

Flipping Coins in the War Room: Skill and Chance in the NFL Draft*

Cade Massey

Yale School of Management

January 7, 2011

Overview

People often judge owners and managers on their ability to pick players. Reputations have been made on the basis of one or two successful picks, and careers ruined by spectacularly unsuccessful ones. Consistent with this, teams spend millions of dollars and thousands of hours on draft preparation. This *seems* appropriate given how much impact drafted players can have on a team. But does that effort and investment make a difference? Is there any evidence of player-picking skill?

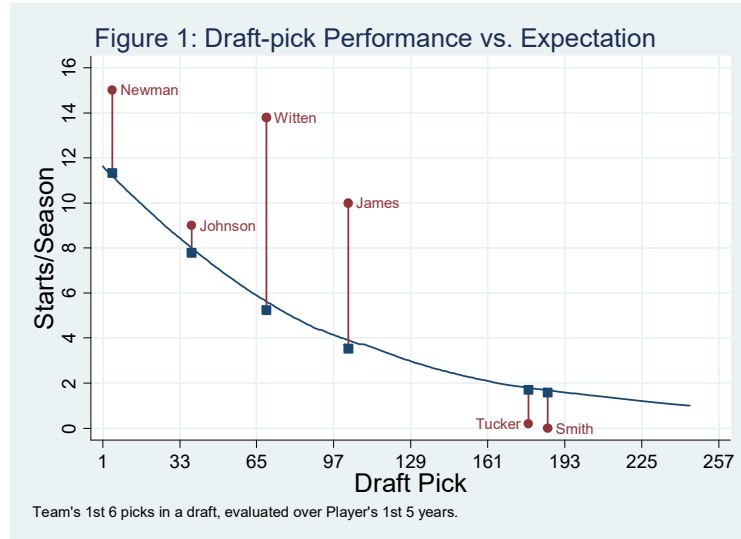
We hypothesize that success in player selection is largely chance. Previous research has shown there is significant uncertainty in forecasting the NFL careers of college players, and that teams underestimate this uncertainty (Massey & Thaler, 2010). Indeed across many domains, people underestimate the role of chance in outcomes (Fischhoff, 1982). One consequence of this is “outcome bias”, a tendency to judge a decision by its eventual outcome rather than by inputs available at the time (Baron & Hershey, 1988). Because of these tendencies, management is overly praised when players succeed and overly criticized when they do not. Management errs as well, investing too much in the “skill” of picking players and too little on the “pick management” that the role of chance suggests.

* Selected for the research paper competition at the 5th-annual MIT Sloan Sports Analytics Conference. Presented at the conference on March 4, 2011.

Data & Method

We investigate these issues using 18 years of NFL data (1991-2008). We consider how well each pick performs using very simple performance statistics – such as games started or pro bowls played – in the five years after being drafted. We first establish the relation between performance and draft order, finding the downward slope one would expect – by all measures, performance *is* related to draft order. We then evaluate a team's drafting ability by measuring deviations from this expectation. If there is team-specific skill in player selection, these deviations should be related within a team. For example, a team good at selecting players might have positive deviations in successive years, or in successive picks.

Figure 1 provides an example from the Dallas Cowboy's 2003 draft. The figure shows the team's first six draft picks, the pick with which they were taken (horizontally), and the average number of games started (0-16) during their first five years on (vertically). The downward sloping blue line is the expected performance for each draft position, estimated from the entire dataset and adjusted for the particular draft year. We measure the success of a draft pick as the performance *relative to* this expectation.



In this example the first four picks outperformed expectations, while the last two underperformed (slightly). Using this approach, the Cowboys get less “credit” for picking Terence Newman 5th than for Jason Witten 69th because, even though Newman started slightly more games, players drafted that high are *supposed* to start many more games than those drafted as low as Witten. We use this process to value every draft pick in the dataset.

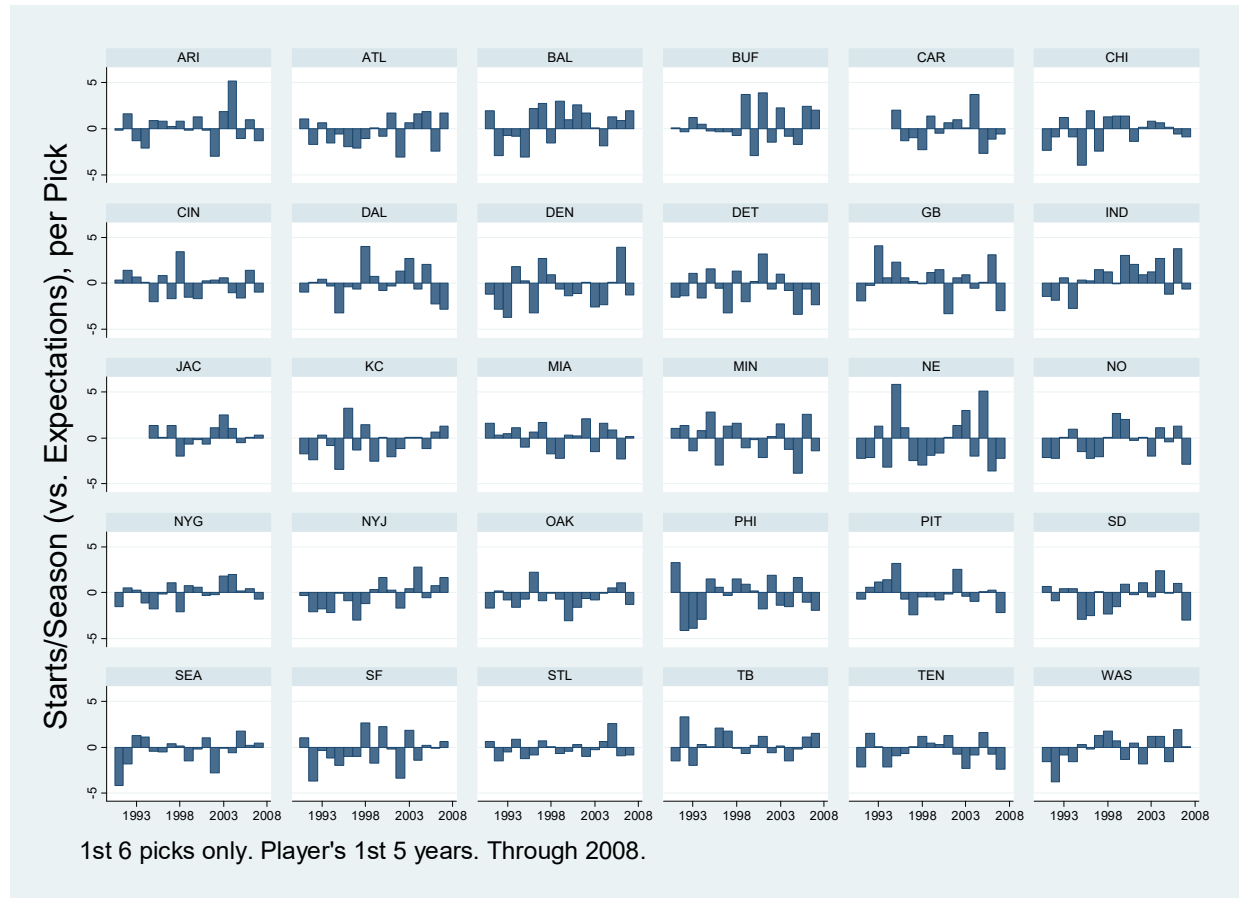
Result 1: Within Year, Across Organizations

We look for evidence of persistent ability – good or bad – in drafting players in multiple ways, two of which we report here. First, we consider whether a team’s picks from a particular year’s draft are correlated with each other. If draft success is due to skill rather than chance, the performance of a player should be positively related to the performance of other players chosen by the same team in the same year. We test for this by estimating the intraclass correlation coefficient (ICC). The ICC reveals whether the within-group variance is reliably different than the between-group variance, where group for us is a team’s pick in a given year.

We estimate the ICC for each of the 18 years in our sample, finding remarkably little evidence of within-team variation. The average intraclass correlation across is .01. The maximum correlation was in 1995, when the ICC was estimated to be .12 (though even here the confidence interval included 0). In 13 of the 18 years the coefficient is estimated to be 0, the minimum value.

Result 2: Within Organization, Across Years

A second way we look for persistent drafting ability is in the relation between a team's draft success in consecutive years. If a particularly good or bad year is due purely to chance, performance one year should be unrelated to performance the next year. To run this test we first average the performance of each set of picks within a team-year. As above, performance is average games started in the player's first five years, relative to expectations for his draft-pick. We then simply regress a year's average performance on the previous year's average performance. We conduct these tests for each team using our entire 18-year sample, as well as for every consecutive 10-year sample within it. The figure below depicts these data.



As one might surmise by inspection alone, we find no evidence of persistent performance across consecutive years. Using the 18-year sample, draft performance is not reliably predicted by the previous year's performance for any of the teams. Looking at the narrower 10-year windows (to mitigate the impact of turnover in organizations), we still find no reliably positive relations between years. That is, in no 10-year period between 1991 and 2008 did a team's draft performance reliably predict the following year's draft performance. Indeed, across these 256 (team-decade) tests the average coefficient for the previous year's performance was -0.14 (the average SE was 0.36).

Discussion

These data strongly suggest there is no persistent player-selection skill in the NFL. Indeed, it is hard to imagine a more resounding failure to find evidence of persistent drafting ability. While in aggregate draft order is clearly related to player ability, individual teams do not show the same talent. Teams appear akin to mutual fund managers – collectively helping create an efficient market, while not able to individually outperform it.

Teams do not appreciate the amount of uncertainty involved in the draft. As a result, they misallocate time and resources to improving their draft *skill* rather than optimally managing that uncertainty. Recognizing the role of uncertainty, and adapting appropriately, is a critical task for organizations and individuals alike, in a wide range of domains. We need to better explore the barriers to acknowledging uncertainty, and improve tools for navigating it. This process involves better understanding the broader objectives served in organizational decision-making. If teams are only out to find the best players, their current approach is flawed. If they also care about objectives like commitment to the player, identifying and reinforcing their preferences and values along the way, etc., there may be something adaptive about it. Our results suggest these questions deserve greater exploration.

References

- Baron, J., & Hershey, J. (1988). Outcome bias in decision evaluation. *Journal of Personality and Social Psychology*, 54(4), 569-579.
- Fischhoff, B. (1982). For those condemned to study the past: Heuristics and biases in hindsight. *Judgment under uncertainty: Heuristics and biases*, 4, 79-93.
- Massey, C., & Thaler, R. (2010). The loser's curse: Overconfidence vs. market efficiency in the National Football League draft. *Under review*.