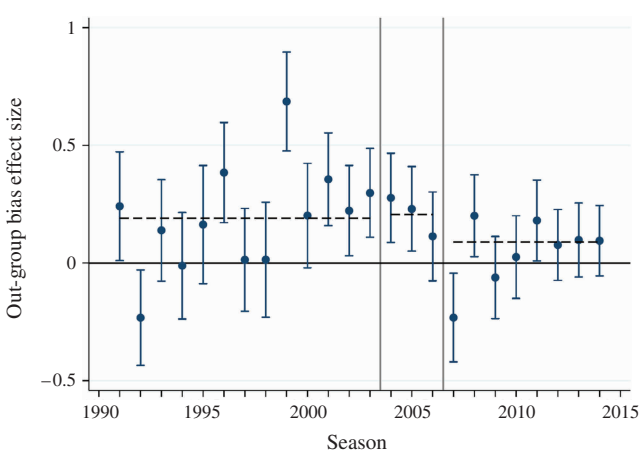

Notes. The point estimates (and accompany 95% confidence intervals) are the bias estimate for each season individually. The first vertical line represents the publication of the original manuscript. The second vertical line represents the timing of media attention. The horizontal dashed lines indicate the average point estimate within each of the three time periods.

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 posttreatment 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 makeup of the referee crew, and we make no regression adjustments. As such, we do not make anything of the raw differences in fouls called against white versus black players, as they tend to reflect the fact that white players are more likely to play positions closer to the basket where more fouls occur. Of greater interest is how this black–white difference varies with the composition of the refereeing crew.

The results in the first two columns of Table 2 show that the more white referees are 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. The fourth and fifth columns of Table 2 show no such differences for the posttreatment 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 in-group favoritism 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 about 5% of our sample, and so this finding should not be overstated).

As is the case with event-study designs, our identification strategy relies on the assumption that nothing happened between the 2006 and 2007 basketball seasons that would cause the racial bias that we find to disappear other than the publicity associated with the Price and Wolfers (2010) study. One candidate alternative explanation is that a prominent referee (Tim Donaghy) was accused of cheating and was fired after the 2006–2007 season. It is possible that the publicity of the Donaghy cheating scandal led the NBA to take actions to improve overall referee quality and that referee bias in general (including racial bias) was reduced because of these actions.[3]

To test for evidence of an overall reduction in referee bias (not just racial bias), we examined the temporal pattern of three other referee biases that have been

Table 2. How Foul Calling Varies by Player Race and Racial Makeup of Referee Crew

	Pretreatment 2004–2006			Posttreatment 2007–2010			Change post vs. pre (%)	
Referees	White players	Black players	Difference	White players	Black players	Difference	White players	Black players
0 white	5.17	4.44	0.73***	4.53	4.05	0.48***	−12.38	−8.78
	(3.88)	(3.35)	(0.03)	(3.54)	(3.04)	(0.02)		
N	17,420	64,168		26,209	99,884			
1 white	4.95	4.41	0.54***	4.51	4.09	0.42***	−8.89	−7.26
	(3.76)	(3.29)	(0.01)	(3.48)	(3.15)	(0.01)		
N	108,040	412,671		153,636	584,718			
2 white	4.81	4.37	0.44***	4.49	4.04	0.45***	−6.65	−7.55
	(3.69)	(3.27)	(0.01)	(3.46)	(3.12)	(0.01)		
N	164,781	621,523		218,400	843,883			
3 white	4.70	4.28	0.42***	4.44	3.97	0.47***	−5.53	−7.24
	(3.63)	(3.21)	(0.02)	(3.58)	(3.12)	(0.01)		
N	62,677	235,309		74,133	294,968			

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 pretreatment 2004–2006 period, for the posttreatment 2007–2010 period, and by the racial makeup of the referee crew (zero to three white referees). The last two columns provide the percentage differences between the 2007–2010 and the 2004–2006 periods.

$^{***}p < 0.01$; $^{**}p < 0.05$; $^{*}p < 0.1$.

shown to exist. Using NBA data from 2004–2010, Price et al. (2012) showed that NBA referees make calls that are biased toward the home team, the team that is losing in a given game, and teams that are losing the game count in a playoff setting. We reanalyzed updated data from Price et al. (2012) to see if these biases decreased in 2007–2010 relative to 2004–2006. The details of this analysis are reported in the appendix. We find no statistically significant difference in any of the three biases. These results suggest that referee bias in general did not decrease after 2007, which lends credibility to the argument that racial bias decreased specifically due to the media exposure generated by the Price and Wolfers (2010) study.

3. Mechanism

Our findings suggest that awareness can reduce the degree of racial bias. While this overall conclusion is important independent of the mechanism, attempting to shed light on the channel through which this reduction in racial bias is achieved can be informative about how these results may generalize to other settings. 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 (Parsons et al. 2011), closer physical proximity to individuals of the other group (Marmaros and Sacerdote 2006, Dovidio et al. 2011), exposure to multicultural education (Richeson and Nussbaum 2004, Springer et al. 1996), and exposure to situations that contradict the particular bias (Plant et al. 2005). Additional studies provide examples of how incentives, pressure, and transparency can change racial biases (Gneezy et al. 2012, Zitzewitz 2014, Garicano et al. 2005).

One obvious mechanism that could lead to the reduction in bias that we find is that the NBA or referee union took explicit steps to remove racial bias by adjusting the assignment of referees across games. To test this, we first reanalyzed our data at the refereeing level, regressing the fraction of fouls called on black players in each game on indicators for each referee with no constant. This yields a referee-specific measure of bias. First, we test whether referees who showed a large amount of in-group favoritism were less likely to continue officiating after the media coverage in 2007. There were 94 referees who officiated more than 100 games between 1991 and 2006. Of these 94 referees, 54 continued and officiated more than 100 games between 2007 and 2010. Of the 40 referees who stopped or significantly reduced the amount of officiating they did starting in the 2007 season, 27 were white and 13 were black. Our estimates of the extent of bias from 1991–2006 among these referees who subsequently left the NBA yield no indication that they were any more favorable to their own race than the average referee of their race during that data period (white referees who stopped officiating were slightly more in-group favoritismed, and black referees who stopped were slightly less in-group favoritismed). Neither difference was significant.

Similarly, we test whether the new referees who started officiating games in 2007 or later were in some way less biased on average. Of the 12 new referees who officiated at least 100 games between 2007 and 2010 (but did not do so in 2004–2006), eight were white and four were black. There is no evidence that these new referees were any less in-group favoritismed 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 in-group favoritism relative to their colleagues).

Another way in which the NBA could have made explicit institutional changes to reduce the bias is by systematically changing 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 2, we document the fraction of games that were officiated by mixed-race crews for each of the seasons between 1991 and 2014. 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 of 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.

While we are unable to find any evidence of explicit changes made by the NBA to impact referee decision making, we cannot completely rule this out as a possible mechanism. At the time when the Price and Wolfers paper was receiving substantial media attention, the NBA publicly challenged the results of the study. Thus, it may be unlikely that they would simultaneously admit to making changes to counteract a bias that they were on the record as claiming did not exist. We are thus open to the possibility that the NBA or referee union took undisclosed steps to encourage referees to be less biased in their decision making.

Figure 2. (Color online) Fraction of Mixed Crews by Season

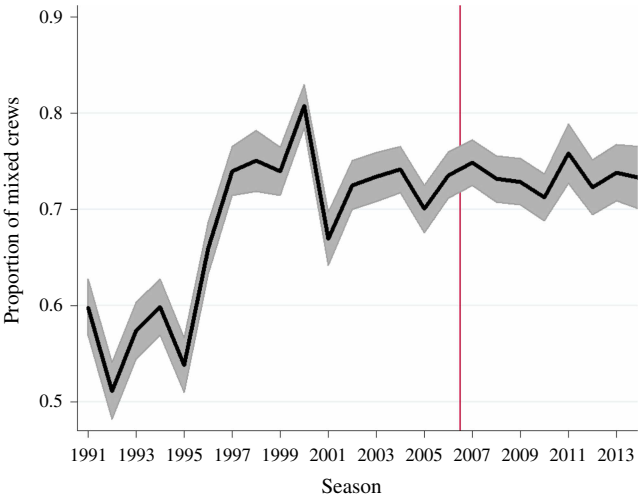

Notes. This figure contains a data point for each basketball season (connected by a line) for the fraction of crews that involved at least one referee who was black and one referee who was not black. The 95% confidence intervals are shaded. The red line represents the treatment date when the original Price and Wolfers study received national media attention.

Another possible mechanism that could be driving the reduction in bias we find is that players and coaches (not referees) reacted to the information from the original Price and Wolfers study, and that this caused a reduction in the bias. For example, it is possible that when a white (black) player is faced with a refereeing crew consisting of three black (white) referees, he may choose to be more passive on the court.

A final possible mechanism for the elimination of the out-group bias that we find is a reduction in out-group bias by individual referees after they became aware of their own racial bias. Given the quick decisions that referees must make in this setting, implicit discrimination may be important (Bertrand et al. 2005). Implicit discrimination suggests that people have certain mental associations between a group (such as African Americans) and a given attribute. Laboratory research that uses the implicit association test (Greenwald et al. 1998, Greenwald and Banaji 1995) 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 (Monteith et al. 2001, Bargh 1999), that awareness of racial bias can potentially be channeled into ways to decrease prejudice (Devine et al. 1991), that contextual variations can impact implicit evaluation (Mitchell et al. 2003), and that being motivated to control one's own bias can moderate automatic attitudes (Devine et al. 2002). Thus, one potential mechanism for our results that has grounding in the experimental literature is that simply making individuals more aware of their own bias leads them to have decreased prejudice.

4. Conclusion

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 publicity that provided evidence of racial bias was enough to bring about meaningful change. 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.

Acknowledgments
All errors are the authors' own.

Appendix
Price et al. (2012) use NBA play-by-play data from the 2002–2008 to test for referee bias. The strategy in that paper for separating referee bias from player behavior was to compare discretionary calls by the referees (traveling, three-second violation, offensive goaltending, and offensive foul) with those that are nondiscretionary (bad pass, lost ball, and shot-clock violation). The play-by-play data allowed the authors to code up all of those specific types of turnovers. The authors found significant evidence of home bias (bias that helps the home team), close bias (bias that helps the team that is losing), and playoff bias (bias that helps the team that is losing the game count in a playoff setting).

For this paper, we updated the data from Price et al. (2012) to include data through 2010. We then use a parallel structure as Table A.1 in this paper to look for evidence of a difference in bias that occurred starting in 2007. Appendix Table A.1 reports the estimates from the same regression used in table 3 from Price et al. (2012) but separated out for the two time

Figure A.1. Publicity of Original Price and Wolfers Study

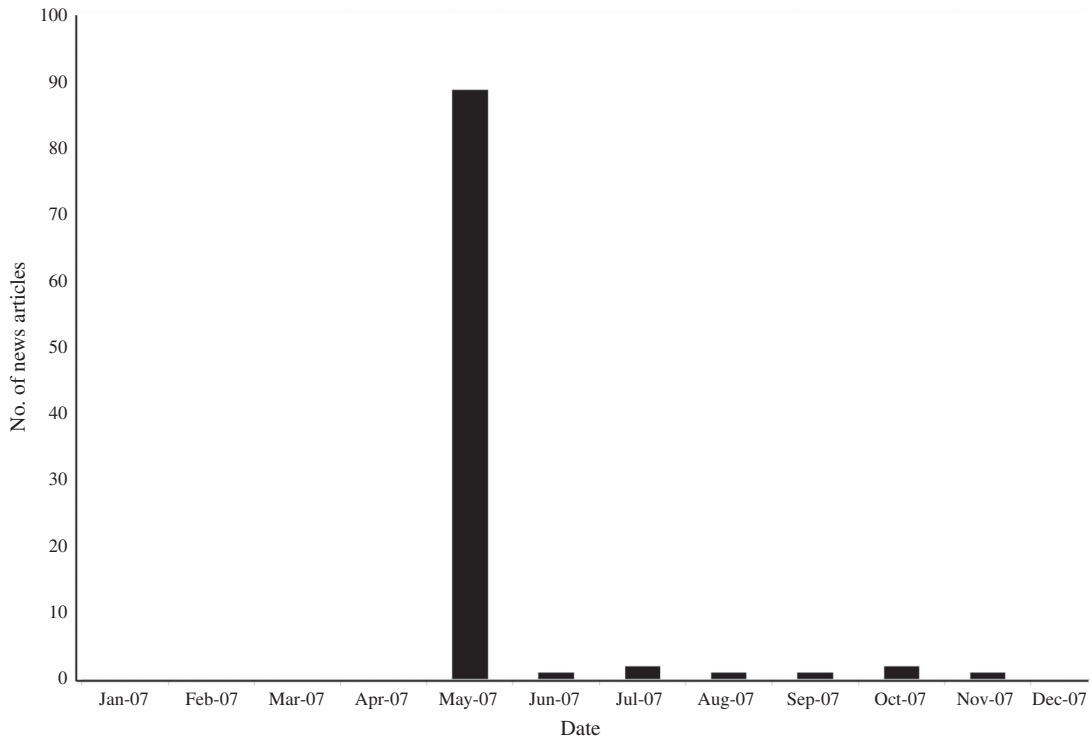

Notes. This figure illustrates the number of news articles produced about the Price and Wolfers study by month of the year in 2007. The data counts come from using the search terms "Justin Wolfers," "Joseph Price," and "NBA" using Nexis Lexis, NewsLibrary.com, and Google News search.

Table A.1. Home and Close Bias Poisson Regression Results

	2004–2006			2007–2010			
	Discretionary	Nondiscr.	Diff.	Discretionary	Nondiscr.	Diff.	Diff.-in-diff.
Home	−0.101***	−0.023***	−0.078***	−0.091***	−0.008	−0.082***	0.004
	(0.0136)	(0.009)	(0.016)	(0.015)	(0.010)	(0.019)	(0.024)
Lead < −10	−0.154***	0.013	−0.167***	−0.014	0.168***	−0.182***	0.015
	(0.024)	(0.015)	(0.028)	(0.027)	(0.016)	(0.031)	(0.042)
Lead >10	0.177***	0.179***	−0.002	0.030	−0.006	0.036	−0.038
	(0.023)	(0.015)	(0.027)	(0.025)	(0.015)	(0.030)	(0.040)
N	323,235	324,000		338,421	338,646		

Notes. These results provide an extension of the results in table 3 of Price et al. (2012). The first row provides the bias in favor of the home team. The second row provides the bias toward the team losing by more than 10 points during the game, and the third row is the bias against the team winning by more than 10 points during the game. All models include matchup (team-opponent-season) fixed effects and quarter fixed effects. The unit of observation is the game-team-minute. Robust standard errors are in parentheses.

$^{***}p < 0.01$; $^{**}p < 0.05$; $^{*}p < 0.1$.

Table A.2. Playoff Bias Poisson Regression Results

	2004–2006			2007–2010			
	Discretionary	Nondiscr.	Diff.	Discretionary	Nondiscr.	Diff.	Diff.-in-diff.
Series difference	0.035***	0.020	0.015	−0.011	−0.008	−0.003	0.018
	(0.019)	(0.014)	(0.024)	(0.020)	(0.012)	(0.023)	(0.033)
Home game	−0.053	0.027	−0.080	−0.171**	0.021	−0.192**	0.112
	(0.084)	(0.057)	(0.102)	(0.080)	(0.053)	(0.096)	(0.140)
Attendance × Home	−0.041***	−0.005	−0.036	−0.018	−0.007	−0.011	−0.025
	(0.025)	(0.019)	(0.031)	(0.031)	(0.021)	(0.037)	(0.048)
Score difference < −10	−0.026	0.021	−0.047	−0.057	0.160***	−0.217***	0.170
	(0.104)	(0.065)	(0.123)	(0.101)	(0.058)	(0.116)	(0.169)
−10 ≤ score difference ≤ −4	0.010	−0.004	0.014	0.005	0.142***	−0.137	0.151
	(0.085)	(0.059)	(0.103)	(0.089)	(0.046)	(0.100)	(0.144)
4 ≤ score difference ≤ −10	0.022	−0.020	0.042	0.030	−0.047	0.077	−0.035
	(0.084)	(0.053)	(0.099)	(0.094)	(0.060)	(0.112)	(0.149)
10 < score difference	0.114	0.087	0.027	0.005	−0.039	0.044	−0.017
	(0.105)	(0.070)	(0.126)	(0.104)	(0.064)	(0.122)	(0.175)
Observations	13,950	13,950		17,536	17,536		

Notes. These results provide an extension of the results in table 4 of Price et al. (2012). Series difference = own games won in series thus far minus opponent's games won in series thus far. Attendance is de-meaned and measured in thousands. Matchup (team-opponent-season) regular season means of the dependent variable, quarter fixed effects, and a constant included as right-hand-side variables in all models. Score difference = start of minute own score minus opponent score (dummy variables for difference being less than −10, greater than −11, and less than −3, etc.). Home = dummy for home game. robust standard errors clustered by game are in parentheses.
$^{***}p < 0.01$; $^{**}p < 0.05$; $^{*}p < 0.1$.

periods of interest in our paper (2003–2006 and 2007–2010). As can be seen in the table, there is significant evidence of home bias and close bias in both data periods. The final column reports the difference in bias across the two periods and finds no statistically significant evidence that a difference exists. The point estimates for home bias and close bias are both small and not close to significant.

The third bias discussed by Price et al. (2012) is playoff bias. This bias uses the game count in the playoffs to test for bias against the underdog. The variable series difference are values based on the difference in the game count. For example, if the series is 3 − 1, then the team ahead in the series would take on a value of 2 and the other team −2. Table A.2 provides the results for playoff bias, once again separating the data into the relevant time periods. Playoff bias is not as robust as home and close bias, and statistically significant evidence of playoff bias is not found in either time period. The point estimates suggest that playoff bias may have been stronger in the 2003–2006 period, but the difference is far from statistically significant.

Figure A.2. (Color online) Out-Group Bias Effect Size by Season and Specification

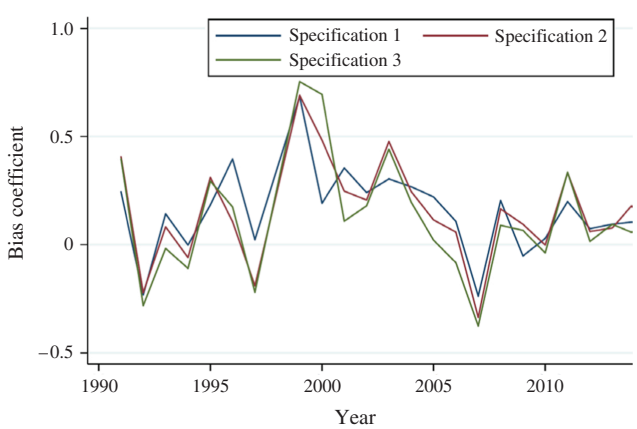

Notes. The points (connected by a line) are the bias estimates for each season individually. Each series of points represents a different specification. Specifically, specification 1 is drawn from column 1 of table 4 of the original Price and Wolfers (2010) paper, and specifications 2 and 3 are from columns 2 and 3, respectively.

Endnotes

[1] Box-score data from all regular-season games during the 1991–2010 seasons (where the year refers to the year the season starts) are used in this analysis. 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 who are white while also controlling directly for each of these measures. 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. Each regression includes player and referee fixed effects and controls for being on the home team or being one of the starters. As in the original paper, each observation is weighted by the number of minutes played by each player so that our analysis gives less weight to players who spend few minutes on the court during the game.

[2] Appendix Figure A.2 replicates Figure 1 using the three different specifications used by Price and Wolfers in their original study (see their table 4).

[3] In the Tim Donaghy case, the cheating scandal was about a referee betting on games and then trying to influence the point spread.

The most obvious action that the NBA could take to eliminate this type of behavior would be to more closely monitor referees off the court and ensure that referee calls are not biased in favor of one team over another (as opposed to one race over another). However, it is feasible that this scandal led the NBA to crack down on all forms of referee bias—including racial bias.

References

Bargh JA (1999) The cognitive monster: The case against the controllability of automatic stereotype effects. Chaiken S, Trope Y, eds. *Dual-Process Theories in Social Psychology* (Guilford Press, New York), 361–382.

Bertrand M, Chugh D, Mullainathan S (2005) Implicit discrimination. *Amer. Econom. Rev*. 95(2):94–98.

Devine PG, Monteith MJ, Zuwerink JR, Elliot AJ (1991) Prejudice with and without compunction. *J. Personality Soc. Psych.* 60(6): 817–830.

Devine PG, Plant EA, Amodio DM, Harmon-Jones E, Vance SL (2002) The regulation of explicit and implicit race bias: The role of motivations to respond without prejudice. *J. Personality Soc. Psych.* 82(5):835–848.

Donohue J, Levitt S (2001) The impact of race on policing and arrests. *J. Law Econom.* 44(2):367–394.

Dovidio JF, 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.

Garicano L, Palacios-Huerta I, Prendergast C (2005) Favoritism under social pressure. *Rev. Econom. Statist.* 87(2):208–216.

Giuliano L, Levine D, Leonard J (2011) Do race, gender, and age differences affect manager-employee relations? An analysis of quits, dismissals, and promotions at a large retail firm. *J. Human Resources* 46(1):26–52.

Gladwell M (2007) *Blink: The Power of Thinking Without Thinking* (Back Bay Books, New York).

Gneezy U, List J, Price M (2012) Toward an understanding of why people discriminate: Evidence from a series of natural field experiments. NBER Working Paper 17855, National Bureau of Economic Research, Cambridge, MA.

Greenwald A, Banaji M (1995) Implicit social cognition: Attitudes, self-esteem, and stereotypes. *Psych. Rev.* 102(1):4–27.

Greenwald A, McGhee D, Schwartz J (1998) Measuring individual differences in implicit cognition: The implicit association test. *J. Personality Soc. Psych.* 74(6):1464–1480.

Marmaros D, Sacerdote B (2006) How do friendships form? *Quart. J. Econom.* 121(1):79–119.

Mitchell JP, Nosek BA, Banaji MR (2003) Contextual variations in implicit evaluation. *J. Experiment. Psych. Gen.* 132(3):455–469.

Monteith MJ, Voils CI, Ashburn-Nardo L (2001) Taking a look underground: Detecting, interpreting, and reacting to implicit racial biases. *Soc. Cognit.* 19(4):395–417.

Parsons CA, Sulaeman J, Yates MC, Hamermesh DS (2011) Strike three: Discrimination, incentives, and evaluation. *Amer. Econom. Rev.* 101(4):1410–1435.

Plant EA, Peruche BM, Butz DA (2005) Eliminating automatic racial bias: Making race non-diagnostic for responses to criminal suspects. *J. Experiment. Soc. Psych.* 41(2):141–156.

Price J, Wolfers J (2010) Racial discrimination among NBA referees. *Quart. J. Econom.* 125(4):1859–1887.

Price J, Remer M, Stone DF (2012) Subperfect game: Profitable biases of NBA referees. *J. Econom. Management Strategy* 21(1):271–300.

Richeson JA, Nussbaum RJ (2004) The impact of multiculturalism versus color-blindness on racial bias. *J. Experiment. Soc. Psych.* 40(3):417–423.

Springer L, Palmer B, Tenenzini PT, Pascarella ET, Nora A (1996) Attitudes toward campus diversity: Participation in a racial or cultural awareness workshop. *Rev. Higher Ed.* 20(1):53–68.

Stoll M, Raphael S, Holzer H (2004) Black job applicants and the hiring officer's race. *Ind. Labor Relations Rev.* 57(2):267–287.

Zitzewitz E (2014) Does transparency reduce favoritism and corruption? Evidence from the reform of figure skating judging. *J. Sports Econom.* 15(1):1–30.