

RACIAL DISCRIMINATION AMONG NBA REFEREES

JOSEPH PRICE AND JUSTIN WOLFERS

The NBA provides an intriguing place to assess discrimination: referees and players are involved in repeated interactions in a high-pressure setting, with referees making split-second decisions that might allow implicit racial biases to become evident. We find that more personal fouls are awarded against players when they are officiated by an opposite-race officiating crew than when they are officiated by an own-race refereeing crew. These biases are sufficiently large so that they affect the outcome of an appreciable number of games. Our results do not distinguish whether the bias stems from the actions of white or black referees.

I. INTRODUCTION

Does race influence our evaluation of others? We provide new evidence on racial biases in evaluation by examining how the number of fouls awarded against black and white National Basketball Association (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 are a particularly expert group, with substantial experience, continual feedback, and large incentives to be accurate. 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.”

NBA referees are effectively randomly assigned to each game. Moreover, the number of games played is large, so we can assess both a very clear baseline rate at which individual players commit fouls and a clear baseline for the number of fouls called by different referees. Against these baselines, we find systematic evidence of an own-race bias. Players earn up to 4% fewer fouls or score up

*This paper has benefited from helpful comments from David Berri, Mark Cuban, John Donohue, Scott Drewnianka, Ronald Ehrenberg, Todd Elder, Gary Fields, Joseph Gyourko, Kevin Hallock, Christine Jolls, Lawrence Kahn, Lawrence Katz, Alan Krueger, Lars Lefgren, David Levine, Janice Madden, Betsey Stevenson, Matthew White, and seminar participants at Chicago Law School, Cornell, Dartmouth, Delaware, Missouri, the National Academy of Science, Northwestern Law School, NYU, Stanford Law School, UCLA, and Wharton and the annual meetings of the American Economic Association, the American Law and Economics Association, the Conference on Empirical Legal Studies, the Society of Labor Economists, and the NBER Summer Institute. We are also grateful to 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. jpp34@cornell.edu.

© 2010 by the President and Fellows of Harvard College and the Massachusetts Institute of Technology.

The Quarterly Journal of Economics, November 2010

to 2¹/₂% more points when they are the recipients of a positive own-race bias, rather than a negative opposite-race effect.

We find similar results when aggregating to the team level, with the racial composition of the refereeing crew having an appreciable effect on the probability of a team winning. In an average game, one team plays around 15% fewer minutes with black players than their opponents. For this team, the chance of victory under an all-black refereeing crew versus an all-white crew differs by about three percentage points.

The simplest interpretation of our findings is that they reflect own-race bias, with either black or white referees (or both) favoring players of their own race, or disfavoring those of other races, though we are unable to make strong statements about which type of bias is occurring. Even so, we explore several other interpretations. Because our unit of analysis is the refereeing crew, we explore whether these findings can be explained by changes in crew-level dynamics, rather than simply reflecting individual referee biases. Alternatively, it may be that the interaction of referee and player race is relevant, not because it affects foul-calling, but because it affects player behavior. We also assess an omitted-variables interpretation in which players may be disadvantaged by opposite-race referees, but this may be the product of different playing styles of black versus white players interacting with different refereeing styles among black versus white referees.

Although we cannot take a strong stance on the mechanisms involved, the accumulated evidence is most consistent with our findings being driven by own-race bias. Comparing games with all-black and all-white refereeing crews yields findings consistent with the rest of the sample, suggesting that the relationship between foul-calling and the composition of refereeing crews is driven by individual referees favoring players of their own race. We examine a variety of player outcomes, finding little evidence of a rise in aggressive play that might explain the rise in the number of fouls called against players. Our findings are also robust to both aggregation to the team level and the inclusion of a wide range of controls, including rich controls for playing styles and their interaction with referee race.

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 (2009) 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 that allow for reasonably precise inferences.

Applying Beckerian taxonomy to our findings, this own-race preference falls under the banner of taste-based discrimination. Within this, customer-based discrimination is unlikely, as the own-race preference continues to exist even after we hold the stadium (and hence customer base) constant. Additionally, employer discrimination is inconsistent with the formal incentives for accuracy provided by the league. This suggests a referee-specific taste for discrimination. Although explicit animus is unlikely, Bertrand, Chugh, and Mullainathan (2005) describe an emerging literature on implicit discrimination that 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 Our findings may reflect these implicit associations varying with the race of the evaluator.

In addition, a large literature has documented substantial evidence of discrimination within sports (Kahn 1991). This setting has afforded useful insights largely because measures of productivity are easily observable. Although 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, whereas these tests for discrimination typically ask whether wages differ for blacks and whites conditional on observable game statistics, we demonstrate that observable game outcomes are

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; Schanzenbach 2005). A recent study by Abrams, Bertrand, and Mullainathan (2006) uses random assignment of judges to particular cases and finds 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, Lambert, and Jacoby (2002) note that the need to make quick judgments increases one's susceptibility to implicit stereotyping.

influenced by the racial mix of the 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).

II. BACKGROUND: BASKETBALL, THE NATIONAL BASKETBALL ASSOCIATION, AND REFEREES

In any season, the NBA has around sixty referees, with a crew of three referees officiating each game. Assignments of referees to crews are made 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 no referee is allowed to officiate more than nine games for any team, or referee twice in a city within a fourteen-day period. Although 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. For example, Table I shows that for each year in our sample, the number of white referees is unrelated to the number of black starters. Likewise, the Appendix shows that none of our variables have any power in explaining the assignment of referees of each race to particular games within each season.

Every game has an observer who meets with the referee for a pregame discussion, observes the game, and reviews video clips from the game with the referees afterward. These observers report to group supervisors, who provide additional input. The director of officiating also provides biweekly feedback to each referee on his or her performance. There is also an informal network of monitoring by coaches, spectators, sports analysts, and fans.

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 its 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 provide substantial additional compensation on

TABLE I
BLACK STARTERS PER TEAM AND THE DISTRIBUTION OF REFEREEING CREWS BY RACE

Season	Black starters per team				χ^2 test of independence ^a (<i>p</i> -value)
	0 white referees	1 white referee	2 white referees	3 white referees	
1991/1992	4.33	4.33	4.27	4.28	.82
1992/1993	4.20	4.20	4.26	4.25	.03
1993/1994	4.27	4.27	4.31	4.30	.80
1994/1995	4.20	4.27	4.29	4.25	.26
1995/1996	4.35	4.26	4.29	4.23	.60
1996/1997	4.11	4.17	4.19	4.17	.97
1997/1998	4.22	4.18	4.19	4.21	.98
1998/1999	4.05	4.13	4.10	4.14	.99
1999/2000	4.26	4.25	4.14	4.25	.07
2000/2001	4.15	4.19	4.22	4.18	.99
2001/2002	4.12	4.08	4.11	4.15	.82
2002/2003	4.16	4.20	4.11	4.20	.79
2003/2004	4.03	4.05	4.03	4.04	.12
Sample size (% of all player-games)	668 (2.7)	4,928 (20.1)	11,580 (47.2)	7,350 (30.0)	<i>n</i> = 24,526

Notes. Each observation is a team × game observation. Sample includes all regular season NBA games from 1991/1992–2003/2004, excluding referee strikes.

^aFinal column tests: H_0 : #white referees is independent of #black starters.

top of the referees’ base salary. Leading referees can earn several hundred thousand dollars per year.

III. PLAYER-LEVEL ANALYSIS

Our data contain box score information from all regular season NBA games played from the 1991–1992 season through to the 2003–2004 season, yielding over a quarter of a million player–game observations. For each player–game, we observe all performance statistics (points, blocks, steals, etc.), as well as minutes played and the number of personal fouls committed. The box score also lists the three referees officiating each game. Although we cannot observe the referee who blows the whistle for each foul, our empirical strategy involves comparing the number of fouls each player earns based on the racial mix of the referee crew.

We coded referees as black or nonblack based on visual inspection of press photographs of referees, supplemented by the able

assistance of a former NBA referee. Our data on player race come from a variety of sources, including Timmerman (2000), Kahn and Shah (2005), and our own coding from past issues of the Official NBA Register and images on 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 nonblacks somewhat imprecisely as “white.” We also draw information about each player’s characteristics (height, weight, and position) from basketballreference.com; characteristics of the game, including the home team and attendance, from the box score; and team characteristics, including the coach’s race, from the NBA Register. We construct a variable for whether a team was out of contention by calculating whether there were fewer games left in the season than the gap between that team’s victories and the record of the eighth best team in its conference. Some of our player-level controls also vary by game, such as whether players were named in the starting five, their age, their experience, and whether they were all-stars that season. Table II 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.

These summary statistics reveal that black players play more minutes per game than white players. Black players receive about the same number of fouls per game (2.55 vs. 2.53) as white players, and hence they 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

However, our focus is on own-race bias, which involves assessing how these differences vary as the racial composition

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

$$\text{Fouls per 48 mins}_{it} = 4.97 - 0.64 \times \text{black player}_i \quad \text{adj. } R^2 = .005 \quad n = 266,984$$

(.016) (.017)

Adding covariates yields

$$\begin{aligned} \text{Fouls rate}_{it} = & -0.017 \times \text{black player}_i + 1.47 \times \text{center}_i + 0.53 \times \text{forward}_i + 0.025 \\ & \quad (.017) \quad (.032) \quad (.021) \quad (.003) \\ & \times \text{height} + 0.010 \times \text{weight} + 0.053 \times \text{age} - 0.086 \times \text{experience}_{it} \\ & \quad (.0004) \quad (.005) \quad (.005) \\ & - 1.366 \times \text{starter} - 0.061 \quad \text{adj. } R^2 = .097 \quad n = 266,984 \\ & \quad (.013) \quad (.252) \end{aligned}$$

TABLE II
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)	3.46***
Fouls	2.55 (1.51)	2.53 (1.54)	0.02***
Points	13.24 (8.37)	11.07 (7.54)	2.16***
Player productivity: stats × 48/minutes played			
Fouls	4.33 (3.20)	4.97 (3.93)	-0.64***
Points	19.76 (10.05)	18.45 (10.11)	1.31***
Free throws made	3.86 (3.90)	3.52 (3.99)	0.34***
Free throws missed	1.33 (1.99)	1.11 (1.99)	0.22***
2 point goals made	6.59 (3.99)	5.96 (4.02)	0.62***
2 point goals missed	7.30 (4.24)	6.42 (4.36)	0.88***
3 point goals made	0.91 (1.63)	1.00 (1.78)	-0.09***
3 point goals missed	1.71 (2.36)	1.70 (2.50)	0.01
Offensive rebounds	2.52 (2.78)	2.70 (3.09)	-0.18***
Defensive rebounds	5.77 (4.10)	6.27 (4.42)	-0.50***
Assists	4.57 (4.08)	4.22 (4.30)	0.35***
Steals	1.66 (1.88)	1.48 (1.93)	0.18***
Blocks	1.00 (1.75)	1.17 (2.06)	-0.18***
Turnovers	2.97 (2.54)	2.83 (2.74)	0.14***
Game information			
Attendance (1,000s)	16.71 (3.69)	16.80 (3.62)	-0.09***
Out of contention	0.06 (0.24)	0.06 (0.24)	0.00

TABLE II
(CONTINUED)

	Black players	White players	Difference
	Mean (SD)	Mean (SD)	
Black coach	0.24 (0.43)	0.20 (0.40)	0.04***
	Player characteristics		
Age	27.90 (4.02)	28.00 (3.87)	-0.09
NBA experience (yrs)	6.19 (3.74)	5.78 (3.73)	0.41**
All Star this year	0.13 (0.34)	0.09 (0.29)	0.04***
Center	0.11 (0.32)	0.34 (0.47)	-0.22***
Forward	0.44 (0.50)	0.35 (0.48)	0.09*
Guard	0.45 (0.50)	0.31 (0.46)	0.13**
Starter	0.69 (0.46)	0.59 (0.49)	0.10***
Height (in.)	78.4 (3.62)	80.54 (4.14)	-2.13***
Weight (lb.)	211.5 (26.5)	223.2 (29.5)	-11.7***
	Referees		
0 white referees	0.03 (0.16)	0.03 (0.17)	-0.00
1 white referee	0.20 (0.40)	0.21 (0.41)	-0.00
2 white referees	0.47 (0.50)	0.47 (0.50)	0.00
3 white referees	0.29 (0.46)	0.29 (0.46)	0.00
# white referees	2.04 (0.78)	2.03 (0.78)	0.01
Sample size			Total
Players	889	301	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

***, **, and * Differences statistically significant at 1%, 5%, and 10%, respectively.

TABLE III
DIFFERENCES IN DIFFERENCES: FOUL RATE (=48 × FOULS/MINUTES PLAYED)

	Black players	White players	Difference: black–white foul rate	Slope: Δ(black–white)/ Δ%white refs
0% white refs (<i>n</i> = 7,359)	4.418 (0.043)	5.245 (0.094)	−0.827 (0.106)	
33% white ref (<i>n</i> = 54,537)	4.317 (0.016)	4.992 (0.035)	−0.675 (0.038)	0.455 (0.331)
67% white refs (<i>n</i> = 126,317)	4.335 (0.010)	4.989 (0.023)	−0.654 (0.025)	0.064 (0.137)
100% white refs (<i>n</i> = 78,771)	4.322 (0.013)	4.897 (0.029)	−0.574 (0.032)	0.240** (0.121)
Average slope: Δfouls/Δ%white refs	−0.022 (0.027)	−0.204*** (0.066)		Diff-in-diff 0.182*** (0.066) (<i>p</i> = .006)

Notes. Sample = 266,984 player–game observations, weighted by minutes played. Standard errors in parentheses.

***, **, and * Statistically significant at 1%, 5%, and 10%.

of the refereeing crew changes. Table III shows an illustrative differences-in-differences analysis. Reading down the columns illustrates the 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 number of fouls earned by black players is, on average, roughly the same whether the refereeing crew is predominantly white or black. In contrast, white players earn fewer fouls under predominantly white refereeing crews. As such, the “difference-in-difference” suggests that a player earns 0.18 fewer fouls per 48 minutes played when facing three referees of the same race than when facing three opposite-race referees.

This analysis reveals that the bias we document primarily affects white players.⁴ This is a departure from more standard accounts of discrimination that involve whites actively discriminating against blacks, although our setting is unusual in that black players are the majority group. In turn, this may reflect either white players being favored by white referees or disfavored

4. The Online Appendix includes regressions confirming that this finding is robust to including a broad set of control variables—although one cannot simultaneously explore this aspect of the result, and control for referee or game fixed effects.

by black referees, although our identification strategy (which relies on random assignment of refereeing crews) does not allow us to sort out which group of referees is responsible for this bias.

The richness of these 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. Consequently, in Table IV we report the results from estimating

$$\begin{aligned}
 (1) \text{ Foul rate}_{igrt} = & \beta_1 \% \text{white referees}_g \times \text{black player}_i \\
 & + \beta_2 \% \text{white referees}_g + \beta_3 \text{black player}_i \\
 & + \beta_4 \text{observable player}_i, \text{ game}_g, \text{player-game}_{ig}, \\
 & \text{team-game}_{ig}, \text{referee}_r \text{ characteristics} \\
 & + \text{player fixed effects}_i + \text{referee fixed effects}_r \\
 & + \text{season fixed effects}_g \\
 & [+ \text{observable controls}_{itg} \times \% \text{white referees}_g \\
 & + \text{black player}_i \times \text{stadium}_g \text{ effects} \\
 & + \text{player}_i \text{ effects} \times \text{year}_g \text{ effects} \\
 & + \text{game}_g \text{ effects} + \text{game}_g \text{ effects} \\
 & \times \text{team}_t \text{ effects}] + \varepsilon_{igrt},
 \end{aligned}$$

where the subscripts denote a player i playing for a team t in a specific game g officiated by referees r . The dependent variable is the number of fouls earned per 48 minutes, and all of our estimates weight player–game observations by the number of minutes played. The coefficient of interest is β_1 , which we interpret as the effect of opposite-race referees on a player's foul rate (relative to own-race referees), or the differential impact of the racial composition of the refereeing crew on black players relative to white players.

In the first column of Table IV, we control for time-varying player characteristics such as age, all-star status, and whether the player was a starter, and team-level variables such as whether the team is playing at home, attendance, whether they are out of contention, and whether the coach is black. These coefficients are reported in subsequent rows. We also control for player fixed effects (which account for both observable differences across players, including height, weight, and position, and unobservable differences), as well as referee fixed effects that measure the differential propensity of each referee to call more or less fouls. We also control for season fixed effects to account for the fact that

TABLE IV
EFFECTS OF OPPOSITE-RACE REFEREES ON FOUL RATES

Independent variables	Dependent variable: foul rate (= 48 × fouls/minutes) (mean = 4.43; SD = 3.34)		
	(1)	(2)	(3)
Black player × %white refs	0.197** (0.061)	0.203** (0.072)	0.181** (0.080)
Control variables			
Age	-0.728*** (0.047)	-0.729*** (0.049)	
All-Star	-0.383*** (0.026)	-0.429*** (0.063)	
Starting five	-0.988** (0.016)	-1.004** (0.040)	-0.775*** (0.044)
Home team	-0.125*** (0.012)	-0.213*** (0.033)	
Attendance (1,000s)	0.008*** (0.002)	0.004 (0.005)	
Out of contention	-0.127** (0.027)	-0.136* (0.071)	
Black coach	-0.107*** (0.017)	-0.080** (0.040)	
<i>R</i> ²	.18	.18	.28
Other controls			
Referee, year, and player fixed effects	✓	✓	✓
Player characteristics × %white refs		✓	✓
Full set of fixed effects			✓

Notes. Sample = 266,984 player–game observations, weighted by minutes played (standard errors in parentheses). Each column reports the results of a separate regression. All specifications control for the observable variables shown (and missing coefficients reflect the fact that some controls are unidentified in the presence of perfectly collinear fixed effects.) The second and third columns add further controls to account for a player’s on-court role, including height, weight, position, experience, and sample averages of 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 the percentage of free throw, two-point and three-point shots made, minutes played, and indicators for missing values. Each of these controls is also interacted with %White referees. The third column also includes a full set of player × year, home team × player race, and team × game fixed effects (including the relevant direct terms).

***, **, and * Statistically significant at 1%, 5%, and 10%.

the racial composition of the refereeing crew is idiosyncratic only within each season. These control variables are all highly significant, but nonetheless, the estimated own-race bias is similar to that estimated in Table III.

Although our player and referee fixed effects take account of the different styles of individual referees and the 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. The second column addresses this by including a series of controls for the share of white referees in a game, interacted with variables describing a player's on-court role. This set of controls that is interacted with %white referees (and also included as direct terms) includes not only all of the controls listed above, but also non-time varying player characteristics such height, weight, and position; we also use our sample data to construct measures describing each player's on-court role by taking sample averages of each of the statistics we track (assists, blocks, defensive rebounds, fouls, offensive rebounds, steals, turnovers, free throw attempts, two-point attempts, and three-point attempts—all measured per 48 minutes played—plus free-throw percentage, two-point percentage, and three-point percentage, minutes played, and indicators for missing values). Although the full set of these 29 interactions is jointly statistically significant (although not in the more complete specification in column (3)), their inclusion does not change our estimate of the extent of own-race bias. The Online Appendix shows these interactions, few of which are individually significant. Moreover, the interaction of %white referees with player race yields the largest partial and semipartial correlation coefficient of all of these interactions.

The final column augments this specification with a large number of fixed effects, which further controls for a range of competing explanations. This specification includes around 5,000 fixed effects for each player in each year, as well as home team \times player race effects that control for different race effects in each stadium. Importantly, we saturate the model, allowing for over 25,000 team \times game fixed effects (which subsume team \times home, team \times year, and team \times refereeing crew and many other effects). These controls ensure that these results are identified only by the differential propensity of teammates to earn extra fouls when the refereeing crew is not of their race. Across each of these specifications, we find that black players receive around 0.18–0.20 more fouls per 48 minutes played (or 4–4½%), relative to white players, when the number of white referees officiating a game increases from zero to three.

Our dependent variable in these regressions—fouls per 48 minutes—is appropriate if fouls are a linear function of playing time, which is unlikely given that the six-foul limit is less likely to

be a constraint for those playing only minor roles. In the extreme case, a player might be sent into a game with the express purpose of committing fouls in order to stop the clock in a close game. Thus, we ran several variants of our baseline regression, finding similar results when analyzing the foul rate only among starters; controlling for a quartic in minutes played; or estimating a count model that includes (log) minutes played as an independent variable. These results are reported in the Online Appendix.

Table V moves beyond fouls to analyze the consequences of opposite-race referees for a number of other measurable player outcomes. Specifically, we measure various box score statistics per 48 minutes played and reestimate equation (1) with that statistic as the dependent variable. Five main points are evident from this table. First, we find suggestive evidence of similar effects operating on flagrant and technical fouls. Although the point estimates are quite large relative to the rarity of these incidents, they are also quite imprecise, and only the effect on flagrant fouls is ever statistically significant. This imprecision reflects the fact that we have data on these two measures only for 1997/1998–2003/2004, whereas 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 often involve incidents when players dispute an on-court ruling.

Second, the propensity to “foul out” appears unaffected by the race of the refereeing crew, with the 4% rise in the foul rate partly countered by a 1%–2% decline in playing time. This suggests that team performance may also be affected by composition effects as opposite-race referees affect the distribution of playing time.

Third, important effects of own-race bias are evident throughout the box score. For instance, increasing the share of opposite-race referees leads to a decline in points scored and a rise in turnovers committed. The pattern of results across all of these box score measures—including results that are statistically insignificant—indicates that player performance appears to deteriorate at nearly every margin when officiated by a larger fraction of opposite-race referees. (Note that measured turnovers include offensive fouls.) Some outcomes may also reflect the role of the race of the potential “victim” rather than the “offender” in shaping foul calls. Specifically, these data yield suggestive evidence of a decline in free throw attempts under opposite-race referees, suggesting that defensive fouls are less likely to be

TABLE V
EFFECTS OF OPPOSITE-RACE REFEREES ON PLAYER PERFORMANCE
(MEASURED PER 48 MINUTES)

Dependent variable	Mean (SD)	Coefficient on black player \times % white referees		
		(1)	(2)	(3)
Personal fouls	4.44 (3.34)	0.197*** (0.061)	0.203*** (0.072)	0.181** (0.080)
Flagrant fouls	0.012 (0.17)	0.006 (0.005)	0.010* (0.006)	0.009 (0.006)
Technical fouls	0.08 (0.38)	0.007 (0.010)	0.016 (0.013)	0.015 (0.014)
Minutes	30.13 (10.1)	-0.408*** (0.136)	-0.503*** (0.160)	-0.403** (0.158)
Fouled out	0.025 (0.16)	-0.000 (0.003)	0.001 (0.004)	0.002 (0.004)
Points	19.54 (10.1)	-0.395** (0.176)	-0.300 (0.206)	-0.482** (0.226)
Free throw attempts	5.09 (4.90)	-0.102 (0.090)	-0.018 (0.106)	-0.041 (0.114)
Free throw %	0.75 (0.23)	0.002 (0.006)	0.000 (0.007)	0.001 (0.008)
Blocks	1.02 (1.81)	-0.057* (0.030)	-0.011 (0.036)	-0.009 (0.039)
Steals	1.63 (1.89)	-0.062* (0.036)	-0.067 (0.043)	-0.078* (0.047)
Turnovers	2.95 (2.57)	0.112** (0.050)	0.153*** (0.058)	0.121* (0.064)
Net effect (win score)	8.36 (9.09)	-0.528*** (0.170)	-0.599*** (0.199)	-0.509** (0.218)
Referee, year, and player fixed effects		✓	✓	✓
Player char. \times %white referees			✓	✓
Full set of fixed effects				✓

Notes. Each cell reports results from a separate regression. See notes to Table IV for specification details. Regressions analyzing shooting percentages are weighted by attempts, rather than minutes. $n = 266,984$, except flagrant and technical fouls $n = 136,509$ (available only 1997–2003).

called against one's opponents when opposite-race players have possession.

Fourth, the key exception to the general pattern of declining player performance under opposite-race referees is that 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, thus

servicing as a natural “placebo” measure. Unfortunately field goal percentage reflects whether the referee assigns blame for physical contact during the shot to the offense or defense, and hence is not a useful placebo.⁵

The final row analyzes a summary measure of a player’s contribution to his team’s winning margin,⁶ which suggests that own-race bias may lead an individual player’s contribution to his team’s winning margin to vary by up to half a point per game. Moreover, the finding that playing time is reduced suggests that there may be additional consequences due to substitutions.

IV. TEAM-LEVEL ANALYSIS

One shortcoming of our analysis of foul propensities in Table IV is that it only reflects the role of own-race bias in determining the guilt of an offender, whereas it may also shape whether a referee is sympathetic to a player as a victim. By aggregating to the team level, we can analyze both the number of fouls awarded against a team, and the number awarded to that team, and see how these vary with the racial composition of each team and the refereeing crew. The cost is that aggregating to the team level substantially reduces the available variation and leads to more imprecise estimates. Our key estimating equation is

(2)

$$\begin{aligned}
 \text{fouls}_{grto} = & \beta_1 \%white\ referees_g \times \%black\ minutes\ played_{gt} \\
 & + \beta_2 \%white\ referees_g \\
 & \times \text{opponent } \%black\ minutes\ played_{go} \\
 & + \beta_3 \%white\ referees_g + \beta_4 \%black\ minutes\ played_{gt} \\
 & + \beta_5 \text{opponent } \%black\ minutes\ played_{go} \\
 & + \beta_6 \text{observable game}_g, \text{ team-game}_{gt}, \text{ and opponent} \\
 & - \text{game}_{ot} \text{ characteristics} + \text{team}_t \text{ fixed effects} \\
 & + \text{opponent}_o \text{ fixed effects} + \text{referee}_r \text{ fixed effects} \\
 & + \text{season}_g \text{ fixed effects} + \text{observable controls} \\
 & \times \%white\ referees_g + \text{opponent observable controls}
 \end{aligned}$$

5. A score is recorded only if the shooter commits no fouls, whereas a miss is not recorded if he is awarded a foul.

6. Berri, Schmidt, and Brook (2006) call this index the “Win Score,” and calculate it as Win Score = (Points – Field goal attempts – 1/2 Free throw attempts) – Turnovers + Rebounds + Steals + 1/2 Blocks + 1/2 Assists – 1/2 Fouls. We analyze this productivity index per 48 minutes played.

$$\begin{aligned} & \times \% \text{white referees}_g + \% \text{black minutes played}_{gt} \\ & \times \text{stadium}_g \text{ effects} + \text{opponent } \% \text{black minutes played}_{go} \\ & \times \text{stadium}_g \text{ effects} + \text{team}_t \times \text{season}_g \text{ effects} + \text{opponent}_o \\ & \times \text{season}_g \text{ effects}] + \varepsilon_{grto}, \end{aligned}$$

where subscript g refers to a particular game, t a particular team, o their opponent, and r an individual referee. We report standard errors clustered at the game level.

The extent to which the fouls earned by a team are driven by their greater racial dissimilarity to the refereeing crew than their opponents' is measured by $\beta_1 - \beta_2$. Note that this estimate incorporates both the direct effect of the referee's propensity to call fouls based on the race of the offender (β_1) and the race of the victim (β_2). The net effect on the foul differential (fouls conceded—fouls awarded) is $\beta_1 - \beta_2$.

More generally, a shortcoming of the analysis in Table V is that it only analyzes the effects of refereeing decisions to the extent that they are captured in individual player box score data. Indeed, Oliver (2003) notes that a key problem with basketball statistics is that individual-level box score statistics paint a rich picture of a player's offensive production, but they do not reveal much about either his defensive contribution or general teamwork. Yet any useful contribution a player makes will be reflected in the scoring of his team or his opponents, and so we can capture these contributions by analyzing aggregate team performance. Consequently we also reestimate equation (2) but analyze points scored as the dependent variable.

This approach also yields an alternative interpretation that is particularly useful: changing a team's racial composition has a direct effect on the team's scoring, measured by the coefficient β_1 on $\% \text{white referees} \times \% \text{black minutes played}$. The same change in a team's racial composition also affects its opponents' expected scoring, and for the opponent, this effect is measured by β_2 , the coefficient on $\% \text{white referees} \times \% \text{opponent black minutes played}$. Thus, β_1 measures the effects of own-race bias on a team's offensive production, whereas β_2 measures the effects on defensive production, with $\beta_1 - \beta_2$ measuring the net effect on the winning margin.

Thus in Table VI we ask whether we see better team outcomes—fewer fouls committed, more fouls earned, more points scored, fewer points conceded, and more games won—when a

TABLE VI
EFFECTS OF OPPOSITE-RACE REFEREES ON TEAM PERFORMANCE

	Mean (SD)	% Black playing time × % white referees		
		(1)	(2)	(3)
A. Total fouls by team (mean = 22.4)				
Total effect	22.4	2.154**	1.899**	1.687
(β ₁ - β ₂)	(4.65)	(0.965)	(0.940)	(1.052)
Of which				
Direct effect (β ₁)		1.135	1.384*	1.192
(fouls committed)		(0.768)	(0.737)	(0.817)
Indirect effect (β ₂)		-1.019	-0.515	-0.495
(fouls awarded)		(0.793)	(0.762)	(0.845)
B. Points scored by team (mean = 98.4)				
Total effect	98.4	-5.733***	-3.836**	-6.185***
(β ₁ - β ₂)	(12.4)	(2.011)	(1.953)	(2.245)
Of which				
Direct effect (β ₁)		-2.073	-2.339	-3.202
(points scored)		(1.924)	(1.792)	(2.012)
Indirect effect (β ₂)		3.660*	1.496	2.983
(points conceded)		(1.914)	(1.800)	(2.013)
C. I (home team wins game)				
% white refs × (%black ^{home} - %black ^{away})		-0.195**	-0.160*	-0.226**
		(0.085)	(0.084)	(0.092)
% white refs × (black coach ^{home} - black coach ^{away})		-0.045	-0.055**	-0.052*
		(0.028)	(0.028)	(0.028)
Control variables				
Observables; year, referees, team, and opponent fixed effects		✓	✓	✓
Full set of fixed effects			✓	✓
Model		OLS	OLS	IV

Notes. Sample = 24,526 team-game observations in Panels A and B and 12,263 game observations in Panel C. Each cell reports results from a separate regression. (Standard errors in parentheses, clustered by game for top two panels.) "Direct" effect refers to coefficient on %black playing time × % white referees. "Indirect" effect refers to coefficient on opponent %black playing time × % white referees. The total effect is reported in the top row as the difference.

IV: The endogenous variables %black minutes played, Opponent %black minutes played, and the interaction of both variables with %white referees are instrumented using Average %black playing time over previous ten games calculated for both teams, and the interaction of each variable with %white referees. Unreported "observable" controls include home, attendance, number of overtimes, out-of-contention, and black coach, with separate control variables recorded for each team.

***, **, and * Statistically significant at 1%, 5%, and 10%.

larger fraction of minutes are played by players who are of the same race as the refereeing crew. Our initial specification includes observable controls such as whether each team is playing at home, is out of contention, and has a black coach, game attendance, and the number of overtimes played; this specification also includes controls for team, opponent, referee, and season fixed effects. The full specification also includes the interaction of the observable control variables with %white referees, as well as separate season effects for each team, and separate race effects for each stadium; in each case, each variable is defined for both the team and its opponent. The number of minutes played by black players may respond endogenously to the racial composition of the refereeing crew assigned to a particular game. Consequently, we also present instrumental variables results in which our variables of interest—the proportion of each team's minutes played by blacks, and that proportion interacted with the racial mix of the referees on that night—are instrumented with the average share of each team's minutes played by black players over that team's previous ten games, included both as direct terms and interacted with the racial mix of the referees on that night. Because team line-ups are persistent, these are very strong instruments.

For continuity with our earlier analysis, Table VI initially presents results on the number of fouls awarded against a team. Although the imprecision in these estimates cautions against a strong interpretation, we find that the estimated *direct* effect of own-race bias on the total number of fouls earned by a team is roughly five times larger than our estimates of the effect of own-race bias on the number of fouls earned by an individual player per 48 minutes. The indirect effect, due to the referee's racial similarity to a team's opponent, is also of a magnitude roughly similar to that of the direct effect, suggesting that the analysis of individual data understates the effects of own-race bias by up to one-half.

Naturally, basketball production is measured not in fouls but in points scored and conceded. Thus, the second set of results in Table VI focus on points scored. These estimates again point to a roughly equal role of own-race bias in shaping a team's offensive production as in shaping 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 its opponent.

The last rows in Table VI examine the effect of racial bias on whether a team wins. Because one team's win is its opponents' loss

and equation (2) controls symmetrically for the characteristics of each team, this specification is equivalent to a game fixed-effects specification or home-versus-away difference regression.⁷ For simplicity, we show this equivalent presentation, analyzing 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, controlling for home-away differences in the independent variables. These results show quite large and statistically significant impacts of the mismatch between the racial composition of the referees and the players.⁸ In addition, it is generally believed that coaches have some influence over the decisions of referees. The bottom row of Panel C provides suggestive evidence of bias against opposite-race coaches, with the magnitude of the coach effect being roughly equivalent to the effect of the race of a single player.

V. QUANTITATIVE INTERPRETATION

The results in Table VI suggest that own-race bias may be an important factor in determining game outcomes. Figure I 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 shows that difference in playing time by black players is correlated with winning margins. This 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 change, depending upon the racial compositions of the refereeing crew.

It is worth pausing to assess the quantitative importance of these results and their consistency with our earlier findings. In order to fix an initial scaling, note that the variable measuring racial mismatch between players and referees, $(\%Black^{home} - \%Black^{away}) \times \%White$ 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 to three

7. The home-away difference specification we show yields coefficient estimates that are exactly half those from estimating equation (2), or the game fixed-effects specification.

8. Although we report results from a linear probability model, a probit model yielded similar estimates. For example, whereas the linear probability model in the first column of Table VI yields a coefficient of -0.196 (with standard error 0.084), the equivalent probit specification yielded a marginal effect of -0.216 (standard error 0.091).

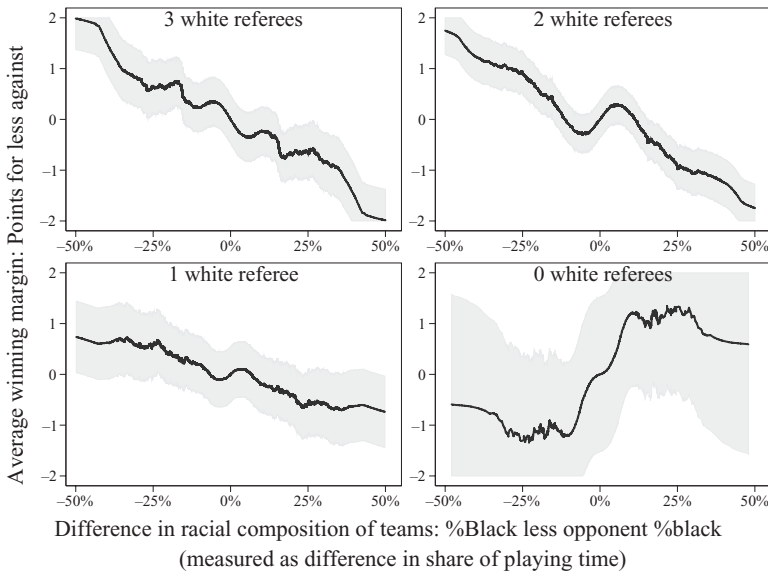


FIGURE I
Effects of Own-Race Bias on Winning Margins

Line shows running mean calculated using Epanechnikov kernel with bandwidth set to 0.4. Shading shows symmetric 95% confidence intervals (if within scale).

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 less black starter). For this team, the chances of victory under an all-black refereeing crew versus an all-white crew differ by around three percentage points ($= 0.196 \times 0.15$). Thus, changing the race of just one referee typically changes the chances of winning by around one percentage point.

Throughout our sample, the refereeing crew was, on average, 68% white, whereas the teams were 83% black (weighted 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 their games; our estimates suggest that a race-normed refereeing panel would raise number by 1.5 percentage points.⁹

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 to 2000:

$$\text{win percentage}_{\text{team,year}} = 0.21 + 0.29 \times (\text{team wage bill}_{\text{team,year}} / \text{league average wage bill}_{\text{year}}).$$

Interpreting this as a causal relationship suggests that a 1.5–percentage point rise in a team’s winning percentage could alternatively be achieved by raising the *aggregate* wage bill of an average team by $1.5\%/0.29 \approx 5\%$. In turn, consider the modal game in our sample: a team with five black starters playing four black starters and one white starter (which occurs in 33% of the games). The team with the one white starter could maintain its winning percentage under a shift to race-normed referees either by upgrading the quality of the team by spending an extra 5% on player salaries, or by simply exchanging the white starter for a similar quality black starter. This exercise suggests that the racial composition of the refereeing pool influences the market value of white versus black players.

The thought experiment also yields interesting player-level implications. Given that the large majority of players, on both the winning and losing sides, are black, race-norming the referee pool could change a lot of game outcomes but still yield only small effects on games won by black players (it would rise from 49.8% to 50.0%, as only a few more players would gain than lose). But the effects on white players would be more dramatic: in our sample, white starters win around 51.3% of their games; our estimates suggest that race-norming the refereeing crew would lower this winning percentage by 1.2 percentage points.

Although these estimates of the number of game outcomes determined by own-race bias may seem large, a simple example illustrates that they are consistent with the player-level analysis in Table V. Consider a game involving five black starters against four blacks and one white. Any team-level differences will be

9. 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 -0.195 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-white versus all-black crew. Further scaling by the magnitude of the proposed change in the proportion of white referees (17%–68%) yields 1.5 percentage points.

driven by the differential treatment of the fifth player, who is black for the home team and white for their rival. Using the coefficient on Berri, Schmidt, and Brook's (2006) "Win Score" metric in Table V, the black player's overall contribution to the team's winning margin will rise by about one-fourth of a point under a race-normed refereeing crew. These individual-level estimates are consistent with the estimates of the "direct" effects measured in Table VI, but that table also shows 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 and points conceded. Consequently race-norming the refereeing crew would, on average, change the winning margin by around half a point, which is what we found in the team-level analysis in Table VI.¹⁰ These apparently small impacts of own-race bias easily yield important effects on win percentages in a league in which around 6½% of games go to overtime, and around 4½% of game outcomes are determined by only one point. That is, when game outcomes are typically very close, even fairly small differences in player performance can yield large differences in how frequently each team wins.

VI. BEHAVIORAL INTERPRETATION

Thus far our analysis has established that player and team performance varies with the racial composition of the refereeing crew. Unfortunately, our framework is not well suited to sorting out whether these results are driven by the actions of black or white referees, because this would require establishing a "no-discrimination" baseline. Although we can control for enough observable features of the game so that perhaps our regression models may establish a reasonable "no-discrimination" benchmark, it is worth emphasizing that this involves substantially stronger assumptions than our earlier analysis.

To illustrate this, we analyze our data at the level of the referee. We use our player–game level data and collect all of the observations 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 earlier: height, weight, age,

10. To see this, multiply the points scored regression coefficient in Table VI ($\beta_1\beta - \beta_2 = 5.733$ points) 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.




FIGURE II
Distribution of Racial Bias, by Referee Race

Each point represents an estimate of the number of extra fouls per 48 minutes an individual referee calls on black versus white players; the bars represent the 95% confidence interval around these estimates. Specifically, we run separate regressions for each referee, regressing the number of fouls earned per 48 minutes for each player–game observation in which the referee participated, against an indicator variable for whether the offending player is black, controlling for year fixed effects and the full set of player, team–game, and player–game controls and career statistics listed in the notes to Table IV. All regressions are weighted by minutes played. The figure only reports results for referees with at least 100 games in our data set.

experience, all-star status, position, and sample averages of various box-score statistics (including their usual foul rate). Figure II plots this estimate for each referee of the degree to which he or she calls more or fewer fouls on blacks, showing those referees with at least 100 games in our sample.

This figure illustrates four important features of our analysis. First, the influence of player race on foul-calling is, on average, different for white and black referees, with each typically favoring players of their own race; the magnitude of the difference is consistent with the estimates reported in Tables III and IV. Second, there are no individual referees whose racial biases are particularly notable. (Although 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: the vast majority of black referees have a greater propensity to call fouls against white players than the majority of white referees. Indeed, despite the imprecision of each referee-specific estimate, only 9 of 28 black referees have an estimated pro-white bias stronger than the game-weighted average among white referees; similarly, only 15 of 52 white referees have a weaker pro-white bias than the game-weighted average among black referees. These findings suggest that statistically significant evidence of own-race bias persists, even when our analysis is aggregated to the level of each individual referee's record. Fourth, because these regressions are estimated separately for each referee, they control for referee-by-referee differences in refereeing "style."

The simplest interpretation of these results is an own-race bias on the part of referees. However, there are a few alternative explanations for our results. First, our results may come from players changing their behavior in response to the racial mix of the refereeing crew. Specifically, players would need to play more aggressively when officiated by more opposite-race referees. However, although fouls rise under opposite-race crews, Table V yields no evidence that other measures of aggression, such as steals or blocks, also rise. 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. This type of strategic response will lead to an attenuation bias, making it harder to discern any effects of own-race bias in the data.

Another possible explanation follows a variant of the usual "omitted-variables" interpretation of race differences. This alternative suggests that white and black referees have different focus areas on the floor, or are trying to penalize different types of behavior. The omitted variable in this interpretation is the differential propensity for white or black players to make these 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.

Some of these possibilities can be addressed by aggregating to the team level, as in Table VI. For instance, if certain on-court roles are typically filled by black players, and these roles are more harshly penalized by white referees than by black referees, this would yield a correlation between foul calls and player race in the individual data. However, aggregating to the team level 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 to a white player filling it. That is, the team-level regressions reflect the net impact of changing the racial composition of playing time, but eliminate variation due to which players have which roles. The fact that we find roughly consistent effects in our individual and team-level analyses speaks against this omitted-variables interpretation.

We also test the sensitivity of our results to various proxies for the omitted variable by attempting 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 rates of free-throw attempts, two-point attempts, three-point attempts, fouls, assists, steals, blocks, turnovers, and offensive and defensive rebounds earned per 48 minutes played, as well as free-throw, two-point, and three-point shooting percentages. Interestingly, these variables do successfully pinpoint an identifiably black playing style quite successfully—a probit model (not shown) attempting to predict a player's race from these "style" variables yielded a pseudo- R^2 of .35, and 11 of 21 variables are individually statistically significant at a 5% level. Even so, the addition of these variables to our main regressions (interacted with %white referees 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 (compare columns (1) and (2) of Table IV). Indeed, these player style \times %white referees control variables are jointly significant only in some specifications, but are insignificant when controlling for game \times team fixed effects.

A third explanation is that black and white referees differ along a number of dimensions (experience, age, birthplace, etc.) and it is these differences, rather than race, that explain our results. For 82% of the games in our sample, we know the NBA referee experience of all three officials. When we include the average experience of the crew interacted with the player's race as an additional control in our model, the coefficient is both small and insignificant, and its inclusion has almost no effect on our estimated own race bias. In addition, for 24% of the games in our sample, we also know the age of all three referees and how many of them were born in the South. We interact these additional crew-level measures (along with average experience) with each player's race and again find that the coefficients on these additional referee characteristics are small and insignificant, and do not have a large effect on our estimate of the own-race bias. Full details of these regressions are provided in the Online Appendix.

Our analysis largely proceeds at the player–game level, and so contrasts the behavior of different refereeing crews, rather than individual referees. Although this is appropriate in the context of arbitrary assignment of refereeing *crews* to games, it admits the possibility that our findings reflect social interactions within refereeing crews. That is, perhaps the relative disadvantages conferred by an increasingly opposite-race refereeing crew reflect referees exhibiting less own-race bias in the presence of referees not of their race. In order to isolate the direct influence of individual referees exhibiting own-race bias from these social interactions, we reran our analysis of the foul data, focusing only on the contrast between games refereed by all-black or all-white crews. Comparing the first and fourth rows of Table III gives a sense of this analysis, but a more complete analysis—available in the Online Appendix—shows that even in this restricted set of games we obtain statistically significant and quantitatively similar estimates of own-race bias. An alternative regression controls for these crew composition effects by including dummies for both mixed race crews, and their interaction with player race; this also yields similar results to our central findings in Table IV.

As additional support for our main findings, two recent papers provide evidence of own-race bias of officials in baseball. In research stimulated by an early draft of this paper, Parsons et al. (2008) find that a strike is more likely to be called when the pitcher and umpire are the same race, and Chen (2007) finds that white umpires provide a larger strike zone to white pitchers and a smaller strike zone to white batters. These papers further demonstrate that this own-race bias is influenced by the amount of monitoring that is in place. In both cases, the own-race bias completely disappears in stadiums with a QuesTec system (devices that provide nearly perfect monitoring of the umpire's decisions about whether the pitch was a strike). In addition, Chen (2007) finds that the own-race bias on the part of a white home-plate umpire is reduced when the umpire works with a racially diverse crew of officials.

VII. CONCLUSIONS

Using a unique data set 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 their race. The richness of our data allows us to control for a host of relevant factors that influence the number of fouls called and thereby to focus specifically on the

racial interaction between players and referees. We find that players have up to 4% fewer fouls called against them 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 so 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. Although 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, although 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.

APPENDIX: FURTHER RANDOMIZATION TESTS

Independent variables	Dependent variable: number of white referees in each game (each cell reports <i>p</i> -values from <i>F</i> -tests of significance)				
	(1)	(2)	(3)	(4)	(5)
Year fixed effects	.00	.00	.00	.00	n.a.
#black starters (home)		.57	.653	.75	.87
#black starters (away)		.41	.40	.72	.42
Attendance			.21	.49	.83
Out of contention (home)			.98	.94	.60
Out of contention (away)			.70	.81	.97
Home team fixed effects				.48	.97
Away team fixed effects				.97	.71
Home team × year fixed effects					1.00
Away team × year fixed effects					1.00
<i>F</i> -test: variables not in prior column		.61	.63	.89	1.00
<i>F</i> -test: all variables except year effects		.61	.74	.92	1.00
Adj. <i>R</i> ²	.0495	.0494	.0493	.0483	.0358

Note. Sample includes 12,263 regular-season games.

BRIGHAM YOUNG UNIVERSITY AND NATIONAL BUREAU OF ECONOMIC RESEARCH
WHARTON, UNIVERSITY OF PENNSYLVANIA, BROOKINGS, CENTRE FOR ECONOMIC
AND POLICY RESEARCH, CENTER FOR ECONOMIC STUDIES AND IFO INSTITUTE FOR
ECONOMIC RESEARCH, INSTITUTE FOR THE STUDY OF LABOR, AND NATIONAL BUREAU
OF ECONOMIC RESEARCH

REFERENCES

- Abrams, David, Marianne Bertrand, and Sendhil Mullainathan, "Do Judges Vary in their Treatment of Race?" University of Chicago Mimeo, 2006.
- Antonovics, Kate, and Brian Knight, "A New Look at Racial Profiling: Evidence from the Boston Police Department," *Review of Economics and Statistics*, 91 (2009), 163–177.
- Berri, David, Martin B. Schmidt, and Stacey L. Brook, *The Wages of Wins: Taking Measure of the Many Myths in Modern Sport* (Stanford, CA: Stanford University Press, 2006).
- Bertrand, Marianne, Dolly Chugh, and Sendhil Mullainathan, "Implicit Discrimination," *American Economic Review*, 95 (2005), 94–98.
- Bodvarsson, Orn, and Raymond Brastow, "A Test of Employer Discrimination in the NBA," *Contemporary Economic Policy*, 17 (1999), 243–255.
- Bushway, Shawn, and Ann Piehl, "Judging Judicial Discretion: Legal Factors and Racial Discrimination in Sentencing," *Law and Society Review*, 35 (2001), 733–764.
- Chen, Jowei, "What Does Baseball Teach Us about Reducing Racial Discrimination? Evidence from Two Natural Experiments," Stanford University Mimeo, 2007.
- Donohue, John, and Steven Levitt, "The Impact of Race on Policing and Arrests," *Journal of Law and Economics*, 44 (2001), 367–394.
- Greenwald, Anthony, and Mahzarin Banaji, "Implicit Social Cognition: Attitudes, Self-Esteem, and Stereotypes," *Psychological Review*, 102 (1995), 4–27.
- Hamilton, Barton, "Racial Discrimination and Professional Basketball Salaries in the 1990s," *Applied Economics*, 29 (1997), 287–296.
- Kahn, Lawrence, "Discrimination in Professional Sports: A Survey of the Literature," *Industrial and Labor Relations Review*, 44 (1991), 395–418.
- Kahn, Lawrence, and Malav Shah, "Race, Compensation and Contract Length in the NBA: 2001–2," *Industrial Relations*, 44 (2005), 444–462.
- Kahn, Lawrence, and Peter Sherer, "Racial Differences in Professional Basketball Players' Compensation," *Journal of Labor Economics*, 6 (1988), 40–61.
- Koch, James, and C. Warren Vander Hill, "Is There Discrimination in the 'Black Man's Game'?" *Social Science Quarterly*, 69 (1988), 83–94.
- Oliver, Dean, *Basketball on Paper: Rules and Tools for Performance Analysis* (Washington, DC: Potomac Books, 2003).
- Parsons, Christopher, Johan Sulaeman, Michael Yates, and Daniel Hamermesh, "Strike Three: Umpires' Demand for Discrimination," IZA Discussion Paper No. 3899, 2008.
- Payne, Keith, Alan Lambert, and Larry Jacoby, "Best Laid Plans: Effects of Goals on Accessibility Bias and Cognitive Control in Race-Based Misperceptions of Weapons," *Journal of Experimental Social Psychology*, 38 (2002), 384–396.
- Schanzenbach, Max, "Racial and Sex Disparities in Prison Sentences: The Effect of District-Level Judicial Demographics," *Journal of Legal Studies*, 34 (2005), 57–92.
- Spohn, Cassia, "The Sentencing Decisions of Black and White Judges: Expected and Unexpected Similarities," *Law and Society Review*, 24 (1990), 1197–1216.
- Stauffer, Joseph M., and M. Ronald Buckley, "The Existence and Nature of Racial Bias in Supervisory Ratings," *Journal of Applied Psychology*, 90 (2005), 586–591.

- Stoll, Michael, Steven Raphael, and Harry Holzer, "Black Job Applicants and the Hiring Officer's Race," *Industrial and Labor Relations Review*, 57 (2004), 267–287.
- Szymanski, Stefan, "The Economic Design of Sporting Contests," *Journal of Economic Literature*, 41 (2003), 1137–1187.
- Timmerman, Thomas, "Racial Diversity, Age Diversity, Interdependence, and Team Performance," *Small Group Research*, 31 (2000), 592–606.
- Welch, Susan, Michael Combs, and John Gruhl, "Do Black Judges Make a Difference?" *American Journal of Political Science*, 32 (1988), 126–136.