

Journal of Quantitative Analysis in Sports

Volume 4, Issue 2

2008

Article 7

Racial Bias in the NBA: Implications in Betting Markets

Tim Larsen* Joe Price[†]

Justin Wolfers[‡]

*Brigham Young University - Utah, tlarsen@byu.net

[†]Brigham Young University - Utah, joe_price@byu.edu

[‡]The Wharton School, University of Pennsylvania/ CEPR/ IZA/ NBER,
jwolfers@wharton.upenn.edu

Copyright ©2008 The Berkeley Electronic Press. All rights reserved.

Racial Bias in the NBA: Implications in Betting Markets*

Tim Larsen, Joe Price, and Justin Wolfers

Abstract

Recent studies have documented the existence of an own-race bias on the part of sports officials. In this paper we explore the implications of these biases on betting markets. We use data from the 1991/92 - 2004/05 NBA regular seasons to show that a betting strategy exploiting own-race biases by referees would systematically beat the spread.

KEYWORDS: betting markets, basketball, own-race bias

*The authors would like to thank the Wharton Sports Business Initiative for generous research support.

I. Introduction

In a recent study, Price and Wolfers (2007) find that referees in the NBA call relatively fewer fouls on players of their own race and that this own-race bias is sufficiently large to affect the outcome of an appreciable number of games. In subsequent research, Parsons et al. (2007) provide similar evidence of own-race bias among baseball umpires, finding that the propensity for an umpire to call strikes rather than balls is higher if he is of the same race as the pitcher.

In this paper, we extend this earlier research by analyzing whether own-race bias also changes the outcomes of point-spread bets on NBA basketball. That is, we analyze data from the 1991/92 to 2004/05 NBA seasons to test whether using information about the racial mix of the players on each team and the racial mix of referees could be used to systematically beat the spread.

We find that in games where the majority of the officials are white, betting on the team expected to have more minutes played by white players always leads to more than a 50% chance of beating the spread. The probability of beating the spread increases as the racial gap between the two teams widens such that, in games with three white referees, a team whose fraction of minutes played by white players is more than 30 percentage points greater than their opponent will beat the spread 57% of the time.

These results yield three interesting implications. First, tests of this nature are typically reported as tests of market efficiency. In this case we learn that the betting market failed to appropriately price the implications of own-race bias on game outcomes. Indeed, whereas Price and Wolfers (2007) find that game outcomes are systematically affected by the interaction of the racial composition of players and referees, we do not find any relationship between these factors and the betting spread. This is a useful contrast to a long literature typically finding that betting markets are extremely efficient aggregators of information (Sauer 1998). Second, given that point spreads control for nearly every other observable factor influencing the outcome of the game, the fact that information about the racial mix of the team and referees systematically leads to a higher chance of beating the spread provides additional evidence of own-race bias on the part of NBA referees. Whereas Price and Wolfers (2007) had relied initially on random assignment and subsequently on regressions with rich controls in order to account for the way in which other factors may influence game outcomes, in this paper we use the betting spread as a summary statistic of the influence of all of these other factors. Third, this paper yields direct evidence that variation across refereeing crews is sufficiently large that advance knowledge of the refereeing crew could be valuable to gamblers.

This is perhaps the most interesting result of our work—the effects of own-race bias on game outcomes are large enough that our betting strategy would not only win a majority of bets, but also win large returns when betting on games with strong differences in the teams' racial composition. It is not clear whether this would continue in the future since Price and Wolfers' (2007) previous study received considerable media attention in May 2007, but winning large amounts of money in our simulated bets further strengthens the evidence for the existence of an underlying racial bias in NBA officiating.

II. Background: The NBA Betting Market and Betting Data

Wagering on NBA basketball games typically involves spread betting, in which one bets \$11 to win \$10 (for a return of \$21) if a team wins the game, relative to a spread. For example, when the Utah Jazz visited the Philadelphia 76ers on December 20, 2004, the spread was Philadelphia $-6\frac{1}{2}$, meaning that a bet on the Jazz would win if either they won the game or if they lost by 6 or fewer points. Thus, the point spread can be considered a market-based forecast of the outcome of a game.

We have collected the betting line (or spread) on all regular season games running from 1991/92 through the 2004/05 seasons from www.covers.com. It is worth emphasizing that these are closing lines and hence represent the set of bets available to gamblers subsequent to the announcement of the refereeing crew. (These announcements occurred 90 minutes prior to tip-off during our sample but were recently changed to the morning of the game. A 90 minute window may seem like a brief time to make bets using information on the officials, but bettors have long kept statistics on several aspects of referees' performances and have used these to make bets after the officiating crews are announced.)

As Figure 1 shows, point spreads have historically been an unbiased forecast of likely game outcomes. Point spreads are in fact so accurate that they have been shown to have greater predictive power than sports experts and economists' models (Song et al., 2007), and it is somewhat rare to find any way to systematically exploit them for a profit (Woodland & Woodland 1994; Woodland & Woodland 2001.).

However, market efficiency is a more demanding standard. One definition of efficiency requires the absence of profit opportunities for gamblers, given available information. Since one must bet \$11 to win \$10, any betting strategy that predicts the winner (against the spread) more than 52.38% of the time amounts to a potentially profitable trading strategy and hence shows evidence of market inefficiency. A stricter efficiency standard also considers the potential profit opportunities available to bookmakers. Under this "strict efficiency" standard, the expected return to accepting bets against any team must be equalized, or else a bookmaker in a competitive market would find it profitable to deviate from the existing spread. Thus, strict efficiency requires that the probability that a team beats the spread is 50%, and more importantly, that it is orthogonal to existing information.

At this point it is worth noting that our definition of efficiency requires that the point spread be a median forecast (deviations above are as likely as deviations below), and not that it forecasts average game outcomes, although the latter is certainly plausible if forecast errors are symmetric. And indeed, Figure 2 shows that the distribution of forecast errors is symmetric and approximately normal.

Our specific interest is in testing whether the interaction of information about the racial composition of teams and referees helps predict whether a team covers the spread. Our data on the racial composition of both players and referees are drawn from Price and Wolfers (2007) who provide a more complete discussion of the data collection procedures. Our racial coding is somewhat coarse—we categorize players and referees as black or non-black (and somewhat imprecisely refer to the latter group as "white"). Thus each game is categorized according to whether it involves 0, 1, 2 or 3 white

referees. The average number of white referees in each game is 2.01. 28.1% of games have three white referees and 3.0% of games have no white referees.

Characterizing the racial composition of a team is somewhat more difficult. We want to be careful to use a metric that is observable prior to tip-off so that it could form the basis of a feasible betting strategy. Thus, for each team, we use information from the last five games to construct a measure of the fraction of minutes that were played by black players. (The first five games of each season were dropped from the dataset since values could not be created for those contests.) This proxy measure is highly correlated with the proportion of minutes played by blacks in the current game ($\rho=0.92$). Using this measure, the difference in the racial mix between the two teams for the games in our sample can be as high as 71 percentage points with a median of 12 percentage points and a 90th percentile of 31 percentage points.

Our original sample includes 15,250 games. We exclude 383 games for which we do not have information on the spread and another 302 for which the final score difference was the same as the spread (called a push, which leads to a cancellation of all of the bets). We also excluded 69 games for which we did not have information on the team's racial composition over the last 5 games. These exclusions leave us with a final sample of 14,496 games.

III. Analysis

We begin by showing some very simple evidence consistent with the efficiency of these betting markets. Table 1 shows that home teams beat the spread 49.43 percent of the time, which is statistically indistinguishable from 50 percent. The rate of return earned from always betting on home teams (at odds of 10/11) is -5.63% . The favorite beat the spread in 50.01% of all games, yielding a rate of return of -4.52% . Further disaggregating these data into home favorites, home underdogs, away favorites, and away underdogs reveals that in all four cases the probability of beating the spread is statistically indistinguishable from 50%. Thus, it appears that the betting market prices these factors appropriately. Many other factors were also tested for their accuracy in predicting winning bets and none yielded results that varied from 50%. Other betting literature concurs with these findings that nearly all observable factors are accounted for in making the point spreads and it is thus difficult to find simple betting strategies that would win more than 50% of the time.

Each of the rows in Table 1 can be re-cast as simple regression equations. This will be important as we assess the role of a player's race—a continuous variable. In Table 2, we estimate a linear probability model, attempting to forecast whether the home team beats the spread, as a function of the difference in the racial composition between the two teams (measured using the average over the last five games played by each team):

$$I(\text{Home team beats spread}) = \alpha + \beta * (\%Black^{home} - \%Black^{away}) \quad [1]$$

The constant term, α , captures any home-team bias in these markets, while β captures whether the relative racial composition of the two teams impacts a team covering the spread ($\%Black$ is the proportion of playing time that a team has given to black players over the preceding five games). We estimate this regression equation separately for games according to the racial composition of the refereeing crew. If this were an efficient market then we expect to see $\alpha=0.5$ and $\beta=0$ in each sub-sample. Our specific interest in own-race bias leads us to ask whether the refereeing crew mediates the relationship between player race and outcomes and hence whether β varies with the racial composition of the refereeing crew.

The results in Panel A of Table 2 yield some striking violations of the efficient markets hypothesis, suggesting that the racial composition of the team has a significant impact on whether a team beats the spread. In particular, if the home team is represented by more black players than are their opponents, then they are much less likely to cover the spread under an all-white refereeing crew. In contrast, the impact of player race on outcomes is much more muted with mixed-race crews (it is insignificant for crews involving two black referees and one white). The final column shows results for all-black crews, in which case the team represented by more black players is more likely to cover the spread, although it should be emphasized that this is a particularly small sample of games, and these estimates are quite imprecise. Probit models yielded virtually identical results.

Figure 3 shows the variation underlying these findings more directly. The upper left panel shows that under an all-white refereeing crew, the chances that the home team covers the spread declines when greater playing time is expected to be given to black players (relative to their opponent). As the number of white referees declines, this pattern becomes substantially more muted; under all-black refereeing crews, giving greater playing time to black players makes a team more likely to beat the spread.

In order to further probe these results, Figure 4 turns to evaluating the margin by which teams beat the spread, analyzing the home team's winning margin relative to the market-based forecast. The advantage of this measure is that it exploits information not just on whether a team beats the spread, but also by how much, yielding more precise estimates. The disadvantage is that stronger assumptions are required (symmetric forecast errors) before one can interpret these as tests of market efficiency. Nonetheless, Figure 4 shows qualitatively similar patterns to Figure 3, a point reinforced by the formal analysis in Panel B of Table 2.

At this point we have established that the race of the referees is a key mediating variable in the relationship between the team that covers the spread and the racial gap between the two teams. That is, the slope in each panel of Figure 3 and Figure 4 differs according to the racial composition of the refereeing crew. In turn, these outcomes are a function of the team's actual winning margin (the focus of Price and Wolfers (2007)) and the betting spread. Thus in Figure 5 and Figure 6, we isolate the separate influences of the winning margin and the betting spread. Figure 5 shows that the relationship between the racial composition of competing teams and the winning margin is strongly affected by the racial composition of the refereeing crew. However, while Figure 6 shows that the betting line varies with the racial composition of the two teams, this relationship is *invariant* to the the racial composition of the refereeing crew. That is, we can infer from

this analysis that the spread is set as if there is no own-race bias, even though game outcomes do in fact reflect an own-race bias. The formal analysis in Panels C and D of Table 2 also confirms this result.

Thus far our analysis has proceeded by examining games involving 0, 1, 2 or 3 white referees separately. A more direct way to test for own-race bias involves pooling these data and running:

$$I(\text{Home team beats spread}) = \alpha + \beta * (\%Black^{\text{home}} - \%Black^{\text{away}}) + \gamma * \%White\ referees + \delta * (\%Black^{\text{home}} - \%Black^{\text{away}}) * \%White\ referees \quad [2]$$

As before, α measures the home team bias, γ measures whether this home team bias varies with the racial composition of the refereeing crew (at least for games in which $\%Black^{\text{home}} - \%Black^{\text{away}}$ is zero, which is also approximately its sample mean), and β measures the baseline impact of the racial composition of the two teams—for games involving zero white referees. Our coefficient of interest, δ , measures the extent to which the relationship between betting outcomes and the racial composition of the teams is affected by the racial composition of the refereeing crew. This specification allows us to directly test this interaction effect, but now also imposes a linear effect of $\%white\ referees$.

Our findings, reported in Table 3, largely confirm the earlier analysis. In particular, the interaction of referee and player race has a statistically significant impact on whether a team covers the spread. To interpret the coefficients, consider a roughly typical game, in which one team is expected to deploy black players for about 15% more time than their opponent. This team is $0.161 * 0.15 \approx 2\frac{1}{2}$ percent more likely to cover the spread under an all-black refereeing crew than under an all-white crew. Alternately phrased, were one to swap out one white referee for one black referee in a typical game, this would change the chances of each team covering the spread by nearly one percentage point. The results in column two—which analyzes the winning margin relative to the spread—are roughly consistent, suggesting that changing from an all-white to an all-black crew changes the expected margin by $3.292 * 0.15 \approx 0.5$ points. In around 8% of games the result is within one point of the spread, so these small effects could still have impacts in betting markets. Further, it is worth noting that these estimated effects on whether a team covers the spread and by how much are quite similar to results reported in Table 7 of Price and Wolfers (2007), who analyzed whether a team won the game and by how much.

The third column of Table 3 analyzes the impact of the own-race bias on the team's outright winning margin, finding a significant effect of a roughly similar magnitude to that estimated for the margin relative to the spread. The final column suggests that the betting spread does not systematically respond to the interaction of player and referee race. All told, these results suggest that the own-race bias documented in column 3 is not priced by the betting market (see column 4), yielding significant evidence that the interaction of the racial composition of player and referees yields potential betting opportunities (in columns 1 and 2).

We now turn to testing whether we can use these results to implement a profitable betting strategy. The simplest betting advice that comes from this research is

to bet on the team that shares the greatest racial similarity with the refereeing crew. The first row of Table 4 reports on the proportion of winning bets such a strategy yields—an overall success rate of 51.01%. Equally, one might expect that the own-race bias we have identified would have a larger impact the more extreme the racial composition of the refereeing crew is, and indeed, this strategy correctly picks 51.37% of games involving all-white crews, and 52.53% of games involving all-black crews. Similarly, one might expect that own-race bias will have a larger impact the greater the difference in the racial composition of the two teams. Table 4 reports the winning percentages when restricting attention to sub-samples of games involving progressively sharper racial contrasts, and we find that the subset of games involving teams which are more racially different tend to yield stronger results.

These data appear to suggest that a profitable betting strategy may exist and that such a strategy would emphasize games involving two teams of quite different racial composition, refereed by either all-white or all-black crews. In Table 5, we assess a couple of fairly simple betting strategies that one might implement based on our analysis. We make assessments against the two alternative notions of market efficiency described previously: whether a betting strategy allows one to predict more than 50% of all games correctly, and whether a strategy yields a positive rate of return (measured here as returns per \$1 bet, taking account of the 10/11 odds usually offered). Given these hypotheses, we implement one-sided hypothesis tests.

Our initial strategy uses the regression equation presented in Panel A of Table 2. In the first case we simply bet on whichever team yields the greatest forecast of beating the spread (ensuring one bet per game), while the second case is more selective, only betting on those games where the chances of beating the spread are forecast to be greater than the 11/21 needed to break even. The former predicts significantly more than 50% of outcomes against the spread (51.35%), but not enough to offer a profit opportunity. The latter yields a more impressive win percentage (53.79%), and a small (albeit statistically insignificant) profit. Betting \$1 on each of the 2,549 games would have returned \$2617.4, for a profit of 2.68%. Given the greater precision available from our earlier analysis of a team's winning margin (relative to the spread), our next strategies exploited the predicted values from Panel B of Table 2. Betting on those games where this prediction margin is greater than half a point yields a small and statistically insignificant profit. Confining our attention only to those games where this formula suggested a one point advantage yields only 661 betting opportunities, but an impressive 55.67% win rate, and a statistically significant 6.29% profit.

These approaches may seem somewhat opaque, and so we subsequently present a more intuitive set of betting strategies, confining attention to the 4,493 games involving all-black or all-white refereeing crews and suggesting simply that one bets on the team with greater racial similarity to the refereeing crew. Overall this yields a 51.48% success rate, which is not quite sufficient to be profitable. In subsequent rows we further focus on those games where the racial contrast between the two teams is starkest, and find that successively tighter rules yield both fewer bets and increasingly higher (and significantly positive) rates of return—when betting on games in which one team averages two more players on the floor at a time that are the same race as the referees (a difference in black players of greater than 40 percentage points), returns would have been over 18% per bet.

IV. Conclusion

This paper assesses the role of own-race bias by NBA officials in shaping outcomes in betting markets. We show that exploiting information on the racial mix of the players and referees would have allowed us to systematically beat the spread. We also show that by restricting one's bets to those games where the bias is likely to be the largest, sizable profits are possible. This finding is notable in that it systematically beats the spread using only information on the race of the players and the officials. It is worth contrasting this with the existing literature showing that simple strategies are almost never profitable, and if they are they generally yield small returns (Sauer 1998).

Equally, our approach is somewhat unique in that the existing literature assessing either sports betting in particular or behavioral finance more generally, has rarely focused on the role that discrimination may play in yielding systematically mistaken beliefs and hence profit opportunities. Wolfers (2006), analyzing the role of CEO gender on stock prices, is another example, albeit from the financial domain.

It is not clear whether the betting strategy we describe in this paper will continue to be profitable in the future. The Price and Wolfers (2007) study received considerable media attention in May 2007 and attracted the criticism of David Stern. It is possible that, since the release of their results, the behavior on the part of referees has changed or that the betting market has begun incorporating information about how the racial composition of the refereeing crew may differentially affect the outcome of the game based on the racial mix of each team.

Table 1: Betting Outcomes for Home v. Away and Favorite v. Underdog

	Bets	Wins	Win %	Rate of return
All teams	28,992	14,496	50%	-4.55%
Home	14,496	7,166	49.43% (0.42)	-5.63% (0.79)
Away	14,496	7,330	50.57% (0.42)	-3.47% (0.79)
Favorites	14,212	7,108	50.01% (0.42)	-4.52% (0.80)
Underdogs	14,212	7,104	49.99% (0.42)	-4.57% (0.80)
Home favorites	10,150	5,036	49.62% (0.50)	-5.28% (0.95)
Home underdogs	4,062	1,990	48.99% (0.78)	-6.47% (1.50)
Away favorites	4,062	2,072	51.01% (0.78)	-2.62% (1.50)
Away underdogs	10,150	5,114	50.38% (0.50)	-3.81% (0.95)

Notes: (Standard errors in parentheses)

Rate of return refers to the return from betting \$1 on each team meeting the specified criterion.

Table 2: Probability of Beating the Spread as a Function of Team and Referee Race

	3 white referees	2 white referees	1 white referee	0 white referees
Panel A				
Dependent Variable: $I(\text{Home Team Beats Spread})$ (Linear Probability Model)				
$\text{Points}^{\text{home}} - \text{Points}^{\text{away}} > E_{q50}[\text{Points}^{\text{home}} - \text{Points}^{\text{away}}]$				
β : % Black ^{home} - % Black ^{away}	-0.122 ^{***} (0.042)	-0.056 [*] (0.033)	-0.034 (0.048)	0.114 (0.125)
α : Constant	0.477 ^{***} (0.008)	0.498 ^{***} (0.006)	0.502 ^{***} (0.009)	0.542 ^{***} (0.024)
F-test ($\alpha=0.5$; $\beta=0$)	$F_{2,4071}=8.26^{***}$ (p=0.0003)	$F_{2,6853}=1.47^*$ (p=0.230)	$F_{2,3122}=0.27$ (p=0.760)	$F_{2,433}=1.83$ (p=0.161)
Panel B				
Dependent Variable: Home Team Winning Margin – Betting Spread				
$\text{Points}^{\text{home}} - \text{Points}^{\text{away}} - E_{q50}[\text{Points}^{\text{home}} - \text{Points}^{\text{away}}]$				
β : % Black ^{home} - % Black ^{away}	-3.151 ^{***} (0.983)	-1.547 ^{**} (0.752)	-1.261 (1.109)	1.241 (2.978)
α : Constant	-0.505 ^{***} (0.182)	-0.038 (0.138)	0.167 (0.209)	0.617 (0.571)
F-test ($\alpha=0$; $\beta=0$)	$F_{2,4071}=8.87^{***}$ (p=0.0001)	$F_{2,6853}=2.15$ (p=0.116)	$F_{2,3122}=0.96$ (p=0.382)	$F_{2,433}=0.64$ (p=0.527)
Panel C				
Dependent Variable: Home Team Winning Margin				
$\text{Points}^{\text{home}} - \text{Points}^{\text{away}}$				
β : % Black ^{home} - % Black ^{away}	-6.128 ^{***} (1.112)	-5.140 ^{***} (0.853)	-3.673 ^{***} (1.256)	-0.903 (3.217)
α : Constant	3.167 ^{***} (0.206)	3.541 ^{***} (0.157)	3.631 ^{***} (0.236)	4.591 ^{***} (0.618)
Panel D				
Dependent Var.: Betting Spread (Forecast home winning margin)				
$E_{q50}[\text{Points}^{\text{home}} - \text{Points}^{\text{away}}]$				
β : % Black ^{home} - % Black ^{away}	-2.976 ^{***} (0.546)	-3.593 ^{***} (0.419)	-2.413 ^{***} (0.603)	-2.144 (1.587)
α : Constant	3.671 ^{***} (0.101)	3.580 ^{***} (0.077)	3.464 ^{***} (0.114)	3.974 ^{***} (0.305)
Sample size	4,073	6,854	3,124	435

Notes: (Standard errors in parentheses). ^{***}, ^{**} and ^{*} denote statistically significant at 1%, 5% and 10%, respectively.

Table 3: Testing for Own-Race Bias

	Dependent Variable			
	I(Home team beats spread)	Home team winning margin, relative to spread	Home team winning margin	Betting spread (Home team expected winning margin)
δ : (%Black ^{home} - %Black ^{away}) * %White referees	-0.161* (0.085)	-3.292* (1.955)	-4.098* (2.210)	-0.807 (1.081)
β : %Black ^{home} - %Black ^{away}	0.043 (0.061)	0.345 (1.398)	-2.236 (1.581)	-2.581*** (0.773)
γ : %White referees	-0.045*** (0.016)	-1.038*** (0.366)	-0.907** (0.414)	0.131 (0.203)
α : Constant	0.524*** (0.011)	0.588** (0.263)	4.093*** (0.297)	3.504*** (0.145)
F-test: Market efficiency	$\alpha=0.5$; $\beta=\gamma=\delta=0$ $F_{4,14482}=5.33$ *** (p=0.0003)	$\alpha=\beta=\gamma=\delta=0$ $F_{4,14482}=6.13$ *** (p=0.0001)		

Notes: (Standard errors in parentheses).

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

Sample $n=14,482$ games.

Table 4: Proportion of Bets Won by Betting on Team with Greatest Racial Similarity to the Refereeing Crew

	3 white referees	2 white referees	1 white referee	0 white referees	All games
All Games	51.37% (0.78) [4,059]	50.99% (0.60) [6,829]	50.39% (0.90) [3,114]	52.53% (2.40) [434]	51.01% (0.42) [14,436]
$\% \text{Black}^{\text{home}} - \% \text{Black}^{\text{away}} > 0.1$	51.73% (1.04) [2,318]	51.17% (0.80) [3,903]	48.86% (1.17) [1,811]	52.71% (3.11) [258]	50.87% (0.55) [8,290]
$\% \text{Black}^{\text{home}} - \% \text{Black}^{\text{away}} > 0.2$	54.36% (1.49) [1,113]	50.83% (1.18) [1,798]	49.88% (1.71) [858]	54.20% (4.37) [131]	51.74% (0.80) [3,900]
$\% \text{Black}^{\text{home}} - \% \text{Black}^{\text{away}} > 0.3$	56.76% (2.34) [451]	52.36% (1.86) [720]	49.86% (2.62) [365]	52.63% (6.67) [57]	53.04% (1.25) [1,593]
$\% \text{Black}^{\text{home}} - \% \text{Black}^{\text{away}} > 0.4$	62.50% (4.05) [144]	54.47% (3.11) [257]	48.48% (4.37) [132]	56.25% (12.80) [16]	55.19% (2.12) [549]

Notes: (Standard errors in parentheses). [Sample size in brackets]

Table 5: Tests of Simple Betting Rules

	Bets	Wins	Win % $H_0: p < 50\%$	Rate of return $H_0: \text{Return} < 0\%$
All teams	28,992	14,496	50%	-4.55%
Based on forecasts of whether a team will beat the spread (Col. A of Table 2)				
$\hat{p} > \frac{1}{2}$	14,496	7,444	51.35%*** (0.41)	-1.96% (0.79)
$\hat{p} > \frac{11}{21}$	2,549	1,371	53.79%*** (0.99)	+2.68% (1.89)
Based on forecasts of how much a team will beat the spread by (Col. B of Table 2)				
$\widehat{\text{Margin}} - \text{spread} > 0$	14,496	7,448	51.38%*** (0.42)	-1.91% (0.79)
$\widehat{\text{Margin}} - \text{spread} > 0.5$	3,605	1,936	53.70%*** (0.83)	+2.52% (1.59)
$\widehat{\text{Margin}} - \text{spread} > 1$	661	368	55.67%*** (1.93)	+6.29%* (3.69)
Simple Rules: Only bet if all-black or all-white refereeing crew				
<i>All games</i>	4,493	2,313	51.48%** (0.75)	-1.72% (1.42)
<i>If $\%Black^{home} - \%Black^{away} > 0.1$</i>	2,576	1,335	51.82%** (0.98)	-1.06% (1.88)
<i>If $\%Black^{home} - \%Black^{away} > 0.2$</i>	1,244	676	54.34%** (1.41)	+3.74% (2.70)
<i>If $\%Black^{home} - \%Black^{away} > 0.3$</i>	508	286	56.30%*** (2.20)	+7.48%** (4.21)
<i>If $\%Black^{home} - \%Black^{away} > 0.4$</i>	160	99	61.88% (3.85)	+18.13%*** (7.35)

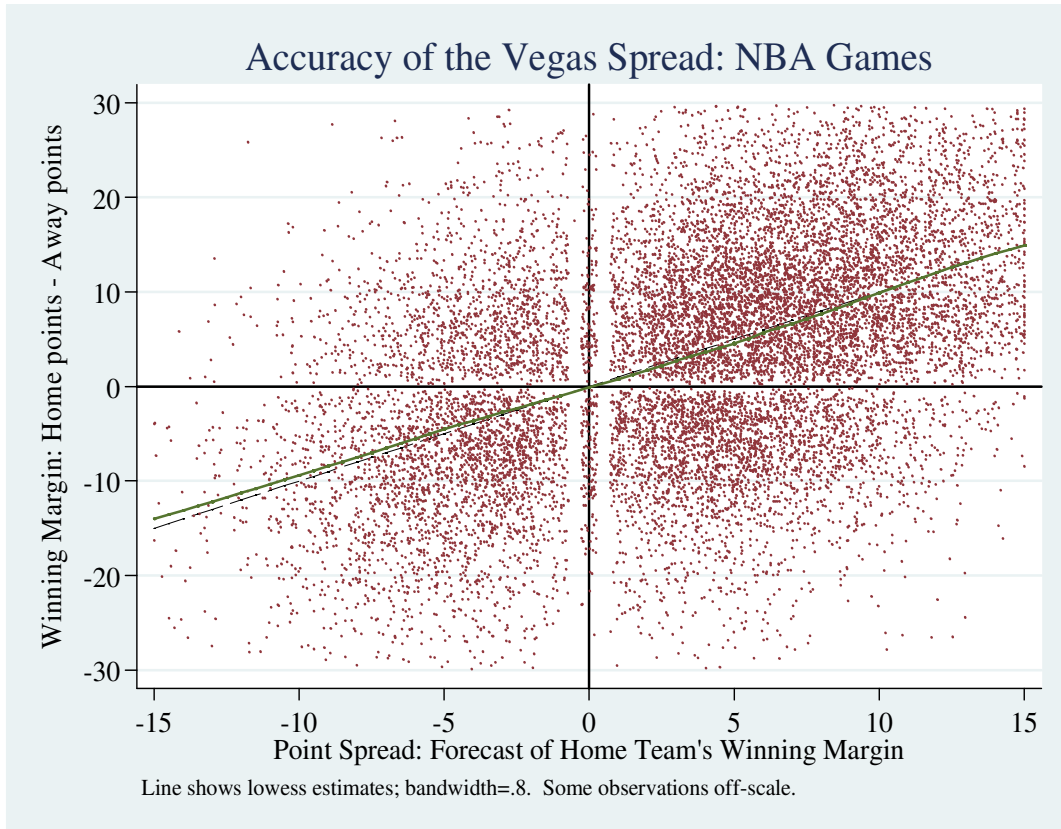
Notes: (Standard errors in parentheses)

Rate of return refers to the return from betting \$1 on each team meeting the specified criterion.

***, ** and * indicate rejection of null hypothesis at 1%, 5% and 10%, respectively.

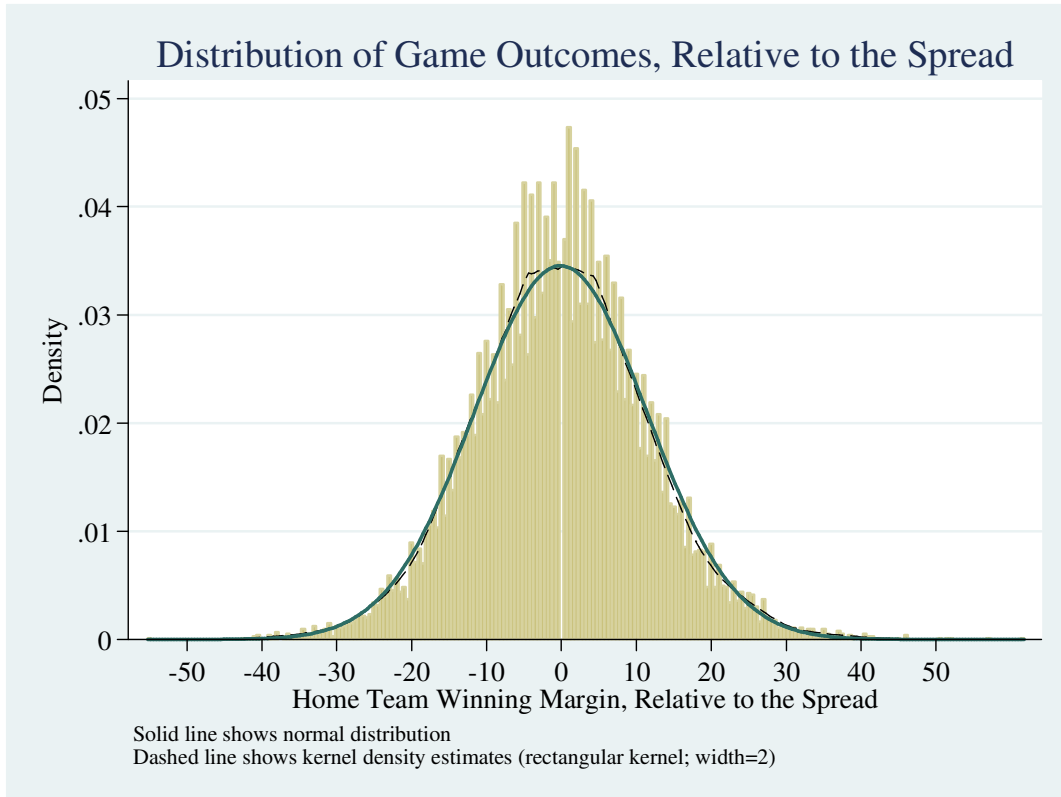
Note that these are one-sided hypotheses.

Figure 1: Point Spreads and Game Outcomes



Source: Authors calculations, based on data from www.covers.com

Figure 2: Histogram of Forecast Errors



Notes: Histogram shows density at both half-point and whole-point margins; because half-point spreads are far more common, this creates the jagged pattern. Also note that we have dropped from the sample all games involving a “push” (when the outcome is exactly equal to the spread).

Figure 3: Betting Outcomes

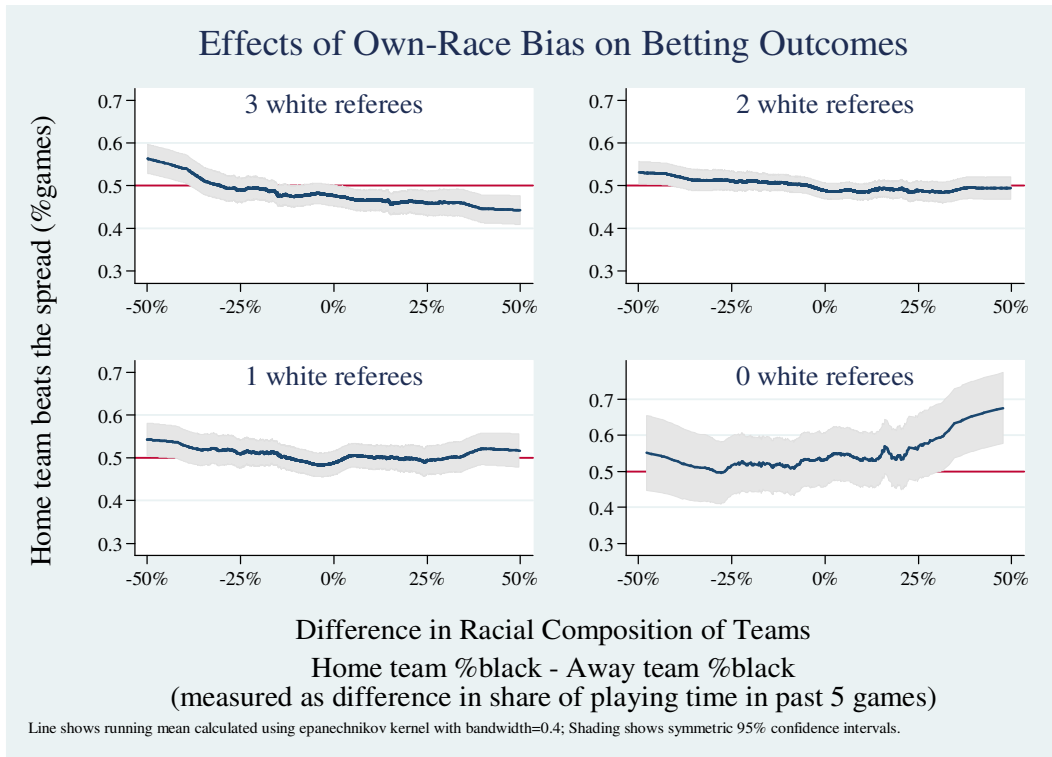


Figure 4: Winning Margin, Relative to the Spread

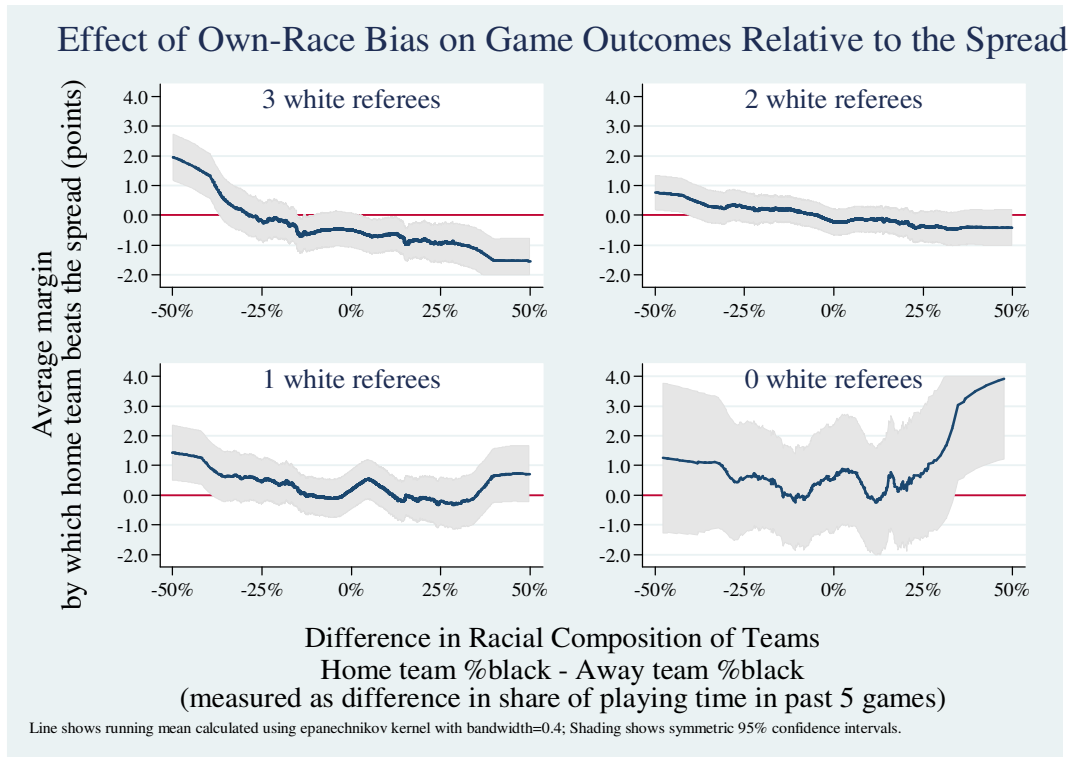


Figure 5: Winning Margin

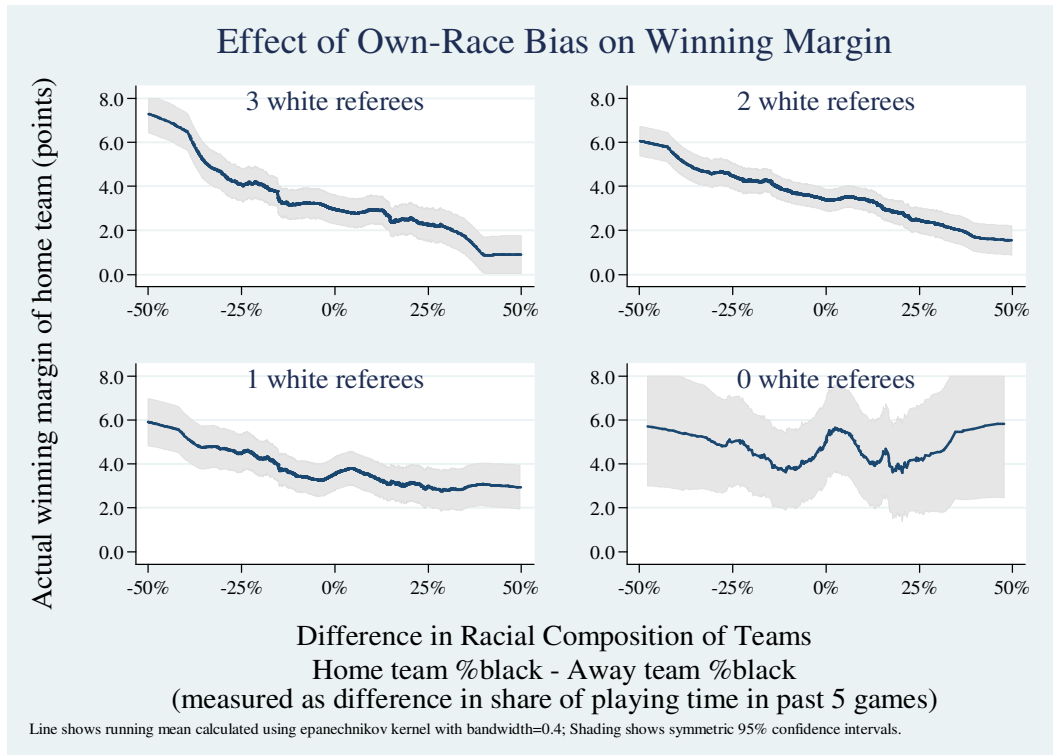
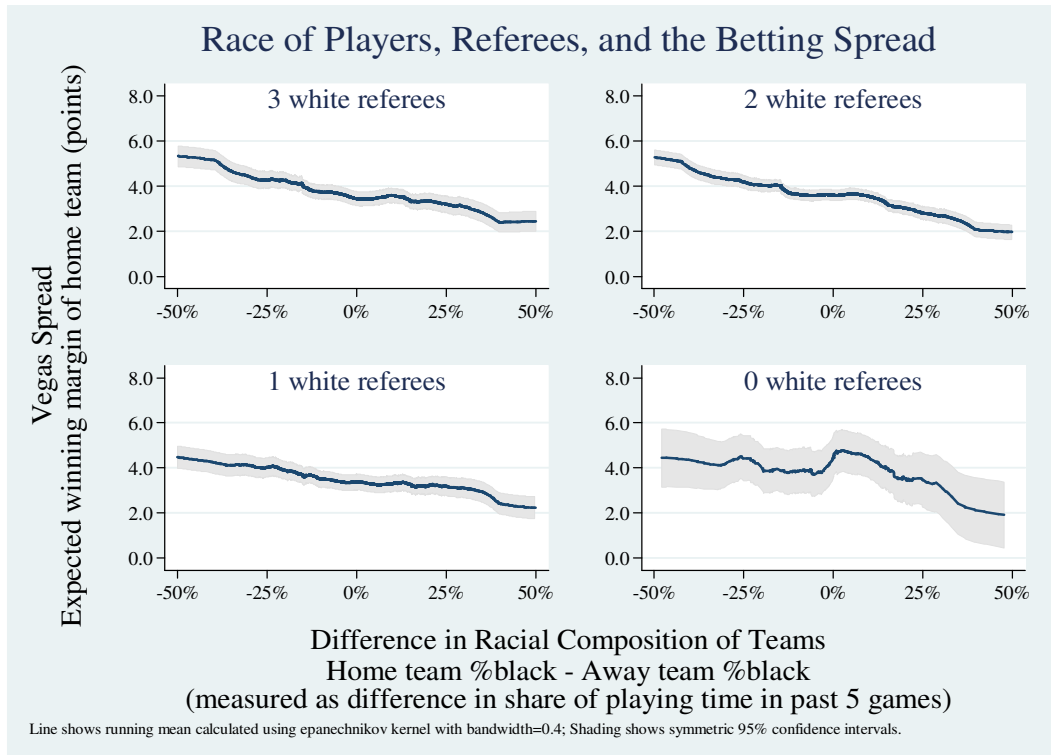


Figure 6: The Betting Spread



V. References

Price, Joseph, and Wolfers, Justin. "Racial Discrimination Among NBA Referees." Working Paper, National Bureau of Economic Research, Cambridge, 2007.

Parsons, Christopher A., Sulaeman, Johan, Yates, Michael C. and Hamermesh, Daniel S. "Strike Three: Umpires' Demand for Discrimination." Mimeo, University of Texas-Austin, 2007.

Sauer, Ramond D. "The Economics of Wagering Markets." *Journal of Economic Literature*, 1998. 36, 4 2021-2064.

Song, Chiung, Boulier Bryan L., and Stekler, Herman O. "The Comparative Accuracy of Judgmental and Model Forecasts of American Football Games." *International Journal of Forecasting*, 2007. 23, 3, 405-413.

Woodland, Linda M. and Woodland, Bill M. "Market Efficiency and the Favorite-Longshot Bias: The Baseball Betting Market." *Journal of Finance*, 1994. 49, 1, 269-279.

Woodland, Linda M. and Woodland, Bill M. "Market Efficiency and Profitable Wagering in the National Hockey League: Can Bettors Score on Longshots?" *Southern Economic Journal*, 2001. 67, 4, 983-995.

Wolfers, Justin. "Diagnosing Discrimination: Stock Returns and CEO Gender." *Journal of the European Economic Association*, 2006., 4(2-3), 531-541.