

Discrimination in Hiring Versus Retention and Promotion: An Empirical Analysis of Within-Firm Treatment of Players in the NFL

Michael Conlin
Michigan State University

Patrick M. Emerson
University of Colorado at Denver

If the costs and benefits of discriminating at the hiring stage differ from those at the retention and promotion stages, as recent evidence suggests, the effect of an individual's race on a firm's hiring decision should differ from its effect on the firm's retention and promotion decisions. This paper presents the first direct empirical test of this proposition. Using data of players drafted into the National Football League (NFL), we show that after controlling for draft selection, position, team, draft year, collegiate division, and team wins the prior season, white players have a 0.13 lower probability of having an active contract and start 1.56 games less than nonwhite players. When compared with the average probability of a drafted player having an active contract, 0.48, and the average number of starts, 2.6 games, these results provide strong evidence that nonwhite players face hiring discrimination in the NFL but are treated more equitably in retention and promotion decisions.

1. Introduction

This paper addresses the question, Is the effect of an individual's race on the selection criteria a firm uses for hiring decisions different than its effect on the criteria used for retention and promotion decisions? There is strong evidence to suggest that the costs and benefits of discriminating at the hiring stage differ from those at the retention and promotion stages.¹ The costs may differ due to the greater uncertainty associated with an individual's ability at the hiring

We are grateful to Michael Waldman, Kaushik Basu, Stacy Dickert-Conlin, Ramsey Shehadeh, David Kaplan, Larry Kahn, Johnny Yinger, two anonymous referees of this journal, participants of the Cornell Applied Microeconomics and Labor Workshops, and participants of workshops at Syracuse University, Vanderbilt University, and the University of Missouri for their comments, insights, and suggestions.

1. This paper is also related to the economics of crime and punishment, beginning with Becker (1968) (for a good overview of the literature see Ehrlich, 1996), which posits that the commission of crime is based on a benefit versus expected cost calculation of a potential perpetrator.

The Journal of Law, Economics, & Organization, Vol. 22, No. 1,
doi:10.1093/jleo/ewj005

Advance Access publication November 2, 2005

© The Author 2005. Published by Oxford University Press on behalf of Yale University.

All rights reserved. For permissions, please email: journals.permissions@oxfordjournals.org

stage. Once employed, an employee develops a performance record from which others can evaluate and judge her effectiveness. If greater uncertainty causes the probability of being perceived as discriminatory to decrease, then the firm's expected cost of discriminating will be less during the hiring decision. In many firms, the individuals making the hiring decision are the owners or human resources personnel, while the individuals making the retention and promotion decisions are the managers. If these groups of individuals have different tastes for discrimination, or if the incentive to respond to consumer discrimination is stronger for one group than for another, the benefits of discrimination will differ between these groups as well. For both cost and benefit reasons, therefore, it is reasonable to expect that a firm's discriminatory behavior is more likely to occur in the hiring stage than in the retention and promotion stages. Though the idea that the costs and benefits to discrimination differ within organizations is not a new one, this is, to the best of our knowledge, the first direct test of this proposition.²

There is a good deal of legal evidence to suggest that the expected cost of discrimination during the hiring stage is less than it is during the retention and promotion stage. For example, Donohue and Siegelman (1991:984) state: "[T]he vast majority of all litigation suits challenge discrimination in discharge." They go on to say (1017) that "the expected penalties for terminating a worker are probably much higher than for failing to hire her."³ Along with the higher expected litigation costs, retention and promotion discrimination may also be more costly due to its effect on other workers. Epstein (1992:60) states:

Although there may be no legal protection against being fired, within the firm dismissal can only be with cause, for reasons that range from financial necessity and business repositioning to employee misconduct. The firm that violates these rules may not be subject to a lawsuit, but it will bear substantial loss of reputation. If other workers regard the dismissal as improper, they will treat it as an amendment of their informal employment arrangement equivalent to a substantial decrease in salary. Their loyalties will diminish, and the rash employer will thus pay an implicit but real price.

Thus employees who perceive inequity in the treatment of others by their firms in retention and promotion decisions can impose costs on their firms through diminished effort, poor work quality, and negative attitudes.

Along with the lower expected costs, there may be greater benefit in discriminating for those individuals who do the hiring compared with those

2. That the labor market within the firm is different than that outside the firm is discussed in the internal labor market literature. See, for example, Baker et al. (1994).

3. Donohue and Siegelman (1991) estimate these costs empirically using federal court data and find that the expected cost of failing to hire a qualified job candidate of a protected class (e.g., a minority applicant) is much less than the expected cost of firing the same person, even under the most extreme assumptions.

making the retention and promotion decisions. The benefit from responding to consumer discrimination is likely greater for individuals making the hiring decision if these individuals are the owners of the firm. In addition, the taste for discrimination may be greater for those making the hiring decision compared with those making the retention and promotion decisions.

This paper uses what Ayres (2001) terms an *outcomes-based* approach to test for employment discrimination. In this approach, outcomes of individuals in a secondary market are compared with their “price” established in a primary market. Using information on football players drafted into the National Football League (NFL), we empirically test (conditional on when a player is drafted) whether race influences retention of a player (i.e., the player makes the team) and the performance opportunities afforded him. The primary market in this environment is the NFL draft, which occurs in early April. After being drafted, a player usually agrees to a contract and in July attends the team’s training camp. The majority of the team’s retention decisions are made during training camp. After training camp, the team plays sixteen regular-season games.

When a player is selected in the draft provides a measure of the team’s valuation for that player while being retained and number of starts are outcome measures in the “secondary” market.⁴ After controlling for draft selection, position, team, draft year, collegiate division, and team wins the prior season, we find that white players have a much lower probability of making the team and start fewer games than nonwhite players.⁵ These results provide strong evidence that hiring discrimination is prevalent among NFL teams.⁶

Once these results are demonstrated, a series of models are estimated to determine whether the level of hiring discrimination depends on factors that influence the costs and benefits of discrimination at the hiring, retention, and/or promotion stages. We find no evidence that the racial differences in the probability of an active contract and the number of starts depend on whether the coach and player are of the same race, and little evidence that they depend on the racial composition of the metropolitan area. However, we do find some evidence that these racially based differences are greater for players with high uncertainty associated with their ability, for players drafted by teams that

4. The “price” a team pays for drafting a player is based on the opportunity cost associated with not being able to select a different player with that draft choice.

5. We use the terms “making the team” and “having an active contract” interchangeably. While the large majority of players with active contracts are those who make the team, this need not be the case. A player’s contract is considered active if the player participates in at least three regular-season games. A contract is considered active if the player plays on the team that drafted him, is on injured reserve, traded, selected off of waivers, or signs with another team under Plan B. In most cases, an active contract occurs when the player makes and plays on the team that drafted him. See Conlin (1999) for a more detailed explanation.

6. These results would also be expected if nonwhite players were being discriminated against in hiring decisions as well as in those of retention/promotion and the discrimination were greater in the former, or if white players were being discriminated against in retention and promotion decisions. Our subsequent results provide no evidence of discrimination against nonwhites or whites in retention and promotion decisions. Therefore, we interpret these results as indicative of nonwhites facing hiring discrimination.

regularly sell out their stadiums, and for players who are drafted by teams of which the head coach has little input into drafting decisions.

Finally, as there should be a price to pay for discriminatory behavior on the part of teams, we estimate the overall effect of discrimination on teams' performance in the league. We find some evidence that teams that discriminate more in the hiring process win fewer games in subsequent seasons.

The rest of this paper is organized as follows: Section 2 summarizes the related literature and discusses how this paper contributes to the research using an outcomes-based approach to test for employment discrimination. Section 3 gives an overview of the NFL labor market and provides summary statistics of the data. Section 4 presents the major findings of the paper, *viz.*, that the probability of having an active contract and the number of starts are much greater if the player is nonwhite (conditional on draft selection, position, draft year, team, collegiate division, and team wins). Section 4 also examines variables associated with the costs and benefits to the teams of discriminating, *viz.*, (i) whether nonwhite players are more or less likely to come from small colleges and what effect this has on the observed differential treatment; (ii) whether teams that regularly sell out their stadiums discriminate less; (iii) whether the coach being of the same race as the player influences the observed differential treatment; (iv) whether more influential coaches mitigate the observed differential treatment in hiring; and (v) whether the racial composition of the fans influences the observed differential treatment. Section 5 examines whether greater uncertainty and different contract structure associated with a particular race can explain the results found in Section 4. Section 6 presents evidence that there is a cost to team performance associated with discrimination. Section 7 summarizes the results and discusses their implications to labor markets outside of professional sports.

2. Related Literature

The majority of the literature that tests for discrimination in employment (i.e., hiring, retention, and promotion) considers the earnings and occupational disparities of individuals of different races (for thorough reviews of the literature, see Darity and Mason, 1998; Yinger, 1998; and Ross, 2003). This research often involves regressing earnings level or occupational status on individual characteristics such as education, race, and age. It is argued that a race coefficient different than zero provides evidence of employment discrimination. The main criticism of this approach is that the race coefficient may be biased due to omitted variables (Ross, 2003; Killingsworth, 1993), and it is very difficult to observe many human characteristics that can be correlated with, for example, employment status and earnings.

A somewhat more direct approach when testing for hiring discrimination based on race involves audit studies such as those conducted by the Urban Institute. These studies involve pairs of testers of different races being matched and trained to minimize the dissimilarities between them. These matched

pairs of testers then apply for similar jobs. While these audit studies provide evidence that minority groups are discriminated against in the hiring process, Heckman and Siegelman (1993) note that the inability to perfectly match testers brings the results into question.⁷ In other words, efforts to minimize dissimilarities between a pair of individuals are not likely to be completely successful in eliminating all of the aspects of human behavior, demeanor, stature, etc., that affect the decision to hire, promote, and compensate individuals.

A third type of approach used to test for discrimination involves what we have referred to in the introduction as outcomes-based tests. This approach has been used to test for discrimination in the bail bond market, motor vehicle searches, journal publications, and the mortgage market. Ayres and Waldfoegel (1994) found that bail bond rates for minority defendants in a particular community were systematically lower than those charged to white defendants, suggesting that the bail amounts set for minority defendants in the courts were systematically too high and thus discriminatory. Knowles et al. (2001) use a similar approach to test for discrimination in motor vehicle searches, comparing the race of drivers that were stopped with the probability of a successful search by police (success being defined as finding contraband that leads to conviction). Laband and Piette (1994) and Berkovic et al. (1998) also used the outcomes approach to test for preferential treatment of certain individuals in the publication of journals and in mortgage markets, respectively. The key element of these types of tests is that they minimize the omitted variable bias problem mentioned above because unobserved individual characteristics that matter in terms of performance (or price) should be taken into account in both the primary and secondary markets.

The main contribution of this paper to the discrimination literature is the application of the outcomes approach to the employment decisions of firms.⁸ The important aspect of the present study is that when considering whether retention decision and performance opportunities afforded to an employee are affected by the employee's race, it conditions on a measure of the firm's private valuation of an employee at the time of hiring (i.e., in which round the player was selected in the draft). This minimizes the problems associated with selection bias, which are common to tests of screening discrimination (for discussions of this issue, see Ross and Yinger, 2002; Ross, 2003).

7. In another application, Ayres and Siegelman (1995) use such an approach to test for discrimination in the new car market.

8. Szymanski (2000) uses a similar methodology to look at the aggregate (external labor market) employment decisions of firms by investigating whether (controlling for their wage payrolls) English professional soccer teams with higher proportions of black players perform better. By using a panel, he was able to focus on within-team variation over a five-year span and found that teams with higher proportions of black players performed better, suggesting that black players were discriminated against.

3. Institutional Background and Data

3.1 The NFL Labor Market and Organizational Structure

The unorthodox nature and structure of the NFL labor market presents a unique opportunity to examine the internal labor markets of firms by providing a measure of a team's private valuation of each player by that player's selection number in the draft. The NFL draft, held every spring, is the primary component of the labor market. The draft is the mechanism through which most football players enter the NFL. From 1986 to 1991 each of the then 28 NFL teams was awarded one draft pick in each of the draft's 12 rounds with which to select a player or to trade. Once drafted by a team, a player is prohibited from signing with another team unless he has been traded to that team by the one that had drafted him. If the player and the team that had drafted him cannot reach a contractual agreement, the player may sit out a year and reenter the draft the following year or the team may trade him (see Conlin, 1999, for a detailed description of the draft and contract negotiation process).

While the labor market may be unorthodox, the organizational structure of the typical NFL team is similar to that of firms in many other industries in that different parties are responsible for the hiring and retention/promotion decisions, respectively. Players entering the league are evaluated and selected primarily by the administrative arm of the teams (the "front office"), but once hired, their future is placed primarily in the hands of the coaching staffs, which make the decision to terminate (to "cut") or to promote (to train players and present them with performance opportunities). There is typically some overlap in these responsibilities—with the front office participating in some retention and promotion decisions and the coaching staff taking part in some hiring decisions—but the primary responsibilities are clearly delineated. (A relatively recent trend of filling both the general manager [front office] and head coach positions with the same person was not prevalent during the time frame of the present study.)

Racial discrimination in the NFL can be observed in the decisions of both the front office and the coaching staff.⁹ Front office discrimination can result in players of a specific race being penalized during the hiring/drafting stage and undercompensated during the contract negotiation stage.¹⁰ Head coach discrimination can lead to players of a certain race making the team ahead of, and receiving more playing time than, other players of equal or greater ability.¹¹

9. The vast majority of the literature testing discrimination among professional athletes focuses on wage discrimination (see Mogull, 1973; Kahn and Sherer, 1988; and Kahn, 1991, 1992). This paper is the first to explicitly test for employment discrimination in professional sports.

10. The front office can also provide coaches the incentive to discriminate. Due to contract-ability issues, it is unlikely that the front office can perfectly align the coaches' incentives with the front office's incentives.

11. Coaches can also discriminate by not providing particular players the opportunity to perform to their capabilities. For example, a coach can discriminate by never calling pass plays where a particular wide receiver is the primary target. In addition, coaches can discriminate by providing less training for players of a certain race.

3.2 Data Description and Summary Statistics

The data employed in this study consist of contract, position, draft, performance, and race information for 1827 of the 2016 players drafted into the NFL between 1986 and 1991. The data set also includes the players' college teams, the number of regular-season games won by each team from 1985 to 1999, the annual attendance and capacity figures of each team's home games from 1985 to 1991, the race of each team's coaches from 1985 to 1993, the 1990 racial composition of each team's metropolitan area, and the head coaches' career winning percentage, career wins, and tenure with the teams. The NFL Players Association (NFLPA) provided initial contract data for 1873 of the 2016 players drafted from 1986 to 1991, along with the players' colleges.¹² Team record, attendance, stadium capacity, player performance, and head coach information were obtained from the 1985 through 1999 NFL Record and Fact Books. College teams' divisions, conferences, and rankings were obtained from the 1995 Official National Collegiate Athletic Association (NCAA) College Football Records Book. Race of players and coaches were obtained from each team's annual media guide from 1986 to 1993, which generally displays individual photographs (head shots without helmet) of the current year's draft picks and coaches. Where pictures were not available, an attempt was made to contact the teams directly and determine the race of their draft picks and coaches. From this procedure, we were able to determine, with a reasonable degree of certainty, the race of all of the coaches and all but 46 of the players. While this method of determining a player's race introduces the possibility of error, no other data on player race are available. By eliminating the players whose race we were unable to determine, our sample size was reduced to 1827. Therefore, our data set contains 90.6% of all players selected in the 1986 to 1991 NFL drafts. Finally, figures for the racial makeup of the metropolitan areas were obtained from the 1990 U.S. Census.

Table 1 presents some summary statistics of the data examined in this paper. Of the 1827 players in this study, 1167 were nonwhite and 660 were white.¹³ Nonwhite players, on average, have a higher probability of an active contract and average more regular-season games started in each of their first three years after being drafted.¹⁴ Overall, NFL draftees average a 0.54 probability of an active contract in their first year, 0.50 in their second, and 0.41 in their third. They also average 1.79 starts in their first year, 2.97 in their second, and 3.10 in their third. The average draft selection number is 154 for nonwhite players and

12. Of the 143 draft choices for which contract information was not available, 15 were selected in supplemental drafts and the remaining 128 either did not sign a contract or did not report the contract to the NFLPA.

13. Almost all nonwhite players in the NFL are black. Only 27 of the 1167 nonwhite players in our sample were nonblack. We therefore refer to *nonwhite* as a single race for ease of discussion but caution the reader that there are a number of different races represented in the nonwhite category.

14. The probability of an active contract and the number of starts for all three contract years generally declines monotonically across draft round. Conditional on draft round and contract year, nonwhite players have a higher probability of an active contract in over 80% of the round-years and start more games in over 90% of the round-years than white players.

Table 1. Means of Variables by Race

Variable	Means		
	Nonwhite	White	All
Active contract 1st year after being drafted	0.60	0.43	0.54
Active contract 2nd year after being drafted	0.54	0.43	0.50
Active contract 3rd year after being drafted	0.46	0.34	0.41
Games started 1st year after being drafted*	2.14	1.17	1.79
Games started 2nd year after being drafted*	3.46	2.10	2.97
Games started 3rd year after being drafted*	3.58	2.24	3.10
Selection number in draft	154	171	160
Small-college football program	0.231	0.199	0.219
Average annual expected compensation (\$1000)	121.63	96.93	112.71
Percent of metropolitan population	21.8	78.2	100.0
Observations	1167	660	1827

*Regular-season games started necessarily equals zero if contract is not active that year.

171 for white players. Slightly more than 20% of the players in the data set had participated in a small-college football program, i.e., were not in a program in the premier NCAA Division 1A. The mean of the average annual expected compensation of a player's contract is \$112,710.¹⁵ The average percentage of the population that is nonwhite in the metropolitan areas of NFL cities is 21.8.

4. Race and NFL Players

To test whether player race has a differential effect on the hiring versus retention and promotion decisions, two types of regressions were performed. The first type is a series of probit regressions on whether a contract is active. Specifically, we initially estimate the following model:

$$ActiveContract_i = \begin{cases} 1 & \text{if } \alpha + \beta\mathbf{X} + \delta(PlayerRace) + \varepsilon > 0 \\ 0 & \text{otherwise.} \end{cases}$$

The dependent variable equals 1 if the contract is active in year i , and 0 if it is not active. Separate models are estimated for whether a player has an active contract in the first, second, and third years after being drafted. The independent variable vector, \mathbf{X} , includes: indicator variables for year of the draft, the player's position, the team that drafted the player, the draft round in which the player was taken, and whether the player participated in a small-college football program, as well as a variable for the selection number of each player and the number of wins the team had in the year prior to the draft in which the

15. Because of antitrust concerns, NFL teams sign all draft choices to "standard form contracts." This allows the value of different contracts to be more easily compared. Almost all drafted players' initial contracts in the NFL are nonguaranteed; therefore, a player who signs a contract after being drafted by a team is entitled to only the signing bonus stipulated in the contract. The player must have an active contract in order to receive the payments in the contract for that year. The average annual expected compensation is calculated by multiplying each payment a player may receive by the probability a player drafted in that round receives the payment, taking the present discounted value of these expected payments and dividing by contract duration.

Table 2. Effects of Selected Variables on Probability of an Active Contract and Number of Games Started in the Player's First, Second, and Third Years. Probit and Ordered Probit Regression Results (SE)

Independent Variables	Active Contract, Year			Games Started, Year		
	1st	2nd	3rd	1st	2nd	3rd
Team wins in prior season	0.003 (0.019)	0.012 (0.019)	0.012 (0.018)	-0.051*** (0.017)	-0.007 (0.016)	0.015 (0.017)
Selection number in draft	-0.017*** (0.006)	-0.020*** (0.006)	-0.019*** (0.006)	-0.014*** (0.005)	-0.012** (0.005)	-0.018*** (0.005)
Small college	-0.007 (0.081)	-0.100 (0.084)	-0.001 (0.084)	-0.046 (0.086)	0.017 (0.083)	0.087 (0.084)
Player's race (= 1 if nonwhite)	0.297*** (0.086)	0.278*** (0.088)	0.230*** (0.088)	0.353*** (0.091)	0.356*** (0.085)	0.350*** (0.085)
Position effects	YES	YES	YES	YES	YES	YES
Round effects	YES	YES	YES	YES	YES	YES
Team effects	YES	YES	YES	YES	YES	YES
Year effects	YES	YES	YES	YES	YES	YES
Pseudo R^2	0.28	0.30	0.27			
F statistic (K, n-K)				9.88 (54,1730)	12.22 (54,1730)	9.79 (54,1730)
Sample size	1827	1827	1827	1784	1784	1784

Note: White's heteroskedastic consistent errors used in all regressions.

*Statistically significant at the 0.10 level; **Statistically significant at the 0.05 level; ***at the 0.01 level (two-tailed tests).

player was selected. This set of variables controls for other factors, besides player race, that are likely to affect whether the player has an active contract. Particularly important is the player's selection number, for this, when controlling for player position, year of the draft, etc., is a measure of the team's private valuation of the player. Small-college football programs, defined as non-NCAA Division 1A, are included as a variable to control for the uncertainty in a player's ability level. This uncertainty is likely to be greater for a player who did not attend a college with a Division 1A football program.¹⁶ *PlayerRace* is an indicator variable that equals 1 if the player is nonwhite and 0 if the player is white. The error term is assumed to be standard normally distributed and independent of \mathbf{X} and *PlayerRace*. A maximum likelihood estimation is used to consistently estimate the parameters.

The results of the probit regressions are shown in the first three columns of Table 2. If a player's race has similar effects upon decisions of hiring (drafting) and of retention and promotion (cutting and starting opportunities, respectively), the coefficient on *PlayerRace* will be zero. However, if there is a difference based on the player's race, the coefficient on *PlayerRace* will be nonzero. The coefficient on *PlayerRace* in Table 2 is positive and significant

16. Uncertainty about the ability of a player can make him more or less attractive to a team depending on the distribution of ability levels and how the team truncates the distribution by cutting players.

non–Division 1A collegiate football programs (i.e., with greater uncertainty associated with their ability) are discriminated against more in the hiring process and whether teams that sell out their stadiums discriminate more in the hiring process.¹⁹ Section 4.2 considers whether the front office’s benefit from discriminating is greater than the coaches’ benefit as the result of their tastes for discrimination or consumer discrimination.²⁰ Specifically, we test whether the differential effect of player race on hiring and retention and promotion decisions depends on the race of the coaches, the input of the head coach on drafting decisions, or the racial composition of a team’s metropolitan area.

4.1 Cost of Discrimination

Once a player participates in training camp, substantial information about the player’s ability is revealed. Discrimination based on race at this stage, then, is likely to be noticed by both the player himself and his teammates. At the time the hiring decision is made, however, there is a great deal of uncertainty associated with a potential player’s ability, and so the likelihood of discrimination being noticed at this stage is considerably less than at stages after employment. Because teams have less information about a player’s ability if he participates in a non–Division 1A collegiate program, there is greater uncertainty at draft time associated with these small-college players (i.e., the signal that the team receives on these players is likely to be noisier than for players who participated in large-college football programs). If this is indeed the case, the front office’s cost of discriminating in its drafting decision is lower when it selects a player from a non–Division 1A program, for the discrimination is harder to detect, both for an outside authority and for other players on the team.

A model in which the cost of hiring discrimination depends on how noisy a signal the team receives on the player’s ability level is presented in the Appendix.²¹ In the model, team quality is a function of the players’ beliefs that their team discriminates. An implication of the model is that the cost of hiring discrimination is lower for noisier signals of the player’s ability level, because the other players on the team are less able to infer whether the team is discriminating when a white player is drafted.

To test whether nonwhite players with greater uncertainty associated with their ability are discriminated against more at the hiring decision, we interact

19. Coaches’ job tenures with a specific NFL team average less than five years and many coaches obtain jobs with other NFL teams after being fired. Therefore, coaches are concerned with how other teams perceive their ability, and this perception is heavily dependent on their coaching record. Furthermore, a large fraction of the NFL’s revenue is obtained from television and merchandising contracts. This revenue is divided equally among the NFL teams. Therefore, if discrimination adversely affects team performance, the cost of this discrimination to the coaches is likely to be greater than to the front office.

20. Owners are often involved in the front office of NFL teams, and coaches are evaluated (both by the team and by the market for league coaches) primarily based on the team’s on-field performance. Therefore, the benefit of responding to consumer discrimination is likely much greater for the front office than for the coaching staff.

21. The model also provides an explanation of why the cost of discrimination is less when it occurs in the drafting decision rather than in the retention and promotion decisions.

SmallCollege (which equals 1 if the player attended a school with a non-Division 1A football program) and PlayerRace and add this interactive variable to the original probit and ordered probit regressions (i.e., the base model). The coefficient estimates associated with PlayerRace and this interaction term are presented in Table 3.²² If players who come from schools with small-college football programs have more uncertainty associated with their ability and if nonwhite players from these schools face greater hiring discrimination, then the coefficients on this interactive variable should be positive, as they are in Table 3. And while almost all have *t*-statistics greater than 1, none are statistically significant at the 10% level. This provides some suggestive evidence that NFL teams discriminate more against players from small-college football programs when hiring.

The cost of hiring discrimination may also vary across teams, depending on their ability to regularly sell out their stadiums. The cost of hiring discrimination to the front office is a function of the effect such a practice has on revenues (since the owners of a team are often part of the front office). As there is likely an adverse effect of discrimination on team performance, the cost of hiring discrimination is likely higher for those teams that do not routinely sell out their stadiums. This is because a team that does not sell out will lose sales at the margin due to the decrease in the team's on-field performance resulting from discrimination.

To test whether selling out affects the level of hiring discrimination, we add to the base specification an indicator variable that equals 1 if the team sells out in the year prior to drafting the player, and we interact this variable with PlayerRace. The interactive term equals 1 if a player is nonwhite and the player's team sells out. The other variables in the specification are the same as in the base specification except for the fact that team indicator variables are not included. If front offices discriminate more in their drafting decisions when the team sells out, then the coefficient on this interactive term should be positive. The results presented in Table 3 indicate that although the coefficients associated with the interactive term are not statistically significant, they are all positive. This lends some support to the hypothesis that teams that sell out their stadiums discriminate more in their drafting decisions.

4.2 Benefit of Discrimination

To test whether coaches' tastes for discrimination contribute to the results in Table 2, we consider whether the probability a player makes the team and the number of games a player starts depend on the race of the coaches. If coaches have some taste for discrimination, one would expect this taste to depend on whether the coach and player's races are different. For example, if coaches discriminate against nonwhite players, one would expect the level of discrimination to be less if the nonwhite player's coach is also nonwhite. If the level of discrimination were found not to depend on whether the race of the coach and the player differ, this would suggest that coaches are not discriminating.

22. While the coefficient estimates associated with the other independent variables are not presented in Table 3 for any of the alternative specifications, they do not change appreciably from those reported in Table 2.

Table 3. Relevant Coefficient Estimates (SE) from Alternate Specifications

	Active Contract, Year			Games Started, Year		
	1st	2nd	3rd	1st	2nd	3rd
Small College–Player Race Interaction						
Player's Race	0.251***	0.218**	0.192**	0.308***	0.320***	0.315***
(= 1 if nonwhite)	(0.095)	(0.097)	(0.096)	(0.098)	(0.091)	(0.092)
(Small College)×	0.210	0.277	0.183	0.285	0.198	0.182
(Player's Race)	(0.174)	(0.185)	(0.180)	(0.210)	(0.183)	(0.187)
Stadium Sell Out						
Player's Race	0.230**	0.216**	0.190*	0.291***	0.318***	0.327***
(= 1 if nonwhite)	(0.107)	(0.107)	(0.110)	(0.109)	(0.105)	(0.110)
Sell Out (= 1 if	0.030	0.066	0.014	-0.245*	-0.061	0.012
stadium is sold out)	(0.118)	(0.121)	(0.119)	(0.134)	(0.117)	(0.119)
(SellOut)*	0.165	0.075	0.029	0.226	0.105	0.034
(Player's Race)	(0.140)	(0.141)	(0.140)	(0.151)	(0.134)	(0.137)
Coach's Race						
Player's Race	0.373	0.634	0.966	0.162	0.565	0.414
(= 1 if nonwhite)	(0.510)	(0.450)	(0.422)	(0.426)	(0.425)	(0.424)
Number of Years	0.140	0.142	0.153	0.055	0.020	-0.040
Head Coach Is	(0.436)	(0.162)	(0.099)	(0.352)	(0.155)	(0.093)
Same Race As Player						
Number of Years	0.052	0.062	0.094	-0.293	0.083	0.092
Coordinator Is	(0.242)	(0.157)	(0.105)	(0.209)	(0.124)	(0.105)
Same Race As Player						
Number of Years	-0.062	-0.024	0.003	0.025	-0.038	-0.022
Position Coach Is	(0.090)	(0.048)	(0.033)	(0.084)	(0.045)	(0.032)
Same Race As Player						
Number of Years	-0.090	-0.005	0.006	0.025	0.043	-0.018
Special Teams	(0.109)	(0.058)	(0.040)	(0.101)	(0.050)	(0.037)
Coach Is Same Race						
As Player						
Head Coach's Winning Percentage						
Player's Race	0.492	0.590**	0.625**	0.430	0.417	0.775***
(= 1 if nonwhite)	(0.308)	(0.291)	(0.292)	(0.287)	(0.271)	(0.272)
Head Coach's Career	-0.745	0.265	0.220	-0.204	-0.811	0.239
Winning Percentage	(0.639)	(0.616)	(0.604)	(0.586)	(0.539)	(0.573)
(Head Coach's	-0.301	-0.521	-0.694	-0.152	-0.106	-0.752
Career Winning	(0.566)	(0.538)	(0.536)	(0.525)	(0.490)	(0.500)
Percentage)×						
(Player's Race)						
Racial Composition of Metro Area						
Player's Race	0.123	0.132	0.020	0.458**	0.606***	0.408**
(= 1 if nonwhite)	(0.189)	(0.193)	(0.192)	(0.196)	(0.189)	(0.190)
MetroArea _{NW}	14.48	5.293	-10.388	15.948*	2.191	0.332
	(10.54)	(9.991)	(9.758)	(9.289)	(9.122)	(8.490)
(MetroArea _{NW})×	0.796	0.671	0.982	-0.491	-1.154	-0.266
(PlayerRace)	(0.784)	(0.792)	(0.802)	(0.814)	(0.750)	(0.771)

Note: White's heteroskedastic consistent errors used in all regressions. Statistically significant at the 0.10 level; *Statistically significant at the 0.05 level; ***at the 0.01 level (two-tailed tests). Indicator variables for year of the draft, the player's position, the team that drafted the player, the draft round in which the player was taken, and whether the player participated in a small-college football program, as well as a variable for the selection number of each player and the number of wins the team had in the year prior to the draft in which the player was selected, are also included as independent variables in the above specifications.

In order to estimate the effect of playing for a coach of the same race, we include a set of independent variables to the base specification that control for the effect of a player and his head coach, coordinator, position coach, and/or special teams coach being of the same race.²³ For the first-year regressions, these variables equal 1 if the player's head coach, coordinator, position coach, and special teams coach are of the same race as the player in the player's first year in the league. Because the effect a coach has on a player's performance measures is cumulative across years, we use a variable that equals the number of years that a player and his respective coaches have been of the same race for the second-year and third-year regressions.²⁴ The coefficients on the coach-player same-race indicator variables should be positive if coaches discriminate against players of races different than their own. The most relevant coefficient estimates from these regressions are given in Table 3.²⁵ The signs of the coefficients on the coach-player same-race indicator variables vary depending on the year and the type of coach. However, none of the coefficients are statistically significant. A test for their joint significance does not reject the null hypothesis that all the coefficients equal zero. These results suggest that coaches' decisions of whom to retain, train, and start do not depend on player's race and, therefore, that coaches are not discriminating.^{26,27}

23. These are the relevant coaches for the player in question. Therefore, a receiver's relevant coaches are the head coach, the offensive coordinator, the receivers coach, and the special teams coach.

24. For example, if a white player plays for a white head coach in his first year, a nonwhite head coach his second year, and a white head coach in his third year, then the variable concerning whether the coach and player are the same race equals 1 the first and second years and 2 the third year.

25. A different specification was also estimated where separate variables were included for white player–white coach and nonwhite player–nonwhite coach. The results from this specification are similar to those when the interactive variables indicate whether the player and coach are of the same race.

26. The proportion of player-years (assuming that all players made the team in their first three years) where the head coach is nonwhite is 0.031 and the proportion of player-years where the coordinator is nonwhite is 0.027 player-years. Nonwhites are better represented as position coaches and special teams coaches, with proportions of 0.212 and 0.180 player-years, respectively.

27. To further test for coach discrimination, we consider another specification that captures the effect of coaching changes that result in the new coach being of a different race than his predecessor. In this specification, the number of games the player starts in the prior season and two sets of indicator variables are added to the base specification. The first set of indicator variables control for any type of coaching change from the prior year. These variables capture any “cleaning house” effect, i.e., the bias of new coaches toward new players due to the lack of an existing relationship with the veterans on the team. The second set of indicator variables concern whether a nonwhite coach replaces a white coach and the player is white and whether a white coach replaces a nonwhite coach and the player is nonwhite. These interactive indicator variables capture the effect on the player's probability of an active contract and regular-season starts of a change in the coach's race away from a player's race. Only players whose contracts were active in the prior year are included in the estimation. The coefficient estimates associated with the interactive indicator variables that capture the effect of a change in coach's race away from player's race vary in sign and are neither economically nor statistically significant. These results are available upon request.

If the benefits from discriminating for NFL teams does in fact come from front office preferences, and if coaches do not share the same taste for discrimination, then we would expect the level of hiring discrimination to diminish in inverse proportion to the influence the coaching staff has on front office decisions. Therefore, we test if the differential effect of player race on hiring and retention and starting decisions depends on the input the coaching staff has on drafting decisions.

To test whether head coach input mitigates the level of discrimination, we add two independent variables to the base model: the head coach's career winning percentage and the interaction term (Head coach's career winning percentage) \times (PlayerRace). If a head coach's level of input increases with his overall success as a coach in the NFL and if this has a mitigating tendency on the front office to discriminate against nonwhites, then the coefficient on the interactive term should be negative. While all of the coefficient estimates in Table 3 associated with (Head coach's career winning percentage) \times (PlayerRace) are negative and several have *t*-statistics over 1, none are statistically significant at the 10% level. Similar results are obtained when the number of seasons the head coach has been with the team and the number of his career coaching wins are used as proxies for the level of input the coaching staff has on drafting decisions. These results provide some support that head coaches' tastes for discrimination are not as strong as those of front offices' and that more head coach input in the hiring decisions can mitigate the level of discrimination.

Another explanation for the results in Table 2, besides taste for discrimination, is that front offices discriminate against nonwhite players in their drafting decisions because of fans' preferences. To test for evidence that front offices discriminate in their hiring practices in response to consumer preferences, we include two additional independent variables to the base specification: MetroArea_{NW} and (MetroArea_{NW}) \times (PlayerRace). MetroArea_{NW} is the percentage of the population of the metropolitan area in which the player's team resides that is nonwhite.²⁸ Assuming that the composition of the team's fan attendance is similar to the racial composition of the metropolitan area in which the team resides, we would expect the coefficient on MetroArea_{NW} to be positive and the sum of the coefficients on MetroArea_{NW} and (MetroArea_{NW}) \times (PlayerRace) to be negative if the front office is responding to consumer discrimination.²⁹ This expectation comes from the fact that there should be less

28. During our sample period, the Green Bay Packers played several regular-season home games in Milwaukee. When determining the Packers' MetroArea_{NW}, a weighted average of the percentage of the population that is nonwhite in both metropolitan areas was calculated (where the weights were based on the fraction of games played in Green Bay and Milwaukee). The Cardinals moved from St. Louis to Phoenix in 1988. The Cardinals' MetroArea_{NW} values for 1986 and 1987 are based on St. Louis' metropolitan area and on Phoenix's metropolitan area in subsequent years.

29. The effect on nonwhite players of playing in an area with a higher concentration of nonwhite residents is obtained by summing the coefficients on MetroArea_{NW} and (MetroArea_{NW}) \times (PlayerRace).

consumer discrimination against nonwhite players if the team is located in a metro area with a particularly high concentration of nonwhite residents.³⁰

The most relevant coefficient estimates from these probit and ordered probit regressions are given in Table 3. The signs of the $\text{MetroArea}_{\text{NW}}$ coefficient and the sum of the $\text{MetroArea}_{\text{NW}}$ and $(\text{MetroArea}_{\text{NW}}) \times (\text{PlayerRace})$ coefficients depend on the year. The only statistically significant coefficient is the positive coefficient associated with $\text{MetroArea}_{\text{NW}}$ when the dependent variable is games started the first year. If the composition of the team's fan attendance is similar to the racial composition of the metropolitan area, these results provide little evidence that front offices are responding to consumer discrimination when drafting nonwhite players after white players of comparable ability.

5. Alternative Explanations

5.1 Uncertainty Associated with Player's Ability

We argue in Section 4.1 that the uncertainty about players' abilities during the hiring phase makes it less costly to discriminate at this stage. At the same time, however, if the quality of information about players' abilities varies systematically with race, it could be this difference that is causing the results in the base specification. In Table 2, the coefficient on the SmallCollege variable is not significant in any of the specifications. The SmallCollege coefficient being close to zero and not statistically significant suggests that the uncertainty associated with a player's ability is not an important consideration when a team drafts a player.³¹

The search literature suggests that the effect of the variance of player's ability level on draft selection should depend on where the distribution is truncated (see Mortensen, 1986, for a review of the search literature; Hendricks et al., 2003, consider the effect of uncertainty on the NFL hiring process). The location of the truncation varies significantly across selection number. (The probability that a player drafted in the first round has an active contract is 0.96, while this probability is 0.25 for a player drafted in the twelfth round.) Therefore, as further evidence that the uncertainty about a player's ability level does not significantly affect draft selection, in another specification we include as an independent variable the interaction term $(\text{SmallCollege}) \times (\text{SelectionNumber})$.

30. A high concentration of nonwhite residents mitigates the tendency to draft white players earlier than equally productive nonwhite players. Therefore, *ceteris paribus*, a white player drafted by a team in a metro area with a relatively high concentration of nonwhite residents would be more likely to have an active contract and start more games than a white player drafted by a team in a metro area with a low concentration of nonwhite residents. However, *ceteris paribus*, a nonwhite player drafted by a team in a metro area with a relatively high concentration of nonwhite residents would be less likely to have an active contract and start more games.

31. To test for robustness, we also include as independent variables the number of times in the past 10 years that the player's collegiate team was ranked in the top 20 at the end of the season as well as indicator variables for the college's conference (e.g., Big Ten, Southeastern, Atlantic Coast, Pacific 10). All of the coefficients associated with these variables were close to zero and not statistically significant. More importantly, including these additional control variables does not appreciably change the magnitude or statistical significance of the PlayerRace coefficient.

The coefficients associated with this interactive term are neither economically nor statistically significant. These results are available upon request.

5.2 Compensation

Though the results in Table 2 provide strong evidence that race affects the employment of NFL players, we cannot attribute it to discrimination until we explore the possibility that a player's contract structure varies systematically with player race. While expected compensation may influence the front office and coach's decision in similar manners, the structure of this compensation (i.e., guaranteed compared with nonguaranteed payments and duration of contract) is likely to have different effects on the decisions of when to draft the player and whether to retain/promote the player, respectively.³² If contract structure is correlated with race, the effect of a contract on the incentives of the front office, coaches, and players could explain the race-based differences in the probability of making the team and number of starts that we observe in Table 2.

We perform three regressions to test for racially based differences in contract structure. Specifically, we regress the log of the average annual expected compensation of a player's initial contract, signing bonus, and the number of years covered by a player's initial contract on \mathbf{X} and PlayerRace. As average annual expected compensation and signing bonus vary systematically with contract duration, we include the number of years of the contract as an independent variable in these regressions.³³ Contracts range from one to six years in duration, with early draft selections usually signing contracts of longer duration than late selections.

The coefficient estimates associated with PlayerRace is slightly negative and statistically significant at the 10% level when the dependent variable is log of the average annual expected compensation. The average marginal effect of a player being nonwhite is to decrease his annual expected compensation by 1.2%. The PlayerRace coefficient is positive but not statistically significant when the dependent variable is signing bonus. When the dependent variable is contract duration, the PlayerRace coefficient is positive and statistically significant at the 10% level. Evaluated at the mean of the dependent variable (which is 2.74 years), nonwhite players agree to contract with an average duration of 0.05 years more than white players. These results, which are available upon request, suggest that while contract structure and compensation may differ slightly between white and nonwhite players, the differences are not large enough to explain the results in Table 2.

32. See Conlin and Emerson, 2003, for a thorough empirical investigation of the double moral hazard issue associated with NFL contracts.

33. While contract duration is likely to be endogenous, there is no obvious instrumental variable. The PlayerRace coefficients do not change appreciably if contract duration is not included as an independent variable.

Table 4. Effect of Discrimination on Number of Regular-Season Games Won from 1991 to 1999. OLS Regression Results (SE)

Independent Variables	Coefficients					
	Active Contract, Year			Games Started, Year		
	1st	2nd	3rd	1st	2nd	3rd
Degree of Discrimination (DEGDIS) Estimates from:						
DEGDIS	-6.18 (11.32)	-6.29 (10.99)	-7.74 (15.51)	-1.38 (2.69)	-1.53 (6.73)	-3.39 (5.59)
Constant	73.7** (4.8)	73.7** (5.0)	73.8** (6.0)	72.8** (2.6)	73.5** (13.2)	76.1** (8.8)
R^2	0.03	0.04	0.05	0.01	0.01	0.07
Sample size	28	28	28	28	28	28

OLS = ordinary least squares.

*Statistically significant at the 0.10 level; **Statistically significant at the 0.05 level (two-tailed tests).

Note: Standard errors are calculated based on the estimation procedure in Murphy and Topel (1985).

6. The Performance Cost of Discrimination

If, as our results suggest, NFL teams discriminate against players based on race in their hiring decisions, then it is reasonable to expect there to be a cost for such discriminatory practices. The cost to NFL teams is in on-field performance, for teams that discriminate settle for less productive workers of the preferred race. Therefore, teams that discriminate more should win fewer games.

In order to test for this effect we alter the base specification by interacting the indicator variable for each team with the PlayerRace variable.³⁴ The coefficients on these interactive variables provide a measure for each team of the difference between hiring and retention/promotion decisions based on player race. (The coefficient estimates on these interactive variables are negative for only a few of the teams. A large coefficient suggests that the team discriminates more in its hiring practices.) We then regress the total number of regular-season games won by each team from 1991 to 1999 on these coefficient estimates from each of the six regressions (using degree of discrimination [DEGDIS]). If teams' on-field performance is adversely affected by discriminatory practices, then teams that practice more discrimination should win fewer games, resulting in a negative DEGDIS coefficient. The results of regressing total team wins on DEGDIS are given in Table 4. Because DEGDIS is estimated in the first stage, the standard errors in Table 4 are corrected for this using the methodology in Murphy and Topel (1985).

The coefficient estimates, while not statistically significant, are negative for all six regressions. Given that there are only 28 data points, these results are quite striking and provide some evidence that discrimination hurts an NFL team's

34. For those regressions where the dependent variable is number of starts, we estimate a linear regression instead of an ordered probit. When estimating an ordered probit, we are unable to obtain standard errors when team interactive terms are included in the specification.

on-field performance.³⁵ If this is the case, the obvious question is, why do teams continue to discriminate? The answer appears to be that teams are willing to pay the price for their own taste for discrimination, and/or that consumers with racially based preferences serve to compensate teams that discriminate.³⁶

7. Conclusion

This paper uses an outcomes-based approach to empirically test for employment discrimination. We test whether an employee's race has a differential effect on a firm's hiring decision versus the firm's retention and promotion decisions using draft, retention, and performance information from the NFL. Our results indicate that white players have a 0.13 lower probability of having an active contract and start 1.56 games less than nonwhite players (after controlling for draft selection, position, team, draft year, collegiate division, and team wins the prior season). These results provide strong evidence that nonwhite players are being discriminated against in the hiring process. Additional specifications provide no evidence that coaches engage in discriminatory activity when making their decisions of whom to retain, start, and train.

The existing literature on employment discrimination consists mainly of tests for earnings and employment disparities of individuals of different races using regression analyses and audit studies that send pairs of individuals and then compare their treatment in the labor market. The outcomes-based approach used in this paper to test for employment discrimination avoids the main criticisms of this regression analysis (i.e., omitted variable bias) and the audit studies (i.e., the inability to perfectly match testers). While this paper examines only players drafted by an NFL team, a differential effect of a person's race on a firm's hiring decision versus the firm's retention and promotion decisions is likely to exist in many industries because the costs and benefits of discriminating are likely to differ in the hiring stage compared with the retention and promotion stage.

35. In another attempt to find evidence of consumer discrimination, we test whether the proportion of the team's metropolitan area that is nonwhite has an effect on the team's degree of discrimination. This is done by regressing DEGDIS on $MetroArea_{NW}$. If team's discrimination is due, at least in part, to consumer preferences, and if the proportion of the metropolitan area that is nonwhite is a proxy for the racial composition of the fans in attendance at the games, then the relationship between DEGDIS and $MetroArea_{NW}$ should be negative. In this regression all the coefficient estimates are positive when $MetroArea_{NW}$ is regressed on DEGDIS coefficient estimates from the probit regressions, but none are significant. However, all of the coefficient estimates are negative when $MetroArea_{NW}$ is regressed on DEGDIS coefficient estimates from the ordered probit regressions, and the coefficient estimate is significant at the 10% level in the second-year regression. The results here are decidedly mixed and provide no evidence of consumer discrimination. However, since $MetroArea_{NW}$ is not necessarily a good representation of the racial composition of NFL teams' paying fans and as these results are mixed, we do not consider this strong evidence for or against consumer discrimination. These results are available upon request.

36. In addition to the results from the regression including the racial composition of the metropolitan area that suggests that consumer discrimination is not a driving force, the revenue-sharing arrangement among NFL teams likely reduces the cost a team incurs from discriminating.

We find strong evidence of race-based hiring discrimination in the NFL. Economic theory (Becker, 1971; Arrow, 1973) suggests that discrimination in the marketplace is likely reduced by competition. The NFL, however, affords teams protection from competition by restricting new entrants and through extensive revenue sharing. In addition, the NFL labor market has an antitrust exemption that enables the NFL to avoid the unwanted side effects of competition, that is, teams bidding against each other for the best new players. Because these protections from competition exist, there is no reason to believe that hiring discrimination will be driven out of the NFL by market forces.

Appendix

This appendix presents a model in which the quality of the team is affected by the players' beliefs of whether the team is discriminatory. The implications of the model are that the team's cost of discriminating is less if the team discriminates in its hiring decisions than if the team discriminates in its retention, training, and starting decisions, and if the signal that the team receives on the player's ability level is noisier.

Let $P(q, Q_T(Q_D, b), w_T(w_D))$ be a team's inverse demand function, where q is the number of tickets sold, Q_T is the quality of the team, and w_T is the fraction of white players on the team. Team quality is a function of the quality of the draft (Q_D) and the players' beliefs of whether the team discriminates (b). Increasing the quality of the draft increases team quality, while increasing the players' beliefs that the team is discriminatory decreases team quality. If players believe that the team discriminates, it will likely affect the players' performance and increase the probability that players will demand a trade or will sign a contract with another team. The fraction of white players on the team is a function of whether the team drafts white players (w_D) and enters the team's inverse demand function because of possible consumer discrimination. A team's cost is a function of the number of tickets sold, and team quality ($C(q, Q_T(Q_D, b))$). $B(w_T(w_D))$ represents the team's taste for white players (i.e., employer discrimination). Thus, the team's payoff function is $P(q, Q_T(Q_D, b), w_T(w_D))q - C(q, Q_T(Q_D, b)) + B(w_T(w_D))$, where q is constrained by stadium capacity.³⁷

Let p be the probability the best available player to draft is white and θ be the probability the team receives the correct signal on the race of the best available player, where $\theta > 0.5$. The probability that the front office receives a signal that the best available player is non-white is $P_r = p(1 - \theta) + (1 - p)\theta$. We expect that if the team drafts a white player, the other players' beliefs that the team is discriminatory (b) would increase with an increase in P_r (perhaps due

37. Discussions with NFL representatives indicate that approximately 80% of a team's revenue during these years came from revenue-sharing sources; most notably television contracts. This source of revenue for the team is not likely to be substantially affected by a team discriminating during the draft. If this is the case, this revenue will enter the front office's payoff function as a constant and not affect its decision to discriminate.

to Bayesian updating). Discrimination in the retention, training, and starting decisions would have a much larger effect on the other players' beliefs because the team and players have much more information on the drafted player's ability level at the time these decisions are made. This is why teams prefer to discriminate in their drafting decisions rather than in their retention, training, and starting decisions.

The probability that the team receives a signal that the best available player is nonwhite increases with θ ($\partial P_r / \partial \theta > 0$) if $p < 0.05$.³⁸ Because teams have more information on a player's ability level if that player participated in an NCAA Division 1A collegiate football program, the signal that the team receives on these players is likely to be less noisy. Therefore, θ is likely to be greater for players whose collegiate experience was in a Division 1A program. If this is indeed the case and team quality is a function of the players' beliefs, the team's cost of discriminating in its drafting decisions is less if it selects a player from a non-Division 1A program.

References

- Arrow, Kenneth J. 1973. "The Theory of Discrimination," in O. Ashenfelter and A. Rees, eds., *Discrimination in Labor Markets*. Princeton, NJ: Princeton University Press.
- Ayres, Ian. 2001. *Pervasive Prejudice? Unconventional Evidence of Race and Gender Discrimination*. Chicago: University of Chicago Press.
- Ayres, Ian, and Peter Siegelman. 1995. "Race and Gender Discrimination in Bargaining for a New Car," 85 *American Economic Review* 304–21.
- Ayres, Ian, and Joel Waldfogel. 1994. "A Market Test for Race Discrimination in Bail Setting," 46 *Stanford Law Review* 987–1047.
- Baker, George, Michael Gibbs, and Bengt Holmstrom. 1994. "The Internal Economics of the Firm: Evidence from Personnel Data," 109 *Quarterly Journal of Economics* 881–919.
- Becker, Gary. 1968. "Crime and Punishment: An Economic Approach," 76 *Journal of Political Economy* 169–217.
- . 1971. *The Economics of Discrimination*, 2d ed. Chicago: University of Chicago Press.
- Berkovec, James A., Glenn B. Canner, Stuart A. Gabriel, and Timothy H. Hannan. 1998. "Discrimination, Competition, and Loan Performance in FHA Mortgage Lending," 80 *Review of Economics and Statistics* 241–50.
- Conlin, Michael. 1999. "Empirical Test of a Separating Equilibrium in National Football League Contract Negotiations," 30 *The Rand Journal of Economics* 289–304.
- Conlin, Michael, and Patrick M. Emerson. 2003. "Multi-Dimensional Separating Equilibria and Double Moral Hazard: A Study of the National Football League Contract Negotiations," 85 *Review of Economics and Statistics* 760–65.
- Darity, William A., Jr., and Patrick L. Mason. 1998. "Evidence on Discrimination in Employment: Codes of Color, Codes of Gender," 12 *Journal of Economic Perspectives* 63–90.
- Donohue, John J., III, and Peter Siegelman. 1991. "The Changing Nature of Employment Discrimination Litigation," 43 *Stanford Law Review* 983–1033.
- Ehrlich, Isaac. 1996. "Crime, Punishment, and the Market for Offenses," 10 *Journal of Economic Perspectives* 43–67.

38. Thirty-six percent of players drafted were white in our data set, which contains over 90% of all players drafted into the NFL from 1986 through 1991. In addition, almost three-quarters of all players in the NFL are nonwhite. We therefore expect p to be much less than 0.5.

- Epstein, Richard A. 1992. *Forbidden Grounds: The Case Against Employment Discrimination Laws*. Cambridge, Mass.: Harvard University Press.
- Heckman, James J., and Peter Siegelman. 1993. "The Urban Institute Audit Studies: Their Methods and Findings," in M. Fix and R. J. Struyk, eds. *Clear and Convincing Evidence: Measurement of Discrimination in America*. Washington, D.C.: The Urban Institute Press.
- Hendricks, Wallace, Lawrence DeBrock, and Roger Koenker. 2003. "Uncertainty, Hiring and Subsequent Performance: The NFL Draft," 21 *Journal of Labor Economics* 857–86.
- Kahn, Lawrence M. 1991. "Discrimination in Professional Sports: A Survey of the Literature," 44 *Industrial and Labor Relations Review* 395–418.
- . 1992. "The Effects of Race on Professional Football Player's Compensation," 45 *Industrial and Labor Relations Review* 295–310.
- Kahn, Lawrence M., and Peter D. Sherer. 1988. "Racial Differences in Professional Basketball Players' Compensation," 6 *Journal of Labor Economics* 40–61.
- Killingsworth, Mark R. 1993. "Analyzing Employment Discrimination: From the Seminar Room to the Courtroom," 83 *American Economic Review* 67–72.
- Knowles, John, Nicola Persico, and Petra Todd. 2001. "Racial Bias in Motor Vehicle Searches: Theory and Evidence," 109 *Journal of Political Economy* 203–29.
- Laband, David N., and Michael J. Piette. 1994. "Favoritism Versus Search for Good Papers: Empirical Evidence Regarding the Behavior of Journal Editors," 102 *Journal of Political Economy* 194–203.
- Mogull, Robert G. 1973. "Football Salaries and Race: Some Empirical Evidence," 12 *Industrial Relations* 109–12.
- Mortensen, Dale T. 1986. "Job Search and Labor Market Analysis," in O. Ashenfelter and R. Layard, eds., *The Handbook of Labor Economics*, vol. 2. Amsterdam: North-Holland.
- Murphy, Kevin M., and Robert H. Topel. 1985. "Estimation and Inference in Two-Step Econometric Models," 3 *Journal of Business and Economic Statistics* 370–79.
- Ross, Stephen L. 2003. "What Is Known About Testing for Discrimination: Lessons Learned by Comparing Across Different Markets," working paper 2003–21R, University of Connecticut Department of Economics.
- Ross, Stephen L., and John Yinger. 2002. *The Color of Credit: Mortgage Lending Discrimination, Research Methodology, and Fair-Lending Enforcement*. Cambridge, Mass.: MIT Press.
- Szymanski, Stefan. 2000. "A Market Test for Discrimination in the English Professional Soccer Leagues," 108 *Journal of Political Economy* 590–603.
- Yinger, John. 1998. "Evidence on Discrimination in Consumer Markets," 12 *Journal of Economic Perspectives* 23–40.