



RACIAL DISCRIMINATION AMONG NBA REFEREES

Executive Summary

Does race color one's evaluation of others? We provide new evidence on racial bias in evaluations, by examining how the number of fouls awarded against black and white NBA players varies with the racial composition of the refereeing crew. Our analysis draws on box-score information for all regular season NBA games from 1991/92 to 2003/04, encompassing over a quarter of a million observations of players in games during which more than 600,000 fouls call were called.

We 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. That is, we find evidence of a pervasive own-race bias. In theory, this may be driven by black or white referees, and either by referees biased toward players of their own race or biased against those of other races.

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. These effects are particularly stark in light of the mismatch of the racial composition of the players and referees: in our sample, blacks account for over 83% of playing time, while white officials account for 68% of officiating time. Had the league instead employed a race-normed refereeing pool, we estimate that this would have changed the outcome of about 1.8% of all games.

The richness of our data allows us to control for a large number of player, team, referee, and game-specific characteristics, allowing us to rule out many alternative explanations for the bias we observe. Were a scientist to attempt to measure these sorts of issues, she might run an experiment, randomly assigning different referee crews to different games, and then comparing the outcomes. In many respects, the NBA process of assigning referees to games approximates this type of experiment, as referee assignments are largely arbitrary, and importantly, referee crews are assigned to games without reference to the characteristics of the relevant teams. Finally, the size of our dataset and the repeated interactions between players and referees allow us to make reasonably precise inferences, taking account of baseline differences in foul-calling of each referee and the playing style of each player.

The own-race preference that we find is not reconcilable with race simply serving as a proxy for other variables (as in theories of statistical discrimination), as this would lead either black or white players to receive more fouls, but not a differential pattern based on the race of the referee. Rather our results suggest that referees interpret similar actions differently, depending on whether the foul was committed by a black or white player, and on whether the referee him or herself is black or white. While explicit animus or malicious intent strikes us as quite unlikely, an emerging literature on implicit discrimination suggests that implicit stereotypes might influence the type of split-second,

high-pressure evaluations required of NBA referees. These sorts of biases may well affect the judgment of even highly-trained professionals, like referees.

We find our results particularly striking given the level of racial harmony achieved along other dimensions in the NBA, the high level of professionalism of NBA referees, and the intense accountability and monitoring under which the referees operate. Thus, while the applicability of our results to other contexts remains an open question, they are at least suggestive that implicit biases may play an important role in shaping our evaluation of others. We study the NBA, because it is a particularly useful laboratory in which to find convincing evidence of own-race bias. While these results may be of interest to those intrigued by the sporting context, we emphasize them instead as potentially suggestive of similar forces operating in a range of other contexts involving rapid subjective assessments.

For instance, just as referees have to evaluate whether or not a foul occurred, teachers must decide whether a student's actions are deserving of disciplinary action, customers decide whether or not to trust proprietors, firms decide who to hire, fire or promote, judges decide who to sentence, and officers decide not only who to arrest, but also make split-second judgments as to whether a suspect is reaching for his gun, or his wallet. The stakes surrounding these decisions are high, and implicit associations may well guide actions beyond the basketball court. Indeed, related evidence suggesting a role for own-race preferences has been documented in a range of other contexts. Donohue and Levitt document evidence that an increase in the number of police of a certain race is associated with an increase in arrests of people of the other race; Antonovics and Knight find that police are more likely to search the vehicle of someone of a different race; Stauffer and Buckley find that supervisors give lower performance ratings for workers of the opposite race; and Stoll, Raphael, and Holzer find that those firms where whites are in charge of hiring are less likely to hire black job applicants than those where blacks control hiring.

We encourage readers to carefully assess our research, and to take full account of the qualifications that necessarily come with this type of research. In particular, we provide greater detail in our research paper (available at: <http://bpp.wharton.upenn.edu/jwolffers/papers/NBARace.pdf>), and encourage a close reading.

We have included one of the tables from the paper in order to illustrate the type of analysis behind our conclusions (next page). The first two columns of this table show the average number of fouls called against black and white players (per 48 minutes played), depending on whether there were 0, 1, 2, or 3 white referees. The third column is the difference in the foul rate between black and white players for each possible racial mix of the officials.

This table shows that black players on average receive fewer fouls than white players. We should emphasize: this observation is not evidence of discrimination—white players are generally taller, heavier, and more likely to play center (all factors that make them more likely to commit fouls). The more striking fact is that the gap in the number of fouls called against black and white players changes as the number of white referees increases (dropping from -0.827 to -0.574). Indeed, this happens even when we make

sure that our comparisons are limited to assessing what happens to the same player in games with different numbers of white referees. The final column reports the change in this gap as we increase the fraction of white referees. The weighted average of these changes indicates that the a black player will receive 0.182 more fouls per 48 minutes played (relative to a white player) as we increase the number of white referees present from zero to three (a 4% increase). The full research paper analyzes a range of other outcomes, as well.

Several people have asked whether the patterns that we observe in our main analytic sample continue to hold in more recent years. In recent days we have put together preliminary data from the 2004/05, 2005/06 and 2006/07 seasons, and find a similar pattern of own race bias. (More information about this supplementary analysis is available on our websites.)

Table 3: Differences in Differences: Foul Rate (= 48*Fouls/Minutes Played)

	Black Players	White Players	Difference: Black – White Foul Rate	Slope: $\Delta(\text{Black-White}) /$ $\Delta\% \text{White Refs}$
0% White Refs	4.418 (0.043)	5.245 (0.094)	-0.827 (0.106)	
33% White Ref	4.317 (0.016)	4.992 (0.035)	-0.675 (0.038)	0.455 (0.331)
67% White Refs	4.335 (0.010)	4.989 (0.023)	-0.654 (0.025)	0.064 (0.137)
100% White Refs	4.322 (0.013)	4.897 (0.029)	-0.574 (0.032)	0.240** (0.121)
Average slope: $\Delta\text{Fouls} / \Delta\% \text{White Refs}$	-0.022 (0.027)	-0.204*** (0.066)		<i>Diff-in-diff</i> 0.182*** (0.066) [p=.006]

Notes: Sample=266,984 player-game observations, weighted by minutes played. The average foul rate in our sample is 4.44. Standard errors in parentheses. ***, **, and * denote statistically significant at 1%, 5% and 10%.

Joseph Price

Ph.D. Candidate
Cornell, Department of Economics
jpp34@cornell.edu
www.people.cornell.edu/pages/jpp34
Office: (607) 255-7622
Cell: (503) 706-5965
Cornell media office: (607) 255-6074

Justin Wolfers

Assistant Professor of Business and Public Policy
The Wharton School, U.Penn
jwolfers@wharton.upenn.edu
www.nber.org/~jwolfers
Office: (215) 898-3013
Cell: (267) 432-4372
Wharton media office: (215) 898-4159

May 2, 2007

ABOUT THE AUTHORS

Joseph Price is a doctoral student in Department of Economics at Cornell University and will defend his doctoral dissertation in July. He will be starting in the fall as an Assistant Professor of Economics at Brigham Young University in Provo, Utah. His research covers labor economics, health economics, and economic analysis of the family.

Justin Wolfers is an Assistant Professor of Business and Public Policy at the Wharton School. He is a visiting scholar with the San Francisco Federal Reserve and a Research Fellow with the National Bureau of Economic Research, and a research affiliate of the Centre for Economic Policy Research in London, and the Institute for the Study of Labor in Bonn. His research covers labor economics, macroeconomics, law and economics, and behavioral economics.