

**LAW AND ECONOMICS SEMINAR
Fall Term 2008**

Professor Polinsky

**Thursday, October 16, 2008
4:15 - 5:45 p.m.
Stanford Law School
Room 272**

**"New Approaches to the Economics of Discrimination: Evidence from NBA
Referees and Female CEOs"**

by

Justin Wolfers

(Wharton School, University of Pennsylvania)

Note to readers: The title of Prof. Wolfers' talk is "New Approaches to the Economics of Discrimination: Evidence from NBA Referees and Female CEOs." His research on NBA referees is complete and is attached. His research on female CEO's is in progress and is not attached, but will be discussed at the seminar.

NBER WORKING PAPER SERIES

RACIAL DISCRIMINATION AMONG NBA REFEREES

Joseph Price
Justin Wolfers

Working Paper 13206
<http://www.nber.org/papers/w13206>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
June 2007

This paper has benefited from helpful comments from David Berri, Mark Cuban, John Donohue, Scott Drewianka, Ron Ehrenberg, Todd Elder, Gary Fields, Joe Gyourko, Kevin Hallock, Christine Jolls, Larry Kahn, Larry Katz, Alan Krueger, Lars Lefgren, David Levine, Janice Madden, Betsey Stevenson, Matthew White, seminar participants at Cornell, Wharton, Northwestern, Chicago, Missouri, UCLA, annual meetings of the Society of Labor Economists and the American Law and Economics Association annual meetings, and various NBA and team officials who provided background information and feedback. We received excellent research assistance from Bryan Elliott. The authors are grateful to the Wharton Sports Business Initiative for research funding. The views expressed herein are those of the author(s) and do not necessarily reflect the views of the National Bureau of Economic Research.

© 2007 by Joseph Price and Justin Wolfers. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Racial Discrimination Among NBA Referees
Joseph Price and Justin Wolfers
NBER Working Paper No. 13206
June 2007
JEL No. J15,J71,K31

ABSTRACT

The NBA provides an intriguing place to test for taste-based discrimination: referees and players are involved in repeated interactions in a high-pressure setting with referees making the type of split-second decisions that might allow implicit racial biases to manifest themselves. Moreover, the referees receive constant monitoring and feedback on their performance. (Commissioner Stern has claimed that NBA referees "are the most ranked, rated, reviewed, statistically analyzed and mentored group of employees of any company in any place in the world.") The essentially arbitrary assignment of refereeing crews to basketball games, and the number of repeated interactions allow us to convincingly test for own-race preferences. We find -- even conditioning on player and referee fixed effects (and specific game fixed effects) -- that more personal fouls are called against players when they are officiated by an opposite-race refereeing crew than when officiated by an own-race crew. These biases are sufficiently large that we find appreciable differences in whether predominantly black teams are more likely to win or lose, based on the racial composition of the refereeing crew.

Joseph Price
Cornell University
414 Uris Hall
Ithaca, NY 14853-7601
United States
jpp34@cornell.edu

Justin Wolfers
Assistant Professor of Economics
Business and Public Policy Department
Wharton School, University of Pennsylvania
3620 Locust Walk
Room 1456 Steinberg-Deitrich Hall
Philadelphia, PA 19104-6372
and NBER
jwolfers@wharton.upenn.edu

Introduction

Does race color our evaluation of others? We provide new evidence on racial biases in evaluation, by examining how the number of fouls awarded against black or white NBA players varies with the racial composition of the refereeing crew. Our setting provides intriguing insights into own-race bias; relative to social, judicial, or labor market settings, the evaluators in our sample (NBA referees) are a particularly expert group, with substantial experience, continual feedback, and they face robust career incentives to be accurate. Indeed, NBA Commissioner Stern has claimed that these referees “are the most ranked, rated, reviewed, statistically analyzed and mentored group of employees of any company in any place in the world.”

The evaluators we study—NBA referees—are effectively randomly assigned to each game. Moreover, the number of games played each year is large, so we can assess both a very clear baseline rate at which individual players commit fouls, and also a clear baseline for the number of fouls called by different referees. Against this baseline, we find systematic evidence of an own-race bias. Notably, players earn up to 4 percent fewer fouls or score up to 2½ percent more points when they are the recipients of a positive own-race bias, rather than a negative opposite-race effect. Player statistics that one might think are unaffected by referee behavior are uncorrelated with referee race. The bias in refereeing is large enough that the probability of a team winning is noticeably affected by the racial composition of the refereeing crew assigned to the game.

These results speak to several literatures. Within the economics of discrimination, this pattern of own-race preference is not reconcilable with efficient statistical discrimination, which would point to a generalized tendency of black or white players to receive more fouls, but not a differential tendency to discriminate correlated with the race of the referee.

Related evidence suggesting a role for own-race preferences has been documented in a range of other contexts. Donohue and Levitt (2001) find 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 (2005) find that police are more likely to search the vehicle of someone of a different race; Stauffer and Buckley (2005) find that supervisors give lower performance ratings for workers of the opposite race; Stoll, Raphael, and Holzer (2004) 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.¹ The advantages of our setting lie in the process for assigning referees to games, which takes no account of player race, thereby ensuring that our findings are not confounded by subjects sorting to preferred evaluators, and repeated interactions which allow for reasonably precise inferences.

Applying a Beckerian taxonomy to our findings, this own-race preference falls under the banner of taste-based discrimination. Within this, we can rule out customer-based discrimination as the cause, as the own-race preference continues to exist even after we hold the stadium (and hence customer base) constant. Additionally, an inference of employer discrimination is inconsistent with our understanding of the formal incentives for accuracy provided by the league.

This suggests a referee-specific taste for discrimination. While explicit animus strikes us as quite unlikely, Bertrand, Chugh and Mullainathan (2005) describe an emerging literature on implicit discrimination which points to the role that implicit associations (such as between blacks and violence) might play in the types of split-second, high-pressure evaluations required of NBA referees.² Frank and Gilovich (1988) have shown that referees tend to call a larger number of fouls on teams wearing black jerseys. Our analysis tests whether such implicit associations vary with the race of the evaluator.

Finally, a large literature has documented substantial evidence of discrimination within sports (Kahn, 1991), a setting which has afforded useful insights largely because measures of

¹ Own-race bias has also been explored in judicial sentencing yielding mixed results (Welch, Combs, and Gruhl 1988; Spohn 1990; Bushway and Piehl 2001; and Schanzenbach 2005). A recent study by Abrams, Bertrand, and Mullainathan (2006) use the random assignment of judges to particular cases and find evidence of racial biases in terms of sentencing but not evidence of own-race bias.

² Greenwald and Banaji (1995) provide an excellent review of implicit social cognition. Payne et al. 2002 note that the need to make quick judgments increases one's susceptibility to implicit stereotyping.

productivity are easily observable. While earlier research suggested that black NBA players suffered substantial wage discrimination (Kahn and Sherer 1988, Koch and Vander Hill 1988), over recent decades, these racial gaps appear to have receded, or even disappeared (Hamilton 1997, Bodvarsson and Brastow 1999). However, while these tests for discrimination typically ask whether wages differ for blacks and whites—conditional on observable game statistics—we demonstrate that these observable game outcomes are themselves the product of biased evaluation by referees. Moreover, in light of the mismatch between the composition of the players (around four-fifths of whom are black) and their evaluators (around two-thirds of referees are white in our sample), an own-race preference may drive an aggregate bias against blacks (or for whites).

While there are important limits to the external validity of extrapolating from our study of NBA referees, it is worth (re-)emphasizing that this is a group that receives far more direct performance feedback than is common in other social and economic evaluation contexts. Thus our finding of own-race preference potentially suggests a role for implicit associations in motivating behavior over a range of other contexts frequently discussed in policy debates. 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 must decide whether or not to trust proprietors; firms who to hire, fire or promote; judges who to sentence; and police officers not only who to arrest, but also in a split-second judgment, 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. Before getting carried away extrapolating from this specific context, we now turn to providing some essential background on the circumstances in which NBA referees make their evaluations and on the incentives they face.

Background: Basketball, the NBA and Referees

In any season, the NBA has around sixty referees, with crews of three referees officiating each game. Assignments of referees to crews are made so as to balance the experience of referees across games, with groups of three referees working together for only a couple of games before being regrouped. According to the NBA, assignments of refereeing crews to specific (regular season) games is “completely arbitrary” with no thought given to the characteristics of the competing teams. Each referee works 70 to 75 games each year and is not allowed to officiate more than nine games for any team or referee twice in a city within a 14 day period. While these constraints mean that assignment of refereeing crews to games is not literally random, the more relevant claim for our approach is that assignment decisions are unrelated to the racial characteristics of either team, and we provide evidence on this point below.

Every game has an observer who meets with the referee for a pre-game discussion, observes the game and reviews video clips from the game with the referees afterwards. These observers report in turn to group supervisors, who provide further input. The director of officiating also provides bi-weekly feedback to each referee on his or her performance. There is also an informal network of monitoring by coaches, spectators, sports analysts, and fans, and feedback from peers, as well as formal channels for feedback from (usually dissatisfied) NBA teams. The league is adamant, however, that while coach input may influence the ratings of individual referees, it has no impact on the assignment of referees to games. These monitoring efforts are supplemented by ongoing referee training, and the league provides ongoing web-mediated training (highlighting controversial calls), videotapes, and summer camps.

The high level of monitoring of referees naturally leads to a high level of accountability for their decisions on the court. The league keeps data on questionable calls made by each referee, and uses this as an input into their internal referee evaluation system. (Unfortunately the NBA refused to share these data with us.) These internal ratings determine which referees will officiate the

playoffs, which provides substantial additional compensation on top of the referee's base salary. While the NBA has not released referee salary data, Jet in 1997 reported that base salary ranged from \$77,000 to \$224,000 per year, and these numbers are widely thought to be accurate. As of 2007, we understand that it is possible for a senior official to earn over half a million dollars per year from officiating-related sources.

Ideally we would like to know how many fouls were called by each referee against each player. However, the NBA boxscore only provides the number of fouls called on each player and the names of the three officials for each game. Thus while we cannot observe the referee who blows the whistle for each foul, our empirical strategy involves comparing the number of fouls each player earns when particular referees are present.

Table 1 provides some initial evidence consistent with the “completely arbitrary” assignment of referees to games, showing that for each year in our sample, the number of white referees is unrelated to the number of black starters. Appendix A also shows that none of our variables have any power in explaining the assignment of referees of each race to particular games within each season.

Player-Level Analysis

Our data contains all boxscore information from all regular season NBA games played from the 1991-92 season through to the 2003-04 season, yielding over a quarter of a million player-game observations. For each player-game, we observe all of their performance statistics (points, blocks, steals, etc) as well as minutes played and the number of personal fouls committed. We coded referees as black or non-black based on visual inspection of press photographs of referees, supplemented by the able assistance of a former NBA referee with a sharp memory. Our data on player race comes from a variety of sources, including Kahn and Shah (2005), Timmerman (2000), and our own coding from past issues of the Official NBA Register, as well as nba.com. In

each case, we simply noted whether a player or referee appeared black, or not. (Hispanics, Asians, and other groups are not well represented among either NBA players or referees, and throughout the paper we refer to non-blacks somewhat imprecisely as “white”.) From separate data sources we draw information about the coach’s race (from the NBA Register), about each player’s characteristics (height, weight and position, from basketballreference.com) and characteristics of each game, including TV coverage. Table 2 provides a list of the variables used in our analysis, as well as a comparison of the mean values between white and black players, weighting all player-level observations by minutes played.

Table 2 shows that, compared to white players, black players play more minutes per game (while Table 2 reports weighted means—30.7 minutes vs 27.2 minutes; while the unweighted means among those with positive playing time are 25.0 vs 20.5). Black players receive about the same number of fouls per game (2.55 vs 2.53) as white players, but receive fewer fouls per 48 minutes played (4.33 vs. 4.97). The differences in foul rates largely reflect the fact that white players tend to be taller, heavier, and more likely to play center than black players.³

Table 3 provides an initial suggestive differences-in-differences analysis. The primary focus of our analysis is the number of fouls called against a player per 48 minutes played, which we refer to throughout as the foul rate. As noted, the foul rate against white players is typically higher than it is for blacks. The number of fouls called against black players is, on average, roughly the same whether the refereeing crew is predominantly white or black. By contrast, white players earn many fewer fouls under white refereeing crews. As such, the “difference-in-difference” suggests that fouls are less likely to be awarded against white players officiated by

³ Note that the large unconditional black-white difference in foul rates is explained by a few observables. First, the unconditional difference:

$$Fouls\ per\ 48\ mins_{it} = 4.97 - 0.64*Black\ player_i \quad Adj.\ R^2=0.005 \quad n=266,984$$

(.016) (.017)

Adding covariates yields:

$$Fouls\ rate_{it} = -0.017*Black\ player_i + 1.47*Center_i + 0.53*Forward_i + 0.025*Height + 0.010*Weight$$

(.017) (.032) (.021) (.003) (.0004)

$$+0.053*Age - 0.086*Experience_{it} -1.366*Starter -0.061 \quad Adj.\ R^2=.097$$

(.005) (.005) (.013) (.252) n=266,984

predominantly white than with predominantly black refereeing crews. The bottom panel shows the full variation in the data, with the difference-in-difference estimate suggesting that a player earns 0.18 fewer fouls per 48 minutes played when facing three referees of his own race than when facing three opposite-race referees.

The richness of our data allows us to extend this analysis to control for the various player, team, referee, and game specific characteristics that might influence the number of fouls called. Thus our estimating equation is:

$$\begin{aligned}
 \text{Foul rate}_{igt} = & \beta_1 \%White\ referees_g * Black\ player_i + \beta_2 \%White\ referees_g + \beta_3 Black\ player_i \\
 & + \beta_4 \text{Observable player}_i, \text{ game}_g, \text{ team}_t, \text{ referee}_r \text{ characteristics} \\
 & [+ \text{Player fixed effects}_i + \text{Referee fixed effects}_i \\
 & + \text{Player characteristics} * \%White\ referees_g \\
 & + Black\ player_i * \text{Stadium}_g \text{ effects} + \text{Team}_t \text{ effects} * \text{Home}_{gt} \\
 & + \text{Team}_t \text{ effects} * \text{Year}_g \text{ effects} \\
 & + \text{Player}_i \text{ effects} * \text{Year}_g \text{ effects} \\
 & + \text{Game}_g \text{ effects} \\
 & + \text{Game}_g \text{ effects} * \text{Team}_t \text{ effects}] + \varepsilon_{igt}
 \end{aligned}$$

where i denotes a player, playing for a team, t , in a specific game, g , officiated by referees, r . All of our estimates weight player-game observations by the number of minutes played. The coefficient of interest is β_1 , and can be interpreted as the effect of opposite-race referees on a player's foul rate (relative to own-race referees). Alternatively phrased, the β_1 coefficient on $\%White\ referees * Black\ player$ captures the *differential* impact of the racial composition of the refereeing crew on black players relative to white players, just as in the difference-in-difference estimates in Table 3.

Table 4 shows our results, beginning with a barebones specification in column one which simply replicates the difference-in-differences specification in Table 3. The next column adds

controls for observable differences including height, weight, position, all-star status, and whether a player is a starter, and whether his team is out-of-contention. These coefficients are reported in subsequent rows. (Not reported are coefficients describing the effect of whether a particular game is played at home, attendance, and the interaction of these variables.) The third column adds player fixed effects, thereby controlling for both observable player-level differences as well as other unobservable differences between black and white players. Similarly we control for referee fixed effects which measure the differential propensity of each referee to call more or less fouls.

While these controls take account of the different styles of individual referees and different roles played by individual players, they do not control for how possible variation in refereeing styles between black and white referees may differentially impact players with different on-court roles. Thus the fourth column adds controls for the share of white referees in a game, interacted with variables describing a player's on-court role, including height, weight, position, age, experience, and whether a player is an all-star that year. We also use our sample data to construct indicators for each player's on-court role, measuring their sample averages on each of the statistics we track (assists, blocks, defensive rebounds, fouls, offensive rebounds, steals, turnovers, free throw attempts, two point attempts, three point attempts—all measured per 48 minutes played—plus free-throw percentage, two-point percentage, and three-point percentage); we then include the interactions of these variables with *%white referees*, as controls. While the full set of these interactions is jointly significant in some specifications (although not in the more complete specifications), their inclusion does not much change our estimated own-race bias.

In order to account for any effects of customer discrimination, the fifth column adds controls each stadium, estimated separately for both black and white players; we also add separate controls for each team, both when at home, and when on the road. The sixth specification adds further controls for team strength, including separate fixed effects for each team-season combination. The seventh column also controls for player-season fixed effects. The eighth

column adds fixed effects for each game, and as such controls for any variation that is common to a game, such as changes in pace, attendance, and the specific refereeing crew. Finally, the last column adds controls for each team-game combination, which means that these results are identified only off the differential propensity of teammates to earn extra fouls when the refereeing crew is of the opposite race.

Across all of these specifications, we find that black players receive around 0.12-0.21 more fouls per 48 minutes played (relative to white players) when the number of white referees officiating a game increases from zero to three (an increase of 2½-4½%).

The time that each player spends on the court is a potential confounder, motivating us to both analyze our dependent variable as a rate (per 48 minutes played), and weight each player's statistics by the number of minutes played. While this is the appropriate specification if the number of foul calls in a game is a linear function of playing time, players who play fewer minutes are less likely to be constrained by the six foul limit,⁴ and hence may be less careful. In the extreme case, a player might be sent in for a few minutes with the express purpose of committing fouls in order to stop the clock in a close game. As such, we re-ran all of these regressions controlling for a quartic in minutes played, or as count data models controlling for a quartic in minutes played, finding similar effects. Alternatively, focusing only on starters yielded similar estimates. We have also re-run our specifications rescaling player statistics to adjust for the "pace" or number of team possessions in each game (as suggested by Kubatko, Oliver, Pelton and Rosenbaum, 2007), again finding similar results. Dropping specific referees or players does not materially affect these results either.

Table 5 moves beyond fouls, to analyze consequences of opposite-race referees on the full range of other measurable player outcomes. As before, we analyze player statistics as rates per 48 minutes played. Five main points are evident from this table. First, beyond the robust

⁴ Once a player earns six fouls, they have "fouled out", and are required to leave the game. Only around 3% of starters actually hit this constraint (and only 1% of bench players).

relationship between personal fouls and opposite-race referees highlighted in Table 4 (and repeated in the first row of Table 5), we find suggestive evidence of similar effects operating on flagrant and technical fouls. While the point estimates are quite large—especially relative to the rarity of these incidents—they are also quite imprecise, and only the effect on flagrant fouls is ever statistically significant, and even then, this varies by specification. (This imprecision reflects the fact that we only have data on these two measures for 1997/98-2003/04; all other measures are available for the full sample). Despite the imprecision of these estimates, they are particularly interesting in that flagrant fouls involve subjective interpretation of physical contact, and technical fouls are often awarded when players dispute an on-court ruling.

Second, the propensity to “foul out” appears unaffected by the race of the refereeing crew, and the 4% rise in the foul rate is partly countered by a 1%-2% decline in playing time. Thus, beyond individual productivity effects, team performance may be affected by composition effects due to effects of opposite-race referees on the distribution of playing time.

Third, beyond fouls, important effects of own-race bias are evident throughout the boxscore. For instance, increasing the share of opposite-race referees leads to an important decline in points scored and a discernible rise in turnovers committed by a player. All told, despite some variability across specifications and imprecision in some measures, the point estimates in Table 5 suggest that player performance deteriorates at nearly every margin when officiated by a larger fraction of opposite-race referees.⁵ Some of these broader responses may directly reflect foul-calling, and indeed, the rise in turnovers suggests that offensive fouls (which are counted as turnovers) may be the key to the effects on personal fouls. A few of these outcomes may also reflect the role of race of the potential “victim” rather than “offender” in shaping foul calls. Specifically, these data yield weakly suggestive evidence of a decline in free throw attempts under

⁵ While the table reports positive effects on the number of rebounds, this is driven entirely by an increase in rebounding opportunities arising from missed shots on offense, and increased shooting by the opposing team. In unreported regressions we measure rebounds as a share of rebounding opportunities, and these positive effects disappear.

opposite-race referees, suggesting that defensive fouls are less likely to be called against one's opponents when opposite-race players have possession. And because a missed shot is not counted as a field goal attempt if a foul is called while shooting, this may help explain the effects on field goal percentage. Alternatively, some responses, such as the decline in steals or blocks, may reflect less aggressive play in response to a player's elevated foul count. There may also be effects of own-race bias mediated through non-foul-related calls, such as a determination of whether the ball is out-of-bounds.

The fourth point speaks to a relatively subtle interpretation issue: while we document a correlation between a player's foul rate and the race of the referee, this may reflect the players responding to the race of the referees, rather than the referees policing opposite-race players more aggressively. *Strategic* responses by players would lead to an attenuation bias: expecting to receive more fouls for a given style of play, the players may play less aggressively, minimizing the impact of referee discrimination on realized fouls. This suggests that our results understate the amount of discrimination. Alternatively, if players exhibit *oppositional* responses, they may play more aggressively when policed by the opposite race. Importantly, such oppositional responses suggest that our findings are driven by changes in player behavior, rather than referee behavior. Yet if this were driving our results, one might expect to see effects not just on the number of fouls called, but on the likelihood of fouling out, as well as other indicators of aggression, including blocks and steals. Instead, we find that blocks and steals actually decline under opposite-race referees.

Fifth, there is only one exception to the general pattern of declining player performance under opposite-race referees: a player's free throw percentage is unaffected by the racial composition of the refereeing pool, and our estimates on this outcome are quite precise. We emphasize this result because this is the one on-court behavior that we expect to be unaffected by referee behavior. As such, free throw percentage serves as a natural "placebo" measure, and it is

reassuring that we find no systematic pattern in this measure. (While this is a useful placebo under the interpretation that our estimates reflect referee behavior; if instead the key is changing player behavior, then theories such as Steele's (1997) "stereotype threat" suggests that the simple presence of opposite-race officials may cause lower free throw accuracy. That we find no such effects cuts against this interpretation.)

Are these effects large enough to affect game outcomes? Berri, Schmidt and Brook (2006) provide a simple framework for analyzing the net effect of these various changes. They note that possession alternates between the competing teams, and hence each will have an equal number of possessions with which to score. Moreover, in expectation, each team scores around one point with each possession. Thus, the contribution of a play to the team's winning margin can be assessed by considering its effects, relative to this baseline. Consequently a missed field goal attempt reduces a team's likely winning margin by one point, while a successful two-point shot increases it by one point, and a three pointer adds two points to the likely margin. Similarly a turnover costs the team possession and hence one point, while winning possession, through either steals or rebounds, adds one point. Assists and blocks are trickier to quantify, but probably change the expected value of a possession by about half a point. Missed free throws only cost half a point, because they only result in a turnover around half the time. Finally, fouls cost a team half a point, which is consistent with a defensive foul changing the game state from one in which the opponent was likely to score one point, to a situation where two free throws are granted, each with a 75% chance of going in.⁶ The bottom row of Table 5 analyzes this composite player productivity metric (measured per 48 minutes played),⁷ suggesting that a team's winning margin would rise by

⁶ Offensive fouls and technical fouls can yield both a change of possession and free throws, and hence may be more costly.

⁷ Berri, Schmidt and Brook call this index the "Win Score", and calculate it as:
 $Win\ Score = (Points - Field\ Goal\ Attempts - \frac{1}{2} Free\ Throw\ Attempts) - Turnovers + Rebounds + Steals + \frac{1}{2} Blocks + \frac{1}{2} Assists - \frac{1}{2} Fouls$. Using a regression to predict whether a team wins as a function of boxscore statistics yields roughly similar weights on these statistics.

up to half a point if they could simply change the race of a player so that it matched that of the refereeing crew.

Team-Level Analysis

One shortcoming of the analysis in Table 5 is that it only analyzes the effects of refereeing decisions to the extent that they are captured in individual player boxscore data. However, a player may make many contributions to their team's performance, only some of which are directly observable. For instance, our analysis of individual personal foul propensities only reflects the role of own-race bias in determining the guilt of an "offender", while it may also shape whether a referee is sympathetic to a player as a "victim".⁸ Moreover, Oliver (2003) notes that a key problem with basketball statistics is that individual-level boxscore statistics paint a rich picture of a player's offensive production, but they do not tell us much about either much of his defensive contribution or general "teamwork".

Ultimately any useful contribution a player makes will be reflected in the scoring of one's team or one's opponents, and hence these contributions can be captured by aggregating our statistics up to the team-game level. Thus, we now turn to analyzing team performance in each game, asking whether we see better team outcomes when a larger fraction of minutes are played by players who are of the same race as the refereeing crew. Naturally, aggregating to the team level substantially reduces the available variation, and so the downside of this approach is more imprecise estimates. (Indeed, recall that results in the final column of Tables 4 and 5 are identified off variation within a team-game, while all of this variation will be lost in this following aggregate analysis.⁹)

⁸ While free throw attempts provide a useful, albeit noisy, measure of the "victim" of defensive fouls, they provide no such detail on offensive fouls.

⁹ Similarly, note that the nature of the identifying variation also changes, as the team-level analysis examines the consequences of alternative refereeing crews on the race of the marginal player (or marginal minutes played), while the individual-level analysis compares outcomes across the average black and white players under alternative refereeing crews.

Thus our key estimating equation in Table 6 is:

$$\begin{aligned}
Fouls_{grio} = & \beta_1 \%White\ referees_g * \%Black\ minutes\ played_{gt} \\
& + \beta_2 \%White\ referees_g * Opponent\ \%Black\ minutes\ played_{go} \\
& + \beta_3 \%White\ referees_g + \beta_4 \%Black\ minutes\ played_{gt} + \beta_5 Opponent\ \%Black\ mins\ played_{go} \\
& + \beta_6 Team_t, opponent_o, game_g, referee_r characteristics_{grio} \\
& [+ Team_t\ fixed\ effects + Opponent_o\ fixed\ effects + Referee_r\ fixed\ effects \\
& + Black\ coach_g * \%White\ referees_g + Opponent\ Black\ coach_g * \%White\ referees_g \\
& + Stadium_g\ effects * \%Black\ minutes_{gt} + Stadium_g\ effects * Opponent\ \%Black\ minutes_{gt} \\
& + Team_t\ effects * Home_{gt} + Opponent_o\ effects * Home_{gt} \\
& + Team_t * Season_g\ effects + Opponent_o * Season_g\ effects \\
& + Game_g\ effects] + \varepsilon_{grio}
\end{aligned}$$

We report standard errors clustered at the game level.

The difference $\beta_1 - \beta_2$ measures the extent to which a team is penalized more often than their opponent because they exhibit greater racial dissimilarity with the refereeing crew than their opponent. Note that this estimate incorporates two causal pathways: the *direct* effect of the referee's propensity to call fouls *against* an opposite-race team (thus β_1 measures the mediating role played by the race of the offending team), as well as an *indirect* effect due to referee's propensity to *protect* a team's opponents by awarding them more fouls if they are of the same race (and hence β_2 measures the role played by the race of the team of the "victim" of the infraction). The net effect on the foul differential (fouls conceded – fouls awarded) is $\beta_1 - \beta_2$. Thus $\beta_1 - \beta_2$ is also the estimate that would result from simply adding game fixed effects to each regression.

This also yields an alternative interpretation that is particularly useful when the dependent variable is points scored. Changing a team's racial composition has a direct effect on the team's scoring, or offensive production, measured by the β_1 coefficient on *%white referees * %black minutes played*. The same change in a team's racial composition also affects their opponent's

expected scoring, and for the opponent, this effect is measured by β_2 , the coefficient on *%white referees * %Opponent black minutes played*. Thus, β_1 measures the effects of own-race bias on a team's offensive production, while β_2 measures the effects on defensive production, with $\beta_1 - \beta_2$ measuring the net effect on winning margin.

Because the number of minutes played by black players may endogenously respond to the racial composition of the refereeing crew assigned to a particular game, we also present instrumental variables results, in which the proportion of a team's minutes played by blacks is instrumented with the average black share of playing time over the team's previous ten games. (Formally we instrument for *%black minutes played * %white referees*—and the equivalent variable for the opponent—with *average %black minutes played in previous ten games * %white referees*, including the direct terms as controls.)¹⁰ Not surprisingly, these are very strong instruments.

For continuity with our earlier analysis, Table 6 initially presents results on the number of fouls awarded against a team. While the imprecision in these estimates cautions against a strong interpretation, as expected we find the estimated *direct* effect of own-race bias on the total number of personal fouls called against a team is roughly five times larger than our estimates of the effect on the number of fouls called against individual players per 48 minutes. Thus the results in Table 4 and 6 are roughly consistent. Importantly, these regressions allow us to test (and reject) an alternative interpretation of the player-level results: that they reflect the referee simply redistributing blame for fouls from own-race members of a team to their opposite-race teammates. Such within-team redistribution would lead the significant individual-level effects to cancel each other out in the team-level regressions.

¹⁰ Random assignment of referees to games allows us to directly test for this endogenous response of black playing time to the racial composition of the refereeing crew. We find little evidence of such a response:

$$\%black\ minutes\ played = 0.00003 * \%White\ referees + Year\ Fixed\ Effects$$

(.00113)

As such, it is not surprising that our IV estimates are quite similar to our OLS estimates, albeit less precise. (While this stands in apparent contrast to the individual-level regressions which suggested that individual players received less playing time under opposite-race referees, a likely reconciliation is that reducing the playing time of one black player increases the likelihood that another black player gets playing time, and the individual-level estimates gave no weight to observations in which zero minutes were played.)

Equally, the indirect effect—due to the referee’s racial similarity to a team’s opponent—is also of a roughly similar magnitude to the direct effect, suggesting that the analysis of individual data understated the effects of own-race bias by about one-half. That is, the race of both the “victim” and “offender” teams are *roughly* of equal importance in shaping own-race bias in foul-calling, while the player-level regressions in Table 4 only emphasized the effects of bias in one’s role as an “offender”.

Naturally, basketball production is measured not in fouls, but in points scored and conceded. Thus the next rows focus on points scored. The estimates again point to a roughly equal role of own-race bias in shaping a team’s offensive production as its defense: the effect of a team’s racial composition is roughly as large on points scored as it is on the points scored by one’s opponent. In further unreported regressions, we look for the proximate sources of these effects, analyzing other boxscore outcomes. These effects appear to be driven by a team’s reduced ability to score once in possession of the ball (there are large effects on field goal percentage, although reassuringly no effects on free throw percentage); there are few effects driven by turnovers.

The estimated effects on offensive and defensive production suggest that own-race bias may have important implications on final game outcomes. To see this, we turn to analyzing team victories in Table 7. Because one team’s win is necessarily their opponent’s loss, the direct and indirect effects will be necessarily equal. Thus, in this analysis we aggregate up to the game level and analyze whether the home team won as a function of the home-versus-away difference in playing time by black players, interacted with the fraction of white referees, plus controls:

$$\begin{aligned}
 I(\text{home team wins}) = & \beta_1 \%White\ referees_g * (\%Black^{home} - \%Black^{away}) \\
 & + \beta_2 \%White\ referees + \beta_3 (\%Black^{home} - \%Black^{away}) \\
 & + \beta_4 (Team\ characteristics^{home} - Team\ characteristics^{away}) \\
 & [+ \beta_4 \%White\ referees * (Black\ coach^{home} - Black\ coach^{away}) \\
 & + Home\ team\ fixed\ effects + Away\ team\ fixed\ effects + Referee\ fixed\ effects
 \end{aligned}$$

$$+ Stadium_g effects * (\%Black^{home} - \%Black^{away})$$

$$+ Home Team * Season effects + Away Team * Season effects] + \varepsilon$$

That is, the coefficient β_1 measures whether a team that has a larger fraction of minutes played by black players, relative to their opponent, is more or less likely to win when more of the referees are white. Noting that the racial mix of the referees might influence the playing time of black and white players, we instrument for each team's racial mix using the average proportion of playing time played by blacks over each team's previous ten games.

The top panel of Table 7 shows quite large and statistically significant impacts of the mismatch between the racial composition of the refereeing crew and that of the players. While Panel A reports the results from a linear probability model, a probit model yielded similar estimates. Panel B turns to analyzing the home team's winning margin instead. The estimates across the two panels are generally rather comparable (in terms of their implications for game outcomes) although the analysis of winning margins yields more precise estimates.

In addition, it is generally believed that coaches have some influence over the decision of referees. If the own-race bias of the referees extends to the race of the coach then we would expect a coach of a particular race to have more influence when a larger fraction of their referees are of his race, especially when facing a coach of the opposite race. The third row of Table 7 shows some weakly suggestive evidence of bias against opposite-race coaches; the magnitude of the coach effect is equal roughly equivalent to the effect of the race of a single player, but quite imprecisely estimated.

Quantitative Interpretation

The results in Table 7 are quite striking, suggesting that own-race bias may be an important factor in determining game outcomes. Figure 1 provides a particularly straightforward representation of the data underlying these findings, plotting local averages of team winning

margins against the proportion of playing time given to black players, relative to the opponents. The slope of these running averages (which show that differences in playing time by black players are correlated with winning margins) is not in itself evidence of bias, as there may be differences in ability. Instead, our analysis highlights the fact that the slope of this relationship appears to be a function of the racial composition of the refereeing crew, and it is this observation that is driving our formal analysis in Table 7.

It is worth pausing to assess the quantitative importance of these results, and their consistency with earlier findings. (For the purposes of this section, we will focus on interpreting the instrumental variables results in Table 7¹¹).

In order to fix an initial scaling, note that the variable measuring racial mismatch between players and referees, $(\%Black^{home} - \%Black^{away}) * \%White\ referees$, has a standard deviation of 0.14, suggesting that a one-standard-deviation rise in mismatch reduces a team's chances of winning by around two percentage points. Of course, this one-standard deviation shock reflects a combination of changes in the racial composition of each team, and changes in the racial composition of the refereeing crew.

We can also use our estimates to assess the sensitivity of game outcomes to changes in just the racial composition of the refereeing crew. For instance, in an average game, one team plays around 15% fewer minutes with black players than their opponent (which roughly corresponds with that team having one fewer black starter). Thus, for this team, the chances of victory under an all-black refereeing crew versus an all-white crew differ by about $0.15 * 0.226$, or around 3.4 percentage points. As such, changing the race of just one referee typically changes the chances of winning by around one percentage point (and the chances of their opponent winning must also change by an offsetting amount).

¹¹ Note that this specification also happens to yield a fairly large estimate of own-race bias; to the extent that other specifications yield smaller coefficient estimates, these magnitudes will decline proportionately.

Throughout our sample, the refereeing crew was on average 68% white, while the teams were 83% black (weighting by playing time). A different thought experiment considers the consequences of race-norming the referee pool so that it matches the racial composition of the player pool. In our sample, the team with a greater share of playing time accounted for by black players won 48.6% of games, which is close to our regression-predicted value of 48.7%. Our estimates suggest that a race-normed refereeing panel would lead this number to rise by 1.8 percentage points, to 50.5%.¹²

In order to translate these magnitudes into payroll consequences, consider the following equation from Szymanski (2003), estimated using team-by-season NBA data from 1986-2000:

$$\text{Win Percentage}_{team, year} = 0.21 + 0.29 * (\text{Team wage bill} / \text{League average wage bill})_{team, year}$$

Interpreting this as a causal relationship suggests that a 1.8 percentage point rise in a team's winning percentage could alternatively be achieved by raising the *aggregate* wage bill of an average team by 6 percent. In turn, consider the modal (and indeed, roughly typical) game in our sample: a team with five black starters playing a team of four black starters and one white. The team with the one white starter could maintain its winning percentage under a shift to race-normed referees by either upgrading the quality of the team by spending an extra 6 percent on player salaries, or by simply exchanging the white starter for a similar quality black starter. As such, this exercise suggests that the racial composition of the refereeing pool has *substantial* consequences for the market value of white versus black players.

The thought experiment also yields interesting player-level implications, and we now turn to analyzing the winning percentage of black and white starters. Given that the large majority of players—on both the winning and losing sides—are black, race-norming the referee pool can change a lot of game outcomes, but still yield only small effects on games won by black players (it would rise from 49.7% to 49.9%, as many black players would gain but nearly as many would lose). But the effects on white

¹² To see this, note that the average absolute difference in the proportion of playing time by blacks is around 15%; multiplying this number by the coefficient of 0.226 yields an estimate of the change in the likelihood of the team with more minutes played by black players winning the game under an all-black versus all-white crew. Further scaling by the magnitude of the proposed change in the proportion of white referees (17%-68%) yields -1.8%. The calculation is less straight-forward for a non-linear model such as a probit, but simulations yielded similar estimates.

players are more dramatic: in our sample, white starters win around 51.8% of their games, but race-norming the refereeing crew would likely lower this winning percentage to 50.4%.

While these estimates of the number of game outcomes determined by own-race bias may seem large, it is worth emphasizing that this is to a large extent a reflection of the high degree of competitive balance within the NBA. Simply put, when game outcomes are typically very close, even fairly small differences in player performance can yield large differences in how frequently each team wins. Indeed, it is this observation that is the key to reconciling what appear to be quantitatively important consequences on game outcomes, with relatively small player-level estimates.

A rough reconciliation goes as follows. Again, consider a game involving five black starters against four blacks and one white. Thus any team-level differences will be driven by the differential treatment of the fifth player, who is black for the home team, and white for their rival. The coefficients in Table 4 suggest that race-norming the refereeing crew would lead the black player to commit around 0.1 fewer fouls per 48 minutes (relative to the change for the white player). Table 5 suggests that he would also score around 0.2 more points and commit 0.05 extra turnovers. Alternatively, using Berri, Schmidt and Brook's (2006) "Win Score" metric, the black player's overall contribution to the team winning margin will rise by about one-quarter of a point under a race-normed refereeing crew (relative to his white rival's contribution). These individual-level estimates are consistent with the estimates of the "direct" effects measured in Table 6. But recall that Table 6 showed that these "direct" effects on fouls committed and points scored are roughly matched by an equal-sized (and opposite-signed) "indirect" effect on fouls awarded, turnovers lost, and points conceded. That is, the away team's boxscore statistics also change in a way that further extends the home team's winning margin by another quarter point.

Thus, race-norming the refereeing crew would, on average, change the winning margin by around half a point, which is what we found in Panel B of Table 7.¹³ Stated this way, it becomes easier to see that the consistency between these game-level results and our earlier player-game-level estimates. Equally, these apparently small impacts of own-race bias have big effects on game-level outcomes in a league in which around 3½ percent of games go to overtime, and around 4 percent of game outcomes are determined by only one point. Indeed, given that the winning margin has a standard deviation of about 12 points and is approximately normally distributed, it is not surprising that only a half-point shift in average winning margins would be sufficient to yield the substantial changes in the winning chances of one or the other team winning.

Behavioral Interpretation

Thus far our analysis has established a robust difference between a player's performance when officiated by an own-race versus opposite-race refereeing crew. Moreover, the consequences of these patterns are quantitatively important, suggesting that the racial composition of the refereeing pool substantially reduces the number of games won by strongly black teams.

Thus, this analysis yields intriguing evidence that the league's historical tendency to hire white referees has a *disparate impact* on black NBA players. While disparate impact may be the relevant legal standard under Title VII of the Civil Right Act,¹⁴ the more interesting social science question remains: what behavior is causing this disparate impact?

Redistribution of Fouls, or Additional Fouls?

While our individual-level results suggest that players earn more fouls under opposite-race referees, it may be that referees simply *redistribute* fouls across team members, charging opposite-

¹³ To see this, multiply the regression coefficient in Panel B of Table 7 by the difference in playing time given to blacks (20% in this example), and further multiply by the difference in the share of white referees (17%-68%), yielding the implication that race-norming referees would lead the winning margin to change by around half a point.

¹⁴ In *Griggs*, the Supreme Court rules that Title VII of the Civil Rights act "proscribes not only overt discrimination but also practices that are fair in form, but discriminatory in operation... [G]ood intent or the absence of discriminatory intent does not redeem employment procedures or testing mechanisms that operate as 'built-in headwinds' for minority groups and are unrelated to measuring job capacity." Equally, it is unclear that the appointment of specific types of refereeing crews qualifies as a relevant "employment practice."

race players for fouls committed by their own-race teammates. Under this view, the patterns we describe are not allocative, and hence of limited interest. The fact that our results are robust even when aggregated at the team level directly falsifies this interpretation. Indeed, the fact that we find large effects on which team wins the game suggests that our estimates point to behavior that is allocative, and not just redistributive.

Player Behavior or Referee Bias?

The observed correlation between player outcomes and referee race could be generated either by referees treating players of the opposite race differently, or by players changing their behavior in response to the refereeing crew. If players responded *oppositionally* to the racial composition of the refereeing crew, then this would increase the number of fouls called under opposite-race officials. However, while fouls rise under opposite-race crews, we find no evidence of that other measures of aggression also rise. While the cost of aggression is a larger number of fouls, one might expect the benefit to be seen elsewhere in the boxscore (such as steals or blocks), and our analysis fails to find this benefit.

On the other hand, it is plausible that players respond *strategically*: Aware of the possibility of earning more fouls under an opposite-race refereeing crew, a player may play more carefully to avoid fouls. Indeed, even if players are unaware of an own-race bias by referees, they are aware of their own foul count, and responding to this alone will yield more careful play under opposite-race referees, and if anything, we find evidence of less aggressive play at other margins. These strategic responses will lead to an attenuation bias, making it harder to discern any effects of own-race bias in the data.

Racial Bias or Different Styles?

One simple explanation of our results is that referees (presumably unconsciously) discriminate toward players of their own race, making split-second evaluations of physical contact in light of their own implicit biases. An alternative explanation follows the usual “omitted variables” interpretation of race differences, but because we are analyzing own-race biases, it is subtly different. This alternative suggests that white and black referees have different focus areas on the floor, or different types of behavior that they are trying to penalize. The omitted variable in this interpretation is the differential propensity for white or black players to make those types of plays, and it may be the interaction of different refereeing styles with different on-court roles that creates the pattern we see in the data. Under either interpretation, a racially unbalanced refereeing crew has a disparate impact on black players relative to white players.

Some of these possibilities can be addressed by aggregating to the team level, as in Table 5. For instance, if certain on-court roles are typically filled by black players, and these roles are more harshly penalized by white referees than black referees, this would yield a correlation between foul calls and player race in the individual data. However, aggregating to the team level effectively aggregates out the differential sorting of blacks and whites to these roles—particularly if the absence of a black player to fill that role would lead a white player to do it instead. As such, the team-level regressions reflect the net impact of changing the racial composition of playing time, but eliminate the influence of which roles are played by which individuals on the team. The fact that we find roughly consistent effects between our individual- and team-level analyses speaks against this omitted variables interpretation.

An alternative approach to understanding what is driving our estimates of discrimination is to test the sensitivity of our results to various proxies for the omitted variable. Thus, we attempt to capture a player’s “style” through variables measuring his height, weight, age, experience, all-star status, and position. We also use each player’s playing history to describe his “style” in terms of

the sample average rate at which free-throw attempts, two point attempts, three point attempts, fouls, assists, steals, blocks, turnovers, offensive and defensive rebounds were earned per 48 minutes played, as well as free throw, two point and three point shooting percentage. Importantly, these variables do successfully pinpoint an identifiably black playing style quite successfully: A probit model attempting to predict a player's race from these "style" variables yielded a pseudo-R² of 0.2, and 12 of 19 variables are individually statistically significant at a 5 percent level.

Even so, the addition of these variables to our main regressions (interacted with *%white referees*, so as to take account of the different response of white referees to the different style of black players) does not appreciably change our estimates of own-race bias (Table 4, columns 3 and 4). Indeed, across the various specifications tested in Table 4, these *player style * %white referees* control variables are jointly significant in some specifications, but insignificant in others (particularly those controlling for game- or team*game fixed effects).

Which Referees?

We now turn to analyzing our data at the level of the referee. We begin with data at the player-game level, and collect all of the observation associated with a particular referee. For each referee, we regress the foul rate against player race, controlling for the full set of player characteristics noted above: height, weight, age, experience, all-star status, position and sample averages of various boxscore statistics. (Not surprisingly, when we disaggregate our data by referee, our statistical power becomes somewhat limited.¹⁵) Even so, Figure 2—which shows each of these point estimates for those referees who have officiated at least 50 games in our sample—highlights four important features of our analysis. First, this figure highlights the intuition of our main result: the influence of player race on foul-calling is, on average, different for white and black referees, and the relevant magnitude appears to be around 0.2 fewer fouls per 48 minutes

¹⁵ Calculating unconditional black-white fouls differences referee-by-referee yields roughly similar results.

called against own-race players. Second, there are no individual referees whose racial biases are particularly notable. (While a few observations are individually statistically significantly different from zero, we do not emphasize this fact, due to the number of referees we test.) Third, the finding of own-race bias is pervasive across all of our referees: nearly all black referees have a greater propensity to call fouls against white players, than nearly all white referees. Fourth, because these regressions are estimated at the referee level, they control for referee-by-referee differences in refereeing “style”, and our main findings appear to remain robust. (That is, these regressions allow the coefficients on player characteristics to be estimated separately for each referee.)

These results also speak to the appropriate measure of the precision of our estimates. Thus far, we have treated the error term as independent across player-game or team-game observations, which may be appropriate under the null that there are no systematic refereeing errors. However a concern about possible interdependencies might lead to concerns that standard errors may be correlated by referee. While the presence of three referees per game makes standard clustering adjustments to our standard errors infeasible, an alternative is to estimate black-white differences in fouls call referee-by-referee. The data in Figure 2 show that 21 out of 29 black referees have a below-average bias in foul-calling against black players, while 34 of 55 white referees have an above-average bias in foul-calling against black players. (Using race-specific norms as our comparison, we find that 22 of 29 black referees show a smaller bias in foul calling against black player than the court-time weighted average among white referees. Similarly, 40 of 55 white referees show a larger bias against black players than the court-time weighted average among black referees.) That is, these results suggest that treating individual referees as the relevant “experiments” still yields statistically significant evidence of own-race bias.

Who Discriminates?

Unfortunately our framework is not well-suited to sorting out whether these results are driven by the actions of black or white referees. To see why, note that the essence of the comparisons underlying our difference-in-difference approach is that the data generated by predominantly white refereeing crews provide a useful baseline for the predominantly black crews (and visa-versa), while asking which group is doing the discriminating or who is discriminated against, requires establishing a “no-discrimination” baseline. While we can control for enough observable features of the game that perhaps our regressions models may establish a reasonable “no-discrimination” benchmark, it is worth emphasizing that this requires substantially stronger assumptions than our earlier analysis.

Figure 2 provides illustrative results: in each of the referee-specific regressions, we control for player characteristics, and assess the referee-specific bias in foul-calling relative to this baseline. While the difference-in-difference framework highlights the systematic differences in this bias between black and white referees, a direct assessment of bias highlights the divergence of these coefficients from zero. Roughly speaking, there is slight evidence of pro-black (or anti-white) bias by black referees (18 of 29 referees have negative coefficients), and somewhat stronger evidence of pro-white (or anti-black) bias by white referees (evident in 43 of 55 cases).

Who is Discriminated Against?

There are also two ways in which these own-race biases may emerge: they may reflect referees favoring players of their own race, or alternatively disfavoring those of the opposite race. The arbitrary assignment of referees to games means that we can test whether our estimates reflect an influence of referee race on black players, or on white players. Table 3 is instructive, showing that the rate at which fouls are called against black players is largely invariant to the racial composition of the refereeing crew. By contrast the rate at which fouls are called against white

players responds quite strongly to referee race. Further regression-based tests yield a similar pattern (see in particular the coefficient on *%white referees* in Table 4), suggesting that the impact of the biases we document is on white players, who are either favored by white referees, or disfavored by black referees.

Conclusion

Using a unique dataset on NBA games, we test whether players of a given race receive fewer fouls when more of the referees present in the game are of the same race. The richness of our data allow us to control for a host of relevant factors that influence the number of fouls called and focus specifically on the repeated interactions between players and referees. We find that players earn up to 4% fewer fouls and score up to 2½% more points on nights in which their race matches that of the refereeing crew. Player statistics that one might think are unaffected by referee behavior are uncorrelated with referee race. The bias in foul-calling is large enough that the probability of a team winning is noticeably affected by the racial composition of the refereeing crew assigned to the game.

These results are striking given the level of racial equality achieved along other dimensions in the NBA and the high level of accountability and monitoring under which the referees operate. Thus, while the external validity of these results remains an open question, they are at least suggestive that implicit biases may play an important role in shaping our evaluation of others, particularly in split-second high-pressure decisions. That is, 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.

References

- Abrams, David; Marianne Bertrand, and Sendhil Mullainathan (2006). "Do Judges Vary in their Treatment of Race?" *mimeo*, University of Chicago.
- Antonovics, Kate and Brian Knight (2004). "A New Look at Racial Profiling: Evidence from the Boston Police Department." NBER working paper 10634, July 2004.
- Berri, David, Martin B. Schmidt and Stacey L. Brook (2006). *The Wages of Wins: Taking Measure of the Many Myths in Modern Sport*, Stanford University Press.
- Bertrand, Marianne; Dolly Chugh, and Sendhil Mullainathan (2005). "Implicit Discrimination." *American Economic Review*, vol. 95, no. 2, 94-98.
- Bodvarsson, Orn and Raymond Brastow (1999). "A Test of Employer Discrimination in the NBA." *Contemporary Economic Policy*, vol. 17, no. 2, 243-255.
- Burdekin, Richard; Richard Hossfield, and Janet Smith (2005). "Are NBA Fans Becoming Indifferent to Race?" *Journal of Sports Economics*, vol. 6, no. 2, 144-159.
- Bushway, Shawn and Ann Piehl (2001). "Judging Judicial Discretion: Legal Factors and Racial Discrimination in Sentencing." *Law and Society Review*, vol. 35, no. 4, 733-64.
- Donohue, John and Steven Levitt (2001). "The Impact of Race on Policing and Arrests." *Journal of Law and Economics*, vol. 44, 367-394.
- Frank, Mark and Thomas Grilovich (1988). "The Dark Side of Self- and Social Perception: Black Uniforms and Aggression in Professional Sports." *Journal of Personality and Social Psychology*, vol. 54, no. 1, 74-85.
- Greenwald, Anthony; and Mahzarin Banaji (1995). "Implicit Social Cognition: Attitudes, self-esteem, and stereotypes." *Psychological Review*, vol. 102, no.1, 4-27.
- Hamilton, Barton (1997). "Racial Discrimination and Professional Basketball Salaries in the 1990s." *Applied Economics*, vol. 29, no. 3, 287-296.
- Jet (1997). "Blacks Shine as Referees in the NBA." *Jet*, May 5, 1997.
- Kahn, Lawrence (1991). "Discrimination in Professional Sports: A Survey of the Literature." *Industrial and Labor Relations Review*, vol. 44, 395-418.
- Kahn, Lawrence and Malav Shah (2005). "Race, Compensation and Contract Length in the NBA: 2001-2." *Industrial Relations*, vol. 44, no. 3, p. 444-462.
- Kahn, Lawrence and Peter Sherer (1988). "Racial Differences in Professional Basketball Players' Compensation." *Journal of Labor Economics*, vol. 6, 40-61.
- Koch, James and C. Warren Vander Hill (1988). "Is There Discrimination in the 'Black Man's Game'?" *Social Science Quarterly*, vol. 69, 83-94.

- Kubatko, Justin, Dean Oliver, Kevin Pelton and Dan Rosenbaum (2007). "A starting Point for Analyzing Basketball Statistics", *mimeo*, UNC-Greensboro.
- Oliver, Dean (2003). *Basketball on Paper: Rules and Tools for Performance Analysis*, Potomac Books.
- Payne, Keith; Alan Lambert; and Larry Jacoby (2002). "Best laid plans: Effects of goals on accessibility bias and cognitive control in race-based misperceptions of weapons." *Journal of Experimental Social Psychology*, vol. 38, 384-396.
- Schanzenbach (2005). "Racial and Sex Disparities in Prison Sentences: The Effect of District-level Judicial Demographics." *Journal of Legal Studies*, vol. 34, no. 1, 57-92.
- Spohn, Cassia (1990). "The Sentencing Decisions of Black and White Judges: Expected and Unexpected Similarities." *Law and Society Review*, vol. 24, no. 5, 1197-1216.
- Stauffer, Joseph M. and M. Ronald Buckley (2005). "The Existence and Nature of Racial Bias in Supervisory Ratings." *Journal of Applied Psychology*, vol. 90, no. 3, 586-591.
- Steele, Claude M. (1997). "A Threat in the Air: How Stereotypes Shape Intellectual Identity and Performance." *American Psychologist*, vol. 52, no. 6, 613-629.
- 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.
- Timmerman, Thomas (2000). "Racial diversity, age diversity, interdependence, and team performance." *Small Group Research*, vol. 31, 592-606.
- Welch, Susan; Michael Combs and John Gruhl (1988). "Do Black Judges Make a Difference?" *American Journal of Political Science*, vol. 32, no. 1, 126-136.

Table 1: Black Starters per Team and the Distribution of Refereeing Crews, by Race

Season	Black Starters per Team				χ^2 -test of independence(a) [p-value]
	0 White Referees	1 White Referee	2 White Referees	3 White Referees	
1991/92	4.33	4.33	4.27	4.28	p=.82
1992/93	4.20	4.20	4.26	4.25	p=.03
1993/94	4.27	4.27	4.31	4.30	p=.80
1994/95	4.20	4.27	4.29	4.25	p=.26
1995/96	4.35	4.26	4.29	4.23	p=.60
1996/97	4.11	4.17	4.19	4.17	p=.97
1997/98	4.22	4.18	4.19	4.21	p=.98
1998/99	4.05	4.13	4.10	4.14	p=.99
1999/00	4.26	4.25	4.14	4.25	p=.07
2000/01	4.15	4.19	4.22	4.18	p=.99
2001/02	4.12	4.08	4.11	4.15	p=.82
2002/03	4.16	4.20	4.11	4.20	p=.79
2003/04	4.03	4.05	4.03	4.04	p=.12
Sample size (% of all player-games)	668 (2.7%)	4,928 (20.1%)	11,580 (47.2%)	7,350 (30.0%)	n=24,526

Notes: Each observation is a team*game observation.

(a) Final column tests: H_0 : #White referees is independent of #black starters

(b) Sample includes all regular season NBA games from 1991/92-2003/04, excluding referee strikes.

Table 2: Summary Statistics (Weighted by Minutes Played)

	Black players		White Players		Difference
	Mean	SD	Mean	SD	
Raw Player Statistics					
Minutes played	30.71	9.98	27.25	10.33	3.464 ^{***}
Fouls	2.547	1.505	2.526	1.542	0.021 ^{***}
Points	13.236	8.366	11.074	7.542	2.163 ^{***}
Player Productivity: Stats*48/Minutes Played					
Fouls	4.330	3.196	4.970	3.933	-0.640 ^{***}
Points	19.759	10.046	18.447	10.106	1.312 ^{***}
Free throws made	3.859	3.897	3.519	3.991	0.340 ^{***}
Free throw missed	1.328	1.987	1.106	1.987	0.223 ^{***}
2 point goals made	6.586	3.988	5.964	4.024	0.623 ^{***}
2 point goals missed	7.3	4.239	6.416	4.361	0.884 ^{***}
3 point goals made	0.909	1.629	1.000	1.780	-0.091 ^{***}
3 point goals missed	1.711	2.363	1.698	2.501	0.014 ^{***}
Offensive rebounds	2.519	2.784	2.696	3.094	-0.177 ^{***}
Defensive rebounds	5.767	4.095	6.271	4.420	-0.504 ^{***}
Assists	4.567	4.077	4.221	4.303	0.346 ^{***}
Steals	1.657	1.885	1.475	1.927	0.182 ^{***}
Blocks	0.995	1.748	1.172	2.065	-0.177 ^{***}
Turnovers	2.968	2.54	2.832	2.737	0.136 ^{***}
Game Information					
Attendance	16,706	3,687	16,798	3,625	-92 ^{***}
Televised game?	0.126	0.332	0.128	0.334	-0.002
Out of contention	0.062	0.241	0.060	0.237	0.002
Black coach	0.241	0.428	0.198	0.398	0.043 ^{***}
Player Characteristics					
Age	27.90	4.02	28.00	3.87	-0.094
NBA experience (yrs)	6.189	3.739	5.783	3.728	0.407 ^{**}
All Star this year	0.134	0.340	0.091	0.287	0.043 ^{***}
Center	0.114	0.318	0.336	0.472	-0.222 ^{***}
Forward	0.440	0.496	0.350	0.477	0.090 [*]
Guard	0.446	0.497	0.315	0.464	0.131 ^{**}
Starter	0.690	0.462	0.588	0.492	0.102 ^{***}
Height (inches)	78.41	3.62	80.54	4.14	-2.13 ^{***}
Weight (lbs)	211.5	26.5	223.2	29.5	-11.7 ^{**}
Foreign-born	0.034	0.181	0.270	0.444	-0.236 ^{***}
Referees					
0 White referees	0.027	0.163	0.028	0.163	-0.001
1 White referee	0.204	0.403	0.207	0.403	-0.003
2 White referees	0.474	0.499	0.471	0.499	0.002
3 White referees	0.295	0.456	0.294	0.456	0.001
# White referees	2.036	0.779	2.031	0.782	0.005
Sample size					
Players	889		301		<u>Total</u> 1,190
Games	13,326		13,130		13,326
Player-games	214,291		52,693		266,984
Player-minutes	5,347,290		1,082,047		6,429,337

Notes: ^{***}, ^{**} and ^{*} denote differences that are statistically significant at 1%, 5% and 10%, respectively.

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

	Black Players	White Players	<i>Difference: Black – White Foul Rate</i>	
Majority White Refs	4.330 (0.008)	4.954 (0.018)	-0.623 ^{***} (0.020)	
Majority Black Refs	4.329 (0.015)	5.023 (0.032)	-0.694 ^{***} (0.036)	
<i>Difference: Majority White - Majority Black</i>	0.001 (0.016)	-0.069 ^{***} (0.040)	<i>Diff-in-diff</i> 0.070 [*] (0.041) [p=.084]	

	Black Players	White Players	<i>Difference: Black – White Foul Rate</i>	<i>Slope: $\Delta(\text{Black-White}) /$ $\Delta\% \text{White Refs}$</i>
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)
<i>Average slope: $\Delta\text{Fouls} / \Delta\% \text{White Refs}$</i>	-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.
(Standard errors in parentheses).
***, **, and * denote statistically significant at 1%, 5% and 10%.

Table 4: Effects of Opposite-Race Referees on Foul Rates

Dependent Variable: Foul Rate (=48*Fouls / Minutes) [Mean=4.43; SD=3.34]									
Independent Variables	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Black player * %White refs	0.182*** (0.066)	0.122* (0.063)	0.197*** (0.061)	0.215*** (0.071)	0.202*** (0.071)	0.201*** (0.071)	0.211*** (0.071)	0.205*** (0.074)	0.188** (0.078)
Control Variables									
Black player	-0.763*** (0.048)	-0.079* (0.046)							
%White refs	-0.204*** (0.060)	-0.137** (0.057)							
Forward		-0.943*** (0.021)							
Guard		-1.426*** (0.032)							
Height (inches)		0.029*** (0.005)							
Weight (lbs)		-0.059*** (0.005)							
All-star		-0.752*** (0.019)	-0.388*** (0.026)	-0.442*** (0.062)	-0.421*** (0.062)	-0.429*** (0.064)			
Starting Five		-1.247*** (0.014)	-0.998*** (0.016)	-1.025*** (0.039)	-1.005*** (0.039)	-0.980*** (0.039)	-0.814*** (0.041)	-0.797*** (0.042)	-0.782*** (0.043)
Out of contention		-0.202*** (0.026)	-0.126*** (0.027)	-0.126*** (0.027)	-0.085*** (0.028)	-0.048* (0.029)	-0.059** (0.029)		
R²	0.01	0.10	0.18	0.18	0.19	0.19	0.22	0.26	0.28
Other Controls									
Observables^(a)		✓	✓	✓	✓	✓	✓	✓	✓
Referee fixed effects			✓	✓	✓	✓	✓	✓	✓
Player fixed effects			✓	✓	✓	✓	✓	✓	✓
Player characteristics *%White refs				✓	✓	✓	✓	✓	✓
Stadium*Black player fixed effects					✓	✓	✓	✓	✓
Team*Home fixed effects					✓	✓	✓	✓	✓
Team*Year fixed effects						✓	✓	✓	✓
Player*Year fixed effects							✓	✓	✓
Game fixed effects								✓	✓
Game*Team fixed effects									✓

Notes: Sample=266,984 player-game observations, weighted by minutes played. Each column reports the results of a separate regression. (Standard errors in parentheses). ***, **, and * denote statistically significant at 1%, 5% and 10%.

Missing values reflect the fact that some controls are unidentified in the presence of perfectly collinear fixed effects.

(a) Observable variables not reported include controls for: home, attendance, home*attendance.

Table 5: Effects of Opposite-Race Referees on Player Performance (Measured per 48 minutes)

Dependent Variable	Mean (SD)	Coefficient on <i>Black Player * % White Referees</i>								
		(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
<i>Personal Fouls</i>	4.44 (3.34)	0.182*** (0.066)	0.122* (0.063)	0.197*** (0.061)	0.215*** (0.071)	0.202*** (0.071)	0.201*** (0.071)	0.211*** (0.071)	0.205*** (0.074)	0.188** (0.078)
<i>Flagrant fouls</i>	0.012 (0.17)	0.005 (0.005)	0.006 (0.005)	0.007 (0.005)	0.010* (0.006)	0.011* (0.006)	0.011* (0.006)	0.011* (0.006)	0.012** (0.006)	0.009 (0.006)
<i>Technical Fouls</i>	0.08 (0.38)	0.006 (0.011)	0.007 (0.011)	0.007 (0.010)	0.017 (0.013)	0.016 (0.013)	0.016 (0.013)	0.018 (0.013)	0.017 (0.013)	0.016 (0.014)
<i>Minutes</i>	30.13 (10.1)	-0.396 (0.199)	-0.510*** (0.146)	-0.427*** (0.137)	-0.613*** (0.159)	-0.565*** (0.159)	-0.581*** (0.157)	-0.319** (0.149)	-0.330** (0.149)	-0.403*** (0.155)
<i>Fouled out</i>	0.025 (0.16)	0.000 (0.003)	-0.002 (0.003)	0.000 (0.003)	0.001 (0.004)	0.001 (0.004)	0.000 (0.004)	0.002 (0.004)	0.003 (0.004)	0.002 (0.004)
<i>Points</i>	19.54 (10.1)	-0.463** (0.200)	-0.356* (0.190)	-0.409** (0.176)	-0.309 (0.205)	-0.339* (0.205)	-0.383* (0.204)	-0.400** (0.202)	-0.450** (0.210)	-0.430* (0.221)
<i>Free Throw Attempts</i>	5.09 (4.90)	-0.119 (0.097)	-0.120 (0.095)	-0.107 (0.09)	-0.035 (0.105)	-0.039 (0.105)	-0.056 (0.105)	-0.077 (0.105)	-0.039 (0.107)	-0.034 (0.112)
<i>Free Throw %</i>	0.75 (0.23)	-0.003 (0.006)	0.004 (0.006)	0.002 (0.006)	0.004 (0.007)	0.005 (0.007)	0.003 (0.007)	0.005 (0.007)	0.004 (0.007)	0.004 (0.008)
<i>2 point attempts</i>	13.63 (6.49)	-0.136 (0.128)	-0.160 (0.121)	-0.026 (0.106)	0.096 (0.123)	0.066 (0.123)	-0.020 (0.121)	0.002 (0.117)	-0.070 (0.117)	-0.064 (0.126)
<i>2 point %</i>	0.48 (0.19)	-0.005 (0.004)	-0.006 (0.004)	-0.007* (0.004)	-0.007 (0.005)	-0.008* (0.005)	-0.007 (0.005)	-0.008 (0.005)	-0.008 (0.005)	-0.007 (0.005)
<i>3 point attempts</i>	2.63 (3.43)	-0.010 (0.068)	0.097 (0.061)	-0.047 (0.049)	-0.086 (0.057)	-0.082 (0.057)	-0.063 (0.055)	-0.048 (0.052)	-0.039 (0.054)	-0.061 (0.056)
<i>3 point %</i>	0.35 (0.27)	-0.014* (0.008)	-0.013 (0.008)	-0.013 (0.008)	-0.014 (0.009)	-0.014 (0.009)	-0.013 (0.009)	-0.014 (0.010)	-0.015 (0.010)	-0.011 (0.012)
<i>Assists</i>	4.51 (4.12)	-0.156* (0.082)	0.027 (0.068)	-0.003 (0.062)	-0.005 (0.072)	0.017 (0.072)	0.023 (0.071)	0.038 (0.071)	0.027 (0.073)	0.046 (0.076)
<i>Blocks</i>	1.02 (1.81)	0.010 (0.036)	-0.056* (0.032)	-0.058* (0.030)	-0.020 (0.035)	-0.019 (0.035)	-0.013 (0.035)	-0.042 (0.035)	-0.029 (0.037)	-0.022 (0.038)
<i>Defensive rebounds</i>	5.85 (4.16)	0.102 (0.082)	0.030 (0.073)	0.019 (0.070)	0.041 (0.081)	0.050 (0.081)	0.054 (0.081)	0.063 (0.081)	0.057 (0.084)	0.056 (0.088)
<i>Offensive rebounds</i>	2.55 (2.84)	0.195*** (0.056)	0.101** (0.051)	0.038 (0.049)	0.029 (0.057)	0.021 (0.057)	0.017 (0.057)	0.014 (0.057)	0.014 (0.059)	0.010 (0.061)
<i>Steals</i>	1.63 (1.89)	-0.043 (0.038)	-0.022 (0.037)	-0.061* (0.037)	-0.038 (0.043)	-0.039 (0.043)	-0.039 (0.043)	-0.039 (0.043)	-0.057 (0.044)	-0.053 (0.046)
<i>Turnovers</i>	2.95 (2.57)	0.043 (0.051)	0.049 (0.051)	0.109** (0.050)	0.132** (0.058)	0.141** (0.058)	0.138** (0.058)	0.128** (0.058)	0.102* (0.060)	0.091 (0.063)
<i>Net Effect (Win Score)</i>	8.36 (9.09)	-0.212 (0.180)	-0.249 (0.174)	-0.524*** (0.170)	-0.522** (0.198)	-0.514*** (0.197)	-0.474** (0.197)	-0.512*** (0.198)	-0.513** (0.207)	-0.449** (0.214)
<i>Observables</i>			✓	✓	✓	✓	✓	✓	✓	✓
<i>Referee and Player</i>				✓	✓	✓	✓	✓	✓	✓
<i>Player char*%W. refs</i>					✓	✓	✓	✓	✓	✓
<i>Stadium*Black player</i>						✓	✓	✓	✓	✓
<i>Team*Home</i>						✓	✓	✓	✓	✓
<i>Team*Year</i>							✓	✓	✓	✓
<i>Player*Year</i>								✓	✓	✓
<i>Game</i>									✓	✓
<i>Game*Team</i>										✓

Notes: Each cell reports results from a separate regression. See notes to Table 4 for specification details. Regressions analyzing shooting percentages are weighted by attempts, rather than minutes. n=266,984, except flagrant and technical fouls n=136,509.

Table 6: Effects of Opposite-Race Referees on Team Performance

		Coefficient on %Black Playing Time * % White Referees						
		(1)	(2)	(3)	(4)	(5)	(6)	(7)
Mean (SD)	Dependent Variable: Total Fouls by Team							
Total effect ($\beta_1 - \beta_2$)	22.4 (4.65)	1.911* (0.995)	1.927** (0.982)	2.148** (0.965)	2.063** (0.968)	2.152** (0.970)	1.906** (0.942)	1.671 (1.043)
Of which:								
Direct effect (β_1) (fouls committed)		0.816 (0.503)	0.789 (0.497)	0.980 (0.775)	1.082 (0.781)	1.372* (0.778)	1.330* (0.738)	1.098 (0.812)
Indirect effect (β_2) (fouls awarded)		-1.095** (0.503)	-1.138** (0.497)	-1.168 (0.797)	-0.981 (0.803)	-0.780 (0.803)	-0.576 (0.762)	-0.573 (0.837)
		Dependent Variable: Points Scored by Team						
Total effect ($\beta_1 - \beta_2$)	98.4 (12.4)	-4.094** (2.152)	-4.777*** (2.041)	-5.726*** (2.010)	-5.425*** (2.021)	-5.640*** (2.017)	-3.832** (1.952)	-6.287*** (2.222)
Of which:								
Direct effect (β_1) (points scored)		-0.415 (1.101)	-0.882 (1.044)	-2.366 (1.988)	-1.768 (2.002)	-2.139 (1.991)	-2.302 (1.789)	-3.303* (1.994)
Indirect effect (β_2) (points conceded)		3.678*** (1.093)	3.894*** (1.041)	3.360* (1.975)	3.657* (1.987)	3.501* (1.966)	1.531 (1.798)	2.984 (1.998)
Observables		✓	✓	✓	✓	✓	✓	✓
Referee			✓	✓	✓	✓	✓	✓
Team / Opponent			✓	✓	✓	✓	✓	✓
Blk coach * %W. refs				✓	✓	✓	✓	✓
Team/Opp*Home					✓	✓	✓	✓
Stadium*%Black						✓	✓	✓
Team/Opp*Year							✓	✓
Model		OLS	OLS	OLS	OLS	OLS	OLS	IV

Notes: Sample=24,526 team-game observations. Each cell reports results from a separate regression.

(Standard errors in parentheses, clustered by game.) ***, **, and * denote statistically significant at 1%, 5% and 10%.

“Direct” effect refers to coefficient on %Black playing time * % white referees; “Indirect” effect refers to coefficient on Opponent %Black playing time * % white referees. The net effect is reported in the top row as the difference.

IV: %minutes played by blacks instrumented using average over previous ten games.

Unreported “observable” controls include home, attendance, home*attendance, game duration, out-of-contention, and black coach. All control variables included both for the team, and for their opponent. Similarly, “Team” and “Team*Year” and “Team*Home” fixed effects include “Opponent”, “Opponent*Year” and “Opponent*Home” fixed effects, respectively.

Table 7: Effects of Opposite-Race Referees on Game Outcomes

Panel A: Dependent Variable: <i>I(Home Team wins game)</i>						
<i>% White refs</i> * (<i>%Black^{home} - %Black^{away}</i>)	-0.139 (0.087)	-0.167* (0.086)	-0.230*** (0.085)	-0.218** (0.085)	-0.160* (0.084)	-0.226** (0.092)
<i>% White refs</i> * (<i>Black coach^{home} - Black coach^{away}</i>)				-0.045 (0.028)	-0.055** (0.028)	-0.052* (0.028)
<i>Adjusted R²</i>	0.002	0.038	0.091	0.091	0.184	0.184
Panel B: Dependent Variable: <i>Home Team's Winning Margin</i>						
<i>% White refs</i> * (<i>%Black^{home} - %Black^{away}</i>)	-4.020* (2.162)	-4.790** (2.115)	-6.210*** (2.069)	-5.5929*** (2.077)	-4.256** (2.007)	-6.574*** (2.196)
<i>% White refs</i> * (<i>Black coach^{home} - Black coach^{away}</i>)				-1.056 (0.684)	-0.850 (0.666)	-0.766 (0.667)
<i>Adjusted R²</i>	0.003	0.047	0.118	0.118	0.244	0.244
<i>Observable controls</i>		✓	✓	✓	✓	✓
<i>Home team fixed effects</i>			✓	✓	✓	✓
<i>Away team fixed effects</i>			✓	✓	✓	✓
<i>Stadium*(%Black^{home} - %Black^{away})</i>			✓	✓	✓	✓
<i>Home team * year fixed effects</i>					✓	✓
<i>Away team * year fixed effects</i>					✓	✓
<i>Model</i>	OLS	OLS	OLS	OLS	OLS	IV

Notes: Sample = 12,263 home game observations (IV regressions: n=12,247)

Each column in each panel represents a separate regression.

(Standard errors in parentheses.) ***, **, and * denote statistically significant at 1%, 5% and 10%.

Coefficient on *%White refs* * (*%Black^{home} - %Black^{away}*) measures own-race bias.

%Black measured as share of minutes played by black players.

Observable controls include home-away differences in: out-of-contention, home*attend and black coach.

IV: Instrumenting for *%Black^{home} - %Black^{away}* and its interaction with *%white referees*, using the average *%Black* for the home team over the preceding ten games less the average *%Black* for the away team over the preceding ten games, interacted with *%white referees*.

Figure 1


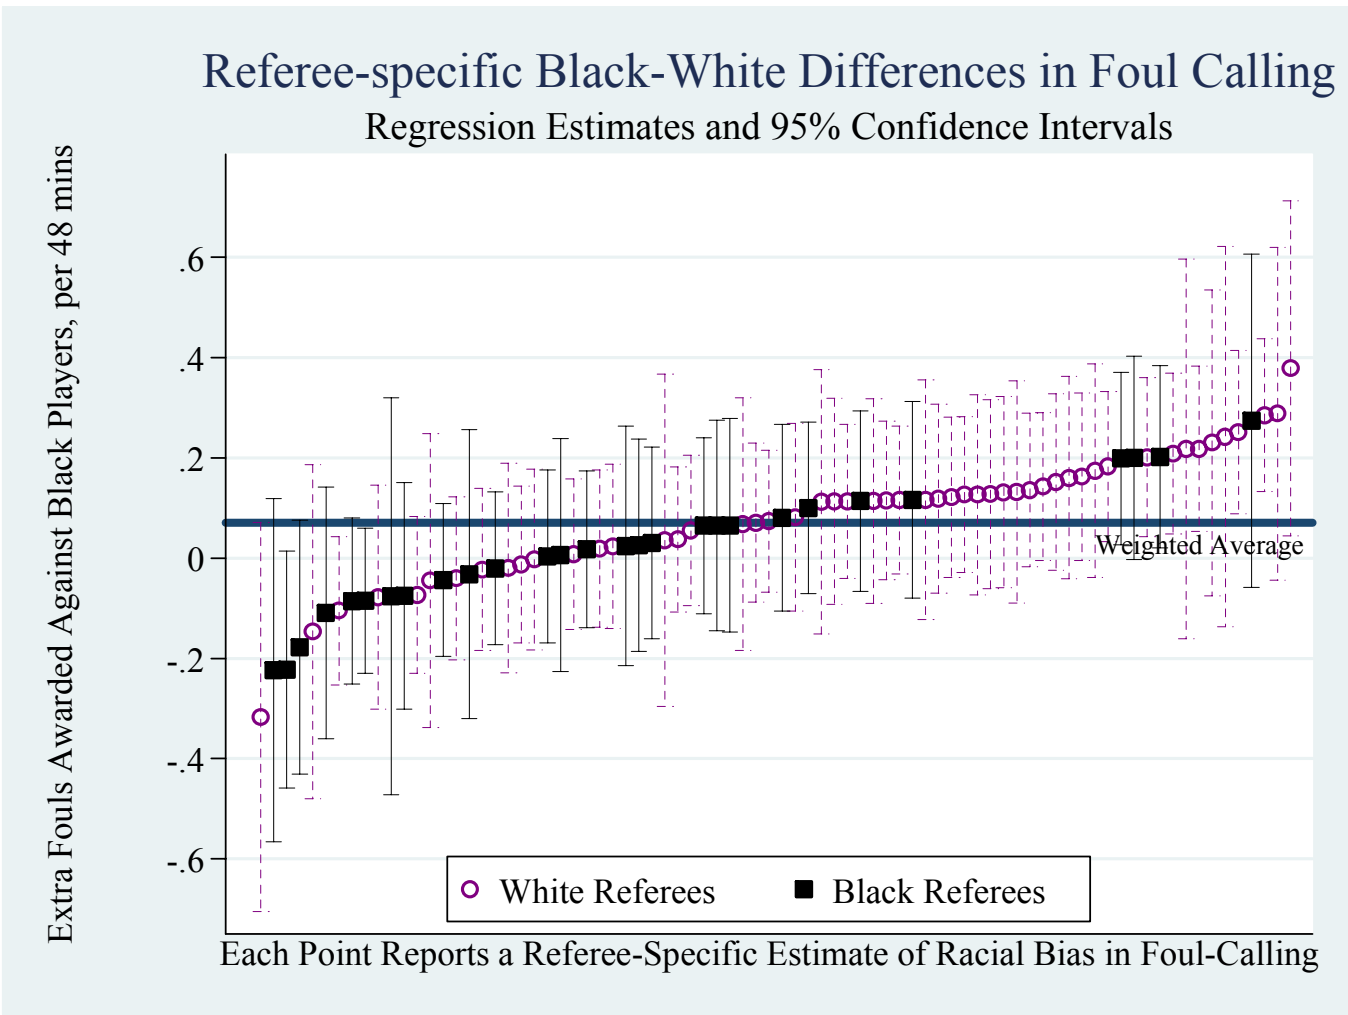


Figure 2: Distribution of Racial Bias, by Referee Race



Notes: Figure shows referee-by-referee estimates of black-white differences in fouls called per player per 48 minutes played. More specifically we run separate regressions for each referee, regressing the number of fouls called per 48 minutes for each player-game observation in which the referee participated, against an indicator variable for whether the offending player is black. These regressions control for player characteristics such as height, weight, age, experience, position, whether he is a starter, and sample-average playing statistics (assists, blocks, offensive and defensive rebounds, steals turnovers, fouls, all per 48 minutes played, and free-throw, two-point and three-point shooting percentage), and weight by minutes played. The figure only reports results for referees with at least 50 games in our dataset.

Appendix A: Further Randomization Tests

Dependent Variable: Number of White Referees in each game					
<i>Each cell reports p-values from F-tests of significance</i>					
<u>Independent Vars</u>	(1)	(2)	(3)	(4)	(5)
<i>Year fixed effects</i>	0.00	0.00	0.00	0.00	0.00
<i>#Black starters (home)</i>		0.72	0.69	0.98	0.81
<i>#Black starters (away)</i>		0.41	0.40	0.72	0.42
<i>Attendance</i>			0.22	0.50	0.83
<i>Out-of-contention (home)</i>			0.98	0.94	0.60
<i>Out-of-contention (away)</i>			0.70	0.80	0.97
<i>Home team FE</i>				0.48	0.97
<i>Away team FE</i>				0.97	0.69
<i>Home team * year FE</i>					0.99
<i>Away team * year FE</i>					0.99
<i>F-test: Variables not in prior column</i>		0.67	0.64	0.88	0.99
<i>F-test: All variables except year effects</i>			0.67	0.78	0.92
<i>Adj. R²</i>	0.0495	0.0494	0.0493	0.0483	0.0358

Notes: Sample includes 12,263 regular-season games.