OVERCONFIDENCE VS. MARKET EFFICIENCY
IN THE NATIONAL FOOTBALL LEAGUE

Cade Massey
Richard H. Thaler

Overconfidence vs. Market Efficiency in the National Football League
Cade Massey and Richard H. Thaler
NBER Working Paper No. 11270
April 2005
JEL No. FD21, J3, G1

ABSTRACT

A question of increasing interest to researchers in a variety of fields is whether the incentives and experience present in many "real world" settings mitigate judgment and decision-making biases. To investigate this question, we analyze the decision making of National Football League teams during their annual player draft. This is a domain in which incentives are exceedingly high and the opportunities for learning rich. It is also a domain in which multiple psychological factors suggest teams may overvalue the "right to choose" in the draft – non-regressive predictions, overconfidence, the winner's curse and false consensus all suggest a bias in this direction. Using archival data on draft-day trades, player performance and compensation, we compare the market value of draft picks with the historical value of drafted players. We find that top draft picks are overvalued in a manner that is inconsistent with rational expectations and efficient markets and consistent with psychological research.

Cade Massey
Duke University
cade.massey@duke.edu

Richard H. Thaler
Graduate School of Business
University of Chicago
1101 East 58th Street
Chicago, IL 60637
and NBER
richard.thaler@gsb.uchicago.edu

Two of the building blocks of modern neo-classical economics are rational expectations and market efficiency. Agents are assumed to make unbiased predictions about the future and markets are assumed to aggregate individual expectations into unbiased estimates of fundamental value. Tests of either of these concepts are often hindered by the lack of data. Although there are countless laboratory demonstrations of biased judgment and decision making (for recent compendiums see Kahneman & Tversky , 2000 and Gilovich, Griffin, & Kahneman, 2002;) there are fewer studies of predictions by market participants with substantial amounts of money at stake. Similarly, tests of market efficiency are plagued by the inability to measure fundamental value. (Even now, in 2005, there is debate among financial economists as to whether prices on Nasdaq were too high in 2000.)

In this paper we offer some evidence on both of these important concepts in an unusual but interesting context: the National Football League, specifically its annual draft of young players. Every year the National Football League (NFL) holds a draft in which teams take turns selecting players. A team that uses an early draft pick to select a player is implicitly forecasting that this player will do well. Of special interest to an economic analysis is that teams often trade picks. For example, a team might give up the 4th pick and get the 10th pick and the 21st pick in return. In aggregate, such trades reveal the *market value* of draft picks. We can compare these market values to the *surplus value* (to the team) of the players chosen with the draft picks. We define surplus value as the player's *performance value* – estimated from the labor market for NFL veterans – less his compensation. In the example just mentioned, if the market for draft picks is rational then the surplus value of the player taken with the 4th pick should equal (on average) the combined surplus value of the players taken with picks 10 and 21.

The rate at which the value of picks declines over the course of the draft should, in a rational world, depend on two factors: the teams' ability to predict the success rate of prospective players, and the compensation that has to be paid to drafted players. If, to take an extreme

2

example, teams have no ability to forecast future value, then early picks are worth no more than later picks: they are all equally valued lottery tickets. A steeply declining price, on the other hand, implies that performance is highly predictable. Compensation matters because if early picks are (for whatever reason) paid more, then the surplus to the teams who select them is reduced. Indeed, if (hypothetically) early picks had to be paid more than later picks, and performance predictability were zero, then early picks would be *less* valuable than late picks.

To illustrate the basic idea of the paper, consider one high-profile example from the 2004 draft. The San Diego Chargers had the rights to the first pick and used it to select a promising quarterback, Eli Manning. Much was expected of Eli: he had a very successful collegiate career, and his father (Archie) and older brother (Peyton) were NFL stars. Peyton had also been the first player selected in the 1998 draft and had become one of the best players in the league. The New York Giants, picking 4[th], were also anxious to draft a quarterback, and it was no secret that they thought Manning was the best prospect in the draft. It was reported (King, 2004) that during the 15 minutes in which they had to make their selection the Giants had two very different options under consideration. They could make a trade with the Chargers in which case the Giants would select Philip Rivers, considered the second-best quarterback in the draft, and then swap players. The price for this "upgrade" was huge: the Giants would have to give up their third-round pick in 2004 (the 65[th] pick) and their first- and fifth-round picks in 2005. Alternatively, the Giants could accept an offer from the Cleveland Browns to move down to the 7[th] pick, where it was expected that they could select the consensus third best QB in the draft, Ben Roethlisberger. The Browns were offering their second-round pick (the 37[th]) in compensation. In summary, Manning was four picks more expensive than Roethlisberger.

As we will show below, the offer the Giants made to the Chargers was in line with previous trades involving the first pick. They were paying the market price to move up, and that price is a steep one. In addition, they knew that Manning would cost them financially. Historically the first pick makes about 60 percent more during his initial (4- to 5-year) contract

3

than the seventh pick. Thus both in terms of pick value and monetary cost, the market prices imply that performance must be highly predictable. But history is full of anecdotes that seem to tell the opposite story. Success in the NFL, especially for quarterbacks, has been notoriously difficult to predict. Indeed, the year in which Eli's brother Peyton was taken with the first pick there was much speculation about whether another quarterback, Ryan Leaf, was a better prospect than Manning. Leaf was taken with the very next pick (by the Chargers) and was as spectacularly unsuccessful as Manning was successful; after trials with several teams he eventually left the league a declared flop.

In this paper we systematically investigate these issues. Our initial conjecture in thinking about these questions was that teams did not have rational expectations regarding their ability to predict player performance. A combination of well-documented behavioral phenomena, all working in the same direction, creates a systematic bias: teams overestimate their ability to discriminate between stars and flops. We reasoned that this would not be eliminated by market forces because, even if there are a few smart teams, they cannot correct the mis-pricing of draft picks through arbitrage. There is no way to sell the early picks short, and successful franchises typically do not "earn" the rights to the very highest picks, so cannot offer to trade them away.

Our findings suggest the biases we had anticipated are actually even stronger than we had guessed. We expected to find that early picks were overpriced, and that the surplus values of picks would decline less steeply than the market values. Instead we have found that the surplus value of the picks during the first round actually increases throughout the round: the players selected with the final pick in the first round on average produces more surplus to his team than than the first pick, and costs one quarter the price!

The plan of the papers is as follows. In section I we review some findings from the psychology of decision making that lead us to predict that teams will put too high a value on picking early. In section II we estimate the market value of draft picks. We show that the very high price the Giants paid in moving from the 4[th] pick to the 1[st] one was not an outlier. Using a

4

data set of 276 draft day trades, we find that the implicit value of picking early is very high. The first pick is valued as much as the 10th and 11th picks combined, and as much as the sum of the last four picks in the first round. We also find that teams discount the future at an extraordinary rate. In section III we examine the relation between draft order and compensation during the initial contract period. We find that compensation also declines very steeply. So, very high draft picks are expensive in two ways: high picks can be traded for multiple lower picks, and high picks are paid higher salaries. In the following sections we ask whether these expensive picks are too expensive. In section IV we analyze the performance of drafted players. We find that performance is correlated with draft order, as expected. Players taken in the first round are more likely to be successful (be on the roster, start games, make the all-star game) than players taken in later rounds. However, performance does not fall as steeply as the implicit price of draft picks. Still, these analyses do not answer the economic question of whether the early picks are mis-priced. We address this question in Section V. We do this by first estimating player performance value using market prices of free agents as an indication of true value. We use compensation data for the sixth year of a player's career since by that stage of their careers players have had the opportunity to test the free-agent market. We regress total compensation on categorical performance data, including position fixed-effects. We use the results of this model to value individual player's year-by-year performances over the first 5 years of their career. Combining this performance value with the player's compensation costs allows us to compute the surplus value to the team drafting each player. We find that surplus value increases throughout the first round, i.e., late-first-round picks generate more value than early-first-round picks. We conclude in section VI.

I. RESEARCH HYPOTHESIS

The NFL draft involves two tasks that have received considerable attention from psychological researchers – predicting the future and bidding competitively. This research

suggests that behavior in these tasks can deviate systematically from rational models. In this section we draw on these findings to develop our research hypothesis. In doing so we have an embarrassment of riches – research on non-regressive predictions, overconfidence, the winner's curse, and false consensus suggests our hypothesis is over-determined. While this means we will not be able to pin the blame on any one underlying cause, it strengthens the case for our overarching hypothesis: teams overvalue the right to choose.

In their early work on representativeness, Kahneman & Tversky (1973) compare "intuitive" predictions to those by normative models. One of their chief findings is that intuitive predictions are insufficiently regressive. That is, intuitive predictions are more extreme and more varied than is justified by the evidence on which they are based. They show this in a series of studies in which individuals predict future states (e.g., a student's grade-point average) from current evidence of various forms (e.g., the results of a "mental-concentration" quiz). Normative models require combining this evidence with the prior probabilities of the future states (e.g., the historical distribution of grade-point averages), with the weight placed on the evidence determined by how diagnostic it is. Hence, one can safely ignore prior probabilities when in possession of a very diagnostic evidence, but should lean on them heavily when the evidence is only noisily related to outcomes. In their studies, Kahneman & Tversky show that such weighting considerations are almost entirely ignored, even when individuals are aware of prior probabilities and evidence diagnosticity. Instead, individuals extrapolate almost directly from evidence to prediction. This results in predictions that are too extreme for all but the most diagnostic evidence.

NFL teams face exactly this kind of task when they predict the future performance of college players – they must combine evidence about the player's ability (his college statistics, scouting reports, fitness tests, etc.) with the prior probabilities of various levels of NFL performance to reach a forecast. For example, over their first five years, players drafted in the first round spend about as many seasons out of the league (8%) or not starting a single game (8%)

6

as in the Pro Bowl (9%). To the extent that the evidence about an individual player is highly

diagnostic of a player's NFL future, prior probabilities such as these can be given less weight.

However, if the evidence is imperfectly related to future performance, then teams should

"regress" player forecasts toward the prior probabilities. If teams act as Kahneman & Tversky's

subjects did, they will rely too heavily on evidence they accumulate on college players. Indeed,

to be regressive is to admit to a limited ability to differentiate the good from the great, and it is

this skill that has secured NFL scouts and general managers their jobs.

      Overconfidence is a closely related concept in the psychological literature. Simply put,

people believe their knowledge is more precise than it is in fact. A simple way this has been

demonstrated is by asking subjects to produce confidence limits for various quantities, for

example, the population of a city. These confidence limits are typically too narrow. For example,

in a review of 28 studies soliciting interquartile ranges (25^{th}-75^{th} percentile) for uncertain

quantities, the median number of observations falling within the ranges was only 37.5%

(normatively it should be 50%) (Lichtenstein, Fischhoff, & Phillips, 1982). Performance was

even worse when asked for broader intervals, e.g., 95% confidence intervals. This is related to

non-regressive forecasts in that subjects are not giving sufficient weight to either the limits of

their cognitive abilities nor to the inherent uncertainty in the world.

      An interesting and important question is how confidence depends on the amount of

information available. When people have more information on which to base their judgments

their confidence can rationally be greater, but often information increases confidence more than it

increases the actual ability to forecast the future. For example, in one classic study, Oskamp

(1965) asked subjects to predict the behavior of a psychology patient based on information

excerpted from the patient's clinical files. Subjects received information about the patient in four

chronological stages, corresponding to phases in the patient's life, making judgments after each

phase. Subjects also reported confidence in the judgments they provided. While the accuracy of

their judgments was relatively constant across the four stages, confidence increased dramatically.

As a consequence, participants progressed from being reasonably well calibrated in the beginning to being quite overconfident after receiving more information. Slovic & Corrigan (1973) find a similar pattern in a study of horse-racing bettors (in Russo & Schoemaker, 2002).

NFL teams face a challenge related to Oskamp's experiment – making judgments about players while accumulating information about them. Teams track some players' performance from their freshman year in college, with the intensity increasing dramatically in the year preceding the draft. In the final months before the draft almost all players are put through additional drills designed to test their speed, strength, agility, intelligence, etc. While one might think such information can only improve a team's judgment about a player, the research just described suggests otherwise. Rather, as teams compile information about players, their confidence in their ability to discriminate between them might outstrip any true improvement in their judgment.

Competitive bidding introduces another set of issues. It is well known that in situations in which many bidders compete for an item with a common but uncertain value then the winner of the auction often overpays (for a review see Thaler, 1988). The winner's curse can occur even if bidders have unbiased but noisy estimates of the object's true worth, because the winning bidder is very likely to be someone who has overestimated the actual value of the object. Rational bidders should recognize this adverse-selection problem and reduce their bids, especially as the number of other bidders increases. Instead, increasing the number of bidders results in more aggressive bidding (Kagel & Levin, 1986). The winners curse was first documented in research on oil-lease bids (Capen, Clapp, & Campbell, 1971), and has since been observed in numerous field (cf. Dessauer, 1981; Roll, 1986) and experimental settings (cf. Samuelson & Bazerman, 1985).

Harrison & March (1984) suggest that a related phenomenon occurs when a single party selects from multiple alternatives. If there is uncertainty about the true value of the alternatives, the decision-maker, on average, will be disappointed with the one she chooses. This problem,

8

which Harrison & March term "expectation inflation", has paradoxical consequences. In the context of choosing football players, the implication is that the more players a team examines, the better will be the player they pick, but the more likely the team will be disappointed in the player.

Harrison & Bazerman (1995) point out that non-regressive predictions, the winner's curse, and expectation inflation have a common underlying cause – the role of uncertainty and individuals' failure to account for it. The authors emphasize that these problems are exacerbated when uncertainty increases and when the number of alternatives increase. The NFL draft is a textbook example of such a situation – teams select among hundreds of alternative players, there are typically many teams interested in any given player, and there is significant uncertainty about the future value of the player. Other than trying to reduce the uncertainty in their predictive models (which is both expensive and of limited potential), teams have little control over these factors. If teams recognize the situation, they will hedge their bids for particular players, reducing the value they place on choosing one player over another. But if they are susceptible to these biases, they will "bid" highly for players, overvaluing the right to choose early.

A final consideration is the false consensus effect (Ross, Greene, & House, 1977). This effect refers to a person's tendency to believe that others are more similar to them in beliefs, preferences and behavior than they actually are.   For example, Ross et al asked their student participants to estimate the percentage of students who believed a woman would be named to the Supreme Court within a decade. Students who themselves believed this was likely, gave an average estimate of 63%, while those who did not believe it was likely gave an average estimate of 35%. This effect does not suggest that everybody believes they are in the majority on all issues, but rather that they believe others are more like them than they actually are. In the NFL draft, the presence of a false concensus effect would mean that teams overestimate the extent to which other teams value players in the same way that they do. This has significant consequences for draft-day trades. As we discuss below, most trades are of a relatively small "distance" for the purpose of drafting a particular player. An alternative to making such a trade is to simply wait

and hope that other teams do not draft the player with the intervening picks. False consensus suggests that teams will overestimate the extent to which other teams covet the same player, and therefore overestimate the importance of trading-up to acquire a particular player. Such a bias will increase the value placed on the right to choose.

Together these biases all push teams toward overvaluing picking early. Teams overestimate their ability to discriminate between the best linebacker in the draft and the next best one, and to overestimate the chance that if they wait, the player they are hoping for will be chosen by another team. Of course, there are strong incentives for teams to overcome these biases, and the draft has been going on for long enough (since 1936) that teams have had ample time to learn. Indeed, sports provides one of the few occupations (academia is perhaps another) where employers can easily monitor the performance of the candidates that they do not hire as well as those they hire. (Every team observed Ryan Leaf's failures, not only the team that picked him.) This should facilitate learning. It should also be possible to overcome the false consensus effect simply by comparing a team's initial ranking of players with the order in which players are selected.

The null hypothesis of rational expectations and market efficiency implies that ratio of market values of picks will be equal (on average) to the ratio of surplus values produced. . Specifically, for the i-th and i-th+k picks in the draft,

$$\frac{MV_i}{MV_{i+k}} = \frac{E(SV_i)}{E(SV_{i+k})}, \tag{1}$$

where MV is the market value of the draft pick and $E(SV)$ is the expected surplus value of players drafted with the pick.

We hypothesize that, in spite of the corrective mechanisms discussed above, teams will overvalue the right to choose early in the draft. For the reasons detailed above, we believe teams will systematically pay too much for the rights to draft one player over another. This will be

reflected in the relative price for draft picks as observed in draft-day trades. Specifically, we predict

$$\frac{MV_i}{MV_{i+k}} > \frac{E(SV_i)}{E(SV_{i+k})}, \tag{2}$$

i.e., that the market value of draft picks will decline more steeply than the surplus value of players drafted with those picks.[1]   Furthermore, we expect this bias to be most acute at the top of the draft, as three of the four mechanisms we've highlighted will be exaggerated there. Regression-to-the-mean effects are strongest for more extreme samples, so we expect the failure to regress predictions to be strongest there as well.[2] Players at the top of the draft receive a disproportionate amount of the attention and analysis, so information-facilitated overconfidence should be most extreme there. And the winner's curse increases with the number of bidders. Although not always the case, we expect more bidders for the typical player near the top of the draft than for the players that come afterwards. False consensus is an exception, as there is typically more true consensus at the top of the draft. We suspect false consensus plays a stronger role with the considerable amount of trading activity that takes place in the middle and lower rounds of the draft. On balance, though, we expect overvaluation to be most extreme at the top of the draft. That is, at the top of the draft we expect the relationship between the market value of draft picks and draft order to be steeper than the relationship between the value of players drafted and draft order.

More generally, we are investigating whether well established judgment and decision making biases are robust to market forces. The NFL seems to provide almost ideal conditions for overcoming these psychological biases. As Michael Lewis, author of *Moneyball*, said of another

---

[1] Note that this expression, by itself, does not imply which side of the equation is "wrong". While our hypothesis is that the left-hand side is the problem, an alternative explanation is that the error is on the right-hand side. This is the claim Bronars (Bronars, 2004) makes, in which he assumes the draft-pick market is rational and points out its discrepancy with subsequent player compensation. A key difference in our approaches is that we also use player performance to explore which of the two sides, or markets, is wrong.

[2] Similarly, De Bondt & Thaler (1985) found the strongest mean reversion in stock prices for the most extreme performers over the past three to five years.

sport, "If professional baseball players, whose achievements are endlessly watched, discussed and analyzed by tens of millions of people, can be radically mis-valued, who can't be? If such a putatively meritocratic culture as professional baseball can be so sloppy and inefficient, what can't be?" (Lewis, in Neyer, 2003). On one hand we agree wholeheartedly with Lewis – one would be hard-pressed to generate better conditions for objective performance evaluation. Consequently, we consider our hypothesis a rather conservative test of the role of psychological biases in organizational decision making. But on the other hand, ideal conditions may not be sufficient for rational decision making. Considering that any one of the psychological factors discussed in this section could strongly bias NFL teams, we would not be surprised to find the NFL draft as "sloppy and inefficient" as Lewis found major league baseball.

II. THE MARKET FOR NFL DRAFT PICKS

In this section we estimate the market value of NFL draft picks as a function of draft order. We value the draft picks in terms of other draft picks. We would like to know, for example, how much the first draft pick is worth relative to say, the fifth, the fifteenth, or the fiftieth. We infer these values from draft-day trades observed over 17 years.

A. Data

The data we use are trades involving NFL draft picks for the years 1988 through 2004.[3] Over this period we observe 334 draft-day trades. Of these, we exclude 51 (15%) that involve NFL players in addition to draft picks, and 6 (2%) with inconsistencies implying a reporting error. We separate the remaining trades into two groups: 213 (64%) involving draft picks from only one year and 63 (19%) involving draft picks from more than one year. We begin by focusing on trades involving a single draft year and subsequently incorporate the multi-year trades in a more general model.

---

[3] This dataset was compiled from newspaper reports. We are missing the second day (of two) of the 1990 draft.

The NFL draft consists of multiple rounds, with each team owning the right to one pick per round. (The order that teams choose depends on the team's won-lost record in the previous season—the worst team chooses first, and the winner of the Super Bowl chooses last.) During the period we observe, the NFL expanded from 28 to 32 teams and reduced the number of rounds from 12 to 7. This means the number of draft picks per year ranges from 222 (1994) to 336 (1990). We designate each pick by its overall order in the draft. In the 213 same-year trades, we observe trades involving picks ranging from 1st to 333rd. Figure 1 depicts the location and distance for all trades in which current-year draft-picks were exchanged (n=238). While we observe trades in every round of the draft, the majority of the trades (n=126, 53%) involve a pick in one of the first two rounds. The average distance moved in these trades (the distance between the top two picks exchanged) was 13.3 (median=10). The average distance moved is shorter for trades involving high draft picks, e.g., 7.8 (median=7) in the first two rounds.

Insert Figure 1 about here

Trades often involve multiple picks (indeed, the team trading down requires something beyond a one-for-one exchange of picks). The average number of picks acquired by the team trading down was 2.2 (sd=.83), with a maximum of 8. The average number of picks acquired by the team trading up is 1.2 (sd=.50), with a maximum of 5. The modal trade was 2-for-1, occurring 159 times (58%).

B. Methodology

We are interested in estimating the value of a draft pick as a function of its order and in terms of other draft picks. We will take the first pick in the draft as the standard against which all other picks are measured. We assume the value of a draft pick drops monotonically with the

13

pick's relative position and that it can be well described using a Weibull distribution.[4] Our task is then estimating the parameters of this distribution.

Let t_i^r denote the *t*-th pick in the draft, either for the team with the relatively higher draft position (if *r*=H) and therefore "trading down", or the team with the relatively lower draft position (if *r*=L) and therefore "trading up". The index *i* indicates the rank among multiple picks involved in a trade, with *i*=1 for the top pick involved.

For each trade, we observe the exchange of a set of draft picks that we assume are equal in value. Thus, for each trade we have

$$\sum_{i=1}^{m} v(t_i^H) = \sum_{j=1}^{n} v(t_j^L) \quad ,$$
(3)

where *m* picks are exchanged by the team trading down for *n* picks from the team trading up. Assuming the value of the picks follow a Weibull distribution, and taking the overall first pick as the numeraire, let the relative value of a pick be

$$v(t_i^r) = e^{-\lambda(t_i^r - 1)^{\beta}} \quad ,$$
(4)

where λ and β are parameters to be estimated. Note that the presence of the β parameter allows the draft value to decay at either an increasing or decreasing rate, depending on whether its value is greater than or less than one. If $\beta = 1$ we have a standard exponential with a constant rate of decay. Also, note that for the first pick in the draft, $v(1) = e^{-\lambda(1-1)^{\beta}} = 1.0$.

Substituting (4) into (3) and solving in terms of the highest pick in the trade, we have

---

[4] A Weibull distribution is a 2-parameter exponential. The single parameter in an exponential indicates the constant rate at which the distribution "decays". The additional parameter in the Weibull allows this decay rate to either increase or decrease. Consequently, the Weibull provides a very flexible distribution with which to estimate the decay of draft-pick value.

$$t_1^H = \left(-\frac{1}{\lambda} \log \left(\sum_{j=1}^{n} e^{-\lambda(t_j^L - 1)^\beta} - \sum_{i=2}^{m} e^{-\lambda(t_i^H - 1)^\beta} \right) \right)^{\frac{1}{\beta}} + 1$$

(5)

which expresses the value of the top pick acquired by the team trading up in terms of the other

picks involved in the trade. Recall that this value is relative to the first pick in the draft. We can

now estimate the value of the parameters λ and β in expression (5) using nonlinear regression.[5]

We would also like to value future draft picks. Modifying (4) to include a discount

rate, ρ, gives us

$$v(t_i^r) = \frac{e^{-\lambda(t_i^r - 1)^\beta}}{(1+\rho)^n}$$

(6)

for a draft pick *n* years in the future. This expression reduces to (4) when *n*=0, i.e., the pick is for

the current year. Substituting (6) into (3) and solving in terms of the highest pick in the trade,

which by definition is in the current year, gives us

$$t_1^H = \left(-\frac{1}{\lambda} \log \left(\sum_{j=1}^{n} \left(\frac{e^{-\lambda(t_j^L - 1)^\beta}}{(1+\rho)^n} \right) - \sum_{i=2}^{m} \left(\frac{e^{-\lambda(t_i^H - 1)^\beta}}{(1+\rho)^n} \right) \right) \right)^{\frac{1}{\beta}} + 1$$

(7)

which simplifies to (5) if all picks are in the current draft.

C. Results

Using the 213 same-year trades described above, we find $\lambda =.148$ (se=.03) and $\beta =.700$

(se=.033). The model fits the data exceedingly well, with $R^2 =0.999$. These results are

summarized in Table 1, column 1. A Weibull distribution with these parameters is graphed in

Figure 2. This graph shows the value of the first 100 draft picks (approximately the first 3

rounds) relative to the first draft pick. This curve indicates that the 10[th] pick is worth 50% of the

1[st] overall pick, the 20[th] is worth 31%, the 30[th] is worth 21%, etc.

---

[5] We first take the log of both sides of expression (5) before estimation in order to adjust for lognormal
errors.

Figure 3 provides another means of evaluating the model's fit. This graph compares the estimated values for "both sides" of a trade – the value of the top pick acquired by the team moving up ($e^{-\hat{\lambda}(t_1^H-1)^{\hat{\beta}}}$), and the value paid for that pick by the team moving up net of the value of additional picks acquired ($\sum_{j=1}^{n} e^{-\hat{\lambda}(t_j^L-1)^{\hat{\beta}}} - \sum_{i=2}^{m} e^{-\hat{\lambda}(t_i^H-1)^{\hat{\beta}}}$), where $\hat{\lambda}$ and $\hat{\beta}$ are estimated parameters. The model fits the data quite well. We can also identify on this graph those trades that appear to be "good deals" for the team trading up (those below the line) and those that appear to be "bad deals" for the team trading up (those above the line), relative to the market price.

We also estimated (7), which includes a third parameter for the discount rate. This expression allows us to include trades involving future picks, expanding our sample to 276 observations. Results are presented in Table 1, column 2. The estimated curve is close to the previous one, with λ =.121 (se=.023) and β =.730 (se=.030), though a bit flatter – e.g., the 10$^{\text{th}}$ pick is valued at 55% of the first. The estimated discount rate, ρ, is a staggering 173.8% (se=.141) per year.

Finally, we investigate how these draft-pick values have changed over time. To do this we estimate separate models for the first half (1988-1996, n=131) and second half (1997-2004, n=145) of our sample. Results are presented in Table 1, columns 3 and 4. In the first period we find λ =.092 (se=.029), β =.775 (se=.052), and ρ =159% (se=.237). In the second period

16

λ =.249 (se=.064), β =.608 (se=.042), and ρ =173.5% (se=.126). Differences in the estimates

for the Weibull parameters, λ and β, are statistically significant, while the difference in the

discount rate estimates are not.   We graph the curves for both periods in Figure 4.

Insert Figure 4 about here

The change in the market value of the draft picks is most easily seen by comparing across

time the implied value of various picks. In the first period, for example, the 10[th] pick is worth

60% of the 1[st], while in the second period it is worth only 39% of the 1[st]. The value of the 20[th]

pick dropped from 41% to 22%, and the value of the 30[th] from 29% to 15%. Overall, the draft-

pick value curve is steeper in the second period than in the first for all picks in the first round.

D. Discussion

One of the most striking features of these data is how well ordered they are – it seems

clear there is a well understood market price for draft picks. Indeed, the use of this kind of "value

curve" has caught on throughout the NFL in recent years. A few years ago Jimmy Johnson, a

former coach turned television commentator, discussed such a curve during television coverage

of the draft, and in 2003 ESPN.com posted a curve it said was representative of curves that teams

use.[6] The ESPN curve very closely approximates the one we estimate for the 1997-2004 period.

The close fit we obtain for our model suggests there is wide agreement among teams (or at least

those who make trades) regarding the relative value of picks. This historical consensus may lend

the considerable power of inertia and precedent to the over-valuation we suggest has

psychological roots.

A second striking feature is how steep the curve is. The drop in value from the 1[st] pick to

the 10[th] is roughly 50%, and another 50% drop from there to the end of the first round. As, we

---

[6] An NFL team confirmed that "everybody has one" of these curves. The one they shared with us was very close to the one we estimate for the second half of our sample, with the notable exception that it was not continuous.

report in the following section, compensation costs follow a very similar pattern.   Moreover, the prices are getting steeper with time. In the first period, the value of the 10th pick is 60% of the 1st pick, whereas in the later period, the 10th pick is only worth 39% of the 1st pick. Since we will argue that even the earlier curve was too steep, the shift has been in the "wrong" direction, that is, it has moved further away from rational pricing.

A third notable feature of these data is the remarkably high discount rate, which we estimate to be 174% per year. A closer look at trading patterns suggests that this rate, though extreme, accurately reflects market behavior. Specifically, teams seem to have adopted a rule of thumb indicating that a pick in this year's round n is equivalent to a pick in next year's round n-1. For example, a team trading this year's 3rd-round pick for a pick in next year's draft would expect to receive a 2nd-round pick in that draft. This pattern is clear in the data. Eighteen of the 26 trades involving 1-for-1 trades for future draft picks follow this pattern. Importantly, the 8 trades that do not follow it all involve picks in the 4th round or later, where more than one pick is needed to compensate for the delay, since one-round differences are smaller later in the draft (i.e., the curve is flatter).

This trading pattern means that the discount rate must equate the value of picks in two adjacent rounds. The curve we estimated above suggests that the middle pick of the 1st round is worth .39 (relative to the top pick), the middle pick of the 2nd round is worth .12, and the middle pick of the 3rd is worth .05. The discount rate required to equate this median 1st-round pick with the median 2nd-round pick is 225%, while the rate required to equate the median 2nd-round pick with the median 3rd–round pick is 140%. Given these rates, and where trades occur (the majority involve a pick in one of the top two rounds), our estimate seems to very reasonably capture market behavior.

This huge premium teams pay to choose a player this year rather than next year is certainly interesting, but is not the primary focus of this paper so we will not analyze it in any detail. We suspect that one reason why the discount rate is high is that picks for the following

year have additional uncertainty attached to them since the exact value depends on the performance of the team trading away the pick in the following year.[7] Still, this factor alone cannot explain a discount rate of this magnitude. Clearly teams giving up second-round picks next year for a third-round pick this year are displaying highly impatient behavior, but it is not possible to say whether this behavior reflects the preferences of the owners or the employees (general manger and coach) who make the choices (or both). Regardless, it is a significant arbitrage opportunity for those teams with a longer-term perspective.

III. INITIAL COMPENSATION COSTS

We have shown that the cost of moving up in the draft is very high in terms of the opportunity cost of picks. Early picks are expensive in terms of compensation as well. NFL teams care about salary costs for two reasons. First, and most obviously, salaries are outlays, and even behavioral economists believe that owners prefer more money to less. The second reason teams care about compensation costs is that NFL teams operate under rules that restrict how much they are allowed to pay their players—a salary cap. We will spare the reader a full description of the salary cap rules – for an excellent summary see Hall & Limm (2002). For our purposes, a few highlights will suffice. Compensation is divided into two components: salary and bonus. A player's salary must be counted against the team's cap in the year in which it is paid. Bonuses, however, even if they are paid up front, can be amortized over the life of the contract as long as the player remains active. If the player leaves the league or is traded, the remaining bonus is charged to the team's salary cap.

In addition to this overall salary cap, there is a rookie salary cap, a "cap within a cap". Teams are allocated an amount of money they are permitted to spend to hire the players they selected in the draft. These allocations depend on the particular picks the team owns, e.g., the

---

[7] We simply use the middle pick of the round when calculating the discount rate. Empirically, there is no reliable difference between the average pick position of teams acquiring future picks (15.48) and those disposing future picks (16.35).

19

team that owns the first pick is given more money to spend on rookies than the team with the last pick, all else equal. When teams and players negotiate their initial contracts, the rookie salary cap plays an important role. The team that has drafted the player has exclusive rights to that player within the NFL. A player who is unhappy with his offer can threaten to hold out for a year and reenter the draft, or go play in another football league, but such threats are very rarely carried out. Teams and players typically come to terms, and the rookie salary cap seems to provide a focal point for these negotiations. As we will see briefly, initial compensation is highly correlated with draft order.

The salary data we use here and later come from published reports in USA Today and its website. For the period 1996-2002 the data include base salary and bonuses paid, i.e., actual team outlays. For 2000-2002 we also have the "cap charge" for each player, i.e., each players' allocation against the team's salary cap. The distinguishing feature of this accounting is that signing bonuses are prorated over the life of the contract, meaning that cap-charge compensation measures are "smoother" than the base+bonus measures.[8]

Insert Figure 5 about here

At this stage we are interested only in the initial costs of signing drafted players, so we consider just the first year's compensation using our 2000-2002 cap-charge dataset. Figure 5 shows first-year compensation (cap charges) as a function of draft order for 1996-2002. This pattern holds – though tempered over time – through the players' first five years, after which virtually all players have reached free agency and are therefore under a new contract, even if remaining with their initial teams.[9] The slope of this curve very closely approximates the draft-

---

[8] If players are cut or traded before the end of their contract, the remaining portion of the pro-rated bonuses is accelerated to that year's salary cap. So the real distinction between base+bonus numbers and cap charges is that the base+bonus charges are disproportionately distributed at the beginning of a player's relationship with a team, while the cap charges are, on average, disproportionately distributed at the end.
[9] Players are not eligible for free agency until after their 3rd year in the league. After 4 years players are eligible for restricted free agency. After 5 years players are unrestricted free agents and can negotiate with

pick value curve estimated in the previous section. Players taken early in the draft are thus

expensive on both counts: foregone picks and salary paid. Are they worth it?

IV. ON-FIELD PERFORMANCE

Recall that in order for the high value associated with early picks to be rational the

relation between draft order and on-field performance must be very steep. The market value of

the first pick in the first round, for example, is roughly four times as high as the last pick in the

first round, and the player selected first will command a salary nearly four times higher than the

player taken 30[th]. Does the quality of performance fall off fast enough to justify these price

differences? In the section following this one we will answer that question using a simple

econometric model of compensation, but before turning to that we offer a few direct measures of

on-field performance.

A. Data

Since we want to include players in every position in our analyses we report four

performance statistics that that are common across all positions: probability of being on a roster

(i.e., in the NFL), number of games played, number of games started, and probability of making

the Pro Bowl (a season ending "All-Star" game). We have these data for the 1991-2002 seasons.[10]

Our analysis involves all players drafted between 1991 and 2002. This means that we

observe different cohorts of players for different periods of time – e.g., we observe the class of

1991 for 12 seasons, but the class of 2002 for only one. While we cannot avoid this cohorting

effect, meaning draft classes carry different weights in our analysis, we are not aware of any

---

any team. This timeframe can be superseded by an initial contract that extends into the free-agency period,
e.g., 6 years and longer. Such contracts were exceedingly rare in the period we observe, though they have
become more common lately.

[10] All performance data are from Stats.Inc. 1991 is the earliest season for which the "games started" are
reliable.

systematic bias this imparts to our analyses.[11] An additional methodological issue is how to treat

players who leave the NFL. For our analysis it is important that our sample is conditioned on

players who are drafted, not players who are observed in the NFL during a season. Consequently,

we keep all drafted players in our data for all years, recording zeros for performance statistics for

those seasons a player is not in the NFL.

<div align="center">

Insert Table 2 about here

</div>

Observations in these data are player-seasons. We have 20,874 such observations, which

are summarized in Table 2. In our sample, the mean probability of making an NFL roster is 47%

per year, while the probability of making the Pro Bowl is 2% per year. The mean number of

games started per season is 3.19, and the mean number of games played in per season is 6.0 (NFL

teams play 16 regular season games per year. We do not include play-off games in our analysis).

Panel B of Table 2 shows how these performance measures change over time. The probability of

making a roster peaks in the player's first year (66%). Games played peaks in year 2 (mean=8.2),

starts in year 4 (mean=4.2), and the probability of making the Pro Bowl in year 6 (3.9%). Recall

that the sample is conditioned on the player being drafted, so these means include zeros for those

players out of the league. This panel also highlights the cohort effect in our sample, with the

number of observations declining with experience. One consequence is that our data are weighted

toward players' early years, e.g., 52% of our observations are from players' first 4 years.

B. Analysis

Our main interest is how player performance varies with draft order. Table 3 summarizes

our data by draft order, showing the average performance for players taken in each round of the

draft. Mean performance generally declines with draft round. For example, first-round picks start

an average of 8.79 games per season in our sample, while 7[th]-round picks start 1.21 games per

---

[11] We have done similar analyses on a sample restricted to players drafted 1991-1998, so that we observe a
full five years from all players. The graphs are virtually identical.

season. The table also lists performance in each round relative to the first round, placing all four

statistics on the same scale. We graph these relative performance statistics in Figure 6. We limit

this graph to the first 7 rounds, the length of the draft since 1994. The graph shows that all

performance categories decline almost monotonically with draft round. This decline is steepest

for the more extreme performance measures – probability of Pro Bowl is steeper than starts,

which is steeper than games played, which is steeper than probability of roster. Finally, we

include on the graph the compensation curve we estimated in the previous section. This curve is

steeper than all the performance curves except the Pro Bowl curve, which it roughly

approximates. The fact that performance declines more slowly than compensation suggests that

early picks may not be good investments, just as we report in the next section.

<div align="center">

Insert Table 3 about here

</div>

This analysis shows how performance varies with large differences in draft position – the

average difference across our one-round categories is 30 picks. A complementary analysis is to

consider performance variation with smaller differences in draft position. After all, teams trading

draft picks typically do not move up entire rounds, but rather half rounds (recall that the overall

average move is 13.3 picks), and even less at the top of the draft (the average move is 7.8 picks in

the top two rounds). One way to investigate these smaller differences is to consider whether a

player is better than the next player drafted at his position. Two observations suggest such an

analysis would be appropriate. First, draft-day trades are frequently for the purpose of drafting a

particular player, implying the team prefers a particular player over the next one available at his

position. So a natural question is whether there are reliable differences in the performance of two

"adjacent" players. A second observation suggesting this analysis is the average difference

between "adjacent" draftees of the same position, 8.26 picks, very closely matching the average

move by teams trading picks in the top 2 rounds (7.8).

<div align="center">

</div>

For this analysis we consider a player's performance over his observed career. Using the performance data described above, we observe 3,114 drafted players for an average of 4.8 years. We use two comparisons to determine which player is "better" – the average number of games started per season, and the per-season probability of making the Pro Bowl. Using all players from 1991-2002 drafts, we consider whether a player performs better than subsequent players drafted at his same position. We use two different samples in this analysis – one that includes all rounds of the draft, and one that includes only the first round. Finally, we vary the lag between players – from 1 (i.e., the next player at drafted at his position) to 4. One way to think of this analysis is that we're asking how far a team has to trade up (within a position) in order to obtain a player that is significantly better than the one they could have picked without moving. Note that this analysis is silent on the cost of trading up, focusing exclusively on the benefits.

One methodological challenge in this analysis is dealing with ties. Since we are interested in the probability a player performs better than another player, we would normally expect a binary observation – 1 for yes and 0 for no. But ties are relatively common, and informative. Censoring them would remove valuable information, while grouping them with either of the extreme outcomes would create a significant bias. Hence we code ties as .5.

The results are shown in Figure 7. There are three notable features in the data. First, for the broadest samples, the probabilities are near chance. Across all rounds, the probability that a player starts more games than the next player chosen at his position is 53%. For the same sample, the probability that a player makes more Pro Bowls than the next player chosen at his position is 51%. Second, these probabilities are higher in the first-round sample than in the full sample. For example, the probabilities of the higher pick performing better in 1-player-lagged comparisons

24

are 58% and 55% for starts and Pro Bowls, respectively, for players drafted in the first round. Finally, longer lags improve discrimination for starts but not for Pro Bowls. For example, across all rounds, the probability of the higher pick starting more games rises from 53% to 58% as we move from 1-player lags to 4-player lags. The rise is even steeper in the first round, increasing from 58% to 69%. The probability of making the Pro Bowl, however, is quite constant for all lags, in both samples.

C. Discussion

There are three important features of the relationship between on-field performance and draft order: 1) performance declines with the draft round (for all measures and almost all rounds), 2) the decline is steeper for more extreme performance measures, and 3) only the steepest decline (Pro Bowls) is as steep as the compensation costs of the draft picks.

We should note that there are two biases in the data that work toward making measured performance decline more steeply than actual performance. First, teams may give high draft picks, particularly early first round picks, "too much" playing time. Such a bias has been found in the National Basketball Association (by Camerer & Weber, 1999; and Staw & Hoang, 1995). These researchers found that draft-order predicts playing time beyond that which is justified by the player's performance. The explanation is that teams are loath to give up on high draft-choices because of their (very public) investment in them. It seems likely this bias exists in the NFL as well, which has a similarly expensive, high-profile college draft. If so, our performance statistics for high draft-choices will look better than they "should". This is especially true in our sample, which is disproportionately weighted by players' early years. To the extent that such a bias exists in the NFL, these results suggest even more strongly that draft-pick value declines too steeply.

The data on Pro Bowl appearances are also biased in a way that makes the performance-draft-order curve too steep. Selections to the Pro Bowl are partly a popularity contest, and players who were high first round picks are likely to have greater name recognition.

25

The within-position analysis provides a finer-grain look at performance by draft order. Here we see that whether a player will be better than the next player taken at his position is close to a coin-flip. These odds can be improved by comparing against those who are taken 3 or 4 players later (again, within position), or by focusing on first-round comparisons only. But even in those cases the probability that one player is better than those taken after him are relatively modest. Combined these analyses suggest that player performance is quite different across draft rounds, but not very different within rounds. In other words, draft order provides good information "in the large", but very little "in the small".

Overall, these analyses support one of the main premises of this paper, namely that predicting performance is difficult, and that the first players taken are not reliably better than ones taken somewhat later. Still, we have not yet addressed the question of valuation—do the early picks provide sufficient value to justify their high market prices? To answer this question we need a method that considers costs and benefits in terms of utility to the team. We turn to such an approach in the next section.

V. COST-BENEFIT ANALYSIS

In this section we estimate the *surplus value* of drafted players, that is the value they provide to the teams less the compensation they are paid. To estimate the value teams assign to various performance levels we start with the assumption that the labor market for veteran players (specifically, those in their 6[th] year in the league) is efficient. By the time a player has reached his sixth year in the league he is under no obligation to the team that originally drafted him, and has had the opportunity to sign (as least one) "free agent" contract. Players at this point have also had five years to establish their quality level, so teams should have a good sense of what they are buying, especially when compared to rookie players with no professional experience. Using these data on the compensation of sixth-year players we estimate the value teams assign to various performance categories for each position (quarterback, linebacker, etc.). We then use

these estimates to value the performance of all drafted players in their first five years and compare these performance valuations to the players' compensation in order to calculate the surplus value a team receives. Our interest is the relationship between surplus value and draft order.

A. Data

We use the performance data described in the previous section and summarized in Table 2. For this analysis, the sample we use is all players who were drafted 1991-1998 so that we observe five years of performance for every player. We place each player-season into one of five mutually exclusive categories: 1) not in the league ("NIL"), 2) started 0 games ("DNS"), 3) started 8 or fewer games ("Backup"), 4) started more than 8 games ("Starter"), 5) selected to the Pro Bowl ("Pro Bowl"). While the first and last performance categories are obvious boundaries, the middle three can be created in a variety of ways. We chose this particular division to emphasize "starters", defined here as players who started more than half the games (each team plays 16 games in the NFL's regular season). This scheme also has the empirical virtue of creating three interior categories of roughly equal size.

For player i in his t-th year in the league, this scheme produces five variables of the form

$$Cat\_n_{i,t} = \{0,1\}$$
(8)

indicating qualification for performance category n according to the criteria discussed above. We also calculate the average of these performance variables,

$$Cat\_\bar{n}_{i,t} = \frac{1}{t}\sum_{j=1}^{t} Cat\_n_{i,j}$$
(9)

over the first t years in player i's career.

Insert Table 4 about here

The data are summarized in Table 4. In Panel A, observations are player-seasons. The first category – players who are not in the league – is easily the largest, with 43% of the observations. Categories 2 through 4 are roughly equal in size, with 18%, 17% and 19%, respectively. The 5[th] category, Pro Bowls, is the smallest, with 2% of the player-seasons. In Panel B, observations are aggregated over a player's first five seasons. While the averages for each category remain the same, the complete distribution provides a bit more information. For example, we see that the median drafted player is out of the league for two of his first five years, and never starts a game.

B. Analysis

We are interested in the market value of different levels of player performance – backups, starter, Pro Bowl, etc. To do this we investigate the relation between a player's 6[th]-year compensation and his performance during his first five years after being drafted. Specifically, we estimate compensation models of the form:

$$Log(Comp_{i,6}) = \alpha + \beta_1 Cat\_\overline{1}_{i,5} + \beta_2 Cat\_\overline{2}_{i,5} + \beta_3 Cat\_\overline{3}_{i,5} + \beta_4 Cat\_\overline{4}_{i,5}$$
$$+ \beta_5 Cat\_\overline{5}_{i,5} + \mathbf{BI}_i + \varepsilon_{i,t},$$

(10)

in which **I** is a vector of indicator variables for the player's position (quarterback, running back, etc.) and $Cat\_\overline{n}_{i,5}$ is player i's relative frequency in performance category n over his first five years. We take the model's predicted values as the estimated market value of each position-performance pair. This general approach is similar to that of previous research on NFL compensation (Ahlburg & Dworkin, 1991; Kahn, 1992; Leeds & Kowalewski, 2001) though we rely more heavily on performance *categories* instead of individual statistics. We omit the bottom two performance categories for our estimation.[12] We use tobit regression for our estimates since our compensation measure is left-censored at the league minimum.[13]

---

[12] This of course collapses them into a single category for this estimation. We do this because we will assume, in the second stage of our analysis, that the value of a player's performance is zero if he is not on a roster. Collapsing these two categories is a means of imposing the zero-value assumption for the category.

Insert Table 5 about here

Table 5 summarizes the compensation data we use as our dependent measures. We show the results of this estimation in Table 6. Model 1 uses the player's cap charge as the dependent measure. Coefficient estimates are in log terms so are difficult to interpret directly – below we turn to a table of transformed values to see model implications in real terms. The coefficients are ordered monotonically, as we would expect. Categories 4 and 5 are significantly different than the omitted categories (1&2) and from each other. In Model 2 we include a variable for the draft-pick value of each player to test whether, controlling for our performance categories, draft position explains 6^{th}-year compensation. To examine this we add a variable equal to the estimated value (from section 3) of the draft pick used to select a player. The coefficient on this variable is not significantly different from zero, and the estimated coefficients on the other variables are essentially unchanged. This is important because it tells us that draft order is not capturing some other unobserved measure of quality. In other words, there is nothing in our data to suggest that former high draft picks are better players than lower draft picks, beyond what is measured in our broad performance categories.

Insert Table 6 about here

In the remaining models we use base+bonus compensation as the dependent measure. Because we have a longer history of compensation in these terms, using this variable increases the size of our sample. This also serves as a robustness test of our model. In model 3 we restrict the estimates to the same sample we use for the cap-charge models so we can compare the results directly. The results are broadly similar, with coefficients ordered monotonically and the top two

There are relatively few player-seasons in the bottom category among the players receiving compensation in year 6 (77 out 1370), but their inclusion presumably biases downward the estimates for Category-2 performance.
[13] These models are censored at $300,000.

categories significantly different from the bottom two and from each other. Draft-pick value is positive in this model, though not significantly so. We expand the sample for model 4, using all seven years of data rather than just the three for which cap-charge information is available. All patterns and formal tests are the same.

To ease the interpretation of model estimates, we transform the predicted values for each position and performance category. These values are summarized for each model in Table 7. The most distinct feature of these value estimates is that they increase with performance, as expected. For model 1, for example, values range from $0.5 to 1.0 million for category 2 (Starts==0), $0.7 to 1.36 million in category 3 (Starts<=8), $2.1 to 4.2 million for category 4 (Starts>8), and $5.0 to 10.1 million for category 5 (Pro Bowl). A second feature is that while there are small differences across positions, the only significant difference is that quarterbacks are valued more highly than other positions. There are also some differences across models, though these are subtle. Broadly, the models give very similar values despite varying both the dependent variable and the sample period.

Insert Table 7 about here

The second step in our analysis is to evaluate the costs and benefits of drafting a player. To do this we apply the performance value estimates from the previous section to performances in players' first five years. This provides an estimate of the benefit teams derive from drafting a player, having exclusive rights to that player for three years and restricted rights for another two. Specifically, we calculate the surplus value for player i in year t,

$$S\hat{V}_{i,t}^{Cap} = P\hat{V}_{i,t}^{Cap} - C_{i,t}^{Cap}, \tag{11}$$

where $P\hat{V}_{i,t}^{Cap}$, a function of the player's performance category and position, is the predicted

value from the compensation model estimated in the previous section, and $C_{i,t}^{Cap}$ is the player's

compensation costs. Our interest is in the relationship between surplus value and draft order.

Note that we make all calculations using the player's cap charge, both for compensation costs and as the basis for the performance value estimates. We also calculate an alternative measure of surplus value,

$$S\hat{V}_{i,t}^{Base} = P\hat{V}_{i,t}^{Base} - C_{i,t}^{Base},$$

(12)

in which the variables are calculated in the same way but rely on base+bonus compensation rather than cap charges. We prefer the cap charge to base+bonus both because it is smoother and, since the salary cap is a binding constraint, this charge reflects the opportunity cost of paying a player – a cap dollar spent on one player cannot be spent on any other. A downside of using cap charges is that our data span fewer years. For this reason, and to check the robustness of our results, we will use both approaches. We explicate our analysis using the cap-charge calculations, but then include both in the formal tests.

Insert Table 8 about here

Our sample is for the 2000-2002 seasons, including all drafted players in their first five years in the league. The performance value estimates, compensation costs, and surplus value calculations are summarized in Table 8. The mean cap charge is $485,462, while the mean estimated performance value is $955,631, resulting in a mean surplus value of $470,169.[14] We graph all three variables in Figure 8. We are most interested in the third panel, estimated surplus value as a function of draft order. The market value of draft picks suggests that this relationship should be negative – that there should be less surplus value later in the draft. In fact, the market-

---

[14] Our compensation data include only players who appear on a roster in a given season, meaning our cap charges do not include any accelerated charges incurred when a player is cut before the end of his contract. This creates an upward bias in our cap-based surplus estimates. We cannot say for sure whether the bias is related to draft order, though we strongly suspect it is *negatively* related to draft order – i.e,. there is less upward bias at the top of the draft – and therefore works against our research hypothesis. The reason for this is that high draft picks are much more likely to receive substantial signing bonuses. Recall that such bonuses are paid immediately but amortized across years for cap purposes. Thus when a top pick is cut we may miss some of what he was really paid, thus underestimating his costs. Note that this bias does not exist with the base+bonus compensation measure since the bonus is charged when paid.

value curve suggests the relationship should be steeply negative at the top of the draft. By contrast, this graph appears to show a positive relationship between surplus value and draft order, especially at the top of the draft.

Insert Figure 8 about here

To represent the data in a more helpful manner we fit lowess curves to these three scatterplots. We show these curves in Figure 9. It is noteworthy that performance value is everywhere higher than compensation costs, and so surplus is always positive. This implies that the rookie cap keeps initial contracts artificially low, at least when compared to the 6th year players who form the basis of our compensation analysis. More central to the thrust of this paper is the fact that while both performance and compensation decline with draft order, compensation declines more steeply. Consequently, surplus value *increases* at the top of the order, rising throughout the first round and into the second. That treasured first pick in the draft is, according to this analysis, actually the least valuable pick in the first round! To be clear, the player taken with the first pick does have the highest expected performance (that is, the performance curve is monotonically decreasing), but he also has the highest salary, and in terms of performance per dollar, is less valuable than players taken in the second round.

Insert Figure 9 about here

To look more closely at the relation between surplus value and draft order we graph that lowess curve in isolation in Figure 10. The curve shows positive value everywhere, increasing over the first 43 picks before declining for the subsequent 200. Surplus value reaches its maximum of ~$750,000 at the 43rd pick, i.e., the 10th pick in the 2nd round of the draft.[15]

_____

[15] The curve goes back up toward the end of the draft but we do not think much should be made of this. It is primarily due to a single outlier, Tom Brady, the all-pro quarterback for the Patriots who was drafted in the 6th round, position 199!

Clearly, considerable caution should be used in interpreting this curve; it is meant to summarize the results simply. We do not have great confidence in its precise shape. More important for our hypothesis is a formal test of the relationship between the estimated surplus value and draft order. Specifically, we need to know whether this relationship is less negative than the market value of draft picks. Having established in section 3 that the market value is strongly negative, we will take as a sufficient (and very conservative) test of our hypothesis whether the relationship between surplus value and draft order is ever positive. Of course this relationship varies with draft-order, so the formal tests need to be specific to regions of the draft. We are distinctly interested in the top of the draft, where the majority of trades – and the overwhelming majority of value-weighted trades – occur. Also, the psychological findings on which we base our hypothesis suggest the over-valuation will be most extreme at the top of the draft.

As a formal test we regress estimated surplus value on a linear spline of draft order. The spline is linear within round and knotted between rounds. Specifically, we estimate

$$
\begin{aligned}
S\hat{V}_{i,t}^{Cap} = {} & \alpha + \beta_1 Round1 + \beta_2 Round2 + \beta_3 Round3 + \beta_4 Round4 \\
& + \beta_5 Round5 + \beta_6 Round6 + \beta_7 Round7 + \varepsilon_{i,t},
\end{aligned}
\tag{13}
$$

where Round$j$ is the linear spline for round j. In this model, then, β_j provides the estimated per-pick change in surplus value during round j. Estimation results are shown in Table 9, model 1. The estimate for *Round1* is significantly positive (p<.01), with a slope of +$18,500 per draft pick. Five of the subsequent six rounds are negative, though only the estimate for Round2 is significantly so (p<.05).

We also estimate

$$S\hat{V}_{i,t}^{Base} = \alpha + \beta_1 Round1 + \beta_2 Round2 + \beta_3 Round3 + \beta_4 Round4$$
$$+ \beta_5 Round5 + \beta_6 Round6 + \beta_7 Round7 + \varepsilon_{i,t},$$

(14)

using the surplus value calculation, $S\hat{V}_{i,t}^{Base}$, which relies on the base+bonus measure of

compensation. We estimate this model over two time periods. The first is for the same time

period as the cap-charge model (2000-2002), while the second is for the broader sample (1996-

2002). Results are shown in Table 9, models 2 and 3, respectively. Most important, the slope for

the first round is significantly positive in both models (p<.01). Estimated first-round slopes are

actually steeper in these models than in the cap-charge model. As in the cap-charge model, the

slopes of subsequent rounds are almost universally negative, though not significantly so.

Finally, we consider a model for player "careers" rather than seasons. We calculate a

new surplus value measure,

$$S\hat{V}_{i,t}^{Car} = \frac{1}{t} \sum_{j=1}^{t} S\hat{V}_{i,t}^{Base},$$

(15)

for player i's average surplus value over the first t years of his career.[16] We then estimate the

model for how the average surplus value over a player's first five years varies with draft order,

$$S\hat{V}_{i,5}^{Car} = \alpha + \beta_1 Round1 + \beta_2 Round2 + \beta_3 Round3 + \beta_4 Round4$$
$$+ \beta_5 Round5 + \beta_6 Round6 + \beta_7 Round7 + \varepsilon_{i,5}.$$

(16)

The results of this estimation are in Table 9, model 4. We find the same pattern as in the

previous models – the first round is significantly positive (p<.05) and subsequent rounds are

negative but not significantly so. The estimated slope on the first round is $+14,000/pick.

Insert Figure 11 about here

We graph the predicted values from these four models in Figure 10. The models all peak

between the first and second round, though we cannot conclude much about the precise location

---

[16] We are unable to calculate a cap-charge version of this 5-year average since we our sample includes only
three years of cap-charge data.

of the peak since the splines are identical for each model and were constructed independently of surplus value. The four models produce patterns that are broadly similar – sharp first-round increases followed by gradual, almost monotonic declines. The maximum surplus value is highest for the models based on the 2000-2002 period, peaking near $1 million, while the two models based on broader samples peak near $750,000. The first-round increase is sharpest for the two base+bonus models using player-season data, both of which show surplus value to be approximately $0 at the very top of the draft.

C. Discussion

Let's take stock. We have shown that the market value of draft picks declines steeply with draft order—the last pick in the first round is worth only 25 percent of the first pick even though the last pick will command a much smaller salary than the first pick. These simple facts are incontrovertible. In a rational market such high prices would forecast high returns; in this context, stellar performance on the field. And, teams do show skill in selecting players—using any performance measure, the players taken at the top of the draft perform better than those taken later. In fact, performance declines steadily thoughout the draft. Still, performance does not decline steeply enough to be consistent with the very high prices of top picks. Indeed, we find that the expected surplus to the team declines throughout the first round. The first pick, in fact, has an expected surplus lower than any pick in the second round!

The magnitude of the market discrepancy we have uncovered is strikingly large. A team blessed with the first pick could, though a series of trades, swap that pick for as many as six picks in the middle of the second round, each of which is worth considerably more than the single pick they gave up. Mispricing this pronounced raises red flags: is there something we have left out of our analysis that can explain the difference between market value and expected surplus?

Since both the market value of picks and the compensation to players are easily observable, the only place our measurements can be seriously off is in valuing performance. Two specific sources of measurement error come to mind. 1. Do top draft picks provide superior

performance in ways that our crude performance measures fail to capture? 2. Do superstars (some of whom are early draft picks) provide value to teams beyond their performance on the field, e.g., in extra ticket sales and/or sales of team apparel such as uniform jerseys? We discuss each in turn.

It is certainly possible that a player such as Peyton Manning provides value to a team (leadership?) that is not captured in statistics such as passing yards. However, we do not think this can explain our results. Most of the variance in player surplus is generated by whether a player becomes a regular. Recall that even first-round draft picks are about as likely to be out of the league as playing in the Pro Bowl. The possibility of landing a so-called "franchise" player is simply too remote to be able to explain our results. Furthermore, if high draft-choices had some intangible value to teams beyond their on-field performance this presumably would be revealed by a significant coefficient on the draft-order term in our compensation regressions. Instead, we find that draft order is not a significant explanatory variable after controlling for prior performance.

A more subtle argument is that the utility to the team of signing a high draft pick is derived from something beyond on-field performance. A very exciting player, Michael Vick comes to mind, might help sell tickets and team paraphernalia even if he doesn't lead the team to many victories. We are skeptical of such arguments generally. Few football players (Vick may be the only one) have the ability to bring in fans without producing wins. But in any case, if high draft-picks had more fan appeal this should show up in their 6^{th} year contracts, and we find no evidence for it.

We have also conducted some other analyses not reported here that add support to the interpretation that high draft-picks are bad investments. For example, for running backs and wide receivers we have computed the statistic "yards gained per dollar of compensation" (using total rushing, receiving and return yards). This simple performance statistic *increases* with draft order at the top of the draft in each of the players' 4 years. We have also computed the number of

games started divided by compensation paid for the first two rounds of the draft, and we find similar results: games started per dollar increases throughout the first two rounds. Pro Bowl appearances per dollar of compensation also increases, even more sharply. Finally, we have estimated the surplus-value-by-draft-order curve shown in Figure 10 separately for two groups of players: "ball handlers" (quarterbacks, wide-receivers and running backs) and others. Presumably, if teams are using high draft-picks to generate fan interest for reasons beyond winning games this argument would apply primarily to the ball handlers who generate the most attention. Yet, we find a significantly positive slope of the surplus value curve for non-ball-handlers throughout the first round.

.

VI. CONCLUSION

Psychologists who study decision making are sometimes criticized for devising what are said to be artificial, contrived, laboratory experiments in which subjects are somehow tricked into making a mistake. In the "real world", the critics allege, people learn over time to do pretty well. Furthermore, the critics add, people specialize, so many difficult decisions are taken by those whose aptitude, training, and experience make them likely to avoid the mistakes that are so prevalent in the lab. This criticism is misguided on many counts. For example, we all have to decide whether to marry, choose careers, and save for retirement, whether or not we are experts— whatever that might mean—in the relevant domain. More germane to the topic of this paper, even professionals who are highly skilled and knowledgeable in their area of expertise are not necessarily experts at making good judgments and decisions. Numerous studies find, for example, that physicians, among the most educated professionals in our society, make diagnoses that display overconfidence and violate Bayes' rule (cf. Christensen-Szalanski & Bushyhead, 1981; Eddy, 1982). The point, of course, is that physicians are experts at medicine, not

37

necessarily probabilistic reasoning. And it should not be surprising that when faced with difficult problems, such as inferring the probability that a patient has cancer from a given test, physicians will be prone to the same types of errors that subjects display in the laboratory. Such findings reveal only that physicians are human.

Our modest claim in this paper is that the owners and managers of National Football League teams are also human, and that market forces have not been strong enough to overcome these human failings. The task of picking players, as we have described here, is an extremely difficult one, much more difficult than the tasks psychologists typically pose to their subjects. Teams must first make predictions about the future performance of (frequently) immature young men. Then they must make judgments about their own abilities: how much confidence should the team have in its forecasting skills? As we detailed in section 2, human nature conspires to make it extremely difficult to avoid overconfidence in this task. The more information teams acquire about players, the more overconfident they will feel about their ability to make fine distinctions. And, though it would seem that there are good opportunities for teams to learn, true learning would require the type of systematic data collection and analysis effort that we have undertaken here. Organizations rarely have the inclination to indulge in such time-intensive analysis.   In the absence of systematic data collection, learning will be inhibited by bad memories and hindsight bias.

We began this study with the strong intuition that teams were putting too high a value on choosing early in the draft. We thought it crazy for the Giants to give up so many picks for the opportunity to move up from the fourth pick to the first one (regardless of which player they chose). But we concede that we did not expect the findings to be as strong as those we report. Rather than a treasure, the right to pick first appears to be a curse. If picks are valued by the surplus they produce, then the first pick in the first round is the worst pick in the round, not the best! In paying a steep price to trade up, teams seem to be getting the sign wrong! We have done numerous "reality checks" to see whether these surprising conclusions are robust, and every

38

analysis gives qualitatively similar results. So, suppose our analyses are taken at face value. Can

they be right? This is a big market, after all, with franchises worth perhaps $1 billion or more.

We think that while our results are surprising, they are plausible. We suspect that some

teams have not fully come to grips with the implications of the salary cap, a relatively new

innovation. Buying expensive players, even if they turn out to be great performers, imposes

opportunity costs elsewhere on the roster. Spending $10 million on a star quarterback instead of

$5 million on a journeyman implies having $5 million less to spend on offensive linemen to block

or linebackers to tackle. Some of the successful franchises seem to understand these concepts,

most notably the New England Patriots, but others do not. Whether because they are smart about

these ideas or others, the Patriots have been doing well recently, and so have not had high draft

picks to use. We can only speculate about whether they would trade down if they somehow

ended up with one of the earliest and most overvalued picks. But notice that if a few teams do

learn and have winning records, there is no market action they can take to make the implied value

of draft picks rational. Indeed, the irony of our results is that the supposed benefit bestowed on

the worst team in the league, the right to pick first in the draft, is really not a benefit at all, unless

the team trades it away.[17] The first pick in the draft is the loser's curse.

The loser's curse can persist even in competitive markets for a reason similar to why the

winner's curse can persist: there are limits to arbitrage. If naïve oil companies bid too much for

drilling rights, then sophisticated competitors can only sit on the sidelines and hope their

competitors go broke – or eventually learn. Since there is no way to sell the oil leases short, the

smart money cannot actively drive the prices down. Similarly, since there is no way to sell the

first draft pick short, there is no way for any team other than the one that owns the pick to exploit

the teams that put too high a value on it Finally, now that the draft-pick value curve is widely

used and accepted in the NFL a team that owns a top draft pick and would like to trade it may be

---

[17] We do note that the San Diego Chargers, the team that took Ryan Leaf with the second pick only to have him flop, has now traded the number one pick twice. This year they are headed to the playoffs. Lesson learned?

reluctant to make a trade at less than "full value". So, even trading down will be hard unless there is a buyer willing to pay the inflated but conventional price.

The implications of this study extend beyond the gridiron. Football players are surely not the only employees whose future performance is difficult to predict. In fact, football teams almost certainly are in a better position to predict performance than most employers choosing workers. Teams get to watch their job candidates perform a very similar task at the college level and then get to administer additional tests on highly diagnostic traits such as strength and speed. Finally, once hired, performance can and is graded, with every action visible on film from multiple angles! Compare that to a company looking to hire a new CEO (or an investment bank hiring an analyst, a law firm hiring an associate, etc.). Candidates from outside the firm will have been performing much of their job out of view. Outside observers see only a portion of the choices made, and options not taken are rarely visible externally. And, even once a CEO is hired, the company's board of directors is unlikely to be able measure his or her performance nearly as accurately as a team can evaluate its quarterback. In our judgment, there is little reason to think that the market for CEOs is more efficient than the market for football players. Perhaps innovative boards of directors should start looking for the next Tom Brady as CEO rather than Eli Manning.

References

Ahlburg, D. A., & Dworkin, J. B. 1991. Player compensation in the National Football League. In P. D. Staudohar, & J. A. Mangan (Eds.), *The Business of Professional Sports*. Urbana: University of Illinois Press.

Bronars, S. G. 2004. Bargaining in Professional Football: Why NFL Superstars are Underpaid. *Working paper.*

Camerer, C., & Weber, R. 1999. The econometrics and behavioral economics of escalation of commitment: a re-examination of Staw and Hoang's NBA data. *Journal of Economic Behavior & Organization*, 39(1): 59-82.

Capen, E. C., Clapp, R. V., & Campbell, W. M. 1971. Competitive bidding in high-risk situations. *Journal of Petroleum Technology*, 23: 641-653.

Christensen-Szalanski, J. J. J., & Bushyhead, J. B. 1981. Physicians' use of probabilistic information in a real clinical setting. *Journal of Experimental Psychology: Human Perception and Performance*, 7: 928-935.

De Bondt, W. F. M., & Thaler, R. 1985. Does the Stock-Market Overreact. *Journal of Finance*, 40(3): 793-805.

Dessauer, J. P. 1981. *Book Publishing*. New York: Bowker.

Eddy, D. M. 1982. Probabilistic reasoning in clinical medicine: Problems and opportunities. In D. Kahneman, P. Slovic, & A. Tversky (Eds.), *Judgment Under Uncertainty: Heuristics and Biases*: 249-267. Cambridge: Cambridge.

Gilovich, T., Griffin, D. W., & Kahneman, D. 2002. *Heuristics and biases: the psychology of intuitive judgement*. Cambridge: Cambridge University Press.

Hall, B. J., & Lim, J. P. 2002. A note on incentives in the NFL. *Harvard Business School case*.

Harrison, J. R., & Bazerman, M. H. 1995. Regression to the mean, expectation inflation, and the winner's curse in organizational contexts. In R. Kramer, & D. M. Messick (Eds.), *The Social Context of Negotiation*: 69-94. Thousand Oaks, CA.: Sage Publications.

Harrison, J. R., & March, J. G. 1984. Decision-making and post-decision surprises. *Administrative Science Quarterly*, 29(1): 26-42.

Kagel, J. H., & Levin, D. 1986. The winner's curse and public information in common value auctions. *American Economic Review*, 76(5): 894-920.

Kahn, L. M. 1992. The effects of race on professional football players compensation. *Industrial & Labor Relations Review*, 45(2): 295-310.

Kahneman, D., & Tversky, A. 1973. On the psychology of prediction. *Psychological Review*, 80: 237-251.

Kahneman, D., & Tversky, A. 2000. *Choices, values, and frames*. Cambridge / New York: Cambridge / Russell Sage Foundation.

King, P. 2004. Three-quarterback monte: SI.com.

Leeds, M. A., & Kowalewski, S. 2001. Winner take all in the NFL: The effect of the salary cap and free agency on the compensation of skill position players. *Journal of Sports Economics*, 2(3): 244-256.

Lichtenstein, S., Fischhoff, B., & Phillips, L. 1982. Calibration of probabilities: The state of the art to 1980. In D. Kahneman, P. Slovic, & A. Tversky (Eds.), *Judgment Under Uncertainty: Heuristics & Biases*: 306-333. Cambridge: Cambridge University Press.

Neyer, R. 2003. Examining the art of evaluating: ESPN.com.

Oskamp, S. 1965. Overconfidence in case-study judgments. *The Journal of Consulting Psychology*, 29: 261-265.

Roll, R. 1986. The hubris hypothesis of corporate takeovers. *Journal of Business*, 59(2): 197-216.

Ross, L., Greene, D., & House, P. 1977. The false consensus effect: an egocentric bias in social perception and attribution processes. *Journal of Experimental Social Psychology*, 13: 279-301.

Russo, J. E., & Schoemaker, P. J. H. 2002. *Winning Decisions: Getting It Right The First Time*. New York: Doubleday.

Samuelson, W. F., & Bazerman, M. H. 1985. The winner's curse in bilateral negotiations. *Research in Experimental Economics*, 3: 105-137.

Slovic, P., & Corrigan, B. 1973. *Behavioral problems of adhering to a decision policy*, Institute for Quantitative Research in Finance.

Staw, B. M., & Hoang, H. 1995. Sunk costs in the NBA - why draft order affects playing time and survival in professional basketball. *Administrative Science Quarterly*, 40(3): 474-494.

Thaler, R. H. 1988. The winner's curse. *Journal of Economic Perspectives*, 2(1): 191-202.

Figure 1
Draft-day trades
1998-2004. Excludes trades for which one team exchanged only players or future picks, as well as two picks in rounds 11 & 12. n=238.

Figure 2
Market value of draft picks

The estimated value of picks in the NFL draft, relative to the first pick. Shown is a Weibull distribution with an estimated scale parameter of .148 (se=.03) and an estimated shape parameter of .700 (se=.033). Estimates are based on 213 draft-day trades observed between 1988 and 2004, excluding trades involving players or future picks.

Figure 3
Market value of draft picks, top pick versus rest of bundle, estimated values
A comparison of estimated values for "both sides" of a trade – the top pick acquired, and the net exchange of all other picks in the trade. These equate to the left-hand and right-hand sides of expression (4), respectively, calculated with the estimated Weibull parameters. There are at least two interpretations of this graph. First, it provides an evaluation of the fit of the estimated model. Second, it suggests the relative "bargain" of each trade – those below the line represent trades that cost less (from the perspective of the party trading up) than expected by the model, while those above the line represent trades that cost more (from the perspective of the party trading up) than expected.

Figure 4
Market value of draft picks over time

The estimated value of picks in the NFL draft, relative to the first pick. Shown are separate estimates for the first half (1988-1996) and second half of the sample period (1997-2004). Estimates are based on 276 draft-day trades observed between 1988 and 2004, excluding trades involving players.

Figure 5
1st-year Compensation, by draft order

Compensation is the cap charge. Drafted players only, 2000-2002. Showing only the first 150 picks.

Figure 6
Performance statistics, by draft round

Observations are player-seasons, for the 1991-2002 seasons, for all players drafted 1991-2002. N=20,874. Graphed are averages for players drafted in each round, relative to the averages for players drafted in the first round. Compensation is average base+bonus during the players' first five years, 1996-2002 (N=8,072).

Figure 7
Within-Position Relative Performance, by Draft Order
The probability that a player is better, over his observed career, than subsequent players drafted at the same position. The *x* axis is the distance between the players compared, in terms of players – 1 is a 1-player lag (e.g., is a QB better than the next QB taken), 2 is a 2-player lag, etc. Two performance statistics are used – average number of starts per season, and the probability of making the Pro Bowl per season. These are each shown for two samples – the solid lines are for all rounds in the draft, and the dashed lines are for the first round only. Each comparison is coded for one of 3 outcomes – 1 if the higher pick is better, 0 if the lower pick is better, and .5 if they are equal. Observations are player careers, for the 1991-2002 seasons, for all players drafted 1991-2002. N=3,114.

Figure 8
Value and compensation by draft order

Observations are player-seasons. The sample is for the 2000-2002 seasons, drafted payers in their first five years in the NFL, excluding punters and kickers. A minimal amount of spherical noise is added to better convey the distribution of overlapping observations.

Figure 9
Performance, Compensation & Surplus

Summary lowess curves for players' first 5 years. Estimated from cap charges, 2000-2002.

Lowess curves. Based on cap charges, 2000-2002.

Figure 10
Surplus value by draft order

Lowess curve for the relationship between estimated surplus value and draft order.
Underlying observations are player-seasons. The sample is for the 2000-2002 seasons,
drafted players in their first five years in the NFL, excluding punters and kickers.

Lowess curves. Based on cap charges, 2000-2002.

Figure 11
Spline regressions, predicted values

Regressing estimated surplus value (performance value – compensation cost) on a linear spline of the draft order. The spline is knotted between rounds, so that the x variables reflect the estimated slope, i.e., .the change in surplus value, during a draft round. Surplus value is calculated using two different measures of compensation – cap charge and base+bonus. Sample is limited to all drafted players, excluding punters and kickers. Reported in 2002 dollars.

Table 1
Regression results
Using non-linear regression to estimate parameter values for a Weibull-function model of draft-pick value. Data are draft-day trades, 1988-2004. Excludes trades involving players (n=51), trades for which we have incomplete information (n=6), and trades on the second day of the 1990 draft (for which no information is available). Standard errors are italicized.

Model	(1)	(2)	(3)	(4)
Years	'88-'04	'88-'04	'88-'96	'97-04
Future Picks	No	Yes	Yes	Yes
Estimates				
lambda	0.148	0.121	0.092	0.249
	0.030	*0.023*	*0.029*	*0.064*
beta	0.700	0.730	0.775	0.608
	0.033	*0.030*	*0.052*	*0.042*
rho		1.738	1.590	1.735
		0.141	*0.237*	*0.126*
Market Values (relative to #1)				
10th pick	50%	55%	60%	39%
20th pick	31%	35%	41%	22%
30th pick	21%	24%	29%	15%
40th pick	15%	17%	21%	10%
50th pick	10%	13%	15%	7%
N	213	276	131	145
R-sq	1.00	0.99	0.98	0.99

Table 2
Performance data

Panel A: Summary statistics.

	Pr(Roster)	Games	Starts	Pr(Pro Bowl)
Obs	20,874	20,874	20,874	20,874
Mean	0.47	6.00	3.19	0.02
SD	0.50	7.04	5.62	0.15
Min	0.00	0.00	0.00	0.00
Max	1.00	16.00	16.00	1.00

Panel B: Mean performance by experience

Year	Pr(Roster)	Games	Starts	Pr(Pro Bowl)	Obs
1	0.661	7.420	2.256	0.004	3,114
2	0.660	8.235	3.744	0.019	2,852
3	0.592	7.599	4.112	0.024	2,606
4	0.527	6.981	4.204	0.034	2,352
5	0.466	6.233	3.943	0.037	2,099
6	0.402	5.402	3.494	0.039	1,858
7	0.342	4.724	3.279	0.038	1,618
8	0.293	3.899	2.732	0.035	1,364
9	0.226	2.989	2.083	0.023	1,115
10	0.165	2.261	1.649	0.015	893
11	0.088	1.158	0.743	0.007	669
12	0.063	0.731	0.452	0.006	334

Table 3
Performance data, by draft order

Round	N	Pr(Roster)		Games		Starts		Pr(Pro Bowl)	
		Mean	vs. Rd 1	Mean	vs. Rd 1	Mean	vs. Rd 1	Mean	vs. Rd 1
1	2,281	0.81	100%	11.15	100%	8.79	100%	0.10	100%
2	2,355	0.71	89%	9.68	87%	6.07	69%	0.04	43%
3	2,478	0.62	77%	7.87	71%	3.73	42%	0.03	26%
4	2,437	0.55	68%	6.68	60%	2.91	33%	0.01	9%
5	2,438	0.43	53%	5.21	47%	2.07	23%	0.02	16%
6	2,600	0.36	45%	4.25	38%	1.57	18%	0.01	10%
7	2,796	0.31	39%	3.50	31%	1.21	14%	0.01	6%
8	924	0.30	37%	3.79	34%	1.72	20%	0.02	15%
9	632	0.20	25%	2.32	21%	1.07	12%	0.00	3%
10	644	0.12	15%	1.56	14%	0.83	9%	0.00	0%
11	645	0.10	13%	1.20	11%	0.30	3%	0.00	0%
12	644	0.13	16%	1.51	14%	0.62	7%	0.00	2%
Total	20,874	0.47		6.00		3.19		0.03	

Table 4
Player performance

Summary statistics for performance categories, including performance in a player's first five years, for all players drafted 1991-1998. Performance categories are comprehensive and mutually exclusive. Observations in Panel A are player-seasons. Observations in Panel B are players.

Panel A: Player performance, season-level, years 1-5.

Category	Criteria	N	mean	sd	min	p25	p50	p75	max
NIL	Not in league	10,500	0.43	0.50	0.00	0.00	0.00	1.00	1.00
DNS	Starts=0	10,500	0.18	0.39	0.00	0.00	0.00	0.00	1.00
Backup	Starts<=8	10,500	0.17	0.38	0.00	0.00	0.00	0.00	1.00
Starter	Starts>8	10,500	0.19	0.39	0.00	0.00	0.00	0.00	1.00
Pro Bowl	Pro Bowl	10,500	0.02	0.15	0.00	0.00	0.00	0.00	1.00

Panel B: Mean performance by player over 1st 5 years.

Category	Criteria	N	mean	sd	min	p25	p50	p75	max
NIL	Not in league	2,099	0.43	0.41	0.00	0.00	0.40	0.80	1.00
DNS	Starts=0	2,099	0.18	0.22	0.00	0.00	0.20	0.20	1.00
Backup	Starts<=8	2,099	0.17	0.21	0.00	0.00	0.00	0.40	1.00
Starter	Starts>8	2,099	0.19	0.29	0.00	0.00	0.00	0.40	1.00
Pro Bowl	Pro Bowl	2,099	0.02	0.10	0.00	0.00	0.00	0.00	0.80

Table 5
Player Compensation

1996-2002. Includes only players drafted 1991-1997. Includes only the players' first six years in the league. Cap charges are available only for 2000-2002.

Panel A: Including players no longer in the league

Compensation	N	mean	sd	min	p25	p50	p75	max
Base Salary	6,494	451,273	744,839	0	0	279,346	573,643	9,271,004
Bonuses	6,494	299,814	739,906	0	0	0	224,592	11,127,144
Base + Bonus	6,494	756,759	1,281,164	0	0	320,883	883,246	13,905,222
Cap Charge	1,435	661,699	1,123,345	0	0	0	898,627	8,679,899

Panel B: Including only players still in the league

Compensation	N	mean	sd	min	p25	p50	p75	max
Base Salary	3,337	878,205	839,497	91,011	417,953	569,552	1,019,716	9,271,004
Bonuses	3,337	583,456	948,685	0	23,481	214,160	740,565	11,127,144
Base + Bonus	3,337	1,472,697	1,462,878	91,011	519,456	860,826	1,945,285	13,905,222
Cap Charge	621	1,529,047	1,261,074	79,412	537,606	1,104,316	2,108,367	8,679,899

Table 6
Compensation Models

Tobit regressions of year-6 compensation on performance over the player's first five years, in 2002 $s. Compensation is left-censored at the league minimum (here, log($300,000)). Compensation is either the player's charge against the salary cap (models 1-2) or the sum of his salary and bonuses (models 3-4). Sample is all drafted players, 1991-1997, who are on a roster in their sixth year, excluding kicker and punters. Omitted performance category is Starts=0. Position fixed-effects are included but not shown. Draft-pick value is estimated (above) from draft-day trades, 1988-2004.

	(1)	(2)	(3)	(4)
	Salary Cap		Base + Bonus	
	2000-2002	2000-2002	2000-2002	1996-2002
Starts<=8, avg	0.306	0.316	0.045	0.071
	[0.205]	[0.206]	[0.220]	[0.130]
Starts>8, avg	1.431	1.459	1.329	1.417
	[0.130]**	[0.140]**	[0.149]**	[0.097]**
Pro Bowl, avg	2.308	2.363	1.849	2.107
	[0.222]**	[0.244]**	[0.262]**	[0.175]**
Draft-pick value		-0.098	0.251	0.215
		[0.185]	[0.198]	[0.125]
Constant	13.818	13.824	13.969	13.837
	[0.210]**	[0.210]**	[0.225]**	[0.122]**
Pseudo R2	0.329	0.295	0.253	0.252
Log Likelihood	-230.83	-230.69	-248.68	-602.51
Observations	274	274	274	661
Left-censored	4	4	2	4
Uncensored	270	270	272	657
Censor value	12.61	12.61	12.61	12.61

Table 7
Estimated performance value, by position
Predicted values from the models reported in Table 6. In 2002 $s. Note that predicted values
apply when a performance category *x* variable equals 1. Hence, these values are the model
predictions for the player performance that falls into a category all 5 years. The draft-pick values
used for the model predictions are the averages for each position-performance pairing in the
sample used for the tobit regression.

Cat	Pos	(1) Salary Cap 2000-2002	(2) Salary Cap 2000-2002	(3) Base + Bonus 2000-2002	(4) Base + Bonus 1996-2002
Starts=0	DB	552,215	546,107	650,590	730,439
	DL	687,811	681,808	804,039	807,545
	LB	566,319	562,136	700,788	717,361
	OL	502,176	496,452	634,252	664,817
	QB	1,002,539	1,000,197	1,189,641	1,039,870
	RB	521,700	513,766	675,598	710,231
	TE	520,289	514,538	781,997	681,533
	WR	665,258	661,378	731,956	781,979
Starts<=8	DB	749,718	744,700	690,680	794,396
	DL	933,812	928,259	857,095	881,345
	LB	768,866	765,626	746,290	782,255
	OL	681,782	680,836	663,635	714,100
	QB	1,361,104	1,364,835	1,260,780	1,129,260
	RB	708,289	701,955	713,684	769,148
	TE	706,373	706,954	814,336	729,076
	WR	903,192	900,335	780,497	853,668
Starts>8	DB	2,310,530	2,317,961	2,543,610	3,105,027
	DL	2,877,882	2,888,141	3,159,747	3,447,937
	LB	2,369,542	2,387,840	2,734,457	3,044,277
	OL	2,101,163	2,110,926	2,468,537	2,815,147
	QB	4,194,740	4,200,189	4,780,221	4,525,227
	RB	2,182,854	2,187,005	2,621,899	3,000,045
	TE	2,176,950	2,193,064	3,025,005	2,870,848
	WR	2,783,517	2,790,640	2,905,468	3,367,574
Pro Bowl	DB	5,550,132	5,719,288	4,283,818	6,192,617
	DL	6,912,970	7,021,486	5,526,844	7,103,115
	LB	5,691,886	5,871,529	4,645,843	6,117,273
	OL	5,047,206	5,058,486	4,480,186	5,985,725
	QB	10,076,199	10,268,860	8,241,777	9,208,248
	RB	5,243,436	5,360,122	4,492,080	6,071,787
	TE	5,229,253	5,387,191	5,152,666	5,781,453
	WR	6,686,295	6,855,982	4,947,470	6,779,927

Table 8
Player Costs and Benefits
2000-2002 seasons, in 2002 dollars. Includes all drafted players in their first five years in the league, excluding punters and kickers.

Variable	N	mean	sd	min	p25	p50	p75	max
Cap Charge (*C*)	3,532	485,462	764,060	0	0	311,153	567,369	10,334,483
Estimated Value (*V*)	3,532	955,631	1,314,062	0	0	566,320	1,361,109	10,076,233
Surplus (*S*)	3,532	470,169	992,136	-6,667,058	0	90,637	543,895	9,747,337

Table 9
Spline Regressions

Regressing estimated surplus value (performance value – compensation cost) on a linear spline of the draft order. The spline is knotted between rounds, so that the x variables reflect the estimated slope, i.e., .the change in surplus, during a draft round. Surplus value is calculated using two different measures of compensation – cap charge and base+bonus. Observations are player-season, except for model 4, which is player-career, i.e., a player's average surplus over his first five years. Sample is limited to all drafted players, excluding punters and kickers. Reported in 2002 dollars.

	(1)	(2)	(3)	(4)
Basis for Surplus	Cap	Base+Bonus	Base+Bonus	Base+Bonus
Sample	2000-2002	2000-2002	1996-2002	1996-2002
Observations	Player-Season	Player-Season	Player-Season	Player-Career
Round1	18,553.87	31,896.13	29,413.63	14,492.71
	[3,866.591]**	[5,000.468]**	[3,006.051]**	[5,614.419]*
Round2	-6,912.95	-9,944.56	-4,471.08	-421.256
	[3,035.595]*	[3,925.783]*	[2,350.429]	[4,398.913]
Round3	-5,737.56	-3,903.16	-4,246.59	-5,658.61
	[2,975.638]	[3,848.244]	[2,298.657]	[4,338.640]
Round4	-5,478.00	-3,634.74	-3,124.70	-3,942.32
	[2,974.153]	[3,846.322]	[2,302.944]	[4,364.393]
Round5	-4,472.07	-4,852.06	-3,533.44	-2,104.10
	[2,948.024]	[3,812.532]	[2,281.149]	[4,314.675]
Round6	1,775.00	1,477.54	-227.911	-1,156.03
	[2,793.427]	[3,612.599]	[2,159.746]	[4,074.920]
Round7	-2,186.13	-2,078.01	-1,395.98	-3,188.80
	[1,808.139]	[2,338.376]	[1,456.271]	[2,809.491]
Constant	372,047.78	-21,745.27	-149,781.57	253,100.97
	[86,667.112]**	[112,082.232]	[67,355.703]*	[125,994.991]*
Observations	3502	3502	8021	717
R-squared	0.06	0.03	0.03	0.07