Uncertainty, Hiring, and Subsequent Performance: The NFL Draft

Wallace Hendricks, *University of Illinois*

Lawrence DeBrock, *University of Illinois*

Roger Koenker, *University of Illinois*

In this article, we analyze the impact of uncertainty on the hiring process. We show the connection between models of statistical discrimination where uncertainty can work against groups that have less reliable indicators of future productivity and models of option value where uncertainty about future productivity can be beneficial for these groups. These models generate hypotheses about the relationship between ex ante hiring patterns and ex post productivity. This is applied to the market for NFL football players. We provide various estimates of NFL success, which suggest that statistical discrimination and option value influence choices in this market.

I. Introduction

Information problems are fundamentally important to the outcomes of most markets. This is perhaps most obvious in labor markets. Uncertainty about the true qualities and effort levels of workers and candidates has attracted much attention from economists. When information is somehow limited, employers are unable to acquire perfect information about worker productivity. Given such uncertainty, it is common to turn to group identification as a signal of underlying productivity.

Firms use a variety of sources to gather information to predict future productivity. Interviews, prejob test scores, and letters of recommendation, for example, provide both objective and subjective information, and

[*Journal of Labor Economics*, 2003, vol. 21, no. 4]

0734-306X/2003/2104-0005$10.00

they vary widely in total cost. Given the nature of information, it is not possible for the firm to discover the true ability of each worker prior to actual employment. Firms will therefore look to information gathered from signals that have proven effective in prediction of worker performance. One such signal is group identification. The battery of tests that the firm uses will result in "scores," but these scores may be noisier for some groups than for other groups.

Previous literature in this area has focused on the relationship between the formulation of the signal and the resulting outcomes for groups. The signal itself could be biased due to prejudice (Becker 1957), or the signal could be inefficient. There are two basic approaches to this inefficiency problem. In the first approach, a highly inefficient signal for one group can put that group at a disadvantage (Phelps 1972; Arrow 1973; Akerlof 1976; Aigner and Cain 1977; Borjas and Goldberg 1978; Bloch and Rao 1993; Coate and Loury 1993; Cornell and Welch 1996). This is called statistical discrimination. Job opportunities are different across groups, even though productivity might be the same. In the second approach, selection between applicants is not viewed as a one-shot choice. Instead, firms are allowed to observe new employees over some probationary period and determine their true productivity. In this case, Lazear (1995) shows that applicants from high-risk groups are more valuable to the firm if the firm has some way of capturing rents associated with knowing this value. Thus, uncertainty improves the job prospects of members of these groups. This is called option value.

In this article, our concern is somewhat different. Rather than focus on the outcome of this formation of expectations, we consider the relationship between ex ante expectations and ex post realizations. That is, q1 our concern is the relationship between prior beliefs and realized outcomes. This allows us to generate empirical predictions from both the statistical discrimination and option value models.

As an example of our approach, consider the case of economics departments. This is an attractive market to consider as ex post productivity is clearly observable. However, to explore this issue further, we would need to determine the ex ante rankings of the candidates. Then we could compare these preemployment ratings with the ex post realization of productivity, the career results of these economists. Unfortunately, such a methodology is flawed by the critical "nature versus nurture" problem. By definition, overvalued candidates will be placed in higher-level departments through sorting in the market. These departments will have environments that are necessary to produce high-quality research. The success of undervalued candidates, on the other hand, may be retarded by their placement in schools with less attractive research environments. Thus, it is not easy to determine how much career success is due to the candidate's training (their "nature") and how much is due to professional growth on the job (their "nurture").

Most occupations that lend themselves to such postemployment observability of productivity are subject to this nature versus nurture problem.[1] If we are to investigate what the ex ante–ex post relationship should look like, we need an industry where the higher-valued candidates are not systematically sorted into the better nurturing environments.

Our response to this problem is to examine employment practices in the National Football League (NFL). We chose the NFL to consider screening for several reasons. First, a particularly rich data set is available. All prospective employees in the NFL are hired after successful college football careers. The individual teams spend a large amount of money in scouting players, analyzing the records of their college performance, and administering additional physical and mental examinations. Once each team has processed this information on all players, it is able to determine a complete ranking of the talent available.

Second, the NFL does not allow potential players to choose among the teams. Instead, a regimented system of "player rights" is exercised in the annual draft. Each team selects a player in a predetermined sequence. These draft positions are allocated according to a formula based on the order of finish in the previous NFL season. Therefore, there is much less likelihood that talented players will be matched with teams that provide greater opportunity for a successful career.

In Section II, we present a simple model of uncertainty and employment screening. This allows us to make empirical predictions about the relationship between preemployment evaluation and observed success. In Section III, we provide institutional detail about the NFL as an empirical test of our predictions. Section IV presents the results of our empirical tests. In Section V, we conclude with some observations and extensions.

II. Statistical Discrimination, Employer Risk Aversion, and Option Value

As in the Aigner and Cain model (1977), suppose that there are two groups of potential job applicants. Employers observe the group affiliation of each applicant, and they observe a signal of individual ability, say a test score, because of preemployment screening. We will assume, for the sake of simplicity, that the joint distribution of the signal, x, and the

[1] This problem is specific to that small subset of industries where we have "good" data on productivity, such as our example of academic economists. In most cases, our measures of preemployment quality and postemployment productivity amount to nothing more than the observations "Did they hire them?" followed by "Did they stay at the firm?"

(eventually) observed productivity, y, is bivariate normal for each group, with mean vector $\mu = (\alpha, \alpha)'$ and covariance matrix,

$$\Omega = \sigma_i^2 \begin{pmatrix} 1 & \rho_i \\ \rho_i & 1 \end{pmatrix}, \qquad i = 1, 2.$$

As in the Aigner and Cain formulation, unconditional mean productivity is the same for both groups; only the covariance parameters differ. Conditional mean productivity, given the signal, can then be written as

$$\mu_i(y \mid x) = \alpha + \rho_i(x - \alpha), \qquad i = 1, 2.$$

With $\rho_1 > \rho_2$, applicants from group 1 are preferred when the signal, x, is above the unconditional mean, α, and applicants from group 2 are preferred when the signal falls below α. This occurs because the signal allows better differentiation of applicants from group 1 than applicants from group 2. This is a disadvantage for high-signal group 2 members, but it is an advantage for low-signal group 2 members. In the absence of risk aversion, employers simply choose in order of the signal value from the two groups as long as the respective signals indicate that expected productivity exceeds some designated cutoff.

A. Risk Aversion

With risk aversion, the situation is somewhat more complicated. There will be a penalty paid by members of a group where it is more difficult to evaluate their ability based on the signal. The size of this penalty, however, may not be constant because of systematic differences in the efficiency of the forecasted productivity. We continue to assume that $\rho_1 > \rho_2$, that is, that group 1's signal is more strongly correlated with productivity than is group 2's signal. In addition, we will assume that $\sigma_2^2 > \sigma_1^2$. The conditional variance of productivity, given the signal under our assumptions, is independent of the signal, that is,

$$V(y \mid x) = \sigma_i^2(1 - \rho_i^2).$$

However, this presumes that employers know the parameters of the (x, y) process. If, more realistically, we adopt the view that employers must base their decisions on an estimated version of $\mu_i(y \mid x)$, we may write for applicant j from group i

$$y_{ij} = \alpha + \rho_i(x_j - \alpha) + u_{ij},$$

where $V(u_{ij}) = \omega_i^2 = \sigma_i^2(1 - \rho_i^2)$. Employing standard least squares forecasting, we have

$$\hat{y}_{ij} = \hat{\alpha} + \hat{\rho}_i(x_j - \hat{\alpha}),$$

and

$$V(\hat{y}_{ij}) = \omega_i^2\left[1 + \frac{(x_j - \alpha)^2}{\lambda}\right].$$

For example, if the employer's estimates of (α, ρ_i) had been based on a sample $(x_i, y_i, i = 1, \ldots, n)$, then $\lambda = \sum_{i=1}^{n}(x_i - \bar{x})^2$.

Of course, there are many other possible sources of heterogeneity in predicted performance as a function of the observed signal; our simple least squares forecasting model is among the simplest of these possibilities. The model illustrates that, even with homogeneous underlying uncertainty, decision makers may face heterogeneity. It is perhaps worth emphasizing that the precise form of the heterogeneity produced by the model is not really crucial to the subsequent empirical interpretation of the model. However, the general characteristic that there is more uncertainty for extreme values of the signal will be an inherent feature of any model in which predictions at these extremes are based on smaller sample sizes than those found at more moderate values of the signal.

Now, if we assume that employers compare prospects using the Pratt (1964) approximation of the "certainty equivalent" prospect, with a coefficient of absolute risk aversion γ,

$$\mathrm{CE}(y \mid x) = E(y \mid x) - \frac{1}{2}\gamma V(y + x),$$

then we have

$$\mathrm{CE}(y \mid x) = \alpha + \rho_i(x - \alpha) - \frac{1}{2}\gamma\omega_i^2\left(1 + \frac{(x - \alpha)^2}{\lambda}\right).$$

In figure 1, we plot the risk-adjusted certainty equivalent productivity curves for group 1 and group 2 under the assumption that group 2 has a larger "risk premium."[2] The plot is drawn to illustrate a situation where the risk premium may be in effect. The solid line is the group 1 certainty equivalent line. In this group, productivity is quite accurately predictable from the signal, and, consequently, the curve is almost linear. For group 2, productivity is much more difficult to predict, and the resulting certainty equivalent curve, shown by the dotted line, is less linear. At the mean signal, the two curves coincide. Group 1 candidates are always preferred when signals are above the mean. Conditional on a "good"

[2] The statement that group 2 has a larger risk premium means that the confidence band for conditional mean productivity is wider for group 2, especially at the extremes of the distribution. For exposition, groups 1 and 2 are plotted by assuming arbitrary values of $\rho_1 = .8$; $\rho_2 = .6$; $\omega_1 = 2$; $\omega_2 = 20$.

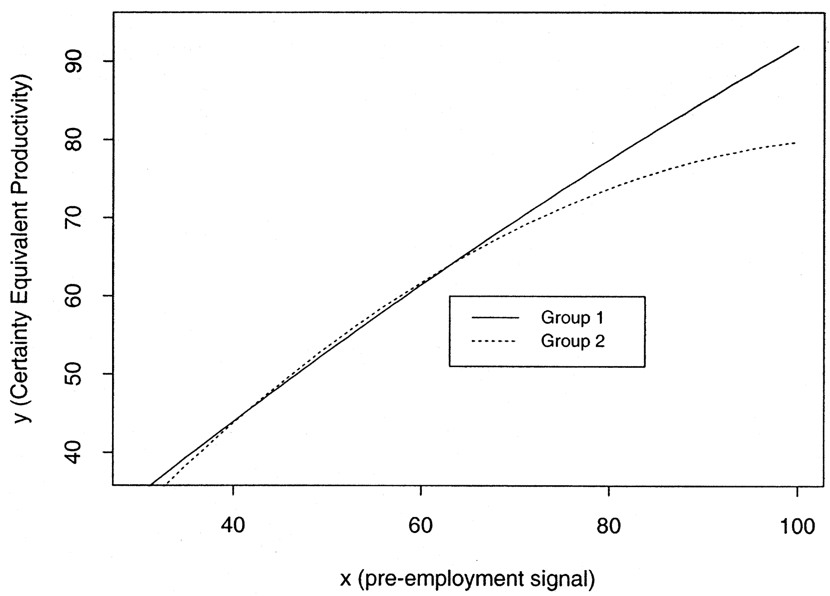

FIG. 1.—Pratt certainty equivalents conditional on preemployment signals for two groups

signal, group 1 applicants have both a higher mean and less risk than group 2 applicants with the same signal. For signals below the mean, the larger risk premium associated with group 2 attenuates their risk neutral (conditional mean) advantage. While we are left with an interval of signals in which group 2 applicants are preferred, for sufficiently poor signals, the risk effect dominates the mean advantage of group 2 and group 1 is again preferred. Therefore, it is possible that the risk aversion discount could cause employers with high risk aversion to again prefer group 1 applicants at low values of the signal. This differs from the prediction without risk aversion. This depends on actual values of the parameters in the market.

If risk aversion by employers exists, we should anticipate that group 1 individuals would dominate choices when a high signal is achieved. Now suppose that the employer ranks applicants from high to low. For scores above the mean, equivalent ranks for group 1 and group 2 individuals would indicate that group 2 individuals are expected to perform better than group 1 individuals since the group 2 individuals pay a risk aversion penalty due to the poor performance of their predictor. This difference in actual job performance should be greatest at high signals. For a given ex ante evaluation of performance, we should expect group 1 individuals to underperform ex post relative to individuals from group 2. For scores below the mean, we will again observe that group 1 individuals will underperform relative to group 2 individuals for a given low ex ante ranking. There is an area in between, however, where the reverse can occur.

B. Option Value

While the above analysis shows that uncertainty about future productivity can work against the choice of individuals from a group where the productivity signal is noisy, as Lazear (1995) shows, uncertainty can also be of benefit for a group. This occurs when the employer can take advantage of information about the worker's productivity that is gathered during an evaluation period.

In the foregoing analysis we presumed that the hiring decision was a one-shot affair. There was no subsequent stage in which productivity was evaluated and a decision to retain or dismiss the employee made. The introduction of such a probationary period, or tenure process, complicates the situation in quite interesting and important ways, which we will now address. In effect, we shall see that the retention decision introduces an option value element into the analysis and that statistical discrimination can be directly incorporated into this model.

The high variance of the group 2 candidates offers a potentially attractive opportunity to the employer who is willing and able to hire candidates from this group and retain them only if they turn out to be successful. While the upper tail of the group 2 productivity distribution is obviously attractive, the employer presumably bears some cost in the probationary period for unsuccessful hires.[3] Suppose that this "tenure review" cost is K. Then the firm can evaluate the expected value of an employee with signal x as

$$W(x) = \pi_0(x)E(y \mid x, y > y_0) - [1 - \pi(x)]K,$$

q2

where $\pi_0(x) = P(y > y_0 \mid x)$ is the probability of success given the signal. In an environment in which employers encounter potential hires from both groups, we can easily compute, for any specified K, the expected

[3] The firm would need to balance the cost of low productivity during the evaluation period plus the replacement cost of these poor workers against the gain it can achieve by paying a salary below the market value of the worker in the event that the worker becomes a "star." Of course, the firm must have some opportunity to capture rent from star performers. Lazear (1995) argues that this requires the employee to have some value that is specific to the firm. However, as he shows in his paper on offer matching (Lazear 1986), if the employing firm can gather information about the employee's productivity that is not readily observed by other firms, then any firm that attempts to raid the employee will face the winner's curse. Thus, high-risk workers can have option value if other firms cannot easily monitor their performance (or there are constraints on their mobility) and if replacement costs are not large. In some circumstances, high costs associated with incorrect hires will negate this option advantage.

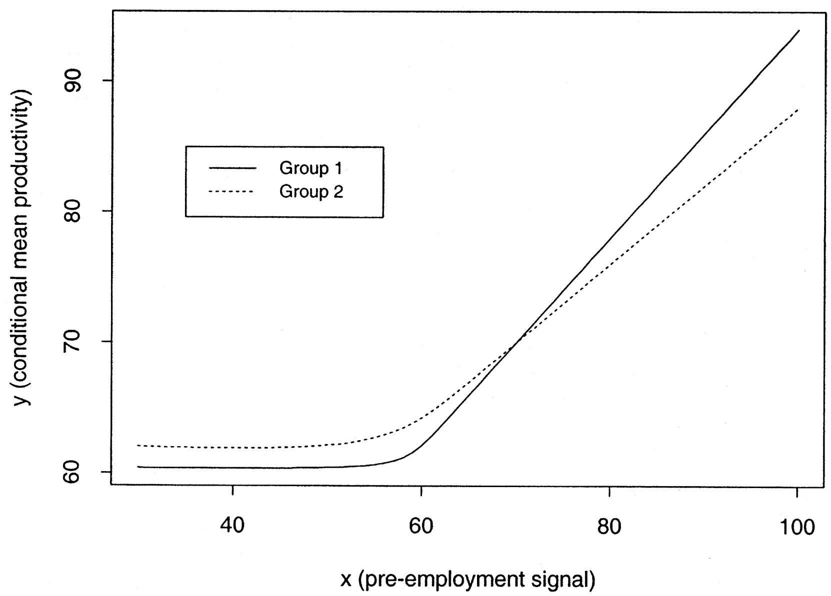

FIG. 2.—Expected productivity conditional on the observed signal with a cutoff $y_0 = 60$ for group 1 applicants and group 2 applicants.

values for the two groups and compare them to the expected productivity of successful applicants from each group.

Suppose that the firm imposes a simple rule: only employees achieving productivity level y_0 will be retained. In figure 2, we plot the expected productivity of the two groups, as well as the computed conditional expectation of productivity given both the signal and that observed productivity exceeds the threshold $y_0 = 60$. We use the same parameter values as for figure 1. The dotted line is the conditional expectation for group 2, and the solid line indicates this conditional expectation for group 1. Again, for sufficiently good signals, group 1 is preferred. But for low and moderate signals there is a clear advantage in group 2. This is due to the fact that the employer is able to truncate the left tail of the productivity distribution.

Both figure 1 and figure 2 plot the relationship between the preemployment signal (x) and the preemployment evaluation of the candidate (y). Qualitatively, above the mean they look quite similar. There is a risk discount that works against group 2 workers with high signal values. For high-signal applicants, the ability to truncate the productivity distribution at the bottom has little effect. However, below the mean in figure 2, there is a clear advantage in the choice of applicants from group 2 workers, while in figure 1, group 1 workers are favored for some of the region

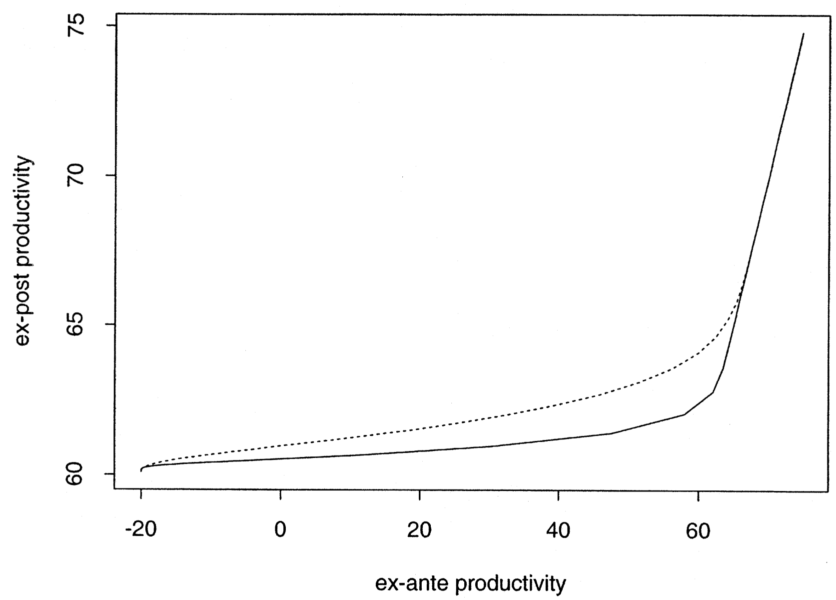

FIG. 3.—Expected productivity by successful applicants against expected value for the two groups' ex ante productivity. Cost of evaluation = K = 20.

below the mean. Why is there a difference? Risk aversion is only reasonable when downside errors generate much more disutility than the benefits gained when value is underestimated. If the downside errors can be truncated at a reasonable cost, then risk aversion can be avoided. Therefore, we can model risk aversion simply as an increase in the fixed cost K of the probationary process. With K extremely high, we expect risk avoidance to be a plausible hiring strategy. We can therefore use the option value model to make predictions about the relationship between expected productivity conditional on meeting some minimum cutoff and actual (realized) productivity on the job. These predictions can be modified for risk aversion by changing the value for the cost of identifying poor workers.

In figure 3, we plot expected productivity by successful applicants (ex post productivity) against expected value (ex ante productivity) for the two groups. The solid line corresponds to group 1 applicants, who have low variability, and the dotted line indicates group 2 applicants, who have higher variability. Here the cost of an unsuccessful applicant is set at $K = 20$. Note that, for high ex ante productivity applicants, the two groups have essentially identical ex post performance, but for lower ex ante applicants, the group 2 applicants have better ex post performance than group 1 applicants. The larger number of failures among group 2

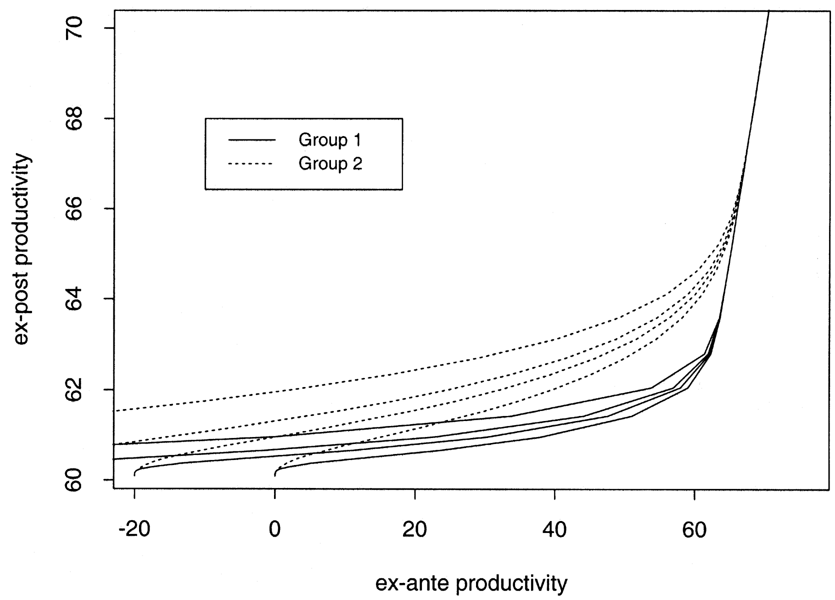

FIG. 4.—Expected productivity of successful applicants from groups 1 and 2 at different costs of evaluation. The four curves for each group depict four distinct levels of cost, $K = \in \{0, 20, 40, 100\}$, imposed on the firm by unsuccessful applicants. Higher curves correspond to higher fixed costs.

applicants is balanced by the higher quality of those who pass the minimum threshold. For a given ex ante ranking, we should observe workers passing the probationary period from group 2 to have higher average productivity than those who survive from group 1, and this difference should be greatest at lower ex ante rankings. However, we should also observe that, conditional on ex ante ranking, more group 2 workers will fail to make the minimum cutoff. This difference is not predicted by pure risk aversion.

In figure 4, we illustrate the impact of increasing the cost of the probationary period to proxy the introduction of risk aversion along with option value. The horizontal axis in figure 4 is truncated to the range 0–70 as compared with the −20–70 range in figure 3. This makes the graphs look different, even though they use the same underlying data. The convergence point is about 65 in both graphs. High cost (K) indicates large risk aversion. Changing this cost does not influence ex post productivity, but it does reduce the ex ante productivity of the two groups differentially. With a low cost of the probationary period, the employer can afford many failures with a few successes. Therefore, there will be a large difference between ex ante and ex post productivity, and this difference will be most evident for a group with high uncertainty (like group 2). This is reflected by the small difference of changing probation costs for group 1 as compared with group 2.

These models predict that the impact of uncertainty will depend on the cost of an incorrect hire. If a mistake in hiring leads to large costs, statistical discrimination is likely to result. If a mistake can be made with low cost, option value will likely favor the group with the most uncertainty. When arraying on mistake costs, we expect to observe hiring differences in the tails of the distribution of productivity that are in opposite directions. Our empirical specification will require measures of productivity at both tails of the distribution.

III. The NFL Draft

Each year, all potential NFL entry-level players must participate in the NFL draft. This is a procedure wherein each team selects a player in a predetermined sequence. The sequence is arranged in reverse order of standings for the previous NFL season in an attempt to improve parity.[4] For the time period of our analysis, there were 12 rounds of the draft. With 28 NFL teams, this meant that 336 players were drafted each year.[5]

All NFL teams have access to an overwhelming array of individual player data compiled during the college career of each player. In addition, the teams jointly operate a common camp where a large number of newly eligible NFL hopefuls gather prior to the draft. This camp is a thorough physical examination of the players' health as well as of their relative abilities.[6] Using all of this combined information about each player, each team must make an estimate of that player's value. It is this estimate and its relationship to the ultimate productivity of the player that form the basis for our analysis.

[4] Of course, since drafted players can be and are traded freely, the imposition of the draft per se has little impact on the final allocation of talent. Teams with higher values will get the player via interteam transactions rather than via the initial allocation.

[5] In fact, more than that number are actually drafted, on average. This is because the NFL holds supplementary drafts at unscheduled times. These drafts handle the allocation of players entering NFL eligibility under extremely special circumstances, such as sudden and unexpected loss of college eligibility. In the cases where NFL rules required teams to give up first-round draft picks to select in the supplemental draft, we assigned first-round draft numbers to the players. In the special case of the U.S. Football League supplemental draft, which occurred when the rival USFL disbanded, we include a dummy variable for all players selected in this draft.

[6] Some players acquire injuries in the period after their college experience but before the draft. Other players may have had unannounced injuries. Finally, some colleges are notoriously inaccurate in their physical descriptions of the players. The physical examinations at this camp reconcile these problems.

To test for statistical discrimination against one group based on the use of a noisy signal of future productivity or for option value in favor of one group based on similar uncertainty, we first need to identify the two groups. For the first group, the preemployment predictor should perform better than for the second group. This translates to the evaluator being able to distinguish the performances of group 1 members from the average for the group better than he can for members of group 2. In the limit, the evaluator would simply predict the mean of group 2 if there were no way to distinguish its members. This will generate a larger variance in the predicted values for the first group. Cornell and Welch (1996) show that the larger variance in predicted outcomes can be expected to lead to a larger selection from this group since there will be more members in the right tail of the distribution. This should be distinguished from larger uncertainty about the predicted values. Thus there should be more variance in the scores for one group, but the uncertainty associated with any individual score should be less.

In our model of Section II, we have assumed for convenience that both groups have identical mean productivity; only the variability of predicted productivity differs between the two groups. In many situations, there can be both mean and variance effects. For example, individuals with less experience are likely to have both lower means and higher variance of outcomes. Mean effects are accounted for in the empirical specification by the presence of several predictors of individual performance available ex ante to the teams and by the inclusion of the group indicator variable. Departures from the conditional mean determined by these predictors found in the interaction effects between the group indicators and the ex ante observables are then attributed to variance effect.

We have selected individuals from the schools that play Division IA football as group 1, with individuals from the remaining schools as group 2. The Division IA schools consist of the major football universities in the NCAA, and they have much higher profile programs than other schools. These Division IA schools can be assumed to attract and train football players better than the other schools.[7]

[7] We have some support for this assumption. In preparation for the 1996 NFL draft, a league reporting service called the NFL Combine did evaluations on approximately 770 players producing Combine ratings which range from 7.40 to 9.89 (mean = 7.76), with 7 indicating a player who probably will not make the NFL and 10 indicating a player who almost surely will start in his first year. These ratings represent the expected value of players as predicted from the Combine tests. Athletes from Division IA schools ($n = 542$) score higher (7.83 vs. 7.59; $t = 5.49$) than non–Division IA athletes ($n = 231$) and have a larger standard error (.62 vs. .40; $F = 73.9$) for their rating. For these data to be consistent with the model, the ratings for non–Division IA athletes should also be more difficult to predict. We have performed regressions that indicate the Combine rating is more predictable for schools that have had traditionally stronger football programs. Scouts make larger

Table 1
Players on 1996 Opening Roster by Draft Round

Round	On Roster	Not on Roster	Total
1	30	0	30
2	31	0	31
3	31	3	34
4	36	1	37
5	24	11	35
6	26	16	42
7	24	21	45
Undrafted	69	464	533

It is possible that option value and risk aversion will dominate in different segments of the draft. If all drafted players in a round far exceed the minimum cutoff, option value is likely to be small. Conversely, if all players are unlikely to meet the minimum, then option value might be quite important while risk aversion is relatively unimportant. To investigate this possibility, table 1 tabulates players on the basis of their draft round and whether or not they made the opening roster of any team in the NFL. A total of 271 new players made NFL rosters in 1996. Every player in the first two rounds of the draft and 67 of 71 players in the next two rounds made opening rosters. On the other hand, only 50 of 87 players in the last two rounds made teams, and only 69 players of 533 additional players, who were evaluated but not drafted, made the rosters. These results suggest that option value may be quite important in the later rounds of the draft and for nondrafted players. Very few candidates will meet the minimum acceptable standards, so that choosing from the higher variance group might be more attractive. A discount for uncertainty about the player' s future performance, if it is important, is likely to occur during the first few rounds, where players receive higher salaries and long-term contracts.

In addition to the 1996 data, we have gathered data on all college players who were drafted into the NFL between 1979 and 1992 (4,765 players) and all NFL players who began their careers in the NFL during the same period but who were not drafted ("free agents," 958 players).[8] For 1989

distinctions in ability for the Division IA athletes than they do for the non–Division IA athletes, based on measured criteria. We also estimated models with interactions between position and weight, height, and time in the 40-yard dash. These models gave the same results as the more parsimonious representation without interactions. The difference in fit for the Division IA and non–Division IA players was even more dramatic since the interactions as a group were significant for the Division IA players but not significant for the non–Division IA players. They therefore increased the adjusted R^2 for the Division IA athletes and reduced the R^2 for the non–Division IA athletes. These results are available from the first author.

[8] See table 5 for details on the sources of our data.

only, we have data on salary for all NFL players. We do not have information on the pool of potential NFL players. These players could come from the senior classes of colleges that play football. They could also come from underclassmen, seniors from the previous years, participants in the Canadian Football League, or players who never attended college. We also have no information on players who were not drafted and who were unsuccessful in their attempt to be employed with an NFL team for these time-series data.

Of the players who were drafted by the NFL, approximately 67% (3,171) actually were employed for at least one game by an NFL team. Of the NFL players who began their careers after 1978, 77% were drafted, with the remaining being free agents. Of the 4,129 NFL players, over these 14 years, approximately 9% were judged to be talented enough to play in at least one Pro Bowl event (the all-star game for the NFL).

In 1992, there were 106 universities that had Division IA football teams. There were approximately 9,000 men playing football at these schools. We estimate that approximately 31,500 men participated in Division IA football over the time period of our data. Of these players, the NFL drafted approximately 12%. The percentage of potential players from all possible sources who are actually drafted in any year is, therefore, likely to be less than 12%. Of the nondrafted players, only 2% were employed by an NFL team, as compared with 69% of the drafted players.

If differences in uncertainty about future performance influence the teams' ex ante evaluation of the players, as given by their draft position, then we should observe differences in career performance between Division IA and non–Division IA players when we condition on draft round. We have gathered several proxies for performance to test this prediction.

Table 2 provides information about career length for NFL players. Although the average career is longer for athletes who come from Division IA schools (4.75 vs. 4.5 years for non–Division IA players), when we hold constant the position where the players where drafted, Division IA players have shorter careers in nine of the 12 rounds and also among free agents. Over the full sample, Division IA players' careers are longer because they are drafted higher. Thus, if we use career length as a measure of productivity, there is some indication that non–Division IA players might be better than Division IA players for the same ex ante evaluation. This could indicate statistical discrimination or option-value effects.

The NFL Pro Bowl is the all-star game that is played once per year. Table 3 shows the probability of appearance in the Pro Bowl during a player's career and the percentage of years in the league when the player appeared in the Pro Bowl. Both statistics show a similar pattern. In eight or nine of the rounds, the non–Division IA athletes have a higher probability of appearing in the Pro Bowl. However, the largest differences appear in the early draft rounds, supporting a risk aversion interpretation

Table 2
Average Years Played for NFL Players by Round of Draft and College Division, 1979–92

Round	Non–Division IA	Division IA	Difference
1	7.78	7.18	.60
2	6.26	5.95	.31
3	5.87	5.12	.75
4	5.00	4.79	.21
5	4.24	4.42	−.18
6	4.79	4.10	.69
7	3.79	3.96	−.17
8	4.88	4.55	.33
9	4.32	4.36	−.04
10	4.90	4.46	.44
11	4.25	3.92	.33
12	3.94	3.75	.19
Free agent	3.80	3.48	.32
Average	4.50	4.75	−.26

NOTE.—The table includes only players who appeared in at least one NFL season. Averages include players with incomplete careers.

but not an option value interpretation if we assume option value is not important in the early rounds. Again, non–Division IA players have productivities that exceed the Division IA players for a given ex ante evaluation.

Table 4 provides information for all players drafted between 1979 and 1992. These data provide a different picture from the results for career length and all-star quality. In eight of the 12 rounds, non–Division IA players were less likely to actually ever play. This indicates that teams were willing to take chances on marginal non–Division IA athletes more

Table 3
Pro Bowl Appearances by Round and Division

	Percentage of Players in Pro Bowl			Percentage of Career in Pro Bowl		
Round	Non–Division IA	Division IA	Difference	Non–Division IA	Division IA	Difference
1	34.5	29.1	5.4	13.1	10.9	2.2
2	20.0	15.6	4.4	6.2	4.8	1.4
3	11.7	8.9	2.8	3.4	2.5	1.0
4	6.6	4.2	2.4	2.5	1.1	1.4
5	5.6	4.5	1.0	2.0	1.3	.7
6	5.5	3.2	2.3	2.1	1.1	1.0
7	3.3	3.9	−.6	2.4	1.1	1.3
8	3.1	1.6	1.6	1.4	.7	.6
9	1.9	2.4	−.6	.6	.9	−.3
10	1.0	2.7	−1.8	.3	1.1	−.8
11	1.7	.4	1.4	1.1	.2	.9
12	1.0	1.4	−.4	.8	.9	−.1
Free agent	4.7	3.6	1.0	1.1	.9	.1
Average	5.4	7.0	−1.6	2.1	2.7	−.7

NOTE.—The table includes players with incomplete careers.

Table 4
Percentage of Draftees Who Actually Played by Round and Division

Round	Non–Division IA	Division IA	Difference
1	92.3	98.9	−6.6
2	91.3	95.9	−4.6
3	89.2	91.2	−2.0
4	89.0	87.1	2.0
5	76.0	80.1	−4.1
6	72.9	69.0	3.9
7	55.9	67.3	−11.4
8	50.0	60.7	−10.7
9	41.5	48.3	−6.7
10	39.6	37.8	1.8
11	38.3	34.1	4.2
12	32.3	37.1	−4.8
Average	70.0	75.1	−5.0

often than they were on Division IA athletes. The evidence from tables 3 and 4 is consistent with the existence of more uncertainty over the expected productivity of non–Division IA players. The table 4 results are consistent with the option value model.

Position in the draft provides significant information about the pre-employment evaluations of the player by the teams in the NFL. Further information is provided by the failure of the teams to draft a player. In this case, the player must have been valued lower than all players in the draft. More players fail among those who are drafted than succeed among free agents. Thus, all free agents are "mistakes" in the sense that they could have been selected in the draft. For the years 1979–92, 34% of free agents came from non–Division IA schools, while players from these schools made up only 17.7% of the drafted players.

These results provide a prima facie case for different evaluations of players who come from schools where the talent of the players is probably harder to evaluate. These results are tested more formally in tables 6–8, where we control for characteristics of the school and the athletes. Table 5 lists the variables that we use to analyze the NFL screening problem. The data include characteristics of the school and the athletes. There are four variables that characterize the athlete's college choice: DIVIA is a measure of the level of competition that the player experienced, YEARSTOP30 is the actual competitiveness of the player's college over the period 1980–92, and BLACK and PRIVATE are dummy variables to capture other characteristics of the school. We have 15 variables that measure the salient characteristics of each athlete in the data set. The variable FREE is a dummy variable representing players who played in the NFL but who were overlooked during the draft ("free agents"). Following Kahn (1992) we use the

Table 5
Variable Names and Sources

Name	Description	Source
DRAFTN	Inverse of position in draft [1/(draft position)]	e, g
DIVIA	Dummy variable indicating membership in Division IA of the NCAA	b
BLACK	Dummy variable indicating membership in traditionally black college	c
PRIVATE	Dummy variable denoting private school	a
YEARSTOP30	Number of top 30 finishes by school in the period 1980–92	b
SUPP84	Dummy variable for the supplemental draft of World League players in 1984	
DFT × TOP30	Interaction term, YEARSTOP30 × DRAFTN	
DRAFT × IA	Interaction term, DIVIA × DRAFTN	
WHITE	Dummy variable indicating the race of the player (one if white, zero otherwise)	f
GRADUATE	Dummy variable indicating that the player graduated from college	e
YEARS	Years played in the NFL	e, g
PLAY	Dummy variable indicating that the drafted player actually appeared in the NFL	e, g
PROBOWL	Percentage of player's active NFL years that he appeared in the Pro Bowl	e, g
SALARY	1989 player salary	h
FREE	Dummy variable indicating that the player was a nondrafted free agent into the NFL	e, g
QB	Dummy variable indicating that the player was a quarterback	e, g
KICKER	Dummy variable indicating that the player was a kicker	e, g
BACKER	Dummy variable indicating that the player was a linebacker	e, g
LINE	Dummy variable indicating that the player was a lineman	e, g
TE	Dummy variable indicating that the player was a tight end	e, g
RECEIVER	Dummy variable indicating that the player was a wide receiver	e, g
DB	Dummy variable indicating that the player was a defensive back	e, g

SOURCES.—(a) Peterson's Guides, *Peterson's Annual Guides to Undergraduate Study: Guide to Four-Year College* (Princeton, NJ, 1992); ACT scores were converted to SAT scores following a chart of SAT-ACT equivalents provided by the University of Illinois. (b) We counted a school as "making" the NCAA championships if their team finished in the top 30 of the final Dunkel Rating (photocopy obtained from author for various years). (c) Office for the Advancement of Public Negro Colleges, *Directory of Traditionally Black Colleges and Universities in the United States*, September 1971. (d) National Collegiate Athletic Association, *1993 NCAA Division I Graduation-Rates Report* (National Collegiate Athletic Association 1993). (e) *The Sporting News Football Register,* various editions; years played and Pro Bowl appearances are measured through the 1996 season for players entering the league between 1979 and 1992. (f) Kahn 1992*; The Sporting News,* various editions; team photos. (g) Neft, Cohen, and Korch (1992). (h) Kahn (1992) supporting data source.

inverse of draft position (DRAFTN) as a measure of the ex ante evaluation.[9] Players selected near the top of the draft have a high DRAFTN. The variable DRAFTN is our best estimate of the ex ante evaluation of the player that was made by the teams. Each free agent was assigned a position that was "one higher" than the number assigned to the last drafted player that year.[10] The variables PROBOWL, PLAY, and YEARS offer different measures of the actual effectiveness of the player as a professional athlete. The variables

[9] Kahn (1992) used the inverse of draft round, while we use the inverse of the draft number.

[10] Thus, if there were 12 rounds for 28 teams in a particular year, all free agents in that year would be assigned draft position 337, one number higher than the last player selected in the twelfth round.

GRADUATE,[11] WHITE, and SUPP84 are additional control variables.[12] Finally, we add a series of dummy variables that control for the position of each player on a professional football team: QB, KICKER, BACKER, LINE, TE, RECEIVER, and DB. The dummy variable coefficients are all estimated with respect to the omitted category of running back.

Every player in the NFL is categorized as drafted (FREE = 0) or undrafted (FREE = 1). Free agent status among actual NFL players is estimated using binomial probit while controlling for school and athlete characteristics. We expect that free agents will be drawn more often from non–Division IA schools and schools that appear less often in the top 30. Our remaining measures of playing status are all contingent on the draft position of the player. In addition to the probability that a drafted athlete actually plays in the NFL, we also estimate the player's expected career length (number of years played) and the percentage of Pro Bowl appearances over the player's career. Persons selected at the beginning of the draft should be expected to have a higher probability of actually playing, play longer, and have a higher percentage of Pro Bowl appearances.

If statistical discrimination or option value influence the relationship between original valuation of the athlete (measured by the position in the draft) and the resulting career, we expect this to be picked up in the main effects of DIVIA and YEARSTOP30 and the interactions of these two variables with DRAFTN. The interaction effect indicates whether the main effect of DRAFTN differs across groups and is thus the prime empirical evidence regarding statistical discrimination and option value. If either statistical discrimination against non–Division IA athletes or option value in favor of non–Division IA athletes is present, we expect the coefficient on the interaction effect to indicate that Division IA athletes have less successful careers. However, if option value is present, we also expect that the interaction effect will indicate that (if drafted) Division IA athletes are more likely to play than are non–Division IA athletes.

Unlike the tests of player status, a test based on years played in the league can be applied to several different samples: (1) all NFL players, (2) draftee NFL players (only NFL players who were drafted), (3) all draftees (thus including draftees who never played and had zero years played), and (4) everyone (free agents and all draftees [sample 3]). Samples 1 and 2 can also be used to analyze the percentage of Pro Bowl appearances.

[11] Graduation information is typically only available for players who actually play in the NFL. This variable is therefore excluded for samples that include unsuccessful draftees. We also estimated NFL player specifications without the graduation variable. There were no changes in significance or sign for the other variables.

[12] The World Football League ceased operations in 1984. The NFL held a special supplementary draft to disperse the players from the WFL to NFL teams.

IV. Estimation and Results

Our model requires us to measure productivity at various positions in the distribution of player talent. To capture this, our dependent variables consist of data on salary of the player in 1989, free agent status or not, whether the player made the team or not, years played, and percentage of years in the Pro Bowl. Whether the player made the team or not and percentage of years in the Pro Bowl are taken from tails of the productivity distribution. Recall the discussion at the end of Section II regarding figure 4. Given the differences in costs incurred from "mistakes" in hiring, we expect to observe differences in the tails of the probability distribution. "Making the team" allows us to distinguish statistical discrimination (SD) from option value (OV) predictions, since the two are different in this case. Probability of a Pro Bowl appearance allows us to focus on production of "stars," although predictions of both SD and OV are the same. Salary and years in the league are better measures of central tendencies of productivity effects, although results at the mean do allow distinction between OV and SD. If salary is a good measure of productivity, we should be able to capture these effects best with our salary measure. Unfortunately, as we detail below, salary data are the least likely to show results. We have therefore used as many measures as possible to give the best picture of total effects.

"Percentage of years in the Pro Bowl" and "years played" exhibit both left censoring and right censoring. The data are left censored because the lowest value of zero years has considerable heterogeneity in the quality of the players. The data are right censored because they include players whose career is incomplete. Because of the discreteness of the data, as well as the problems of left and right censoring, there is no "ideal" estimation scheme. Thus, we report several alternative methods of estimation on different subsets of the data.

The first results in table 6 (those reported in col. 1) are for the dichotomous dependent variable FREE, which indicates that the player was not drafted out of college but ultimately played in the NFL. Players who played outside Division IA or who played for weaker Division IA teams while in college tend to be overrepresented among free agents. Both the DIVIA coefficient and the YEARSTOP30 coefficient are negative and significant. White players are also overrepresented among free agents.

The second set of results in table 6 (those in col. 2) concern the dummy variable PLAY, which indicates whether a player chosen in the draft ultimately played for any NFL team. The specification (and all our subsequent estimations) includes the draft variable (high values of DRAFTN indicate high ex ante evaluation) and interactions between the draft variable and DIVIA and YEARSTOP30. If the impact of uncertainty does not vary with position in the draft, then the dummy variables for DIVIA

Table 6
Probability of Choice as Free Agent, Probability of Draftees Playing in the NFL, and Probability of Choice for Pro Bowl

	Dependent Variable			
	FREE	PLAY	PROBOWL	PROBOWL
Sample	NFL players	Draftees	Drafted NFL players	NFL players
Estimation method	Binomial probit	Binomial probit	Weighted probit	Weighted probit
N	3,838	4,515	3,020	3,838
Log-likelihood	−1,851	−2,349	−3,700	−4,189
Variable:				
QB	.03015	−.19158	.12832	.21006 *
	(.1269)	(.107)	(.07976)	.07289
KICKER	.93210 *	.00371	.32171 *	.30837 *
	(.1217)	(.1295)	(.09581)	.07866
BACKER	.06994	.15042	−.19293 *	−.23261 *
	(.08774)	(.07891)	(.05818)	.05437
LINE	.16976 *	−.16967 *	.18258 *	.20804 *
	(.08249)	(.07086)	(.05178)	.04871
TE	.15780	.12254	.00677	.01165
	(.1095)	(.09875)	(.0802)	.07491
RECEIVER	.18390 *	−.13332	.08864	.05375
	(.08303)	(.07224)	(.05498)	.05303
DB	.31644 *	−.03373	−.01943	−.01426
	(.07316)	(.06752)	(.05009)	.04708
WHITE	.13999 *	−.12635 *	−.33293 *	−.32661 *
	(.05778)	(.05104)	(.04188)	.03887
PRIVATE	.05413	.07752	.03622	.04779
	(.05804)	(.05123)	(.03291)	.03156
GRADUATE	−.08536		.10272 *	.07770 *
	(.05353)		(.03427)	.03210
BLACK	.05403	−.16001	−.02029	−.02079
	(.1028)	(.1016)	(.08772)	.07650
SUPP84		−2.36850 *	.36009 *	.37759 *
		(.1953)	(.09563)	.09563
DRAFTN		76.84700 *	7.21700	7.62080 *
		(7.642)	(1.501)	1.43200
YEARSTOP30	−.06506 *	.05777 *	.01157 *	.01670 *
	(.007749)	(.009907)	(.005326)	.00496
DIVIA	−.27426 *	−.34487 *	.05663	.03906
	(.06934)	(.0982)	(.0636)	.05491
DRAFT × IA		50.61200 *	−4.77220 *	−5.09370 *
		(11.07)	(1.514)	1.44600
DFT × TOP30		−7.66080 *	−.10389 *	−.11399 *
		(1.098)	(.0316)	.03132
Constant	−.61190	−.06204 *	−1.74560 *	−1.77720 *
	(.07265)	(.08447)	(.0623)	.05430

NOTE.—NFL players are players with at least 1 year in the NFL. They include both drafted NFL players and free agents. Draftees are all athletes chosen in the draft from 1979 to 1992. FREE is a dummy variable for free agent status; PLAY is a dummy variable for playing in at least one game in 1 year in the NFL. PROBOWL is the proportion of a player's career when he was chosen for the Pro Bowl game.
* Significant at the .05 level or better (two-tailed test).

and YEARSTOP30 should be significant but the interactions should not be significant.

The coefficient for DRAFTN is strongly positive, as expected. Those with good draft positions are much more likely to actually play in the NFL. Among draftees, white athletes are less likely to play. This might reflect the fact that white players have a much higher probability of graduation than black players. Thus, their employment alternatives are greater, and their effort to make the team could be less than the effort given by black players. However, white players are also overrepresented among free agents. This seems counter to the effort explanation.

If athletes from lesser-regarded football schools have more option value, the teams will take a chance on them. This chance could result in fewer of them actually surviving their first probationary period. In this case, their probability of making the team should be less. If teams are reluctant to take a chance on players with noisy predraft information due to risk aversion, then the players from lesser-regarded football schools should actually have a higher probability of surviving the probationary period. The results for school characteristics present a mixed relationship. The effects of DIVIA and YEARSTOP30 are exactly opposite. The main effect of DIVIA is negative and the interaction effect is positive, while the main effect of YEARSTOP30 is positive and the interaction effect is negative.

The combined effects of Division (DIVIA) and quality within division (YEARSTOP30) are evaluated in table 7. This table shows the predicted *Z*-scores (and their corresponding probabilities for draftees playing in the NFL) from the probit estimation for the probability of a draftee actually playing (see col. 2 of table 6). These are evaluated at three points in the draft distribution for three different school levels. We selected the midpoints of the first draft round (drafting position = 15), the third round (drafting position = 75), and the tenth round (drafting position = 285).[13]

The pattern of the *Z*-scores and their probabilities is consistent with both option value and statistical discrimination explanations. The stronger DIVIA school athletes have the lowest *Z*-scores at the top of the draft (compare 3.434 to 4.935 and 7.511) but the highest at the bottom (compare .223 to .081 and −.055). The smaller probabilities of playing at the top of the draft are consistent with statistical discrimination against athletes from weaker, less visible programs. Since all athletes at the top of the draft are likely to play (all predicted probabilities are essentially 1.0), the actual effect of this is very small even though the differences are significant. The larger probabilities of playing near the bottom of the draft are consistent with oversampling of non–Division IA athletes and Division IA athletes from weaker programs due to their option value. Thus, the PLAY results provide support for both option value and statistical discrimination for this comparison.

[13] The predictions assume that the control variables (all dummies) take the value of the omitted category, which are specified in the footnote to the table. Based on Wald tests for linear combinations of the coefficients in this probit, all differences in the *Z*-scores going down rows are significant. All differences across the columns are also significant, with the exception of the difference between very strong Division IA schools (DIVIA = 1 and YEARSTOP30 = 10) and non–Division IA schools in the 75th and 285th draft position and the difference between weaker Division IA schools (YEARSTOP30 = 1) and non–Division IA schools in the 285th draft position.

Table 7
Predicted Probability of Draftees Playing (Table 6, Col. 2) Evaluated for Different Draft Rounds and Different Player's School Characteristics

Round Drafted (Drafting Position)	Player's School's Characteristics: Non–Division IA (0 Years Top 30)	Division IA (1 Year Top 30)	Division IA (10 Years Top 30)
Round 1 (15):			
Z-score predicted	4.935	7.511	3.434
Predicted probability	1.000	1.000	1.000
Round 3 (75):			
Z-score predicted	.836	1.122	.722
Predicted probability	.798	.869	.765
Round 10 (285):			
Z-score predicted	.081	−.055	.223
Predicted probability	.532	.478	.588

Predicted *Z*-scores are computed for the two divisions, years in the top 30, and the different draft numbers with all control dummies set to the omitted categories (black running back from a public [non–traditionally black] institution and not chosen in the supplemental 1984 draft).

While it is natural to equate the outcome we call "making the team" with the probationary period, this is not exactly correct. NFL teams often keep players who are not expected to help their teams in the near future but who are expected to help later. These players might get the opportunity to play several times before the team releases them. Their probationary period thus includes not only their preseason tryout but also a small amount of playing time with the team. Therefore, it is necessary to evaluate other measures of career success that might better capture the quality of the player's career.

The last two specifications in table 6 (cols. 3 and 4) provide results for the quality of a player's career when "quality" is measured by the percentage of a player's career where he is elected to the "Pro Bowl." Each year, the players in the NFL vote to determine which of their colleagues will play in the Pro Bowl. Election to the Pro Bowl is a sign of dominance at the player's particular position on the field.

Since players' career length ranges from 1 to 14 years, this percentage is based on from 1 to 14 actual observations for each player on whether or not he was selected for the Pro Bowl. If the proportion is used without weighting, the underlying assumption is that we are making a single draw from a distribution with variance $p(1 - p)$, where in fact we are making multiple draws, which have unconditional variance of $p(1 - p)/N$. To account for the more precise estimate of the variance and to treat the heteroskedasticity that is introduced by having varying career length, we treat each proportion as a replication of N draws of size one. Therefore, we use probit as our estimation procedure. We label this "weighted" probit to indicate that career length is used as the weight to replicate observations.[14] We perform this weighted probit on two subsamples of the data. The first limits the data to just those NFL players who were drafted out of college: 3,020 observations. The second subset represents all players that appeared in the NFL: 3,838 observations on successful draftees as well as free agents. In both cases, DRAFTN is once again strongly positive. The effects of both Division IA membership and strength of program (YEARSTOP30) and their interactions with DRAFTN are in the same direction for both samples. All coefficients except the main effect of DIVIA are significant.

We again find a negative coefficient on WHITE in the PROBOWL estimations. Throughout this period, white players represented a minority of the sample (about 38%). Thus, this might be explained by a voting effect. However, it is again consistent with the results for PLAY estimations. Controlling for their playing position, their position in the draft, and the characteristics of their school, white players tend to have less distinguished careers.[15]

The results for the predicted Z-scores and probabilities of playing in the Pro Bowl are evaluated at different draft positions and school categories in table 8.[16] Athletes from non–Division IA schools are much more likely to be chosen for the Pro Bowl when they are drafted in the first round. This result supports both statistical discrimination and option value. However, option value is not supported at the bottom of the draft, since we should observe more distinguished careers among those who actually survive the probationary period. Again, we are constrained by not having a specific probationary period that would provide the best test of option value.

Our next estimate uses the length of a player's tenure in the NFL as the measure of success. Those players with high YEARS should represent very successful careers. They should have been identified as valuable selections in the draft.

The dependent variable is YEARS, a series of integers representing the length of career. This data series is truncated on the right for those players whose career was not complete at the time our data set was compiled, and it is truncated on the left (in principle) for players who are not good enough to even play "zero" years. In addition, the data are not really

[14] See Greene (1995, p. 413) for a discussion of the use of probit with grouped data.

[15] Kahn and Sherer (1988) find evidence consistent with discrimination against black players in the National Basketball Association that is similar to this result. In analyzing draft position, however, they do not find evidence that white players are drafted before black players, ceteris paribus.

[16] All the differences across rows and columns are significant, with the exception of the comparison of athletes from weaker Division IA schools and non–Division IA athletes in the 75th and 285th draft positions.

Table 8
Predicted Probability of NFL Players Playing in the Pro Bowl (Table 6, Col. 4) Evaluated for Different Draft Rounds and Different School Characteristics

	Player's School Characteristics		
Round Drafted (Drafting Position)	Non–Division IA (0 Years Top 30)	Division IA (1 Year Top 30)	Division IA (10 Years Top 30)
Round 1 (15):			
Z-score predicted	−1.596	−1.887	−1.805
Predicted probability	.055	.030	.036
Round 3 (75):			
Z-score predicted	−2.002	−2.016	−1.879
Predicted probability	.023	.022	.030
Round 10 (285):			
Z-score predicted	−2.077	−2.040	−1.893
Predicted probability	.019	.021	.029

NOTE.—Predicted *Z*-scores computed for the two divisions, years in the top 30, and the different draft numbers, with all control dummies set to the omitted categories (black running back from a public (non–traditionally black) institution and not chosen in the supplemental 1984 draft).

continuous. Thus both ordered probit and survival techniques have some drawbacks in the estimation of career length. Finally, we used four different subsets of the data. The first included those players, draftees and free agents, who had played at least one game in the NFL. The next estimation used the entire data set, which included all players drafted (regardless of tenure in the NFL) and free agents. Our next set used the subset of NFL players who were drafted, and it therefore excluded free agents. Finally, our last cut included all drafted players, regardless of their potential success.

In all four samples, DRAFTN had the expected sign: higher-drafted players have longer careers. The main effects of DRAFTN and YEARSTOP30 were all positive and significant in six of the eight estimations. The interaction terms DRAFT × IA and DFT × TOP30 were all negative and significant. In contrast to our other results, white athletes did not have significantly different career lengths than black athletes.

Figure 5 illustrates the estimated probabilities of specific career lengths for Division IA and non–Division IA players selected in different draft rounds.[17] This is not like a survival function. It is like a density function. For example, non–Division IA athletes have an expected probability of about 2% of leaving the league after 1 year if they are drafted in position 15. The same value for Division IA athletes from strong schools is about 4%. We present the probabilities of career lengths in years for two groups:

[17] These estimations are based on the results of estimates for all players and draftees. The differences in expected years in the league for the two groups were evaluated for significance for each round. All four tests yielded significant differences. The full coefficient estimates are available in Hendricks et al. (2001), which is available from the first author.

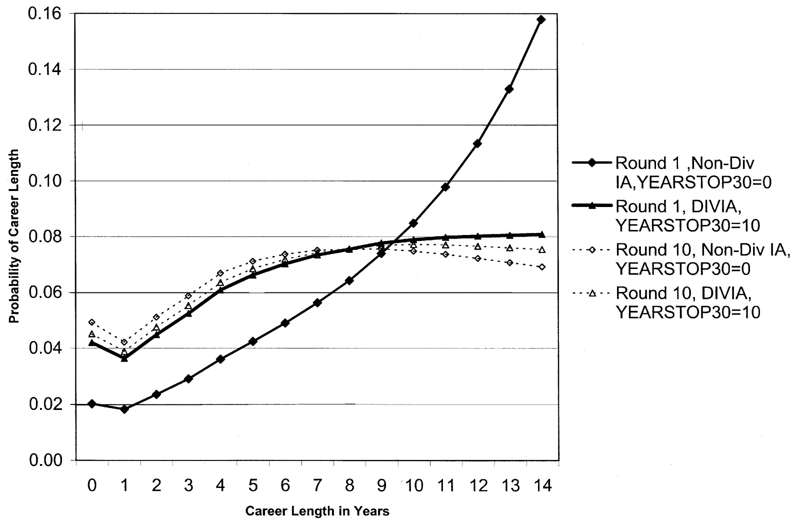

Fig. 5.—Predicted probabilities of career length

strong Division IA school athletes (DIVIA = 1 and YEARSTOP30 = 10) and non–Division IA school athletes (DIVIA = 0 and YEARSTOP30). For each group, we present the pattern for those chosen in round 1 (draft position = 15) and those chosen in round 10 (draft position = 285). There is a quite striking difference in career lengths for athletes from strong Division IA schools and athletes from other schools who are both chosen in round 1 of the draft. Non–Division IA athletes chosen in the first round show a much larger probability of having a long career length and a much lower probability of having a short career as compared with Division IA athletes chosen in the first round. The reverse pattern occurs at the end of the draft. Although the differences are not so striking, athletes from strong Division IA schools have higher probabilities of longer careers than non–Division IA athletes when both are drafted near the bottom (in our example, the tenth round). These results are again consistent with option value and statistical discrimination. Non–Division IA athletes are drafted more often at the bottom of the draft and less often at the top than one would expect based on their realized career lengths.

We had salary data for only one point in the player's careers, the year 1989. The explanatory variables used in our salary regression were drawn from Kahn (1992).[18] We simply added DIVIA and YEARSTOP30 to his

[18] See Kahn (1992) for a full description of these variables.

data and included interactions of these two variables with his draft variable.[19] Altonji and Pierret (2001) show that, under statistical discrimination, the impact of easily observed signals of ability (such as school football program) should fall over time, while the impact of hard-to-observe correlates of productivity should increase through time. Therefore, we should expect the effect of original draft position to decline as the player's career is extended. In Kahn's original study, he found that the interactions of draft position and experience and experience squared with draft position were negative. This is consistent with Altonji and Pierret's (2001) conjecture. The first-year salary of most players is almost entirely based on the team' s ex ante estimate of their value. As the player's tenure in the league lengthens, his true productivity becomes more evident, and the impact of his draft position becomes less important. This is reflected in the negative interaction of draft with experience.

Unfortunately, this makes the interpretation of salary as a measure of true ability problematical. Salary will be a good measure of true ability only after recontracting based on revealed productivity. However, many players sign multiple-year contracts. Since the average career span is only 3–4 years, it will be difficult to find the impact of statistical discrimination or option value using salary data. The relationship between the decrease in the effect of the draft and time in the league should be moderated by the school variables if either influence is associated with ex ante evaluation. This requires interacting the experience interaction terms with DIVIA and YEARSTOP30 to allow a difference in this moderating influence of time for schools with different characteristics. To be consistent with previous results, these new interaction terms should be negative. Players from more visible schools should see their salaries increase less quickly as compared with those of players from less visible schools if either statistical
q3 discrimination or option value is influencing ex ante choices.

The full results are available in our longer working paper (Hendricks et al. 2001). While an *F*-test of the joint effect of DIVIA and YEARSTOP30 and their interactions indicates that these variables are significant at the .003 level, many of the individual double and triple interactions were insignificant. This could be expected given the short panel for most players and the collinearity between the main effects and interactions. The predictions suggest that there is approximately a 5% beginning positive increment for players coming from highly rated Division IA schools. This effect is eliminated after about 2 years.

[19] In the Kahn study, the draft variable was measured as 1/(draft round) rather than 1/(draft number) as we used in our study. We have maintained this former definition in our salary analysis.

V. Conclusion

In this article, we have argued that the literature on statistical discrimination in the labor market and the literature on option value are linked by their connection with uncertainty about future productivity. The literature on statistical discrimination suggests that groups may be at a disadvantage when the reliability of the test instrument used to predict their performance is less than the reliability of this instrument when it is used to predict a competing group's performance. As Cornell and Welch (1996) show, this can result in selection primarily from the majority group when there are many applicants for a few slots. Thus, the effect of uncertainty is to provide fewer opportunities for minority individuals competing for highly competitive positions. They, like Aigner and Cain (1977), also show that this also implies that there should be a spot in the queue where the employer will begin to value minority applicants more highly than majority applicants. This latter prediction has raised some question because researchers have not cited situations where this observation appears to hold. On the other hand, the labor literature on option value (Lazear 1986, 1995) suggests that uncertainty can be of positive benefit to risky groups. As long as the employer can eliminate poor performers without other employers taking advantage of the sorting information, it will pay employers to take members of minority groups. Viewed in this manner, it seems quite possible that employers take chances on risky workers in the hope of finding "stars."

The National Football League provides a good test case for both ideas. First, we do not run into the nature-nurture problem in attempting to evaluate future performance. Second, employees are bound to the employer for a fixed period because league rules do not allow other teams to bid for the player until after a number of years have passed. This allows the employer to earn rents from keeping a risky worker.

Our empirical work used a variety of outcome measures in our tests. One measure, PLAY, indicated that the athlete was good enough to actually play in the league. Another (PROBOWL) indicated that the player was actually very talented relative to his peers. Finally, salary regressions offer more traditional measures of outcomes as rewards. We find broad support for both statistical discrimination and option value effects. Conditional on selection in one of the early rounds of the draft, athletes from less visible programs seem to have better careers. Their salaries are less likely to fall with experience, although they do pay an initial salary penalty for not being from a visible school. It is possible that this reflects option value, but it is more likely the result of statistical discrimination. When teams are choosing between two star athletes at the top of the draft, they seem to act in a risk adverse manner and select the athlete from the more visible football program. In the last rounds of the draft, the reverse appears to be true. Division IA athletes from top programs are undervalued. This supports the option value explanation, but we should see the results for "survivors" of the probationary period reversed. Athletes in this group should have stronger careers. We do not find evidence of this.

This conclusion is also consistent with our results for free agents. However, in contrast to the university tenure process, it is not always clear when the probationary period ceases for football players. In an article in *USA Today* (Bell 1997), one NFL owner was quoted in regard to the reduction in the length of time each team holds exclusive rights to draftees. The owner noted that hastening the time when a draftee is allowed to negotiate with other teams has forced teams to draft players who are likely to play immediately. This could potentially reduce the demand for athletes with more uncertain futures, that is, those with higher option value.

Even though the NFL has an enormous array of information sources at its disposal, we see evidence of systematic effects in the selection of players in their annual draft. Recalling the special nature of our data set only strengthens the significance of these findings. Our goal was to consider the existence of systematic "error" in labor market decisions. Certainly, one would expect such "errors" to be more likely in markets with less information. The greater the uncertainty, the more likely it is to observe such a pattern of choice. We have shown evidence for the existence of such effects in a market with excellent preemployment information. Very few markets have comparable information regarding preemployment performance. The print and electronic media devote tremendous attention to college athletics, in general, and to potential professional talent, in particular. In addition to their own predictions, each team has access to a variety of pundits' predictions. The fact that we are still able to document a systematic group effect is even more striking in such an environment of copious information.

If a test score (or any other characteristic that is used to differentiate candidates) is more unreliable for one group than another (even though it is unbiased), it is likely that few candidates from the less reliable group will be chosen when there is a large number of candidates for a few slots. One remedy for this statistical discrimination is to oversample from the group with the less reliable indicator. This might serve as one basis for affirmative action efforts. The impact of these efforts should be similar to the effect of option value. The employer will be able to choose more accurately from the group with less uncertainty. As a result, their ex ante predicted performance should be higher. However, if we allow the employer to eliminate the weakest performers, the ex post mean of the group with the unreliable indicator can be higher than the mean for the other group. We should expect more failures in the group with large uncertainty. However, these failures will be balanced off with successes that were not anticipated. The option value message is to not focus exclusively on the failures.

References

Aigner, Dennis, and Cain, Glen. "Statistical Theories of Discrimination in the Labor Market." *Industrial and Labor Relations Review* 30 (January 1977): 175–87.

Akerlof, George. "The Economics of Caste and of the Rat Race and Other Woeful Tales." *Quarterly Journal of Economics* 90 (November 1976): 599–617.

Altonji, Joseph G., and Pierret, Charles R. "Employer Learning and Statistical Discrimination." *Quarterly Journal of Economics* 116 (February 2001): 313–50.

Arrow, Kenneth. "The Theory of Discrimination." In *Discrimination in Labor Markets,* edited by Orley Ashenfelter and Albert Rees, pp. 3–33. Princeton, NJ: Princeton University Press, 1973.

Becker, Gary. *The Economics of Discrimination*. 2d ed. Chicago: University of Chicago Press, 1957.

Bell, Jarrett. "Early Returns: NFL Caught in Rookie Blitz as Quick Studies Hit Pay Dirt." *USA Today* (September 17, 1997), p. C1.

Bloch, Francis, and Rao, Vijayendra. "Statistical Discrimination, Identity Selection, and the Social Transformation of Caste and Race." Research Report no. 93-276. Ann Arbor: University of Michigan, Population Studies Center, 1993.

Borjas, George J., and Goldberg, Matthew S. "Biased Screening and Discrimination in the Labor Market." *American Economic Review* 68, no. 5 (1978): 918–22.

Coate, Stephan, and Loury, Glenn. "Will Affirmative-Action Policies Eliminate Negative Stereotypes?" *American Economic Review* 83 (December 1993): 1220–40.

Cornell, Bradford, and Welch, Ivo. "Culture, Information, and Screening Discrimination." *Journal of Political Economy* 104, no. 3 (1996): 542–71.

Greene, William. *Limdep, Version 7.0*. New York: Econometric Software, 1995.

Hendricks, Wallace; DeBrock, Lawrence; and Koenker, Roger. "Uncertainty, Hiring, and Subsequent Performance: The NFL Draft." Working paper. Urbana: University of Illinois, Institute of Labor and Industrial Relations, 2001.

Kahn, Lawrence M. "The Effects of Race on Professional Football Players' Compensation." *Industrial and Labor Relations Review* 45, no. 1 (1992): 295–310.

Kahn, Lawrence M., and Sherer, Peter D. "Racial Differences in Professional Basketball Players' Compensation." *Journal of Labor Economics* 6 (January 1988): 40–61.

Lazear, Edward P. "Raids and Offer-Matching." In *Research in Labor Economics,* vol. 8, pt. A, edited by Ronald Ehrenberg, pp. 141–65. Greenwich, CT: JAI, 1986.

———. *1993 NCAA Division I Graduation-Rates Report.* Overland Park, KS: National Collegiate Athletic Association, 1993.

———. "Hiring Risky Workers." Working Paper no. 5334. Cambridge, MA: National Bureau of Economic Research, 1995.

Neft, D.; Cohen, R.; and Korch, R. *The Sports Encyclopedia: Pro Football*, 10th ed. New York: St. Martin's, 1992.

Office for Advancement of Public Negro Colleges. *Directory of Traditionally Black Colleges and Universities in the United States.* Atlanta: Office for Advancement of Public Negro Colleges, 1971.

Peterson's Guides, *Peterson's Annual Guides to Undergraduate Study: Guide to Four-Year Colleges.* Princeton, NJ: Peterson's Guides, 1992.

Phelps, Edmund. "The Statistical Theory of Racism and Sexism." *American Economic Review* 62 (September 1972): 659–61.

Pratt, J. "Risk Aversion in the Small and in the Large." *Econometrica* 32 (January 1964): 122–36.

Sporting News. St. Louis, MO: Sporting News, various editions.

Sporting News Football Register. St. Louis, MO: Sporting News, various editions.

QUERIES TO THE AUTHOR

1 AU: This journal's style does not italicize foreign phrases that appear as a regular entry in Webster's 11th edition dictionary. Therefore, ex post and ex ante have been changed to roman font. There are also not hyphenated.

2 AU: In the last term of the equation above, is the x being multiplied by 1 - pi or only by pi?

3 AU: The original sentence was grammatically awkward. Is the rewrite OK?

Journal of Labor Economics

Graduate School of Business
1011 East 58th Street
Chicago, IL 60637

Reprint Order Form

Please return this form even if no reprints are desired.

☐ Send 2 free copies of the journal issue and one-year free subscription to the address below

☐ No Reprints Desired

TO BE COMPLETED BY AUTHOR:

Journal of Labor Economics Vol ______ No _____ Month _______________________

Author(s): ____ No of pages in article __________

Title of Article: ____

REPRINT PRICE LIST: Prices include shipping for U.S. and Canadian orders. Non-US and non-Canadian orders are shipped via Airmail at an additional cost of 45% of the total printing charge.

	Total Quantity				add'l
Page s	**50**	**100**	**150**	**200**	**50's**
2-4	$64.00	$76.00	$89.00	$100.00	$11.00
5-8	71.00	91.00	110.00	126.00	19.00
9-12	77.00	111.00	139.00	165.00	28.00
13-16	86.00	123.00	156.00	188.00	34.00
17-20	98.00	146.00	190.00	234.00	44.00
21-24	105.00	161.00	214.00	265.00	53.00
add'l 4 pgs	12.00	24.00	42.00	60.00	16.00
Covers	93.00	105.00	123.00	140.00	19.00

Charges (please compute)

_______ Quantity	$ ___________
Covers	$ ___________
Subtotal	$ ___________
GST (7% for Canadian destinations only)	$ ___________
Non-U.S./non-Canada orders add 45% to subtotal	$ ___________
TOTAL DUE (US $)	**$ ___________**

Shipping Instructions

Name ____

Phone* ________________________ Fax ____________________

Dept ___________________________________Room __________

Institution ____

Street____

City ________________________ State_____ Zip ______________

Country ____

Billing Instructions (Institutional Orders Only)

Institution ____

Street ____

City _______________________ State _____ Zip _____________

Country ____

Phone ____

Email ____

* Please include a phone number in case we need to contact you about your order.

MAKE CHECKS AND PURCHASE ORDERS PAYABLE TO: The University of Chicago Press

Orders need to be received prior to the printing of the journal issue. All orders must be accompanied by one of the three payment options: purchase order, check/money order, or Visa/MasterCard in U.S. currency via an American bank. Terms are net 30 days. Reprints are shipped 2-4 weeks after publication of the journal.

1) Institutional Purchase Order No. ________________________________ Purchase Order attached ☐ to come ☐

order will not be processed without a number

2) ☐ Check or Money Order for total charges is attached **OR** **3)** Charge to: ☐ VISA ☐ MASTERCARD

Cardmember name as it appears on card (please print clearly) ____

Card Number ____ Expiration Date ________________________________

Signature ____ Phone ____

revised 05/27/03