

The Loser's Curse: Decision Making and Market Efficiency in the National Football League Draft

Cade Massey

The Wharton School, University of Pennsylvania, Philadelphia, Pennsylvania 19104,
cadem@wharton.upenn.edu

Richard H. Thaler

Booth School of Business, University of Chicago, Chicago, Illinois 60637,
richard.thaler@chicagobooth.edu

A question of increasing interest to researchers in a variety of fields is whether the biases found in judgment and decision-making research remain present in contexts in which experienced participants face strong economic incentives. 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 monetary stakes are exceedingly high and the opportunities for learning are rich. It is also a domain in which multiple psychological factors suggest that teams may overvalue the chance to pick early in the draft. Using archival data on draft-day trades, player performance, and compensation, we compare the market value of draft picks with the surplus value to teams provided by the drafted players. We find that top draft picks are significantly overvalued in a manner that is inconsistent with rational expectations and efficient markets, and consistent with psychological research.

Key words: overconfidence; judgment under uncertainty; efficient market hypothesis; organizational studies; decision making

History: Received August 9, 2011; accepted August 31, 2012, by Uri Gneezy, behavioral economics. Published online in *Articles in Advance* March 18, 2013.

1. Introduction

Two of the building blocks of modern neoclassical 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 Gilovich et al. 2002, Kahneman and Tversky 2000), there are far fewer studies of predictions by market participants with substantial amounts of money at stake (for a recent review, see DellaVigna 2009). Similarly, tests of financial market efficiency are often plagued by the inability to measure fundamental value.

In this paper we investigate how rational expectations and market efficiency play out in an unusual but interesting labor market: the National Football League (NFL), specifically, its annual draft of young players. Every year the 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 12th pick and the 31st pick in return. In aggregate, such trades reveal the *market value* of draft picks. Although it is not immediately obvious what the rate of exchange should be for such picks, a consensus has emerged over time that is highly regular. One reason for this regularity is that a price list—known in the league circles as “the Chart”—has emerged and teams now routinely refer to the Chart when bargaining for picks. What our analysis shows is that although this chart is widely used, it has the “wrong” prices. That is, the prices on the chart do not correspond to the correct relative value of the players. We are able to say this because player performance is observable.

To determine whether the market values of picks are “correct,” we compare them 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 12 and 31. Thus,

our null hypothesis is that the ratio of pick values will be equal to the ratio of surplus values.

The alternative hypothesis we investigate is that a combination of well-documented behavioral phenomena, all working in the same direction, creates a systematic bias causing teams to overvalue the highest picks in the draft. For example, this is the result we would expect if teams overestimate their ability to determine the quality of young players. Market forces will not necessarily eliminate this mispricing because even if there are a few smart teams, they cannot correct the mispricing 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, and so cannot offer to trade them away.

Our findings strongly reject the hypothesis of market efficiency. Although the market prices of picks decline sharply initially (the Chart prices the first pick at three times the 16th pick), we find surplus value of the picks during the first round actually *increases* throughout most of the round: the player selected with the final pick in the first round, on average, produces more surplus to his team than the first pick! The market seems to have converged on an inefficient equilibrium. As we discuss in §4, both the Chart and a robust rule of thumb regarding the trading of a pick this year for a pick next year have emerged as norms in the league, norms that appear to be difficult to dislodge even though the values these norms imply are demonstrably wrong.

The setting for this study is unusual, but we suggest that the implications are quite general. It is known in financial economics that limits to arbitrage can allow prices to diverge from intrinsic value, but some version of market efficiency remains the working hypotheses (sometimes implicitly) even in markets where there are no arbitrage opportunities. Are competition and high stakes enough to produce efficiency? We show that they are not. We study a domain in which it is arguably easier to predict performance than, say, the market for most employees, even CEOs. Teams have been able to watch prospects for several years play the same game they will play in the pros, and also have administered days of physical and mental tests. Still, we find their ability to predict performance incommensurate with their confidence. Similar judgments about the future undergird many important decisions. Whether deciding to hire a CEO, invest in a new technology, or to use military force, it is critical that one’s confidence level is appropriate.

For the initial step in our analysis, we use a data set of 407 draft-day trades to estimate the market value of draft picks. We then ask whether the highest picks are too expensive, as the relevant psychology predicts. To do so, we first take a nonparametric approach,

comparing the benefit of using a pick to the opportunity cost of foregone trades. We then perform a more detailed cost-benefit analysis of each player selected in the draft, by calculating the surplus value that the player provides to the team, namely, the performance value (estimated by the price of an equivalent veteran player) minus the salary paid. These analyses allows us to test for and reject market efficiency.

2. Background Information

Although it is not necessary to know the difference between an outside linebacker and a cheerleader to follow the analysis in this paper, it is important to have some background regarding the nature of this labor market. There are three essential features. First, new players to the league are allocated to teams via an annual draft. Teams take turns selecting players in an order determined by the previous year’s record. There are seven rounds of the draft, and in each round the worst team chooses first and the champion chooses last (with some minor exceptions). The players selected are then signed to a contract, historically four to six years. Players can only sign with the team that selected them.

Second, the league has adopted a rule setting a maximum amount any team can pay its players in a given year. This is called the salary cap. The cap has increased over time, from \$34.6 million in 1994 to \$128 million in 2009. When players are signed to multiple-year contracts, there is usually a guaranteed up-front bonus payment plus annual salaries. The accounting for the salary cap rule allows the teams to allocate the bonus equally across the years of the contract. Whenever we report player compensation in this paper we are using the official cap charge as reported to the league.¹ The existence of this salary cap makes it easier to draw robust conclusions about market efficiency because all owners face the same upper limit on what they can spend, unlike in professional baseball or European soccer. In those sports, rich owners can buy the rights to star players to suit their own preferences, and it would be impossible to say they are paying “too much” without knowing their utility function. In the NFL, hiring a star to a big salary limits what can be offered to other players, so owners are forced to choose which players they wish to spend their budget on.

Third, there is also a special “rookie salary cap” that limits the amount of money a team can spend on first-year players, both drafted and undrafted (most teams typically sign several of these each year). This rookie salary cap is a “cap within a cap,” meaning

¹For an excellent summary of salary cap rules, see Hall and Lim (2002).

that the money spent on rookies counts toward the overall cap, but is an extra constraint. A key feature of the rookie salary cap is that, unlike the overall cap, it varies by team. Specifically, the team's rookie salary cap depends on the portfolio of picks the team has (subsequent to all trades), and teams with high first-round picks are given larger amounts to spend on rookie salaries. As we shall show, these rookie salary cap allocations largely determine the compensation of draft picks.²

A few other features of the league are worth noting. The teams earn most of their revenue from television contracts, and these revenues are divided equally. Teams also share all revenues from sales of team paraphernalia such as hats or jerseys. Finally, during the period we study, the salary cap is a binding constraint, or nearly so, for most teams.

3. Research Hypothesis

Draft picks are valuable. Because of the rookie salary cap, and teams' exclusive rights to the players they draft, teams spend less on drafted players than they would for veteran players of the same expected quality. (We document this fact below.) The greater this "surplus," the more a team can spend to sign high-quality players in the free-agent market. This suggests that if teams are profit maximizing in their decisions to trade (or not trade) draft picks, the relative value of any two picks will equal the relative expected surplus of the picks.³ Specifically, for the i th and $i+k$ picks in the draft,

$$\frac{M_i}{M_{i+k}} = \frac{E(S_i)}{E(S_{i+k})}, \quad (1)$$

where M_i is the market value of the i th draft pick relative to other picks, and $E(S_i)$ is the expected surplus value (i.e., salary-cap savings) of players drafted with the i th pick. We assume that the player's value can be observed on the field and estimated from the labor market, although we stress test those assumptions in the penultimate section of the paper.

In contrast to this null hypothesis, we predict that teams will overvalue the right to choose early

in the draft. Specifically, we believe teams will systematically pay too much for the right 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 that

$$\frac{M_i}{M_{i+k}} > \frac{E(S_i)}{E(S_{i+k})}; \quad (2)$$

i.e., that the market value of draft picks will decline more steeply than the surplus value of players drafted with those picks.⁴

The bases for our prediction that top picks will be overvalued are rooted in numerous findings in the psychology of decision making. The NFL draft involves predicting the future, a task that has received considerable attention from psychological researchers. This research suggests that behavior can deviate systematically from rational models. In this particular domain, all the psychological biases point in the direction of our central prediction, so we will not belabor our discussion of these various findings. Briefly, the following robust empirical findings support our prediction.

Nonregressive Predictions. One of the earliest findings in this literature is that intuitive predictions are insufficiently regressive (Kahneman and Tversky 1973). That is, intuitive predictions are more extreme and more varied than is justified by the evidence on which they are based. Normatively, one should combine evidence (e.g., player's running speed) with the prior probabilities of future states. Teams should find these prior odds quite sobering. For example, over their first five years, *first-round* draft picks (that is, the top 32 picks) have more seasons with zero starts (15.3%) than with selections to the Pro Bowl⁵ (12.8%). 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. However, to be regressive is to admit to a limited ability to differentiate the good from the great; this perceived ability to differentiate is the very thing that

² The NFL signed a new labor agreement in 2011. The rules we describe are those that were in place during the period we studied. In brief, the pay to the top draft picks has been cut, a change consistent with viewing our results as both accurate and problematic for the league. However, the changes were small relative to the effects we document, and primarily focused on the top three to five picks. The broad patterns we find under the previous labor agreement still exist under the new one.

³ Alternative assumptions, that firms try to maximize profits, winning percentage, or chance of winning the Super Bowl, are conceptually quite similar. Teams do make more money if they win, and the salary cap means that they have to win without spending an unlimited amount on players.

⁴ Note that this expression, by itself, does not imply which side of the equation is "wrong." Whereas 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 (2004) makes, in which he assumes the draft-pick market is rational and points out its discrepancy with subsequent player compensation. The key difference in our approaches is that we appeal to a third, objective measure—player performance—to determine which of the two sides, or markets, is wrong.

⁵ The Pro Bowl is held at the end of each year with the best players selected to play. We use the selection to play in this game as one measure of outstanding performance.

has secured NFL scouts and general managers their jobs. Hence, we suspect NFL decision makers put more weight on scouting evidence than is justified.⁶

Overconfidence. Another robust finding in psychology, similar in spirit to the aforementioned tendency to make excessively extreme forecasts, is that people are overconfident in their judgments (Alpert and Raiffa 1982). Furthermore, overconfidence is exacerbated by information—the more information experts have, the more overconfident they become.⁷ NFL teams face a related challenge—making judgments about players while accumulating increasing amounts of information about them as the draft approaches.

Other Psychological Factors. There are additional factors that could reinforce the tendency for teams to overvalue top picks. The winner's curse suggests that teams will fail to adjust for the fact that the winner among many bidders for an object of uncertain but common value is likely to overpay (for a review, see Thaler 1988).⁸ False consensus suggests that teams will overestimate the need to trade up to acquire a player they value because they will believe, unduly, that other teams value him similarly (Ross et al. 1977). Also, anticipated regret can lead teams to exercise rights to high-profile players because to miss out on a superstar would be particularly painful (Lerner and Tetlock 1999). Together these biases all push teams toward overvaluing picking early.

Of course there are strong incentives for teams to overcome these biases, and the draft has been going

⁶ In unreported analyses we find that team scouts predict exceptional performance by college players in the NFL more frequently than is warranted, and that among these players predicted to be superstars there is *no* relation between ratings and performance.

⁷ Research subjects have included clinical psychologists (Oskamp 1965) and horserace bettors (Russo and Schoemaker 2002, Slovic and Corrigan 1973).

⁸ Although values are not perfectly common—there is certainly some true heterogeneity in the value teams place on players—there are multiple reasons this characterization fits. First, we simply assert that there is far more variation in value across players (i.e., uncertainty) than there is variation in value within players across team (i.e., heterogeneity). Second, any true heterogeneity is muted by the relatively liquid trading market in players. Finally, and most important, almost all the heterogeneity is determined by player *position* due to team needs. This is fine, because the winners curse should apply within position as well. Harrison and March (1984) suggest that a related phenomenon, “expectation inflation,” 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. Harrison and Bazerman (1995) point out that nonregressive predictions, the winner's curse, and expectation inflation have a common underlying cause—the role of uncertainty and individuals' failure to account for it. Harrison and Bazerman (1995) emphasize that these problems are exacerbated when uncertainty increases and when the number of alternatives increases—precisely the conditions of the NFL draft.

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 they do not hire as well as those they do. This could facilitate learning. This same feature, that performance is observable, is what makes this research project possible.

We expect the deviation between market prices and surplus value to be most acute at the top of the draft, because the psychological mechanisms we have highlighted above are most acute there. Regression to the mean is strongest for more extreme samples, so we expect the failure to regress predictions to be strongest there as well.⁹ Players at the top of the draft also receive a disproportionate amount of the teams' attention and analysis, so information-facilitated overconfidence should be most extreme there.¹⁰

A finding supporting our hypothesis implies teams should trade down when endowed with a very high pick. There are of course limits to this strategy, because the roster size constrains the number of new players a team can acquire via trading down. These are not tight constraints though, because teams routinely invite 5 to 10 undrafted free agents to summer training camp. Sufficient for a test of our hypothesis is the possibility of converting (at market rates) one first-round draft pick into just two lower picks. Indeed, this is the modal type of trade we observe, and also matches well our theoretical focus on very high picks.

More generally, we are investigating whether well-established judgment and decision-making biases are robust to market forces. Gary Becker asserts that, “Division of labor strongly attenuates if not eliminates any effects caused by bounded rationality. . . . [It] doesn't matter if 90% of people can't do the complex analysis required to calculate probabilities. The 10% of people who can will end up in the jobs where it's required” (Stewart 2005, p. 41). Romer's (2006) insightful analysis of the decision about whether to punt or “go for it” on 4th down suggests that NFL coaches are not members of Becker's elite 10% (see also Carter and Machol 1978). Here we see whether market forces can help NFL owners and general managers to do better.

4. 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

⁹ Similarly, De Bondt and Thaler (1985) found the strongest mean reversion in stock prices for the most extreme performers over the past three to five years.

¹⁰ The tendency to overweight small probabilities, well documented in the psychological literature, also suggests that overvaluation will be worse at the top of the draft. For example, consider how suspiciously often we hear a college prospect described as a “once-in-a-lifetime player.”

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 10th, 16th, or 32nd. We infer these values from draft-day trades observed over 26 years.

4.1. Data

The NFL draft consists of multiple rounds, with each team owning the right to one pick per round.¹¹ We designate each pick by its overall order in the draft. 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). The data we use are trades of these draft picks from 1983–2008.¹² Over this period, we observe 1,078 draft-pick trades. Of these, we exclude 663 (61%) that involve NFL players in addition to draft picks and 7 (< 1%) with inconsistencies implying a reporting error. We separate the remaining trades into two groups: 314 (29%) involving draft picks from the current year only and 94 (9%) involving draft picks from both the current and future years.¹³ Although we observe trades in every round of the draft, the majority of the trades ($n = 171$, 54%) involve a pick in one of the first two rounds, precisely the domain in which we are predicting the strongest deviations from market efficiency.¹⁴ We observe every team trade both up and down at least once.

4.2. Results and Discussion

We estimate the market value draft picks using a two-parameter Weibull distribution (see the appendix for details). Figure 1 plots this function, showing the value of the first 160 draft picks (the first five rounds) relative to the first draft pick. It does this by comparing the estimated values for “both sides” of a trade: the value of the top pick acquired by the team moving up, and the value paid for that pick by the team moving up, net of the value of additional picks acquired. The model fits the data exceedingly well, in part because of the reliance on the Chart, discussed in detail below.

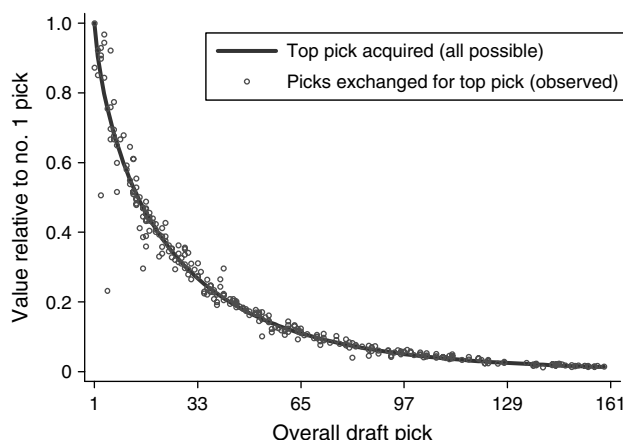
¹¹ The order that teams choose depends on the team’s won–lost record in the previous season—the worst team chooses first, and the winner of the Super Bowl chooses last.

¹² This data set was compiled and cross-checked from a variety of publicly available sources, including newspapers and ESPN.com.

¹³ See the supplemental analysis for a summary table of draft-pick trades. The supplement is available online at http://consafo.com/losers_curse_supplemental_analysis.pdf.

¹⁴ The average number of picks acquired by the team trading down was 2.3 (SD = 0.61), with a maximum of 6. The average number of picks acquired by the team trading up is 1.1 (SD = 0.38), with a maximum of 3. The modal trade was 2-for-1, occurring 212 times (68%).

Figure 1 Estimated Trade Value of Draft Picks



Notes. 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- and right-hand sides of expression (7), 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, whereas those above the line represent trades that cost more (from the perspective of the party trading up) than expected.

A striking feature of these data is how steep the curve is. The drop in value from the 1st pick to the 10th is roughly 50%, and values fall another 50% from there to the end of the first round. As we report in the following section, compensation costs follow a very similar pattern. Although the curve is not as steep as it used to be, this flattening has slowed over time. In an efficient market the curve’s steepness would imply both that player performance falls sharply at the top of the draft, and that teams are highly skilled in their ability to identify these performance differences.

Another notable feature is the remarkably high discount rate, which we estimate to be 136% per year. Although this finding is not the focus of the paper, it is clear that teams who “borrow” picks on these terms are displaying highly impatient behavior. Although it is not possible to say whether this behavior reflects the preferences of the team owners, of their employees who typically make the decisions (general manager, head coach, etc.), or both, it provides a significant opportunity for teams with a longer-term perspective.

Norms. As noted above, one reason why our trading-price estimates fit the data so well is that teams have come to rely on the Chart to help them negotiate the terms of trade. The Chart was originally estimated in 1991 by Mike McCoy, then a part owner of the Dallas Cowboys (McCoy 2006). An engineer, McCoy estimated the values from a subset of the trades that occurred from 1987 to 1990. His goal was merely to characterize past trading behavior rather

than to determine what the picks should be worth. The Chart then made its way through the league as personnel moved from the Cowboys to other teams, taking the Chart with them. In 2003 ESPN.com posted a graphical version of the Chart, reporting that it was representative of curves that teams use (ESPN 2004). McCoy's original curve, as well as the ESPN curve, closely approximates the one we estimate for the 1983–2008 period.

Teams were beginning to agree about the market value of picks by the time McCoy estimated his chart. As the Chart spread around the league, it became standard for teams to openly use it to negotiate the terms of trades.¹⁵ A norm emerged for trades involving future picks as well: "gain a round by waiting a year." For example, this year's third-round pick would bring a pick in next year's second round. McCoy mentioned this heuristic explicitly when discussing his construction of the Chart, and it is clear in the data. This rule of thumb leads to huge discount rates because they must equate the value of picks in two adjacent rounds. It is also surprisingly arbitrary. Consider that it depends on the number of teams in the league (which in fact has changed over time).

Predictably, the emergence of widely accepted prices made trading easier; between 1983 and 2008, the deviation in prices from the Chart dropped by 50% (and the year-to-year volatility of that deviation shrunk considerably). Also, trading activity tripled to over 20 trades per year.¹⁶ Thus, the emergence of consensus—a norm—seems to have lent the considerable power of precedent and conventional wisdom to the overvaluation we suggest has psychological roots.

The very steep curve we document in this section suggests that teams believe they have the ability to distinguish great players from the merely good. Before moving to full cost-benefit analyses, let us consider a simple question: What is the likelihood that a player is better than the next player chosen at his position (e.g., linebacker) by some reasonable measure of performance, such as games started in his first five seasons? After all, this is the question teams face as they decide whether to trade up to acquire a specific player.

The answer is 52%. Across all rounds, all positions, all years, the chance that a player proves to be better than the next-best alternative is only slightly better

than a coin flip.¹⁷ This (overly) simple observation suggests a discrepancy between the teams' perceived and actual ability to discriminate between prospective players. We explore this potential discrepancy in the sections that follow.

5. Opportunity Costs

In our first analysis, we evaluate the benefit of using a draft pick relative to the opportunity cost of trading it for two lesser picks. We focus only on directly observable performance—starts and pro bowls—taking a nonparametric approach free from any monetary calculations. We can avoid the step of estimating the relative value of these two performance measures because the results show that trading down is a dominant strategy; the players acquired by trading down make significantly more starts and just as many pro bowls.¹⁸

To conduct this analysis, we evaluate all possible 2-for-1 trades. We focus exclusively on first-round picks, i.e., each first-round pick and the 2-for-1 trades down that are possible from that position. The possibility of a trade depends on the Chart—we consider all two-pick combinations whose total Chart value is 90%–100% of the value of the first-round pick.¹⁹ For example, a team could trade the top pick in the draft for the 2nd and 181st, for the 14th and 15th, or a number of combinations in between. Every draft-pick position has, on average, 19 of these two-player combinations. We use the 1991–2004 drafts, stopping in 2004 so that we have five years of performance data. We estimate means and standard errors separately for each draft-pick position, clustering on draft year. For each possible trade, we consider the number of starts and pro bowls generated by the players involved over their first five seasons.

We analyze 8,526 potential trades over the 14-year period and find overwhelming evidence that a team would do better in the draft by trading down. The average gain from trading down is 5.4 starts per

¹⁷ The median number of picks between players at the same position is 7.

¹⁸ Of course even this analysis is not assumption free. There could be *some* measure of performance that is sacrificed by this strategy, at least partially offsetting the gains in starts. This is unlikely given our unsuccessful search for elite performance measures, documented in §7.1.1. Moreover, even if present, this kind of performance is by definition exceedingly rare and must be weighted accordingly. It would also require that such performance is tightly related to draft order, a common intuition with scant empirical support.

¹⁹ Hence, the trades we consider always weakly favor the team trading up according to the chart. Compensation costs are very similar, although slightly lower for a single pick. We also estimate compensation-constant portfolios (rather than these chart-value constant) and find the same pattern of results. See the supplemental analysis for details.

¹⁵ In a conversation with the authors, McCoy (2006) stated, "It gave us more confidence. If you just had a sticker—bread is 49 cents—everything would be easier." It also provided cover. "A standard price list also protects you." McCoy added, "Now nobody gets skinned."

¹⁶ By 2008 the average absolute deviation from the Chart was equivalent in value to a mid-4th-round pick, 1/50th the value of the top pick in draft. See the supplemental analysis for a figure showing these trends.

season. We estimate this gain to be reliably positive for 31 of the 32 draft-pick positions in the first round. Indeed, the mean gain is greater than three starts/trade for most of the round (25 of the 32 positions). Importantly, these gains are generated without cost in terms of pro bowls; the net change in pro bowls for most (20 of 32) draft-pick positions is not different than zero, and there are more that are positive (9) than negative (3).

Of course not every possible trade will work out well; sometimes the team with the high pick will trade away a star for two duds, but this strategy has a very high hit rate. For 74% of the trades, a team would have acquired more starts by trading down than by using a pick. Also, it is not the case that these gains come at the expense of giving up the chance at a big hit. In fact, in terms of starts and pro bowls, trading down is a stochastically dominant strategy; 61% of the time the team trading down would have done better in terms of starts without doing worse in terms of pro bowls. This is 2.5 times the risk of trading down and doing worse on pro bowls without doing any better on starts (24%).

We also conducted an alternative version of this analysis by comparing the maximum (as opposed to the sum) of the two players acquired by trading down to the player taken with the original (higher) pick. This is an extremely conservative test of the value of trading down because, (1) the player costs almost 50% less than the original player and (2) it neglects the possibly high value of the second player. However, even setting aside those benefits, trading down is beneficial. Teams would have gained an average of 0.83 starts per season by trading down and keeping only the best of the two acquired players, with a 0.0 change in the number of pro bowls. In fact, this strategy is almost stochastically dominant; 48% of the time the best acquired player is better than the original player on either starts or pro bowls without being worse on the other (versus 40% of the time being worse on one without being better on the other). This means that even a team simply trying to fill a single spot on their roster would be better taking two draws later in the draft than one draw in the first round. This strategy is even more appealing when considering the reduced compensation cost, as well as the option value from the second player. It is difficult to overstate the strength of these results.

6. Cost-Benefit Analysis

The opportunity-cost analysis above indicates that teams make a mistake by holding onto a single first-round draft pick rather than trading it for two lower picks. However, the analysis is silent on the magnitude of the mistake, as well as how it varies at different points in the draft. These finer questions require

a more complicated analysis. For this we turn to a two-stage cost-benefit analysis. In the first stage, we establish the value that teams place on performance by looking at the compensation of veteran players. In the second stage, we apply these values to all drafted players. We estimate the “surplus value” of these players to their teams by subtracting their compensation from these performance values. Our interest is the relation between surplus value and draft order.

6.1. Data

Because we want to include players in every position in our analyses, we rely on three performance statistics that we can use for all positions: whether the player is on a roster (i.e., in the NFL), the number of games he starts, and whether he makes the Pro Bowl (a season-ending “All-Star” game). We have these data for the 1991–2008 seasons.²⁰ (Later we replicate these analyses with more detailed performance data, using only wide receivers.) Using these statistics, we create five comprehensive and mutually exclusive performance categories for each player-season: players elected to the Pro Bowl (“Pro Bowl”), those who start at least 14 of the 16 regular season games (“regular starter”), those who start fewer than 14 games (“occasional starter”), those who do not start any games (“backup”), and those not in the league (“NIL”).²¹ For player i in his t th year in the league, this gives the measure $Cat_{n,i,t} = \{0, 1\}$, indicating qualification for performance category n according to the criteria described above.

We rely on a sample of experienced players to estimate the value that teams place on these performance categories. These are veteran players who have signed at least one free-agent contract. We limit this sample to players drafted in 1991–2001 who are in their sixth, seventh or eighth year in the NFL, and restrict our analysis to the 1996–2008 seasons so we can observe five years of performance for each player in previous years. As shown in Table 1, panel A, this leaves 3,014 players-seasons. These veteran players averaged 16% of their previous five seasons as a backup, 41% as an occasional starter, 33% as a regular starter, and 10% elected to the Pro Bowl. They are paid an average of \$3.4 million per year (median = \$2.7 million, SD = \$2.7 million).

²⁰ Performance data are from Stats. Inc. The 1991 season is the earliest for which the “games started” measure is reliable.

²¹ Most of these category boundaries are obvious. The exception is dividing the two “starter” categories at 14 games. We do this to avoid excluding a player from the top starter category because of very small perturbations due to injury, chance, coaching, etc. Estimation results are robust to moving this cutoff higher or lower. Players elected to the Pro Bowl are assigned to that category regardless of how many games they started, with the exception of special-teams players.

Table 1 Player Samples

Years in league	<i>N</i>	Historical performance (%)					Compensation (\$ mm)
		NIL	Backup	Occasional starter	Regular starter	Pro Bowl	
Panel A: Experienced-player sample							
6	1,169	0	18	40	31	10	3.052
7	993	0	14	42	34	10	3.439
8	852	0	16	41	33	11	3.883
Total	3,014	0	16	41	33	10	3.414
Panel B: Drafted-player sample							
1	3,483	21	38	34	7	1	0.747
2	3,456	23	26	33	15	3	0.881
3	3,438	30	19	29	18	3	0.952
4	3,430	38	15	26	17	4	1.186
5	3,348	45	11	22	17	4	1.472
Total	17,155	31	22	29	15	3	1.044

Notes. Panel A: Includes players drafted 1991–2001 who are in the league during their 6th–8th years. Means for each category over the previous five seasons. Observations are player-seasons. Compensation is the player's salary-cap charge, in 2008 dollars. Panel B: Includes drafted players in their first five seasons, 1994–2008. Limited to players in the first seven rounds of the draft (this is binding only 1991–1993). Observations are player-seasons. Compensation is the player's salary-cap charge, in 2008 dollars.

Ultimately, we are interested in the value of the player to the drafting team. To assess this, we turn to a second sample consisting of players in their first five years after being drafted. We restrict our analysis to the salary cap era, 1994–2008. We also limit our analysis to the first seven rounds of the draft because the draft has included only seven rounds since 1994. As shown in Table 1, panel B, this yields 17,155 player-seasons. Thirty-one percent of the player-seasons are NIL, 22% are backup, 29% are occasional starter, 15% are regular starter, and 3% are Pro Bowl. Note that we avoid survivorship bias by retaining players in our analysis who are not in the league. Players in this sample are paid an average of \$1.04 million per year (median = \$0.606 million, SD = \$1.438 million).

As one would expect, the correlation between compensation and player performance is much higher in years 6–8 (0.73) than in the players' first five years (0.55). Performance becomes easier to predict after a player has played several years in the league, and the market (rather than draft order) is determining compensation. This is the primary motivation for basing the compensation model on the sample of experienced players.

6.2. Analysis and Results

6.2.1. Performance Value. We are interested in the market value of different levels of player performance: backup, Pro Bowl, etc. To do this, we investigate the relation between a player's compensation (salary-cap value) in years 6–8 and his performance during the previous five seasons. Recent years are likely to carry more weight because they are more closely related to future performance. To allow this possibility we use a weighted average

of the player's performance history, estimating the best-fitting "memory" parameter for these weights. Specifically, for player i in year t , we estimate that

$$\log(\text{Comp}_{i,t}) = \alpha + \text{BCat}_{\tilde{n}_{i,t}} + \text{III}_i^P + \text{KI}_{i,t}^T + \varepsilon_{i,t}, \quad (3)$$

where $\text{BCat}_{\tilde{n}_{i,t}}$ is a weighted average of the player's n th performance category over the previous five years, I_i^P is a vector of indicator variables for the player's position (quarterback, running back, etc.), and I_i^T is a vector of indicator variables for the player's year in the league (6th–8th). Weights are given by $w_t = \exp(-\eta(r-1))$ for player performance r years in the past. This model lets "memory" in compensation decay at an exponential rate. The amount of decay is determined by η , which we estimate. The special case of full memory, in which all five years are equally weighted, is given when $\eta = 0$. By construction the weight is 1 for the most recent year.²²

The model's predicted values provide the estimated market value for each position-performance pair.²³ This general approach is similar to that of previous research on NFL compensation (Ahlburg and

²² We allow the memory parameter, η , to vary by player year. This is because we expect the distant past to carry less weight for a history covering years 1–5 (at the beginning of which a player has just entered the league and sometimes does not even play) than for a history covering years 3–7. Let $\eta = \eta_1 \eta_2^{(t-6)t}$ for a player's t th year in the league. We estimate η_1 , which provides the exponential memory parameter, and η_2 , which modifies that parameter by player year. See the supplemental analysis for a depiction of these functions.

²³ Of course this is an approximation, because there is variation in true value within a position-performance pair. For our purposes, these approximations will be adequate as long as they are unbiased relative to draft order. We relax this assumption to test its implications (see Footnote 26).

Table 2 Compensation Models

Coefficient	(1)	(2)
κ_1 Year 7	0.1251 (0.019)	-0.0607 (0.031)
κ_2 Year 8	0.2449 (0.024)	-0.0943 (0.055)
β_2 Starts = 0		0.0725 (0.031)
β_3 Starts \leq 14		0.342 (0.029)
β_4 Starts > 14		0.683 (0.033)
β_5 Pro Bowl		0.886 (0.039)
η_1 Memory		0.360 (0.031)
η_2 Memory (yearly modifier)		0.896 (0.040)
α Constant	14.22 (0.070)	13.22 (0.081)
Observations	3,014	3,014
Adjusted <i>R</i> -squared	0.07	0.59

Notes. Nonlinear regression results for compensation in years 6–8. Compensation is the salary-cap charge. Sample is all players who were drafted 1991–2001 and on an NFL roster 6–8 years later (excluding kickers and punters). Position fixed-effects are included but suppressed for presentation. The omitted player year is 6th. In model (2), the omitted performance category is NIL. Standard errors are shown in parentheses.

Dworkin 1991, Kahn 1992, Leeds and Kowalewski 2001), though, aside from our analysis of wide receivers below, we rely on performance *categories* rather than performance *statistics*. Our modeling compensation as a function of past performance is consistent with these earlier approaches, as well as with industry practice, in football and other professional sports.²⁴ Indeed, player negotiations are often considered “boring” because they are largely a matter of finding comparable players based on historical performance (Burke 2012).

We present the results from this estimation in Table 2. Model (1) provides a baseline, including only indicator variables for player position and player year. Results indicate that compensation increases from year 6 to 7 and from year 7 to 8. This model explains only 6.7% of the variance in player compensation. In model (2) we add player performance categories, including the memory parameters. Most important, we find that values increase monotonically with performance category, with each category statisti-

cally distinct.²⁵ The timing of the performance matters, as well—estimates for the memory parameter indicate that a player’s performance two years before has only 65% as much influence on salary as the most recent year. Comparable values for three, four, and five years past are 42%, 28%, and 18%, respectively. Hence, there is considerable “decay” in memory, providing a more predictive model for future performance. Finally, having controlled for performance and hence survivorship, we no longer find compensation rising with player experience. This model explains a considerable portion of the variance in player compensation, with an adjusted *R*-squared of 0.59.²⁶

In Figure 2 we present the predicted values of this final model for each position and performance category, transformed into dollars. As we saw in the model estimates, values increase with performance. The mean values increase from \$917,000 for the player-seasons without any starts, to \$1.9 million, \$4.8 million and \$8.2 million for occasional starters, full-time starters, and Pro Bowl players, respectively. One striking feature of the results is the variation in compensation for various positions. Most notable is the incremental value of quarterbacks, who are paid more than 50% above the next-highest paid position, defensive end.

6.2.2. Compensation Cost. 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. Second, as we discussed above, the NFL teams operate under rules restricting how much they are allowed to pay their players—the salary cap.

The compensation data we use are from a variety of public sources and have been checked for accuracy by an NFL team.²⁷ Our sample includes the first 15 years of the free-agency era, 1994–2008. We focus on a player’s salary-cap charge each year, which includes

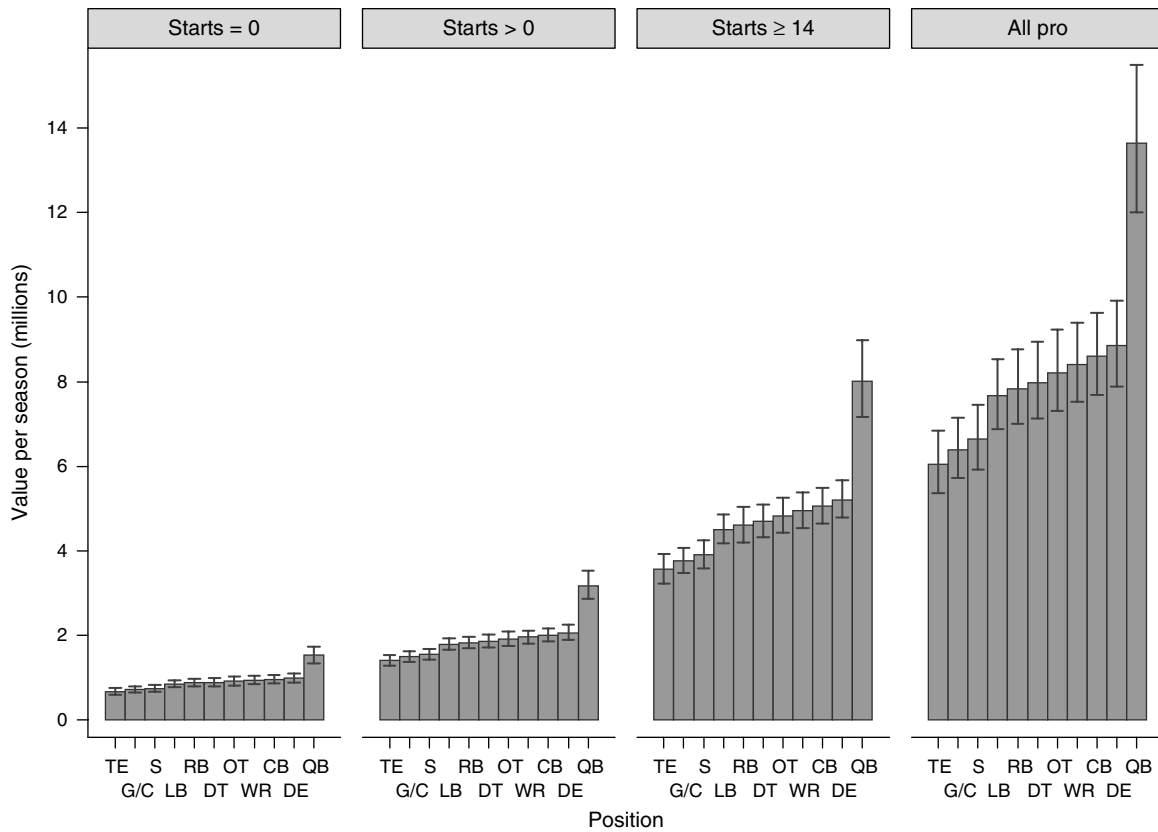
²⁵ Estimates are in log terms and therefore difficult to interpret directly; we transform their values in Figure 3 to see the results in real terms.

²⁶ In unreported analyses we consider two further elaborations of this model. First, we interacted player position with performance categories. This increased the *R*-squared 0.01 (to 0.60) and did not change the results of our subsequent tests. Second, we included the player’s original draft-pick value, using an exponential distribution estimated from the data. This parameter was significantly positive, indicating that model (2) does not completely capture the relation between draft pick and player value as a free agent. However, the effect was quite small, and incorporating it into our subsequent tests did not change any of our results.

²⁷ Player contracts have to be submitted in full to the league, and the details are made available to all the teams and registered player agents. In other words, compensation is common knowledge within the league.

²⁴ We also experimented with models that used predicted performance in years 6–8 as the explanatory variable rather than past performance. Models of this type make more sense theoretically, because teams should be interested in what the players will do in the future rather than how they performed in the past, but explain much less of the variance in salary. Apparently, teams pay based on the past rather than a rational forecast of the future.

Figure 2 Performance Valuation



Notes. The labor market value of a player’s previous five years. These are the predicted values from model (2) in Table 2, estimated from compensation in player years 6–8 for draft classes 1991–2001. Note that these are values for player performance that falls into a category 100% of the five-year history. Reported in 2008 dollars.

his salary and a prorated portion of his bonus.²⁸ There are also minimum salaries, which vary by year and with player experience. In our sample only 12% of players are paid the league minimum.

The data reveal a very steep relation between compensation and draft order at the top of the draft.²⁹ This general pattern holds 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.³⁰ The slope of this curve approximates the draft-pick value curve estimated in the previous section. Thus, players taken early in the draft are thus expensive on both counts: foregone picks and salary paid.

6.2.3. Surplus Value. The third and final 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 the 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 ,

$$\hat{S}_{i,t} = \hat{P}_{i,t} - C_{i,t}, \tag{4}$$

where $\hat{P}_{i,t}$ is the performance value estimated from the compensation model above for his position and

²⁸ Our compensation data include only players who appear on a roster in a given season, meaning that 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, although we strongly suspect it is negatively related to draft order—i.e., that 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 are 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.

²⁹ There is also a distinct discontinuity after pick 32, the last pick in the first round. Compensation shifts down sharply at this point, creating a first-round premium, although of course there is no such discontinuity in performance. See the supplemental analysis for a figure.

³⁰ After four years, players are eligible for restricted free agency. After five years, players are unrestricted free agents and can negotiate with any team. This time frame can be superseded by an initial contract that extends into the free-agency period, e.g., six years and longer. Such contracts were exceedingly rare in the period we observe, although they are becoming more common.

actual performance, and $C_{i,t}$ is the player's actual compensation costs. Our interest is in the relationship between surplus value and draft order.

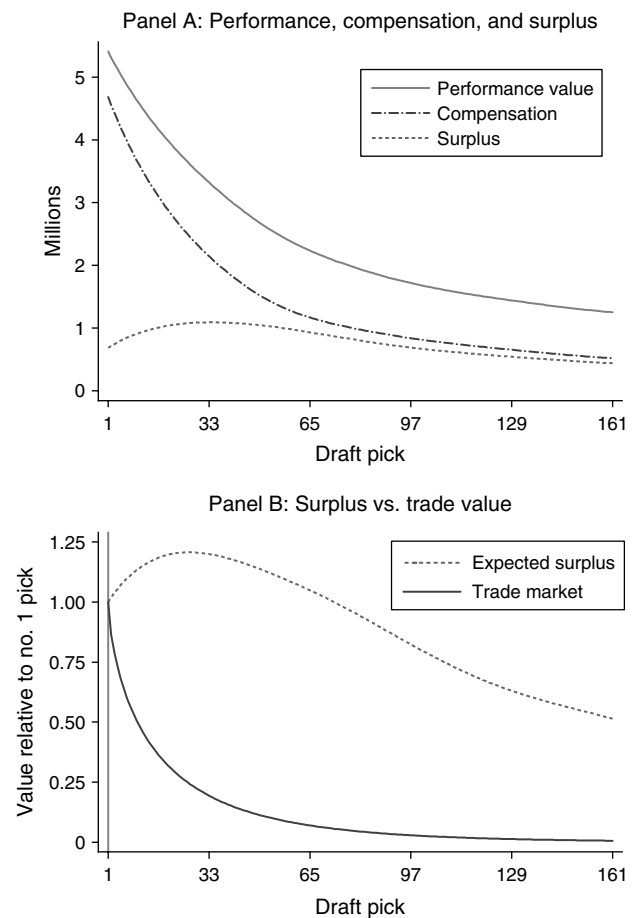
Across all rounds, the mean salary-cap charge is \$1,044,029, whereas the mean estimated performance value is \$1,703,390, resulting in a mean surplus value of \$659,361. For an initial look at the relation between these values and draft order, we consider how the values for first-round players compare to those for second-round players. Surprisingly, we find that the mean surplus value is higher in the second round (\$1,171,834) than in the first round (\$1,016,797). Indeed, the median surplus value is more than 60% higher in the second round (\$762,785) than in the first round (\$462,634). Recall that the draft-pick market values the first pick approximately four times higher than the first pick in the second round.

In Figure 3, panel A, we graph all three variables as a function of draft order, fitting lowess curves to the underlying player-seasons. It is noteworthy that performance value is everywhere higher than compensation costs, and so surplus is always positive. This implies that the rookie salary cap keeps initial contracts artificially low relative to the more experienced players who form the basis of our compensation analysis. More central to the thrust of this paper is the fact that although both performance and compensation decline with draft order, compensation declines more steeply. Consequently, surplus value *increases* at the top of the order, rising to its maximum of approximately \$1,200,000 near the beginning of the second round before declining through the rest of the draft. 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 value curve is monotonically decreasing), but he also has the highest salary, and in terms of performance per dollar, is less valuable than most players taken in the second round.³¹

Clearly we should be cautious in interpreting this surplus curve; it is meant to summarize the results simply. Whereas the general shape is robust to a wide range of modeling decisions, the precise values are not. More important for our hypothesis is a formal test of the relation between the estimated surplus value and draft order. Specifically, we need to know whether

³¹ We also find that the standard deviation of surplus value is strongly *negatively* related to draft order. That is, not only do the top picks have lower mean surplus value than those in the second round, they also have the highest variance in the draft. Of course teams might value variance if it means there is a fat right tail offering the chance of a superstar, but the opportunity-cost analysis in §5 shows this is not the case. See also our discussion of superstars in §7.1.1.

Figure 3 Surplus Value by Draft Order



Notes. Panel A: Summary of lowess curves for player performance value, compensation, and surplus (performance value less compensation) in the player's first five years. Underlying observations are player-seasons, 1994–2008; $n = 16,502$. Reported in 2008 dollars. Panel B: “Expected surplus” is the lowess curve for the relationship between estimated surplus value and draft order (panel A). Observations are player-seasons. The sample is for the 1994–2008 seasons, including all drafted players in their first five years in the NFL, excluding punters and kickers. “Trade market” is the Weibull estimated from draft-day trades (Figure 1).

this relation is less negative than the one between *market value* and draft order. Certainly it appears to be less negative, as shown in Figure 3, panel B. Whereas the market value of draft picks drops immediately and precipitously, the surplus value expected from the draft pick actually increases. Having established in §4 that the market value relationship is strongly negative and measured quite precisely, we will take as a sufficient, and conservative, test of our hypothesis whether the relationship between surplus value and draft order is positive over a substantial part of the draft. This relationship varies with draft order, so the formal tests should be specific to regions of the draft. We are most 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 that the overvaluation will be most extreme at the top of the draft.

6.2.4. Spline Regressions. 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

$$\hat{S}_{i,t} = \alpha + \beta_1 \text{Rd}_1 + \beta_2 \text{Rd}_2 + \beta_3 \text{Rd}_3 + \beta_4 \text{Rd}_4 + \beta_5 \text{Rd}_5 + \beta_6 \text{Rd}_6 + \beta_7 \text{Rd}_7 + \varepsilon_{i,t}, \quad (5)$$

where Rd_j is the linear spline for round j . In this model, β_j provides the estimated per-pick change in surplus value during round j . Estimation results are significantly positive for round 1. This is true whether using ordinary least squares ($M = 0.025$, $SE = 0.003$) or quantile regressions for the 25th ($M = 0.048$, $SE = 0.001$), 50th ($M = 0.017$, $SE = 0.000$), and 75th ($M = 0.012$, $SE = 0.002$) percentiles. In contrast, rounds 2–5 are negative in all models, significantly so in all the models for the second and fifth rounds.³²

6.3. Discussion

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% 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 throughout 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 and is riskier as well. Furthermore, the risk associated with very high picks is mostly on the downside. Because top picks are paid so much, there is little room for a player to greatly exceed expectations, but when top picks turn out to be complete busts, tens of millions of dollars are wasted.

The magnitude of the market discrepancy we have uncovered is strikingly large. A team blessed with the first pick could, in principle, through a series of trades, swap that pick for four or more picks in the top of the second round, *each* of which is worth more

than the single pick they gave up.³³ 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? We turn to this question next.

7. Additional Empirical Evidence

In this section we consider a variety of alternative explanations and provide additional empirical evidence relevant to the most common questions about these results. We also construct a new test of our research hypothesis, based on a very different dependent variable—wins. The objective throughout is to determine whether the main results are robust to alternative empirical formulations.

7.1. Alternative Explanations

7.1.1. Superstars. One might worry that our results could be produced by a failure to capture the true value of superstar players who can singlehandedly transform a team. We are skeptical of this explanation on three counts. First, a football team has so many players (53 on the roster, of which 22 are regular starters, not including specialists such as kickers) that it is difficult for a single player to have such a profound effect (unlike in basketball, for example).

Second, not all great players come from the top of the draft. The two best quarterbacks in recent years, Peyton Manning and Tom Brady, are cases in point. Manning was taken with the first pick in the draft, but Brady was taken 199th. Also, as we showed in the opportunity-cost analysis above, trading down to get *more* players does not reduce the chance of getting *top* players.

Third, we are already valuing the performance of top players quite highly, and the valuation function is extremely convex. We estimate the value of the top (99th) percentile of players at more than twice that of players at the 94th percentile, and, in turn, value those twice as highly as players at the 72nd percentile.³⁴ This is one of the reasons we have found that there is no need for an additional “elite” performance category beyond our all-pro designation. We have estimated a wide range of compensation models

³² The four models produce patterns that are broadly similar; see the supplemental analysis for a complete table and graph.

³³ Theoretically, roster limits constrain the extent to which a team could pursue this strategy. However, as a practical matter, this is not binding. Teams can carry 80 players into summer training camp (versus 53 during the season), a six-week period that provides a much more thorough assessment of the player. Teams usually include 10–18 rookies on this roster, meaning they might have as many undrafted rookies as drafted. Hence, the marginal player displaced by an extra draft pick is an undrafted rookie. A team could certainly trade down enough to double the number of picks it has from 7 to 14 without bumping up against roster constraints.

³⁴ See the supplemental analysis for a more complete summary.

using an additional, sixth performance category for the players elected to some combination of the all-pro teams approved by the collective bargaining agreement (CBA).³⁵ No matter how exclusively or inclusively we construct the super-elite category, the labor market does not appear to distinguish it from our existing top category.

Still, we test the plausibility of this hypothesis by arbitrarily increasing by 50% the performance value of players who are consensus All-Pro, that is, elected to all three all-star teams approved in the CBA. There are, on average, about seven players a season (the top 0.4%) in this elite group of superstars. Despite this increase, which if fully compensated would almost certainly violate the salary cap of every team with one of these players, our estimated surplus value still increases during the first round of the draft according to the spline regressions estimated as in the previous section ($\beta = 0.022$, $t = 7.22$, $p < 0.01$). Indeed, even doubling the value of these elite players does not alter this pattern ($\beta = 0.016$, $t = 4.92$, $p < 0.01$). Thus, it does not appear that undervaluing superstars is a valid explanation for our results. Although this exercise is clearly arbitrary, these results and others from similar exercises demonstrate the robustness of the pattern we observe.³⁶

7.1.2. Other Alternative Explanations. *Off-Field Utility.* 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. The intuition is that a very exciting player might help sell tickets and team paraphernalia in a way his performance statistics do not reflect. Setting aside the fact that paraphernalia sales are shared equally across teams (unlike in European soccer, where jersey sales can yield a team millions of Euros), such arguments are dubious in American football. Very few football players are able to bring in fans without performing well on the field, the value of which we have captured in our analysis. The fans' interest in an exciting

player will not last long if the player does not contribute to the team winning on the field. However, to be certain, we replicated our analysis using only offensive linemen, the very large players who protect the quarterback and create holes for the running backs to run through, but who are forbidden to carry the ball. While the football cognoscenti may tell you they are the most important unit on the field, they attract little fan attention (or jersey sales). However, we find an almost identical relation between surplus value and draft order in this subsample.³⁷

Finer Performance Measures. Our main analysis of player valuation includes all NFL players. This restricts the performance measures we can use to those common across all positions, measures that are admittedly coarse. A question that naturally arises is whether a more fine-grained evaluation of player performance might alter our results. To evaluate this possibility, we estimate a separate valuation model for wide receivers ($n = 304$), the players whose main job is to catch the passes thrown by the quarterback.³⁸ We use the same estimation strategy as in our main analysis, except that instead of using broad categories to measure performance (e.g., starter, Pro Bowl, etc.), we use a continuous measure of performance—receiving yards—and explicitly allow for nonlinearities. As in the general model, we find that surplus values increases sharply through the first round, peaking somewhere in the second before gradually declining.³⁹ This relation is strikingly similar to that which we found in our general model.

Traded-For Players. A conclusion from our analysis is that teams should trade down, not up. A possible objection to this conclusion is that when teams trade up, they might have a special need at the position and/or believe they have particularly good information about the acquired player. To assess this possibility, we compare the performance of players “traded for”—the highest drafted player obtained by a team trading up ($n = 221$)—with the performance of all other players ($n = 3,409$). We calculate the player's performance over his first five years using four measures: probability of being in the NFL, games played, games started, and the probability of making the Pro Bowl. Using Tobit regressions, we estimate a separate model for each performance measure. In each model

³⁵ Pro Football Writers of America, the Associated Press, and the *Sporting News*.

³⁶ We perform an additional robustness check by increasing the value of *all* players by 50%. This addresses the concern—setting aside its validity—that the salary cap artificially reduces the wage that would be paid under a free market, and because this is done by a fixed percentage, the impact is greatest among the best-paid players. If these players are also systematically drafted early, our estimate of the draft-pick-performance relation will be muted. We find that after inflating the performance value of all players by 50%, the spline regression of surplus value on draft-pick order does become reliably negative for the first round. However, the decline (~20% over the first round) is much less steep than the decline of draft-pick values (~75%). We also note that for smaller increases in performance value (e.g., 20%), the surplus value still reliably *increases* over the first round.

³⁷ See the supplemental analysis for the complete analysis.

³⁸ We chose wide receivers over quarterbacks because a wide receiver's contribution is better captured by a single statistic than is a quarterback's. We chose wide receivers over running backs because we know from separate analyses that running backs are the poorest value of any position in the draft and therefore might bias the results in favor of our hypothesis. Still, the results are similar for those positions.

³⁹ R-squared is 0.80. Full results can be found in the supplemental analysis.

we regress player performance on draft order (using both linear and quadratic terms) and a dummy variable for whether the player was “traded for.” Evaluating 14 draft classes (1991–2004) over 18 seasons (1991–2008), we find that “traded-for” players do not perform differently than other players. In each of the four models, the dummy variable for traded-for players is not statistically different than zero.⁴⁰ This means that the players targeted in these trades perform no better than would be expected for their draft position.

7.2. Alternative Test: Wins

In this section we construct a final test of our hypothesis. Specifically, we test the implication that those teams who make wise trades according to these estimates—wittingly or unwittingly—will perform better on the football field.

The discrepancy between the market prices implied by the Chart and the surplus values we estimate suggest that teams who successfully exploit this difference can substantially improve their on-field performance. In this section we investigate whether the teams that make “smart” trades, by our measure, end up winning more games. There are various aspects of the National Football League that make finding a statistically significant relationship unlikely. With so many players on the roster, a single draft-pick trade may not have much effect. Also, teams play only 16 games a season, each with a substantial random component, so there is not much power to detect an increase in the chance of winning. Nevertheless, we test this implication.

Observations are team-years. For each draft year we calculate the net surplus value of the picks exchanged in trades. If teams do not trade any picks in a particular year, their net value is 0. We arbitrarily assign the first pick a value of 1.0 and compare the surplus values of other picks to that.⁴¹ To reflect the net value on a roster at a given time, we accumulate these single-year values into a rolling four-year sum for each team-year. We analyze the period following the introduction of free agency (1993). Given a four-year lag, this means our sample is limited to performance in years 1997–2008. As would be expected because trades are zero sum, the median value for a team-year is 0 ($M = -0.03$). There is wide range, though, with a minimum of -12.06 , a maximum of 11.23 , and an interquartile range of -1.80 to 1.86 ($SD = 3.0$).

We then evaluate the relation between this measure of draft-trade acuity and a team’s winning percentage. To do so, we regress a team’s winning percentage on

Table 3 Wins Analysis

Coefficient	Win %			
	(1)	(2)	(3)	(4)
Win %, lag 1	0.2191 (0.0427)	0.2213 (0.0426)	0.2161 (0.0440)	0.2174 (0.0442)
Win %, lag 2	0.1203 (0.0488)	0.1186 (0.0479)	0.1261 (0.0483)	0.1249 (0.0473)
Win %, lag 3	0.0629 (0.0549)	0.0609 (0.0554)	0.0736 (0.0555)	0.0721 (0.0563)
Win %, lag 4	0.0346 (0.0586)	0.0341 (0.0584)	0.0411 (0.0594)	0.0407 (0.0594)
No. of picks (4-yr. accum.)		0.0015 (0.0026)		0.0009 (0.0026)
Net trade value (4-yr. accum.)			0.0057 (0.0028)	0.0056 (0.0029)
Constant	0.2834 (0.0366)	0.2355 (0.0914)	0.2732 (0.0353)	0.2452 (0.0922)
Observations	393	393	393	393
R-squared	0.0932	0.0940	0.1011	0.1014

Notes. Regression results evaluating the impact of draft-pick trades on team winning percentage. Observations are team-seasons, 1997–2008. No. of picks is a team’s total number of draft picks. The net trade value is the net value exchanged by a team via draft-day trades, as estimated by the surplus curve in the previous section. Both no. of picks and net trade value are accumulated over a four-year window. Standard errors are clustered on team and reported in parentheses.

the four-year trade value accumulated on that year’s roster. We include four years of a team’s lagged winning percentages to control for previous performance. Finally, we cluster standard errors by team.

Regression results are presented in Table 3. Model (1) is simply the time series of winning percentage; it shows that winning is reliably persistent for two years before dropping off. Model (3) adds our trade-value variable. This variable is positive and significant ($p < 0.05$). Additional analyses reveal that the strength of the effect has increased over time and is strongest for the last four years of the sample. During this final period, a one-standard-deviation improvement in draft-pick trading produced an estimated 1.5 wins per year, a huge number in a 16-game season.

These results should be interpreted cautiously. We can say that teams that make the type of draft-day trades that our model evaluates positively tend to subsequently win more games. Of course it is likely that such teams also do other things right, so it is possible that draft-pick trading might be serving as a proxy for overall management intelligence.⁴²

8. Conclusion

Our modest claim in this paper is that the owners and managers of National Football League teams are

⁴⁰ See the supplemental analysis for full regression results.

⁴¹ For example, we estimate that surplus value peaks at approximately 1.2 near the end of the first round.

⁴² In fact, two of the teams that do well in our evaluation of draft trading (New England and Philadelphia) also try more fourth-down conversions than average, a smart strategy as judged by Romer (2006).

subject to the same biased judgments found in countless other domains. Furthermore, market forces have not been strong enough to overcome these human failings. The task of picking players, as we have described in this paper, 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 §3, 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, although 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.

The Chart is an example of an especially interesting social phenomenon in which a bias or wrong belief becomes conventional wisdom and then, eventually, a norm. The NFL draft is a situation with both great uncertainty and the need to coordinate, making a norm—such as the Chart—especially valuable. However, *which* norm is determined by the psychological biases at play. The early trades on which the original Chart was based were priced according to the intuitions of team decision makers. As we have argued, we have reason to expect these intuitions are overconfident. Once distilled as a norm, this overconfidence is self-enforcing via the confirmation bias (Klayman and Ha 1987), status quo bias (Samuelson and Zeckhauser 1988), and regret aversion (Bell 1983). It also changes the incentives that decision makers face, because there may be sanctions from fans, media, and possibly even ownership for deviating from such a widely accepted practice. Because norms exert such power, biases once codified are particularly pernicious. Hence, the Chart appears both a symptom of biased judgment and also a self-perpetuating cause. This dynamic between biases and norms deserves greater investigation.

Our findings are strikingly strong. 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 are paying a lot to acquire a pick that is worth less than the ones they are giving up. We have conducted a wide range of empirical tests, and every analysis gives qualitatively similar results. The same is true under the 2011 labor agreement. The new rookie

salary cap reduced the cost of the very top draft picks, but not enough of them to alter our results.

We think that although our results are surprising, they are plausible. We suspect that some teams, even after 15 years, have not fully come to grips with the implications of the salary cap. Buying expensive players, even if they turn out to be great performers, imposes opportunity costs elsewhere on the roster. Some of the most successful franchises seem to understand these concepts, but others do not. However, 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 only a benefit if the team trades it away. 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. Because there is no way to sell the oil leases short, the smart money cannot actively drive the prices of those leases down. Similarly, because 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 Chart is widely used and accepted in the NFL, a team that owns a top draft pick and would like to trade it may be 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. At its most general, these findings stand as a reminder that decision makers often know less than they think they know. This lesson has been implicated in disaster after disaster, from financial markets to international affairs. Closer to the topic at hand, 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, whether newly minted MBAs or the next CEO.⁴³ In our judgment, there is

⁴³ Kaplan et al. (2012) find that observable traits reliably predict the performance of CEOs in buyout firms and, to a lesser extent, venture-capital backed firms. The explanatory power of these models is understandably modest (*R*-squared ranges from 0.14 to 0.42). The present research suggests decision makers in these environments will overestimate this predictive ability. This has consequences both before the hiring decision (e.g., settling on candidates too soon, paying too much for one over another) and afterward (e.g., insufficient monitoring, misconstruing early performance).

little reason to think that the market for CEOs is more efficient than the market for football players. The problem is not that future performance is difficult to predict, but that decision makers do not appreciate how difficult it is. Or, as Mark Twain supposedly put it, more colorfully: “It ain’t what you don’t know that gets you into trouble. It’s what you know for sure that just ain’t so.”

Acknowledgments

The authors thank Marianne Bertrand, Jim Baron, Rodrigo Canales, Russ Fuller, Shane Frederick, Rob Gertner, Jon Guryan, Tom Hubbard, Rick Larrick, Michael Lewis, Toby Moskowitz, Barry Nalebuff, Devin Pope, Olav Sorenson, David Robinson, Yuval Rottenstreich, Suzanne Shu, Jack Soll, George Wu, and especially David Romer. They also thank workshop participants at the University of California at Berkeley, Carnegie Mellon University, Cornell University, Duke University, the Massachusetts Institute of Technology, the University of Pennsylvania, Stanford University, the University of California at Los Angeles, the University of California at San Diego, the University of Chicago, and Yale University for valuable comments; and Chad Reuter, Al Mannes, and Wagish Bhartiya for very helpful research assistance.

Appendix. Estimating the Trade-Market Value of Draft Picks

A.1. Methodology

We are interested in estimating the value of a draft pick in terms of other draft picks, as a function of its order. We let the first pick be the standard by which we measure other picks. We assume that the value of a draft pick drops monotonically with the pick’s relative position, and that it can be well described using a Weibull distribution.⁴⁴ Our task is then estimating the parameters of this distribution.

Let t_i^r denote the i 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

$$\sum_{i=1}^m v(t_i^H) = \sum_{j=1}^n v(t_j^L), \tag{6}$$

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

$$v(t_i^r) = e^{-\lambda(t_i^r - 1)^\beta}, \tag{7}$$

where λ and β are parameters to be estimated. Note that the presence of the β parameter allows the draft value to

⁴⁴ The Weibull distribution is a two-parameter function that nests the exponential, providing a more flexible estimation than a standard exponential would provide.

Table A.1 Draft-Pick Trade Value: Regression Results

Model:	(1)	(2)	(3)	(4)	(5)
Years:	1983–2008	1983–2008	1983–1992	1993–2000	2001–2008
Future picks:	No	Yes	No	No	No
Parameter estimates					
λ	0.146 (0.027)	0.0996 (0.016)	0.199 (0.086)	0.184 (0.068)	0.0994 (0.021)
β	0.698 (0.030)	0.745 (0.026)	0.642 (0.068)	0.662 (0.060)	0.764 (0.035)
ρ	1.358 (0.084)				
Implied values (relative to the first pick) (%)					
5th pick	68	76	62	63	75
10th pick	51	60	44	45	59
16th pick	38	47	32	33	46
32nd pick	20	28	16	17	25
64th pick	7	11	6	6	9
N	313	407	70	108	135
R -squared	0.99	0.99	0.98	0.99	0.99

Notes. Results from using nonlinear regression to estimate parameter values for a Weibull-function model of draft-pick value. Data are draft-day trades, 1983–2008. Excludes trades involving players ($n = 663$). Standard errors are in parentheses.

decay at either an increasing or decreasing rate, depending on whether its value is greater than or less than 1. 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 (7) into (6) and solving in terms of the high-end pick in the trade, we have

$$t_1^H = \left(-\frac{1}{\lambda} \log \left(\sum_{j=1}^n e^{-\lambda(t_j^L - 1)^\beta} - \sum_{i=2}^m e^{-\lambda(t_i^H - 1)^\beta} \right) \right)^{1/\beta} + 1, \tag{8}$$

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 (8) using nonlinear regression.⁴⁵

A.2. Results

We estimate (8) using the 313 current-year trades only, finding $\lambda = 0.146$ (se = 0.027) and $\beta = 0.698$ (se = 0.030). These results are summarized in Table A.1, column (1). As shown in the bottom half of the table, these values imply a steep drop in the value of draft picks. In short, the 5th pick is valued approximately 2/3 as much as the first pick, the 10th pick 1/2 as much, and the last pick in the first round about 1/5 as much.⁴⁶ We also estimated an extended version of (8) that includes a parameter for the discount rate.⁴⁷

⁴⁵ We first take the log of both sides of expression (8) before estimation to adjust for lognormal errors.

⁴⁶ We drop one trade from our estimation because of its disproportionate influence. We identify this trade by repeatedly estimating this model while dropping one observation at a time. Results are robust to the exclusion of all trades except one, the inclusion of which changes values dramatically. Excluding this observation provides a conservative test of our main hypothesis since the valuation curve is flatter without it.

⁴⁷ See the supplemental analysis for the full derivation.

This expression allows us to include trades involving future picks, expanding our sample to 407 observations. Results are presented in Table A.1, column (2). The estimated curve is close to the previous one, with $\lambda = 0.0996$ (SE = 0.016) and $\beta = 0.745$ (SE = 0.026), although a bit flatter—e.g., the 10th pick is valued at 60% (versus 51%) of the first. The estimated discount rate, ρ , is a staggering 136% (se = 0.084) per year.

Finally, we investigate how these draft-pick values have changed over time, focusing on trades for current picks only. We find that the valuation curve has flattened some in recent years, meaning pick values do not decline as rapidly. This is consistent with teams learning, however slowly and insufficiently, that the top picks are relatively overpriced.

References

- Ahlburg DA, Dworkin JB (1991) Player compensation in the National Football League. Staudohar PD, Mangan JA, eds. *The Business of Professional Sports* (University of Illinois Press, Urbana), 61–70.
- Alpert M, Raiffa H (1982) A progress report on the training of probability assessors. Kahneman D, Slovic P, Tversky A, eds. *Judgment Under Uncertainty: Heuristics and Biases* (Cambridge University Press, Cambridge, UK), 294–305.
- Bell D (1983) Risk premiums for decision regret. *Management Sci.* 29(10):1156–1166.
- Bronars SG (2004) Bargaining in professional football: Why NFL superstars are underpaid. Working paper, Georgetown University, Washington, DC.
- Burke B (2012) Art and analytics of negotiation. Panel discussion, MIT Sloan Sports Analytics Conference, March 2, Boston.
- Carter V, Machol RE (1978) Optimal strategies on fourth down. *Management Sci.* 24(16):1758–1762.
- De Bondt WFM, Thaler R (1985) Does the stock-market overreact. *J. Finance* 40(3):793–805.
- DellaVigna S (2009) Psychology and economics: Evidence from the field. *J. Econom. Literature* 47(2):315–372.
- ESPN (2004) NFL draft-pick value chart. *ESPN Insider* (March 1), <http://sports.espn.go.com/nfl/draft06/news/story?id=2410670>.
- Gilovich T, Griffin DW, Kahneman D (2002) *Heuristics and Biases: The Psychology of Intuitive Judgement* (Cambridge University Press, Cambridge, UK).
- Hall BJ, Lim JP (2002) A note on incentives in the NFL. HBS Case 902–129, Harvard Business School, Boston.
- Harrison JR, Bazerman MH (1995) Regression to the mean, expectation inflation, and the winner's curse in organizational contexts. Kramer R, Messick DM, eds. *The Social Context of Negotiation* (Sage Publications, Thousand Oaks, CA), 69–94.
- Harrison JR, March JG (1984) Decision-making and post-decision surprises. *Admin. Sci. Quart.* 29(1):26–42.
- Kahn LM (1992) The effects of race on professional football players compensation. *Indust. Labor Relations Rev.* 45(2):295–310.
- Kahneman D, Tversky A (1973) On the psychology of prediction. *Psych. Rev.* 80(4):237–251.
- Kahneman D, Tversky A (2000) *Choices, Values, and Frames* (Russell Sage Foundation, Cambridge University Press, Cambridge, UK).
- Kaplan S, Klebanov M, Sorensen M (2012) Which CEO characteristics and abilities matter? *J. Finance* 67(3):973–1007.
- Klayman J, Ha Y (1987) Confirmation, disconfirmation, and information in hypothesis testing. *Psych. Rev.* 94(2):211–228.
- Leeds MA, 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. *J. Sports Econom.* 2(3):244–256.
- Lerner J, Tetlock P (1999) Accounting for the effects of accountability. *Psych. Bull.* 125(2):255–275.
- McCoy M (2006) Personal correspondence with the authors, September 12.
- Oskamp S (1965) Overconfidence in case-study judgments. *J. Consulting Psych.* 29(3):261–265.
- Romer D (2006) Do firms maximize? Evidence from professional football. *J. Political Econom.* 114(2):340–365.
- Ross L, Greene D, House P (1977) The false consensus effect: An egocentric bias in social perception and attribution processes. *J. Experiment. Soc. Psych.* 13(3):279–301.
- Russo J, Schoemaker P (2002) *Winning Decisions: Getting It Right the First Time* (Doubleday, New York).
- Samuelson W, Zeckhauser R (1988) Status quo bias in decision making. *J. Risk Uncertainty* 1(1):7–59.
- Slovic P, Corrigan B (1973) Behavioral problems of adhering to a decision policy. Seminar presentation, May 1, Institute for Quantitative Research in Finance, Napa, CA.
- Stewart S (2005) Can behavioral economics save us from ourselves? *Univ. Chicago Magazine* 97(3):36–42.
- Thaler RH (1988) The winner's curse. *J. Econom. Perspect.* 2(1):191–202.