

Salary Allocation and Risk Preferences in the National Football League:
The Implications of Salary Allocation in Understanding the Preferences of NFL
Owners

Papa Chakravarthy

Presented to the Department of Economics

In partial fulfillment of the requirements

For a Bachelor of Arts degree with Honors

Harvard College

Cambridge, Massachusetts

March 8th, 2012

Abstract

The study of risk preferences in the allocation of the National Football League's salary cap has not seen much academic research. Previous analysis shows that the salary cap improves parity across the NFL and may be partially responsible for the growth of the United States' most popular sports league. Allocating salary to players, however, can reveal a great deal of information regarding the utility function of NFL Owners. This paper illustrates, using data on wide receivers in the NFL from 2005 to 2009, variables predicting future performance in the NFL do not predict future salary, meaning Owners value something in addition to future expected performance when allocating salary. The potential to become a star, leadership or popularity may be the characteristic valued by Owners that is not shown by OLS regression. A comparison of NFL Owners to fantasy football owners shows that while the method of risk aversion differs between the two, it is impossible to rule out risk aversion by NFL Owners trying to retain aged players with high salaries. Finally, it is possible that NFL teams may improve team performance while staying within the salary cap by cutting players more frequently and signing shorter contracts, thus eliminating the need to overpay players and fielding a more competitive team.

Acknowledgements

This thesis would not be possible without my advisor, Professor Christopher Avery, whose enthusiasm, guidance and support were invaluable. I never left a meeting with him without feeling optimistic I could handle the work ahead. I am truly indebted to Dr. Niels Rosenquist who convinced me to step forward to the most satisfying undertaking of my undergraduate experience. Finally, I owe a great deal to Cathy Barerra, without whom this paper would be largely inaccessible; Professor Susan Athey, whose advice was incredibly helpful; and my parents, who may not know exactly what this paper concerns but provided a constant stream of encouragement nonetheless.

Table of Contents

Introduction	4
Background	7
Previous Results	15
Theory and Testable Hypotheses.....	25
<i>Problem I: Player Salary and Salary Growth</i>	26
<i>Problem II: Player Performance and Performance Growth</i>	28
<i>Problem III: Lagged Values</i>	30
<i>Problem IV: Risk Preferences and Cutting Players</i>	33
Methods.....	36
Data.....	43
Results.....	46
<i>Part I: Player Salary and Salary Growth</i>	48
<i>Part II: Player Value</i>	51
<i>Part III: Lagged Values</i>	54
<i>Part IV: Risk Preferences and Cutting Players</i>	56
Limitations	59
Conclusion.....	60
Tables.....	64
Figures.....	71
References	82

Introduction

The recent Academy Award nominated film “Moneyball” brought sabermetrics as evaluation of baseball players into the mainstream. Evaluation of talent in the National Football League (“NFL”), however, has not yet seen the levels of statistical analysis common in Major League Baseball. The study of salary allocation and talent evaluation in the NFL is a valuable pursuit because it focuses on a huge market that has seen little to no academic study. NFL players together make approximately \$3.5 billion per year, but the division of that money and the reason certain players make more than their peers has not been considered from an economic perspective. Often, teams appear to pay players in a different way than their on-field performance would indicate – a possible inefficiency in incentivizing future performance that merits investigation.

This paper concerns the process governing NFL salary allocations with multi-year contracts, the reasons for player payment strategies, as well as the induced risk preferences shown by NFL teams. Implicit in that query is the need to understand exactly the set of preferences espoused by NFL Owners, a major focus of this thesis. For example, perhaps Owners¹ prefer older, more popular players to younger, less experienced athletes because of the profits gained from jersey sales of popular players. If so, those Owners would probably weight different characteristics than on-field production, like popularity or teamwork, and would pay players differently as a result. If we assume that every Owner has the sole goal of winning more games, we can examine

¹ To distinguish from owners in the game of fantasy football, I have capitalized “Owner” when referring to those in the NFL.

the induced risk preferences from on-field characteristics. Studying player contribution to team wins, through on-field statistical results, and relating that value to a player's salary is imperative for a strong analysis. If player's on-field value to his team is largely uncorrelated to player salary, owners probably value characteristics invisible during the football games, rendering my previous assumption invalid. However, such a result would provide great insight into the utility curves of Owners and General Managers, by showing that winning games takes a backseat to some other goal.

My junior research paper on market design could elucidate findings in this study (P. Chakravarthy, 2010). I found risk-averse preferences among owners in fantasy football, a widely played game in which a group of individuals select real football players and compete with each other to earn points based on those players' performance. This risk-aversion is mostly evidenced by an owner spending a larger allocation of his/her limited budget on players likely to be in a bench role than would be expected by on-field contributions and appears to be a result of overestimating the probability of an injury or other performance limitation to a player in a starting role. Bench players are those would not contribute to the team except in the case of replacing a better starting player. Risk-aversion on the part of many owners would increase the price of many players as the market finds equilibrium. Spending a higher allocation on backups forms a type of insurance against injury or performance-risk. Such a strategy may be induced by the type of market – certain owners may derive utility from factors other than their success in a league, perhaps they generate more utility from avoiding a poor rank than attaining a high rank. In the NFL, Owners and General Managers of teams might express similar

behavior. A team that wants to avoid finishing last in its division or conference will take a risk-averse set of choices by allocating money to backups and avoiding the risk of an injury derailing a season. Alternatively, if a team is weak at a given position, they may over spend relative to value contributed to the team to garner top level players. This induced risk preference differs from the risk-averse preference above, but may be a result of scarcity of talent, which would drive up prices for top tier players. While fantasy football usually consists of only ten teams, the National Football League has 32 teams. This means that the same pool of top players is being divided into far more teams, and the relative lack of hugely talented players makes them more valuable in the NFL. The question of how the risk-averse preferences of fantasy football owners apply to their NFL counterparts is an interesting one because of the competing effects described above.

The questions above will be discussed as outlined here. First, I will detail the background of the NFL, including information on how players are allocated and previous findings regarding the salary cap's impact on a team's operations. This is followed by a section detailing my previous research concerning fantasy football and the hypotheses formed from that work. In the Theory and Testable Hypotheses section, I detail exactly which tests will provide the answers to each question, and what each possible result implies about the analysis. Following this breakdown of tests, are Data, Methods and Results sections, explaining test results. The Conclusions section provides summarizing of thoughts and a reiteration of the points learned from testing hypotheses.

Background

The NFL is the most popular sports league in the United States and among the most popular in the world (A. Heller, 2000). The direct consequence of this popularity is the special niche which a team holds in the hearts and minds of its followers, not unlike a religion. Because football teams have such a cult-like following and all cities hosting NFL teams contribute financially to building and maintaining a stadium without taking direct profit,² many people feel that Owners should not treat the team as a financial investment, but rather as an opportunity to provide a public good while still gaining some financial benefit. This general attitude is the foundation for the assumption discussed earlier – that owners have one goal above all others – to field a successful football team as measured by winning percentage. While Owners have the power to do whatever they want with their team, when they are seen as miserly it can have a severe effect on their finances. This inference is based in part on a major movement among the followers of the Cincinnati Bengals franchise, angered by a consistently poor team and an Owner unwilling to change. The anger of the movement is based in large part on the perceived goal of the team's Owner – to make money rather than to win games.³ While an angry group may feel an Owner is responsible for providing a competitive team, the

² City governments contribute to stadiums because of the associated revenue teams bring into towns. Also, the risk that a team may transfer to another city is a major factor. An example is the recent negotiations between the Minnesota Vikings and Minneapolis, in which the Vikings threatened to move to Los Angeles. http://espn.go.com/nfl/story/_/id/7190566/stadium-drama-stoking-fears-minnesota-vikings-move

³ Who Dey Revolution, named for the chant "Who dey think gon' beat them Bengals," believes in destabilizing the organization by affecting the franchise's ticket intake and directly limiting their financial success. Group members routinely boycott the games and advise against purchase of any memorabilia. One notable protest involved a nude man streaking through the stadium while holding a poster reading "Fire Mike [Brown, the Bengals' Owner]." <http://www.whodeyrevolution.com/>

following analysis can show whether that is, in fact, an Owner's primary objective or if there are other motivations in choosing players besides expected on-field performance.

The NFL is made up of 32 clubs, split up into eight different divisions. Each team plays 16 games in a season. For the sake of parity, only a certain number of players may play in any given game. Teams may have 53 players on their roster, but only 45 of those players may "dress for" and participate in any given game (R. Goodell, 2011). Finding and replacing players is difficult because of the particular skill set required for success playing football, making the labor market extremely thin, which could have a significant effect on the relationship between player contribution and salary. The football player's market is governed by a Collective Bargaining Agreement ("CBA") negotiated between Owners and players. All teams have Owners⁴ and most have General Managers or Presidents. Officers such as these run the daily operations of a team, including acquisition of players and interaction with fans or media.

NFL teams may acquire players in three ways: free agency, drafts, and trades. Free agency as we know it now arose during the CBA negotiations of the late 1980's (E. Garvey, 1989). Garvey details the source of the disagreement surrounding those negotiations – player desire for complete free agency with its accompanying increases in payment and owner desire to avoid salary hikes associated with the free market. A free agent is a player who has completed the requirements of his contract and is currently

⁴ The notable exception to this statement is the Green Bay Packers. The Packers are a publicly owned company with a Board of Directors. The reason for this division is a fundraising one – the team needed money to survive in 1950 and sold shares of ownership. There are over four million shares, but they may not be traded or sold. Individuals holding shares have the right to vote for Board members, who in turn elect a President to run the team. <http://www.packers.com/history/birth-of-a-team-and-a-legend.html>

available to negotiate new contract with any team (M. Truelock, 1993). He may be signed by whichever team he selects and has complete control over his decision, be it motivated by finances or other reasons. A result of free agency, aside from increased player salaries, is more sophisticated coaching schemes fueled by the ability to find specific skills in personnel (S.E. Backman, 2002). Free agency has also increased competitive balance and parity (A. Larsen et al., 2006) along with player payment inequality (M.A. Leeds and S. Kowalewski, 2001, G.W. Scully, 2004). Consequently, teams use lucrative, long term contracts used to lure free agents. One important concession of free agency, which is not seen in any other major sports league, is a team's right to "cut" a player before the season begins by removing him from their roster without having to pay any non-guaranteed contractual obligations, such as a bonus or future salary. The ability to cut a player drastically lowers a team's risk and helps limit the damage possible from a poor contract decision, by limiting the effect of a salary on a team's salary cap. It may have an effect on the largesse of contracts – teams give more money than expected for players, knowing that if the contract becomes unwieldy they can cut the player. Teams may also renegotiate contracts with players during the term of the contract, allowing teams to find creative ways to allocate salaries, usually through the use of increased bonuses, while remaining under the league-wide salary cap.⁵ Signing a renegotiated contract puts players at risk of being cut before

⁵ Ben Roethlisberger of the Pittsburgh Steelers recently agreed to such a restructuring for the reasons outlined above. http://espn.go.com/nfl/story/_/id/7609160/pittsburgh-steelers-ben-roethlisberger-restructures-contract

reaching their highest paid years but also allows for the signing of better teammates and being part of a better team, the reason players often agree to such an option.

New players usually enter the NFL through the league's amateur draft. Most players in the draft played in a college as the league does not allow players who are less than three years removed from their high school graduation (R. Goodell, 2011). Teams enter the draft ranked in reverse order of performance and draft for seven rounds in the same order. This means that the worst team in 2011 will draft first for each round, one to seven, in 2012's draft. The team that won the Super Bowl this February, the New York Giants, will draft 32nd in each round. Massey and Thaler (2005) detail the overvaluation of early draft picks with regard to contribution to a team, an effect which reiterates the importance and difficulty of predicting future performance in the NFL – a major focus of this study. Highly touted college players are paid far more than their on-field contribution to the team would indicate, perhaps because of the potential such a player has to become a team's main attraction, thereby generating revenue for the team. This overvaluation may be important to understanding risk preferences as it is a common thread in many sports, which value having a "star" capable of drawing fans to buy tickets and merchandise (S. Rosen, 1981). The motive behind paying for star potential – the idea one great player is worth far more than any number of good players because of associated financial benefit – could be a driving factor in how owner's allocate their funds in acquiring and retaining players.

Trades also improve competitive balance through player movement (A. Nissim, 2004) by allowing teams to find players matching their ideal schemes. A player trade is executed by trading the rights for a player. This means that a player may not be under contract while being traded; only that he has been assigned to a team.⁶ When a player is drafted, the team owns the rights to that player, which means no other NFL team may sign a contract with that player for one calendar year. During that time, a player may still be traded even if he has not signed a contract. The result of these trades is that teams have players who fit within a malleable system and can be maximally utilized according to fit within a team rather than to pure ability, two attributes that do not always correlate directly. Trading rights makes salary figures more accurate predictors of player value to the team instead of simply indicators of player ability, which means there is less noise in comparing salary to player contribution. Moreover, studying the utilization of players can give insight into a player's skill. While this may not always be the case, players who are used more often are generally being used because their coaches believe they give their team the best chance of winning. Measuring utilization is a way of measuring a team's faith in a player.

The most important factor in a club's acquisition of players is the salary cap. The NFL uses a hard salary cap, which means that teams may not exceed a certain number in yearly salary (B. Richard, 2008). The salary cap became a part of the NFL in 1994. It started at just more than \$34 million per team, but has since grown rapidly. Currently,

⁶ This distinction allowed the famous Eli Manning – Phillip Rivers trade in 2004, when neither player had signed a contract and Manning forced a trade because he didn't want to play in San Diego. http://espn.go.com/blog/afwest/post/_/id/39688/poll-eli-manning-trade

the salary cap stands at \$128 million per team (NFL Enterprises, 2010).⁷ Despite the changes in the NFL Collective Bargaining Agreement during the summer of 2011, the basic rules regarding player salary have not been changed a great deal since free agency began. Under the cap, the allocation of salaries between players is largely unregulated. The only limitation is that the NFL has minimum salaries for players depending on years in the league. Players in the league for more than a decade can make no less than the “veteran minimum” of \$810,000 but younger players can be paid as little as \$405,000 (NFL Enterprises, 2010). This means that with a cap of more than \$100 million, a team could theoretically devote a huge amount of money to a single player and spread their remaining salary among lots of players or they could opt to give the same amount to every player. A common assumption is that players are compensated according to expected future performance (C. Massey and R. Thaler, 2005). This is a consistent theme throughout major sports contracts – younger players with similar immediate performance to older players are expected to earn more money because of the likelihood of improved performance in the future.

A great deal of research studies the effects of this cap on the NFL. Richard (2008) shows that teams oriented around paying a single player a huge amount tend to perform worse on average, perhaps because of dissatisfaction from teammates. Other literature suggests that the salary cap exists to keep player salaries evenly distributed among teams (M.J. Redding and D.R. Peterson, 2009), which increases parity and

⁷ The cap has grown because of increases in revenue from television and merchandising, but also because players are demanding a higher ratio of the revenue taken in by clubs. This debate was the center of the 2011 NFL lockout. <http://sports.espn.go.com/nfl/news/story?id=6687485>

therefore consumer interest (A. Larsen, A.J. Fenn and E.L. Spenner, 2006). Alternatively, Heller (2000) focuses on how the salary cap incentivizes owners and players working together to improve everyone's revenue. Allocating that money to players is a difficult task and understanding that allocation is important in the NFL because of the huge salaries involved. The general finding of this research is that the salary cap's limitations require teams to develop an understanding of how best to allocate their funds interest (A. Larsen, A.J. Fenn and E.L. Spenner, 2006). Although different teams have very different strategies they all need to understand how the salary cap works – indeed, this is backed up by anecdotal evidence. Most teams have several employees who work exclusively with cap allocations. However, it is worth noticing that while many have studied the effects of a salary cap on football games and how teams work under a cap (L.M. Kahn, 2000) there is a dearth of information regarding evaluation of players under the cap and understanding how Owners view risk in defining salary allocation.

When a team has a signed contract with a player, they may add him to a roster. To sign a contract, a team must know the expected contribution of a player to an Owner's utility (S. Rosen, 1981). Understanding all the ways a player contributes to his team's performance can help to judge a player's worth to his team. Also important but harder to define is a player's off-field contributions to his team, such as revenue benefits. These attributes may make a player more or less desirable and affect a salary.

The complexity of football complicates the addition or removal of a player from a team. Coaches spend weeks generating playbooks with proprietary information suited

toward the personnel in place on a team. This means there is an additional cost in adding or removing a player. Such a cost is a direct effect of the effort required to train him. Similarly, trading away a player costs more than production lost, it also leads to lost training time. An example of this cost is player Chad Ochocinco, who was largely successful during his tenure with the Cincinnati Bengals but had an exceedingly poor season with the New England Patriots in 2011. There are a variety of possible causes for poor performance this season – including such mundane aspects as the weather of his new home. However, one possible explanation is the complexity of the New England Patriots playbook – widely known as being among the most difficult to understand in the NFL – in conjunction with the 2011 NFL lockout, which drastically shortened the off season time for players to become acclimated to their new teams and new schemes. In the case of Ochocinco, such a limited time to learn an incredibly complex playbook undoubtedly limited his ability to perform at his expected high level.

Previous Results

Most of the following is a restatement of information and conclusions from my junior paper: *Optimizing Draft Strategies in Fantasy Football* (Chakravarthy, 2010).

Fantasy football is an interactive, often online, game played within a group of eight or ten people.⁸ Such a group is termed a “league.” Each of the people involved is considered an “owner” and runs a “team.” Each owner selects ten players in the National Football League for his team. Based on the performance of those players in real games against NFL opponents, the owner receives points for his team. The number of points received varies slightly depending on the league, but as a rule, better players earn their owners more points as a result of better performance. The more points a team receives relative to other teams in the league, the more likely a team is to win its league, which may earn its owner money or simply bragging rights. The strength of fantasy football as a model for statistical analysis of the NFL lies in the ability to isolate performance from induced revenue. In fantasy football, owners do not care about the revenue a player earns and are only influenced by the utility he provides them. Thus, we can isolate the ways a player may generate value for his team. Moreover, acquiring new players does not affect the performance of players on the team – players are independent from each other and qualities like leadership and teamwork.

Even with only a superficial understanding of the game, it is immediately apparent that obtaining the best players provides an owner with the best chance to

⁸ Estimates of the number of fantasy football participants vary widely, because many participants do not use online clients to manage their leagues. The number of total players in the United States may be as high as 50 million people (C. Harris, 2008)

place well in the league. Acquiring players occurs just prior to the beginning of the NFL season in a process known as a draft. Most leagues do this draft online using a draft engine, like ESPN.com or Yahoo's Fantasy Football portal. These drafts are separated into two major categories: snake and auction. A snake draft requires a random number generator to predetermine draft order for every team in the league. Then, each team makes a player selection in that defined order. The process is repeated, but in alternating forward and reverse order until every team has a full roster. This type of draft is currently popular but that popularity is trending downwards in favor of the more complex auction draft, which resembles the NFL free agency market.

In an auction draft, owners are again randomly defined in an order, but instead of selecting players in that order, owners nominate players. Then, each team bids on those players in a slightly modified English auction. Whichever owner bids highest retains that player for the year. In this internet-driven age, every owner has the same information about each player. Moreover, the online draft portal provides an estimated valuation for the player. This English auction has some important distinctions from a standard auction: Because each owner needs to fill his roster, the draft consists of a series of auctions which affect one another; additionally, each owner has a salary cap, a number set by the league to protect parity and prevent one owner with more money from buying all of the best players. Still, because the auction is an English auction, it is efficient with regard to allocating players to the owner with the highest valuation of that player.

The method of allocating players in an auction is very similar to the free agency marketplace in the NFL because each player can be offered varying amounts of money from different teams. While there may be other factors to be considered in choosing a location to play, money is definitely a major part of the decision making process for real players. My previous work was in this aspect of fantasy football – understanding how and why players were priced at certain levels and whether there is a better way to price players or strategically operate in the draft given an owner’s risk preferences. A player’s price can be impacted by two major effects – the expected performance and the possibility of injury or some other risk to the expected performance. The owner may also have his own reasons for wanting or not wanting a certain player – things like team allegiances and favoritism may come into in fantasy football just as they might in the NFL. One important factor to note in defining price is whether the player will earn his team points. Teams are limited in that only first-string players can earn points. Even though a player may be a first-string player in the NFL, he may perform less well than another first-string player and be a fantasy football backup.⁹ Since fantasy backups are less likely to earn the team more points than fantasy first-string players, fantasy backups are rarely utilized. Thus, even if a fantasy backup could earn his owner fifty points, since the first-string player would earn one hundred points, the backup would never be utilized and did not earn the owner any points. In such a situation, the backup’s zero

⁹ An example of this situation would be the wide receivers on the New England Patriots in 2012. Even though both Deion Branch and Wes Welker started in the NFL at wide receiver, Branch performed less well than many other NFL starters and in fantasy football was a backup.
http://espn.go.com/nfl/player/_id/3593/deion-branch

points are worth nothing to an owner aside from possible utility gained from preventing his presence on another owner's team.

Gleaning risk preferences from fantasy data requires a major assumption: all fantasy football owners have the same goal – to win their leagues. Therefore, an owner's utility would be determined exclusively by the probability of winning his or her league. There are two ways to improve probability of winning a league: one may improve his own team or one may try to hoard good players to prevent other teams from being able to score points. The latter strategy would greatly complicate any analysis of player price. Unfortunately tracking individual leagues is impossible; however there is a simple way to detect whether top tier players are being hoarded. Because a team must have five wide receivers only three of those wide receivers are being used at any time, so two wide receivers will not be started and will be second-string players for their team. If top players are being hoarded, we would see some top players who are not being started as often as would be expected. Figure 1, Figure 2 and Figure 3 are repetitive, but reinforce the pattern of starting players over multiple years. Both scatter plots are associated with regressions having statistically significant positive coefficients on points, meaning that better players must, in fact, be started more often than worse players. Unfortunately, I do not have starting data concurrent with my pricing analysis, however we can see the pattern in 2011 and 2012 is sufficiently strong that any other year of fantasy football would likely not have significantly different ways of starting and benching players. Therefore, we can rule out the idea of a systematic attempt to hoard top tier players to prevent other teams from scoring points. Thus, we can infer teams

usually hope to win by improving their own performance rather than hindering their competitors.

The distinction between starting and bench players is an important one to much of this study and particularly in the case of induced risk-aversion. The reason for this is that teams can only reap effective fantasy points, not all fantasy points. An effective point is a fantasy point scored by a player who started. Consequently, regardless of the number of fantasy points a bench player would score, since he doesn't start, his effective fantasy points are zero. Because a starting player could always be easily replaced by a bench player, a starter's value to his team is the difference in performance above a bench player. The direct conclusion of such a situation is that bench players would all be interchangeable and have negligible value, while starting players would have value to teams according to their performance with respect to bench players. Therefore, a team's valuation function for players would look just like Figure 4. An effective point is a measure of points scored multiplied by percent started, or a measure of value to a team. Thus, if a player is not started (i.e. is a bench player) he will not be able to earn any points. We can see in the figure that the value provided to a team by a bench player is zero¹⁰ and then value increases among starting players with increased performance. One would expect that the bids in an auction for players would see a similar distinct break between starters and bench players.

¹⁰ The slope of the OLS regression of effective points on points among bench players is not statistically significant.

To test risk-aversion, I regressed fantasy price on points scored for starters and bench players, and studied the difference in slope of the OLS regression between the two pools. Figure 5 and Figure 6 show that there is no difference in price with performance among fantasy players regardless of whether the player will be a starter or bench player. Although without data to track specific leagues, one cannot completely rule out the possibility of teams hoarding top players on a micro scale. However, given the strength of the pattern of fantasy prices on bench players in 2008 and 2009 as well as the pattern of starting players given earlier, it is reasonable to state that it is extremely unlikely fantasy football owners bid more on bench players than would follow if their only motive was to increase the probability of winning. In fact, one would expect a distinctly piecewise function based on the effective points a player could score for his team. Such a function for price would be flat for bench players at the lowest possible price – in this case \$1 – and then have positive slope for starting players.

This type of function – flat for bench players and increasing for starting players – is not the one seen in practice. There are several reasons that slope would be positive for bench players in practice even though it would be zero in theory even without risk-averse preferences. A possible explanation for an overpaid bench player is a misperception regarding future contributions, although such a misperception problem is unlikely to be systematic because there is very little information asymmetry in fantasy football. However, the most obvious reason is that there is always a risk of a starting player getting injured or performing poorly for some other reason, and having a good backup to replace such a player hedges risk and lowers the standard deviation of point

total. Deciding whether the equality of slopes is justified by these reasons can give insight into fantasy owners' risk preferences.

Risk to players in the NFL is dependent on two major aspects: likelihood of injuries and performance decrease. Likelihood of injury is simply games played by viable players over total games available. This may underestimate injury risk as some players develop injuries before the season begins, making them nonviable despite being drafted. On the other hand performance decrease, which is calculated by comparing performance from one season to the next, may easily overestimate risk. Assuming the previous season's output is a reasonable predictor of performance in the next season, by calculating the percentage of players that see any drop in performance from the previous season, we may find the risk of underperforming during a fluky season. That value fails to include players outperforming their expectations and improving, thus this measure overestimates risk. This overestimation helps to compensate for underestimation from injuries.

Using empirical data, I found that the average risk is quite small, only applying to 10% of all players (Figure 9 Chakravarthy, 2010). Meanwhile, fantasy owners appear to drastically overpay bench players, thinking that such players were likely to contribute to the team despite the low chance they are needed (Figure 10 Chakravarthy, 2010). By assuming that the goal of each owner is to win his league, I could conclude that owners were overestimating the danger of a sudden drop in performance. When a team corrects for that overestimated risk by lowering the degree of induced risk-aversion, it is

able to allocate more money to starting players and therefore win the auction for top free agents. This allows a more risk-neutral team to collect better starting players at the expense of bench players, although such a team would also take on more risk. Whether an owner's induced risk preferences would be risk-averse depends wholly on whether he would want to win his league or simply avoid finishing near the bottom. An owner with a great desire to win his league would have a risk-loving preference because taking on higher risk gives him a greater average mean performance. Continuing in the same vein and defining a threshold for winning a league,¹¹ I was able to generate a direct relationship between player performance and cost. In the NFL, that type of relationship is far more difficult to do because teams cannot merely set a threshold for amount of talent required to win the Super Bowl. However, the results from the previous study are not without usefulness in considering live football.

Fantasy owners often appear induced to have preferences that are very risk-averse, as can be seen by the coefficient between fantasy auction prices on fantasy points. A player is worth his expected total contribution to a team, discounted by the probability he does not contribute to that potential – the risk of that player getting injured. The regression showed that the discount rate of starting players is nearly 40%, four times higher than the empirical value I found for risk (Figure 10 Chakravarthy, 2010). By testing player valuations with a lower discount rate in a simulated auction, I found that such a valuation lends itself to a draft with a mean higher expected fantasy

¹¹ Because I had access to the distribution of fantasy team total scores over a season, I was able to construct a distribution of scores. Teams need to score in the top 10% of their league to win – generating an average threshold of scores.

points, implying a higher expected utility. However, low risk-aversion also increases the standard deviation of expected point total, which is detrimental to utility, suggesting that risk-aversion may not be irrational; rather, it is induced by the preferences faced by the individual. It also appears that having strong positional influences, as well as a value based drafting strategy is the best way to obtain a team with high output. The impact of this study in professional football seems to be ambiguous, but in reality has direct, and intuitively apparent, implications on how to approach the free agent markets.

This study has some limitations but found clear evidence that fantasy football owners were overestimating the probability that players under performed as a result of injury or simply a fluky season. This could be considered induced risk-aversion and it may be rational – owners do not want to lose their leagues and would take a lower mean expected point total in exchange for low standard deviation. There are multiple explanations for this with direct NFL impact. The first is that owners do not always use empirical data when evaluating players. Instead, they rely on other aspects of the player and get utility from other aspects not directly related to winning. Popularity would be the most obvious factor in fantasy football, but in the NFL it could be any form of non-statistical benefit – leadership, teamwork, or other form of self-sacrifice. This explanation is difficult to prove, but runs counter to expectations by fans of the NFL and if systematically evident might spark similar protests to the one in Cincinnati.

Other value gained from this study lies in how best to evaluate players. Understanding the positions that need the most help will provide teams with another

factor with which to weigh skills and determine how best to evaluate players. The indirect effects of a team spending more money on certain players could also generate positive externalities – some veterans may join successful teams in the hopes of winning a championship before retirement. Such an externality was observed most recently in the New England Patriots, who were able to sign Randy Moss and Junior Seau in no small part because of a history of success and a high probability of future championships. Being willing to spend a higher allocation of the salary cap on starting players instead of bench players and have less risk-averse preference could also contribute to helping a team generate the level of talent necessary to be successful.

Theory and Testable Hypotheses

This analysis tests two questions in the NFL: whether teams allocate salary purely based on expectations of future performance and whether NFL Owners exhibit the same type of induced risk-averse preferences shown in their fantasy football peers. The allocation of salary has not been plumbed by research and concerns three major aspects: whether salary is predictable by past performance, whether performance in the past can predict performance in the future, and how salary and performance weight the past. Answering these questions can shed light on the way teams evaluate their players and decide to allocate salary to different players. Such an analysis could even provide ways for teams interested in maximizing on-field talent under the salary cap to do so in new ways.

Induced risk-aversion in the NFL is hard to gauge because it is impacted by a pair of competing effects. The salary cap limits the amount an owner may spend on his players, just as in fantasy football. This would lead such an owner to avoid committing a huge amount to one player. However, all NFL contracts have the unique salary escape clause – cutting players. A team can cut a player before the season and avoid having to pay that salary beyond the guaranteed amount. Even that amount could be waived by a team if a player is given a contract by another team. This ability to release players gives a general manager or owner more flexibility in signing contracts. The effects of being able to remove a contract from the books and being forced to spend within the limits of a cap, make it unclear which effect will be more significant and whether NFL Owners have specific risk preferences induced by their market.

Problem 1: Player Salary and Salary Growth

Earlier, I mentioned the value of understanding why players get paid in a certain manner – how payment is aligned with expectations for the future. If a player is paid exactly according to how he is predicted to perform in the future, clearly the team is willing to take the risk of some injury or weak performance. If, on the other hand, teams are paying different players differently with regard to future expected performance, I hypothesize it would indicate a fear of injury or some other motive in allocating cap space as discussed earlier. Breaking down the relationship between pay and expected future performance can lend insight into how teams view the conflicting impacts of the salary cap and ability to cut players. Regressing past performance on current salary can define that relationship clearly and succinctly.

The question being considered is: What are the best predictors of salary and are they consistent over time? The question must be considered using OLS regressions of previous season performance metrics on future cap value. This would be a series of one year lagged regressions in the form:

$$\text{Salary} = a + B_S * (\text{Previous Year's Performance}) \quad 1$$

$$H_0: B_S^{2008} = B_S^{2007} = B_S^{2006} = B_S^{2005} \quad 2$$

The null hypothesis is that there are certain strong predictors of salary and they have similar coefficients in all the years considered by step down regressions. Alternatively, if salaries do not have consistently strong predictors, as seen in those step down regressions, salary is determined by randomly changing characteristics or, more likely,

salary is not determined by on-field characteristics and the relationship illustrated by the regressions has little significance.

Problem II: Player Performance and Performance Growth

The valuation of players can help to understand team utility. If a team gives players a high salary for reasons that cannot be explained on the field, that team clearly has a different utility function than a team that pays players in direct accordance with expected future performance. Unfortunately, parsing the difference between reasons for differing utility functions is harder than it looks. The difference between systems of the two teams may have a huge impact on salaries offered. Football is such a system specific sport that small disruptions between systems may render a player highly ineffective, as was seen with Chad Ochocinco during the 2012 season. It is also possible that a team cares a great deal about merchandising and signs players based on expected revenue, which would have an additional impact on salary and may not be directly visible in data analysis.

The question being considered is: which aspects of player performance do the most effective job of predicting future performance and whether those aspects are consistent. This is similar to the earlier question asked regarding salary, but this time focuses on whether it is possible to predict performance. Implied by the question is the comparison of these results with the results from the previous problem to consider whether predictions of salary align with predictors of performance. The question will be considered by a series of step down OLS regressions of previous performance metrics on future fantasy point totals.

$$\text{Performance} = a + B_p * (\text{Previous Year's Performance}) \quad 3$$

$$H_0: B_p^{2008} = B_p^{2007} = B_p^{2006} = B_p^{2005} \quad 4$$

The null hypothesis is that the predictors of performance are consistent. Alternatively, predictors of performance could be inconsistent, meaning teams have no good on-field performance method for predicting future performance.

This conclusion, in conjunction with the previous one can shed light on the movement of salary and whether it is in accordance with changes in performance and the most important question of Problems I and II.

$$\text{Salary} = a + B_S * (\text{Previous Year's Performance}) \quad 5$$

$$\text{Performance} = a + B_P * (\text{Previous Year's Performance}) \quad 6$$

$$H_0: B_P = B_S \quad 7$$

If we reject the null hypothesis we imply that teams provide salary not according to future performance but rather due to some other characteristics that cannot be measured by on-field statistics. Such a response would support the conclusion that salary is not given in accordance with expected future performance, but according to some other preferences. If the null hypothesis holds, however, salary is allocated with regard to future performance. The null hypothesis is complicated by the questions previously discussed – whether B_S is constant for predicting future performance and B_P constant for future salary. If those hypotheses do not hold, then we cannot accept the null hypothesis and are forced to infer that salary is not given in accordance with expected future performance, a crucially important conclusion in trying to define Owner preferences.

Problem III: Lagged Values

The performance –salary relationship can also be defined by the relative weights of the past. Studying lagged values will be an important portion of this study because it shows how past performance dictates future performance and future payment. The relative weights of past performance can lend insight into the reasons behind paying players. For example, if performance in year three is mostly strongly predicted by a 2:1 ratio of performance in year two to performance in year one, one would expect the same ratio of performance in years two and one to predict salary in year three. If that ratio changes drastically, it could indicate the motives behind paying players – whether it is in direct accordance with future on-field performance, or whether it is a result of something not moving with expected future performance. Paying players for previous work that is unlikely to be reprised suggests that Owners value characteristics that are not on-field performance, like leadership or merchandising revenue.

Continued testing of this data can provide answers to major questions regarding NFL salaries. Understanding how players are paid and whether that is different from the way players are evaluated can give information regarding the way teams handle contract signings. Fantasy football is an underappreciated tool in NFL circles, but could provide a way for a team trying to maximize talent on the field and expected future production while staying under the NFL’s hard salary cap. In fantasy football, owners have a huge pool of players to choose from,¹² meaning they can afford to be risk-averse

¹² An average fantasy football roster includes fewer than twelve players at six positions. Fantasy owners can choose those players from all 32 NFL teams’ 53 player rosters. Thus, a ten team league selects 120 of the nearly 1700 players available

and hope for the possibility of a player with low expectations outperforming those expectations. In the NFL, however, the ability to find players who cannot compete is rare, and quite noteworthy. Victor Cruz of the New York Giants is a striking example joining the league as a relative unknown quantity and succeeding. Given the number of undrafted players, it is rare to find such a talent and groom him to compete on a high level. And even after finding that player, he must be paid in accordance with his performance in future years, so holding him for the duration of his career at the same low rate is effectively impossible. Thus, we see the selective pressures of the NFL and fantasy football – in the NFL, being highly risk-averse is a far more difficult strategy to justify because of the scarcity of top tier players makes acquiring those players more important and the ability to cut a contract allows for an immediate out if the contract signing was too generous from a team’s perspective.

The question of the lagged values of performance as predictors of future performance and future salary can clarify the motivations of allocating salary and may help define the way players are paid. By regressing lagged performance values on future salary and future performance, we may compare the relative weights of the past performance by observing the ratio between coefficients.

$$Salary = a + B_S^1 * (Previous\ Year's\ Performance) + B_S^2 * (Performance\ Two\ Years\ Previous) \quad 8$$

Performance

9

$$= a + B_p^1 * (\text{Previous Year's Performance}) + B_p^2 * (\text{Performance Two Years Previous})$$

$$H_0: \frac{B_S^1}{B_S^2} = \frac{B_P^1}{B_P^2} \quad 10$$

The null hypothesis is that the relative weights on the past are the same, as determined by the series of OLS regressions, meaning that teams allocate salary based on their expectations for future performance. Alternatively, if the weights are different for future salary and future performance, it would indicate that future salary is not allocated according to expected future performance, implying other reasons for allocating future salaries, perhaps including future potential as a star or characteristics important to the team like teamwork. It may seem more straightforward to compare B_S^1 directly to B_P^1 , and vice versa, but such a comparison is drastically complicated by the magnitude of each coefficient. Such a test would allow us insight into how exactly this ratio diverges, but does not account for relative weights. Rather than seeing the exact magnitude given to each year, the more important question is whether the weights on past years are equal, because that would illustrate a consistency in salary allocation in accordance with performance growth over time.

Problem IV: Risk Preferences and Cutting Players

Studying the risk preferences of NFL teams is an important question and it meshes with my previous work. The question of whether NFL teams display induced risk-aversion is of two parts, and it relies on the fantasy pricing analysis before. A piecewise OLS regression of performance on fantasy price, split at starting and backup players illustrates the risk preference of fantasy team owners. Meanwhile, the same piecewise regression of performance on NFL cap value separated by starting and backup players illustrates any risk-aversion of NFL Owners. My previous research showed that there is little difference between the starters and bench players in fantasy price and that the piecewise regressions have little change in slope between the two pools of players.¹³ This is the case despite expected performance for a fantasy team. Although most of the bench players will earn the team few, if any, points, they are paid as though they would, seemingly indicating very risk-averse preferences by fantasy owners who seem to assume that those players would be necessary. Unfortunately, NFL backups rarely earn any statistics at all because, unlike their fantasy counterparts, they do not often take the field. This is minor problem with measuring the value of backups, but assuming that weighting of certain statistics among backups are not different from the corresponding weights of those statistics among starters, the problem of limited observability is a minor one.

The distinction of NFL players between first-string and second-string is a complex one, because viewing the problem after a season presents a much different picture than

¹³ Figure 13 illustrates this point.

before the season. There is an argument to be made that teams may think that one player is a starter only to find out halfway through the season that another player is better. If this were the case, it would overestimate any overpayment of low performing players. However, this study shows, particularly in Problems II and III, that there are many strong indicators for future performance, to which all teams would have access. Thus, the fact that teams could statistically infer future production mitigates the ex post bias caused by defining starters and bench players after the season is complete.

The strength of this analysis requires understanding the benefit of target as a proxy for a team's faith in the quality of a player combined with a measure of talent. Better players, starters, would receive more targets than bench players. For this case, the definition of starter will be a player above the median in targets – bench players are those who are below the mean in targets. In the Results section Problem II, a breakdown of the validity of using targets as a proxy for talent will be described in full, so splitting the group of wide receivers into pools of starters and bench players by their position in targets follows.

$$Cap\ Hit_{Starter} = a + B_{Starter} * (Performance) \quad 11$$

$$Cap\ Hit_{Bench} = a + B_{Bench} * (Performance) \quad 12$$

$$H_0: B_{Starter} = B_{Bench} \quad 13$$

The null hypothesis is that the piecewise regression of starters and bench players in performance on NFL cap value will exhibit the same pattern as in fantasy football – very similar slopes, indicating the same sort of risk preference as among fantasy owners. The

alternative hypothesis is that the slopes are different in the piecewise regression, suggesting that the risk preference pattern in the NFL is distinctly different than that of the fantasy market. If the different slopes consist of a lower slope in among backups and a higher slope among starters, it would most likely indicate that the NFL player market induces less risk-averse preference because the scarcity of top level players in the NFL limits the ability for teams to be risk-averse in searching for top tier players. Alternatively, if slope is higher among backups, it indicates even greater risk-averse preferences than in fantasy football because teams are anxious to secure backup players.

While understanding the motives of NFL owners may be difficult, the possibility of finding a salary allocation system with the goal of winning games could still be extremely valuable and there are lessons available in fantasy football. The system to be considered later in this paper is one of shorter term contracts with higher, incentivized payment structures and greater reliance on the team's ability to cut players. Such a system would result in a radical change in the way contract negotiations are held, but if it were done across the league rather than just by one or two teams it could lead to a more open market for players. There may, however, be large scale limitations for teams treating players as single-year performers. Those limitations would include the reputational effects of being a team accused of not caring about players. Such a label might stigmatize a team, and could be so strong that the way salary allocation and contract length is handled now might be risk-neutral.

Methods

This analysis is separated into several steps designed to generate a proper understanding of why NFL players are paid in given ways. In order to understand which characteristics of players are most important to generating salary, one must first consider a player's contribution. Leeds and Kowalewski (2001) discuss the 1993 Collective Bargaining Agreement and its impact on skill position players. The term "skill position" refers to the players who will touch the ball most on a team. Quarterbacks, wide receivers, running backs and tight ends are all considered skill players. The methodological focus on the skill players is the result of statistics regarding skill player performance being widely available and is a tactic I must replicate. Wide receivers are among the most tracked players, with a variety of statistical metrics for their performance, which is the reason this study is looking mostly at that position. Leeds and Kowalewski (2001) also focus on players at the ends of the wage spectrum, meaning that generating a model for performance at any wage may require a piecewise breakdown of players based on wage or skill.

I am studying wide receivers because they have a plethora of statistics and despite the variety of routes they run, they face a single event model of performance. This means that although wide receivers face many targets every game, each of these targets is roughly independent, particularly when considered over an entire season. A "target" refers to a pass in a player's direction. This serves to regulate an individual's production and differentiate between players that may seem to be equally productive but have different degrees of opportunity. For example, if player A catches four passes

for forty yards and player B catches four passes for forty yards, they may appear to have been equally effective. However, if player A receives six targets and B receives ten, A is the more efficient player, accomplishing the same amount of team benefit in fewer opportunities. Despite the benefit of using points per target as a variable, targets are endogenous with regard to talent level – a player who receives fewer targets most likely has less ability than his counterparts. Other statistics that are used to determine the value of a wide receiver are yards per catch, which place a premium on “big play ability,” an important characteristic for many teams, particularly those who often run instead of passing on offense. Another benefit for the use of wide receivers is that they rarely have to make crucial decisions, unlike a quarterback or a middle line backer, who serves as a coach on the field. This means that the amount of non-measurable abilities of such a player is limited. On the other hand, certain players are associated with qualities that are important to teams and such an association gives that player extra weight in salary negotiations. The converse is also true-players seen as a “cancer” to their team are less valued and tend to make less money than their peers who are seen as less egotistical.¹⁴

First, to detect which aspects of player performance are most rewarded, I must discover which metrics are most important in determining pay scale. This step involves the Thaler and Massey technique (2005)of using the simplest indicators of performance

¹⁴ There are many famous examples of this type of player, but Randy Moss is among the most significant. A Hall of Fame level player is rarely traded or cut because of the irreproducible performance he brings and the associated interest from fans. Moss, however, was released three separate times in 2010 alone and is currently planning a comeback to the NFL that has been viewed negatively because he is such a divisive character. It is currently unknown whether he will be signed by any teams.
http://espn.go.com/nfl/story/_/id/7616262/report-randy-moss-willing-sign-deal-guarantee-money

to generate an idea of what players are paid for doing. For example, is age a larger factor in salary than yards? If so, which way does that effect trend? To do this, I must use a variety of factors to predict the salary cap portion of contract benefits and determine which is most important. In identifying this, I can then determine the relative validity of these explanatory variables. For example, if touchdowns are a strong predictor of wide receiver pay grades, and the distribution of touchdowns is somewhat stochastic when compared to yards or receptions, knowing that salaries can be based on inconsistent variables is important and suggests that owners may pay players for things other than expected performance on the field. Alternatively, if I find that pay rate is linked closest to a factor without direct impact on production, like age, that may also prove illuminating. Because salaries often do not change by a large factor, measuring whether or not a player experiences radical changes in pay is important. If a player experiences large changes in pay, what could have caused such a change? Which types of performances lead to changes in cap value? How do the raw statistics generate future salary growth? This type of analysis can be fruitful in studying the creation of new contracts, but can be plagued by small sample size since signing new contracts is rare.

Another major input to study is salary growth. Salary, or cap value in particular, can change depending on a variety of factors. What is most significant in generating increases in salary? Age drives up salary because of the NFL's CBA, but if that is the only significant predictor of salary growth, the NFL allocation system may be weak because of an inability to pay better players more money. Similar to the previous portion, lagged values can be very important. How do owners treat the past when allocating salary

changes? To understand this, I must study current salary in concert with a similar variety of factors, but from previous NFL seasons. This analysis would reveal if salaries are most influenced by recent performance, or if they are influenced more by performance from many years past. Studying the way salary allocation is affected by past performance is not valuable without a point of comparison. Performance is also predicted by previous performance, and if future performance is predicted by past performance in the same manner that future salary is predicted by past performance, we can infer that salaries are allotted with an eye to future performance. On the other hand, if the relative weights of the past differ between predicting salary and performance, understanding the difference between salary allocation and what owners “should” be could be illuminating. Unfortunately, the tendency of most NFL contracts to be multiple years in length complicates this analysis, but ostensibly teams would be trying to compensate players in accordance with their expected performance in a given year.

To analyze production, this study makes use of the fantasy football point scheme to determine a player’s value, which could also be called fantasy points:

$$Fantasy\ Value = 6 * Touchdowns + \frac{Yards}{10} - Fumbles \quad 14$$

This formula is a measure of player production and is used in fantasy football to determine how many points a player earns in a given game. Better players score more points, worse players score fewer. This formula breaks performance into multiple pieces – scoring touchdowns, gaining yards, and maintaining possession of the football for the offensive team. In a football game, scoring a touchdown earns a team six points, which

is the reason fantasy football players get six points per touchdown scored. The yardage coefficient is based on the average NFL drive, which gains approximately twenty yards and scores approximately two points (J. Armstrong, 2011). Thus, covering one hundred yards is worth approximately ten points, as shown by the formula above. Similarly, about half of fumbles are lost, and a lost fumble ends a drive, corresponding to the lost opportunity of two points. It's impossible to know which fumbles will cause a change in possession, but on average, a fumble is a negative point for the team.

The reason behind using this type of performance metric is that allows for far greater ease in OLS regressions. By using one response variable, the regression will be far easier to analyze. Unfortunately, fantasy points are not normally distributed, but otherwise it serves as an excellent statistic for academic use. The fact that it is right skew supports my earlier statement regarding the scarcity of top level players. This scarcity explains why top end players make so much money. There are a variety of other metrics that can be used to measure wide receiver performance, like Defense-adjusted Yards Above Replacement (DYAR) and Defense-adjusted Value Over Average (DVOA). The problem with those metrics is that they require play by play analysis for each player being studied, data currently unavailable for every year necessary. In any case, these statistics are highly correlated with each other and with fantasy points, meaning that all three are approximately the same metric of success (S. Walder, 2011). To confirm those findings, I computed the correlation between DYAR and fantasy points (.856) and the correlation between DVOA and fantasy points (.839). Those correlations are sufficiently

large to suggest my analysis would provide similar qualitative, if not necessarily quantitative, results with any of the three metrics.

Similarly to the analysis of payment, analyzing performance growth can also be revealing. Generally, one would expect performance growth to be highest in the first years of a player's career, before he enters his prime. If performance growth and salary growth move together, we may conclude that players are paid in a direct relationship to performance. The growth variables are more complex to analyze because salary growth is hampered by the fact that many NFL contracts are multiple years in length. This presents a caveat on any conclusions gleaned from an analysis on salary growth because teams may be trying to predict performance very far into the future, when extrapolations may grow increasingly hazy. Drawing too heavily on these inferences, then, would lead to an overestimation of any difference between future expected performance and future salary. The benefit of using data ending in 2009, though, allows for corrections in the case of teams restructuring contracts and showing the restructured value rather than the initial value, which may have been a shortsighted decision, and lowers the effect of huge extrapolations.¹⁵ This limitation aside, looking at growth between contracts as well as the effects of cutting a player may provide more information on the decisions made by NFL teams when a player is up for a new contract. Generating these dummy variables relies on finding players who see a very large jump in expected cap value or who did not retire but were not on a team's roster during a

¹⁵ Teams may sign contracts including the possibility of dead money, in which a player is guaranteed money and then cut. The benefit of the older dataset is that such money would be included in a player's cap hit for the season immediately before being cut, so the cost to the team of that player is not lost.

season. These patterns would suggest that a player signed a new contract or had been cut. Studying the effects of performance on new contracts and cuts can help define a prescriptive approach

The remainder of this work lies in identifying the pattern of relationship between player skill and player salary which could reveal the risk preferences induced by the NFL's salary cap. Defining such risk preferences may enable the formation of a direct, prescriptive policy recommendation. Identifying risk-aversion seems straightforward – a more skilled player should be paid more than a less skilled player. Similarly, a player with high expectations of future growth should make more money than a peer with a lower potential for growth. If these relationships follow non linear patterns, that may illuminate our understanding of spending patterns and induced risk preferences. If, for example, a player with a very high ceiling is paid more than a player currently outperforming him but with a lower ceiling, it could suggest that owners and general managers are risk loving – placing greater value on the possibility of high future performance than on the certainty of moderate current performance. Similarly, the inverse could be true. In any case, a focused analysis of salary and performance factors is necessary before making any conclusions and forming any recommendations.

Data

The salary data for this study is coming largely from USA Today's NFL Salaries Database. It is the most accessible data source for National Football League Salaries and has information as recent as 2009. This work uses the five most recent seasons of NFL salary data available, from 2005 to 2009, to ensure that all the salary figures used are from the same CBA, which was signed prior to the 2005 NFL season. USA Today provides the yearly base salary, signing bonus, cap figure, and total payment including bonus and salary for each player in the NFL in dollar figures. Cap figure refers to the amount that counts against a team's total salary limit. Because teams are forced to allocate salaries to stay within a certain spending budget and the owners are very wealthy, the cap allocation is the most important aspect of a player's income.

Thus, if a team is paying a player \$11 million, but the cap value is only \$6 million, the latter figure is of more interest since the remaining \$5 million is coming directly from the owner's pocket and the marginal cost of that remaining payment is essentially inconsequential. Moreover, much of that \$5 million figure will not be paid to the player as it includes performance incentive bonuses that are extremely unlikely to be realized.¹⁶ USA Today collects information for this database from the NFL Players Association, player agents, and their own research. The database uses terms from contracts completed at the beginning of a season. Limitations on salary like the rookie cap may account for some salary inefficiencies, but dummy variables for the rookie

¹⁶ In the remainder of this article I will use the terms "salary" and "cap value" interchangeably as salary and cap figure are identical when bonuses are discounted. When referring to the total payment above, I will use "total salary."

contract are not available. Rarely, a player may extend or sign a contract during the season. Such a contract would not be included in this database until the next season. If a team changes the way the money is distributed in a contract, such a change would only be included in data for the next season, and cap value would account for any such change. The ability of teams to restructure contracts would theoretically improve the ability of cap value to move along with expected performance and underestimate any differences in expected future performance and future salary.

Pro Football Reference is the source for raw on-field data from 2005 to 2009. This website gets data from the ESPN Football Encyclopedia and tracks stats using the ESPN statistics center. ESPN also provides the data for fantasy football auction prices and usage. For wide receivers Pro Football Reference provides most of the data to be studied and includes data on every wide receiver on an NFL roster at any point during a season. If a player played in one season, but not in the next, he is not included for multiyear growth statistics, but a player would be included for single-year statistics, and all regressions using only single-year statistics. This analysis also makes use of data from Fantasy Football Today, which provides the variable “targets.” Fantasy Football Today gives strong target information only for the 2007, 2008 and 2009 seasons. Before those seasons, target information is not available. In an effort to consider off the field effects, like player popularity with the general public, I contacted the NFL shop – the official seller of NFL merchandising – in order to find the ranking system for NFL jersey sales. The shop, however, does not release that data so including external variables such as

popularity is currently impossible. Finding the relationship between popularity and pay would also be helpful in defining an owner's utility function but is a current limitation.

Results

The first important procedural effect was correction of the variables included. Many variables were right skew, including the salary cap figure (“cap value”) and player fantasy points. Fantasy points¹⁷ were created according to the formula given above and shown again here:

$$\textit{Fantasy Points} = 6 * \textit{Touchdowns} + \frac{\textit{Yards}}{10} - \textit{Fumbles} \quad 14$$

This formula is the central characteristic to all portions of the results that deal with player value and merits a discussion of strengths and weaknesses. By combining many aspects of performance into a single number, it allows for a great deal of ease in regressions. It attempts to be somewhat intuitive by converting all performance into expected points added to a team. Moreover, it is a commonly used system with the continued growth of fantasy football, as described earlier. However, such a simplification remains just that – a simplification. It ignores some statistics that may be important, like targets. This formula cannot consider the ability of a player to serve as a leader or be a good teammate, though comprehension of personality cannot be had without extensive interviews. In short, this metric is very useful, but does have some flaws due to a lack of detail.

The target is a valuable explanatory variable because it can serve as a measure of a team’s faith in a certain player. A target is a ball thrown in a receiver’s direction. Better receivers usually have more targets because the team wants to utilize the talent

¹⁷ Value and points are interchangeable terms in the formula above.

of that receiver. Usually, receivers with the most targets have the most production on the field because they are the best receivers on their team, lending some endogeneity to the comparison between targets and performance. However, because targets measure the team's faith in a player, they are a useful method for dividing players between first and second-string, as in Part IV. Figure 7 illustrates the validity of that statement. The strength of this relationship decreases when targets are used as a predictive variable for value, but the relationship is still clearly positive. Fantasy points in previous seasons may also predict targets – Figure 8 is an example of that relationship. Intuitively, this is reasonable – the better a player is in the previous season, the more likely he is to be utilized by his team in the future. These graphs show the strength of targets as a proxy for fantasy points and how similar the two statistics are. Because a player who has more faith will get more targets, this fact is useful in the comparison between a player's fantasy points and his salary, as targets can be used as a metric for worth in the eyes of a coaching staff.

Part I: Player Salary and Salary Growth

A regression of previous performance indicators on salary shows predictors of player cap value appear to be widely scattered dependent on the year. The variables shown as strong predictors ($p < .05$) of the natural log of cap value appears in Table 1. This is a series of regressions for the four years between 2006 and 2009. Each of the variables refers to the variable in the year before the regression. Thus, in the 2006 coefficients, the games played variable refers to games played in 2005.

Table 1¹⁸¹⁹ illustrates the general inconsistency in determining which variables are significant in determining salary. Over the four regressions above, age is the only variable which is consistently significant although each year produces a fairly high R-squared term, forcing us to reject the null hypothesis and accept the idea that salary is not able to be predicted directly from the previous year's on-field performance since the coefficient on Age does not stay within a 95% confidence interval overtime. Despite the lack of precision of comparing coefficients from different regressions, the multicollinearity of age prevents a stronger test, and the difference in magnitude of the coefficients seems to preclude the necessity of a stronger test. In order to test the B_S in Part II of the results, one needs to calculate the average value of each predictor to compare to the average values for B_P . Below is the sum of the average coefficient for each of these variables for the four year period, illustrating the relative strength of each:

¹⁸ This table, and many others, show different variables for each regression. That is because the only variables shown in each regression are the ones significant at $p = .05$. To see the complete regression, including statistically insignificant variables, see Table 2.

¹⁹ The reason for listing Fantasy Points as a square root is to correct the right skew.

$$E(\text{Salary Cap Figure}(t))$$

15

$$= e^{.08 * \text{Age}(t-1) - .03 * \text{Games Played}(t-1) + .01 * \text{Receptions}(t-1) + .093 * \sqrt{\text{Value}(t-1)}}$$

Where t is the season in question and $t-1$ is the previous season. By breaking down this equation, we may see even small changes in each of these variables could have a major effect on the estimated salary cap value for a player because of the exponential nature of the function above. Given the specification with log of salary cap value as the dependent variable, the positive coefficient on performance indicates *increased* marginal return to improvements in performance.²⁰ This type of increased marginal return suggests the benefits of being a “star.” Star power is a focus of all athletics because the financial benefits of having one of the few top players at a position. The analogy used by sportswriters, “never trade a dollar for four quarters,” is an apt description of the issues faced by teams trying to acquire star performers. Thus, a player moving from being simply good to being elite has huge repercussions for his team that go beyond helping his team to win games, supporting the exponential equation above.

While Tables 1 and 2 illustrate the inconsistencies in finding strong predictors for salary in a given season, using salary growth as a proxy for salary does not illuminate the situation. Measuring the salary cap value in year two as a fraction of that hit in year one is one way to measure salary growth. Running the same stepwise regressions for salary growth instead of salary reiterates the lack of consistent predictors for salary. Tables 3 and 4 illustrate several striking results, namely that on-field performance metrics do a

²⁰ This statement is supported by the results of the regression in Table 1a, showing the coefficient on the square root term to be far less significant than the coefficient on the squared term, $p < 0.001$

poor job of predicting salary growth in the future and that they have a wide range of variables that may have been expected to be important are not consistently statistically significant. This inconsistency reinforces rejection of the null hypothesis that future salary and expected performance growth move together and may be a result of the fact that most players sign multiyear contracts that are held for extended periods of time despite changes in player performance.

However, combating the effects of multiyear contracts is difficult. When limiting focus to those players who experience large changes in salary cap values – players who signed new contracts, sample size drastically drops, reinforcing the statements above that predicting future salary is imperfect and that players often sign multi-year contracts instead of short term contracts that would require them to sign new contracts often. For example, there were no significant predictors for salary growth in 2006 or 2008. In 2007 and 2009 the sample size was too low to generate any real conclusions (N=34, 17 respectively). Finding weak regressions, however, does give us insight into the way teams pay players. Teams rarely give players large changes in salary, even when there are large changes in performance. However when a change does occur, the reasons for it differ across individuals. There may be many explanations for this, but the most likely one might be the opportunity cost of changing a player's salary. Terminating a player's contract and renegotiating a salary is certainly not easy and could be a point of frustration for everyone involved. This might be a driving force behind the limited number of contract changes seen.

Part II: Player Value

Predictors of player performance, unlike salary, remain relatively consistent, though these regressions have weaker R-squared values. Table 5, 6 and 7 both show the issues with predicting future salary – although none of the explanatory variables in Table 5 appear in all four regressions it is clear that fantasy points per game is a significant predictor of future performance because the magnitude is relatively constant. In order to look for stronger predictors, Table 7 is the same analysis as above, but extended into the past to find better predictors and over two years instead of one. Thus we can see that because B_p is within a 95% confidence interval over each year tested for fantasy points per game we cannot reject the null hypothesis that B_p is constant across time.

To consider the other question in this portion of the analysis, we must compare B_p and B_s . The average coefficients are shown in the equation below:

$$\begin{aligned} E(\text{Performance}(t)) & & 16 \\ & = [.465 * \text{Fantasy Points per Game}(t - 1) + .016 \\ & * \text{Receptions}(t - 1)]^2 \end{aligned}$$

This allows us to reject the null hypothesis that $B_s = B_p$, which implies future salary grows in direct compensation for expected future performance, because the consistent predictors are different from the inconsistent predictors of salary. Rejecting the null hypothesis complicates the ability to generate a conclusion regarding the exact motives of allocating salary, but does elucidate the point that salary does not move with increases in player value.

Table 6 illustrates the difficulty of predicting future performance. Although R-squared terms appear reasonably high and the variable fantasy points per game is a strong predictor, the other explanatory variables are not significant. Appearing in more games has a negative correlation with future performance, as do age and number of games started. It is unclear what exactly caused this effect, but it may have been the result of declining performance after years of performing in one's prime. Thus, aging seems to be more significant as a factor in lost ability than in gained experience. This is an interesting finding that could be explored in the future by another study.

Another method of measuring the strength of a wide receiver is considering his fantasy points per target. A high point per target figure means that a receiver makes the most of his opportunities, while a low point per target figure suggests weaker production when corrected for opportunities. This would seem to be a strong explanatory variable for performance because presumably a player with high marginal returns per opportunity should see more opportunities in the future. Unfortunately, using points per target as a predictive tool is largely ineffective. In Figure 9, one can see that points per target isn't nearly as useful as a proxy for fantasy points as number of targets, which is strongly correlated both with current production and future production. This is a major disconnection between points per target and its impact on performance and on payment. Figure 10 shows a very strong positive correlation between points per target and future performance, but Figure 11 shows that points per target isn't nearly as strong an explanatory variable of future salary. This effect may be

confounded by defenses focusing on better receivers, making their points per target lower than their less skilled peers.

Another way to think about predicting performance is by considering performance growth. What predicts sharp increases in performance? Tables 8 and 9 show that in the years 2007, 2008 and 2009, changes in performance are very much inconclusive with low R-squared values, despite the large sample size. There is a sample bias in those regressions since performance growth will be overestimated. Players cut after a season are not included and only improving or high level players would be part of the dataset, so the individuals experiencing negative performance growth are less likely to remain in the league and would not be in the dataset. Surprisingly, each step down regression reported that the games played variable was significant with a negative coefficient. This suggests that the number of games played is important and not just because it is negatively correlated with a statistic evaluated on a per game basis.²¹ That is, players who play a surprising amount of games in one year may do so as a fluke and see their production drop back down to the norm as a regression to the mean.

²¹ Fantasy points per game also have a negative coefficient, suggesting that there is something beyond simple confusion of explanatory variables.

Part III: Lagged Values

Table 10 includes the results of two time-lagged regressions to predict performance. Each used the three previous seasons' values as predictors. Both step down regressions eliminated the most distant year and came up with similar ratios of years into the past to predict future performance. This shows that predicting performance based on the past can use a somewhat consistent model. Table 11 shows that predicting salary using lagged values from the previous three seasons generates similar results, in that the third season past is immediately removed from the regression as salary may be predicted using the two most recent seasons only. While these results do only include the two previous seasons, the relationship in weighting the two previous seasons is not only inconsistent; it is markedly different from the relationship between weighting the two previous seasons seen earlier.

$$E (\text{Performance}) \qquad \qquad \qquad 17$$

$$= 8.585 * \sqrt{\text{Previous Year's Fantasy Value}} + 3.967$$

$$* \sqrt{\text{Fantasy Value Two Years Previous}}$$

$$E (\text{Cap Hit}) = 281388.5 * \sqrt{\text{Previous Year's Fantasy Value}} + 308856 \qquad 18$$

$$* \sqrt{\text{Fantasy Value Two Years Previous}}$$

$$\frac{B_p^1}{B_p^2} = 2.161 \qquad \qquad \qquad 19$$

$$\frac{B_s^1}{B_s^2} = 0.912 \qquad \qquad \qquad 20$$

The fact that the ratio of lagged indicators for Performance is not within a 95% Confidence Interval of the same ratio for Salary requires we reject the null hypothesis (Tables 11 and 12). This indicates that weighting of the past, as done by NFL General Managers and Owners in determining salary is different from the weighting of the past in predicting performance – a hugely important conclusion because this means that players are decidedly not being paid in accordance with what their future performance is expected to be. These findings allow us to reject the null hypothesis and embrace the alternative hypothesis and clarify the inferences from Problem II. Because lagged values of performance on salary and performance do not line up, clearly future salary is not allocated on the basis of expected future performance. This explains the conclusion earlier that salary growth does not occur in conjunction with performance growth by showing that future salary is not meant as compensation for future performance.

Part IV: Risk Preferences and Cutting Players

The relationship between fantasy football and the NFL is an interesting one because both are dependent on essentially the same group of people, but fantasy football situates players as if they were all on one year contracts. This simplifies a player's price to equate with expected performance. Essentially, if I expect a player to give me 500 yards and four touchdowns, I'll pay him directly in accordance with that performance unless I already have four wide receivers, in which case the marginal benefit of adding a fifth is small because I only have a limited number of spots for players to be on the field.

Estimating induced risk-aversion in fantasy football is a fairly straightforward problem. While better players are going to receive more "payment" than their worse counterparts, the relationship between the two groups²² is important. Worse players are less likely to play than their peers and bring very little marginal benefit to their owner because a team can very easily limit starting players to their better players. Therefore, while starters should be paid according to performance, and bench players should be paid according to performance, the relationship between rates of payment for performance differs between the two pools depending on the utility function of the owner. An owner with a high level of risk-aversion would have similar relationship between performance and payment in the two pools of players. An owner who is risk

²² The two groups are split according to "targets." Players above the median number of targets are "starters" in the NFL because teams have roughly the same number of starters and bench players at the wide receiver position. Players below the median are "bench" players.

loving would have a very high slope in an OLS regression of performance on payment in the starting pool and a very low slope in the same regression in the bench pool.

In fantasy football, we see that the slopes of the OLS regression in the two pools are almost identical – reiterating the finding of high risk-aversion mentioned earlier. Figure 5 and Figure 6 illustrate the similarities of the two OLS regressions. While this analysis may seem unclear, viewing the difference between bench players and starting players in the NFL in performance regressed on NFL cap value clarifies the distinction (Figure 12, Figure 13, and Figure 13a).²³ The slopes of the OLS regressions of the two pools are similar in fantasy football, but are drastically different in the NFL.²⁴ This supports my earlier hypothesis that teams in the NFL view players very differently than do their fantasy counterparts, perhaps as a result of the scarcity of top end talent. This is similar to the prediction of the efficiency wage theory, perhaps as a result of adverse selection – teams pay the best players more because players differ from each other in terms of talent and talent is the best indicator of performance.

Fantasy football makes use of one year contracts for players in exchange for an English auction. The NFL does not operate in the same way because, although it's certainly possible, teams sign players to extended contracts and rarely cut prominent players. Teams do cut players although the logic with which they do so does not always

²³ Interestingly, there is much greater spread in the NFL regressions than in the fantasy version. This may be a result of other characteristics Owners find valuable, like schematic fit or personality.

²⁴ The parabolic regression in Figure 13 and Figure 13a is of particular interest because it seems to indicate further conclusions about the way Owners pay their players. The players making high amounts of money in the backup pool are all older players, as evidenced by Figure 14. These players may have been given lucrative contracts before experiencing a decrease in performance but for some reason have not been cut.

follow. Figure 15 details the fantasy points and production of players cut in 2009. For some players, the reason for cuts seems obvious – poor production and very high salary. On the other hand, there are several players who seem to produce at a high level but get cut anyway. One would expect Figure 15 to look very different. Performance is a major factor, but often players who are cut are not given opportunities (Figure 16). Breaking down production per dollar should lend more light on the situation, but that is not the case. Players who provide a large amount of value at a low cost can still be cut (Figure 17). Further analysis seems to indicate that the players most likely to get cut are those who receive few targets or convert fewer targets into receptions than do their peers. Figure 18 illustrates that effect in graphic form, in which very few cut players are above the average.

The apparent conclusion from this information is that NFL owners could take a page from the books of fantasy football gurus and start cutting players more aggressively. There are many players who clearly do not have high catch rates (Figure 18) but are not cut. There are also older players who are paid in accordance with their stature rather than with performance (Figure 13a, Figure 14). Cutting more aggressively would allow those players to be signed again in concert with their on-field production. Obviously there are other factors to consider – namely interpersonal relationships crucial to team chemistry – but the relevance of fantasy football and the usefulness of single-year contracts cannot be denied.

Limitations

The most apparent limitation of this study is the position being studied: wide receivers. Because other positions have not been looked at, it seems like a major weakness. However, the conclusions found for wide receivers should hold for any position, even if finding statistical evidence is hard to do. The source of this mispricing is not wide receiver specific, rather it is the result of administrators weighting the past differently than performance track record would recommend. Linemen, for example, are most likely also mispriced with respect to future performance, but because the data tracked for linemen is very limited, sharpening those conclusions is nearly impossible. Weighting of the past would not be different depending on position, and in a position where statistical data is difficult to find pricing may be more ambiguous.

Additionally, there are issues with the pursuit of direct policy recommendations because it is impossible to fully understand the reputational effects of those recommendations. It is quite possible, for example, that a team has considered the position of only offering short term contracts, but that such a position would be untenable in the free agent market and players would never sign with such a team. This issue hurts my ability to make a strong recommendation. Similarly, judging a team's risk-aversion is not particularly easy given the difficulties of knowing whether a team has considered an alternative approach and given that understanding the goals of a team is a nebulous task at best. Those weaknesses aside, there is the capacity to elucidate some conclusions and see how an Owner understanding these conclusions might be able to affect his management style in the future.

Conclusion

The most important findings in this research are the following: NFL teams do not pay players in accordance to future expected performance and they display form of risk preferences with regard to backup players distinctly different from those of their fantasy counterparts. Instead of paying directly for future performance, they weight the past in a distinctly different manner than would seem correct if only paying for performance, as evidenced by the results in Problems II and III, when the predictors of salary and performance were different, as were the lagged predictors of each. The most reasonable inference, therefore, is that teams value aspects other than future performance and are willing to pay for those aspects. An example would include the capacity for “star power.” In Part I of the results section I discussed the increased marginal returns of salary to an increase in value, reinforcing the idea that teams care about things other than future performance and that star power is one of them, because more talented players are more recognized and capable of selling more merchandise. Capacity for becoming a star is not something afforded to every NFL player and the ones that do have such precocious talent are scarce. Teams seize opportunities to get stars perhaps because stars bring the organization a huge amount of revenue.

Despite the appearance that teams in the NFL care most for winning, this sports league, like all others, is driven first and foremost by experiencing financial success. Without that financial success, the league would have to fold. Owners in the NFL today have a wildly popular product that provides a return on investment even during the

worst seasons. As a result of the financial security of owning an NFL team, many fans believe teams must try to win regardless of financial sacrifices. However, the NFL's treatment of salary in comparison with expected performance, as in Problem III, suggests that there are other motives at play in the selection of players and some of those motives could well be financial. This renders my initial assumption that Owners obtain utility only from winning and allocate player salary in accordance with utility function.

As mentioned earlier, it is difficult to judge the induced risk preferences of Owners because knowing their exact utility function is impossible. Instead, this paper shows that the utility function of Owners is distinctly different from the expectation fans have of their Owners – an important piece of knowledge. Moreover, owners are not risk-averse in the same way that fantasy football owners are risk-averse, evidenced by the difference in slopes seen in Problem IV. Fantasy owners pay back up players higher than expected values in case of injuries whereas NFL Owners pay certain backups very little and other backups quite a lot. It is impossible to say whether this choice of payment is, in fact, any different than the risk-aversion in fantasy football. Fantasy football has the limiting factor of single season contracts and NFL teams do not have the luxury of only having to sign single-year contracts. Therefore, Owners may pay certain backups in a very risk-averse manner – over paying on the possibility of a player becoming a star or on the chance such a player will stay on a team throughout his prime. One way to avoid this possible lowered marginal benefit per dollar would be to cut players more frequently and sign shorter contracts, as in fantasy football, and pay

players according to immediate future performance, something much easier to predict than performance several years into the future.

Rosch and Hodgson (2000) use a medical perspective to determine a football player's physical qualifications. This type of analysis requires the ability to work directly with the player and have individuals participating in tests. The option for direct player contact brings to mind using similar data to compare to salaries and in game performance. Would coaches and general managers over estimate the values of conventional performance tests in evaluating players and writing contracts?

Market forces in the NFL are complex, particularly in the advent of statistical analysis in sports. Many teams are aware of the value of using statistics to predict future performance and may have individuals who use all statistical tools available. This fact may indicate the rationality of NFL teams – even while avoiding some of the strategies mentioned in this paper, teams are willing to adopt different pay scales in an effort to generate revenue by capitalizing on well known players. This presents a utility function different from the one fans expect from Owners – one based on aspects not encompassed by on-field performance.

Alternatively, if there was a team interested in winning regardless of financial motives or other personal motives, there appears to be a way forward through the benefit of the NFL's allowances regarding "cutting" of players. Signing shorter contracts would allow teams to be better, albeit with a higher short term cost. Paying purely for expected future performance instead of the hopes that a player remains on a team for a

long time would be another way for teams to gain an on-field competitive advantage.

There are several ways that a team could become more talented than their competitors, even while staying within the salary cap, that are currently being underutilized. Perhaps the reason for this is that teams don't recognize these other strategies, but more likely, most Owners get more utility from increased financial benefit than from higher probability of victory on the field.

Tables

Table 1: Summary of the Significant ($p < 0.05$) Explanatory Variables of Future Salary OLS Regression (Author's Calculations Using: USA Today NFL Salary Database; ESPN; Fantasy Football Today)

Variable	2006 Coeff	2007 Coeff	2008 Coeff	2009 Coeff	Average Coefficient
Age	0.042 (.009)	0.078 (.019)	0.098 (.019)	0.100 (.023)	.080
Fantasy Points Per Game	0.109 (.029)			0.174 (.021)	.071
Receptions		0.021 (.002)			.005
Square Root of Points			0.177 (.024)		.044
N	117	112	122	119	
R- Squared	.542	.559	.607	.611	

Table 1a: Confirmation of Increasing Marginal Returns of Skill on Future Salary (Author's Calculations Using: USA Today NFL Salary Database; ESPN)

Variable	2009 Coeff
Square Root of Fantasy Points	-213545.29 (197847)
Fantasy Points	44988.55*** (12940)
N	117
R- Squared	.479

*** $p < 0.001$

Table 2: All Explanatory Variables of Future Salary OLS Regression (Author's Calculations Using: USA Today NFL Salary Database; ESPN; Fantasy Football Today)

Variable	2006 Coeff	2007 Coeff	2008 Coeff	2009 Coeff
Age	.043* (.021)	.070*** (.020)	.094*** (.021)	.095*** (.024)
Games Played	-0.007 (.035)	-0.008 (.036)	-0.016 (.049)	-0.003 (.034)
Games Started	.062** (.021)	0.028 (.018)	-0.019 (.028)	0.006 (.026)
Receptions	0.013 (.010)	0.014 (.009)	-0.018 (.012)	-0.013 (.013)
Yards	0.0007 (.001)	0.0003 (.001)	0.001 (.001)	0.001 (.001)
Fumbles	-0.033 (.053)	0.066 (.072)	0.013 (.058)	-0.071 (.071)
Fantasy Points	-0.006 (.008)	-0.002 (.009)	-.022* (.008)	-0.002 (.010)
Targets			0.012 (.008)	0.006 (.008)
Square Root of Points	-0.127 (.108)	-0.034 (.095)	0.186 (.127)	-0.029 (.110)
Fantasy Points per Game	0.154 (.082)	0.046 (.635)	0.210 (.118)	.150* (.072)
N	117	112	122	119
R- Squared	.569	.574	.667	.625

* $p < 0.05$; ** $p < 0.01$; *** $p < 0.001$

Table 3: Summary of the Significant ($p < 0.05$) Explanatory Variables of Salary Growth OLS Regression (Author's Calculations Using: USA Today NFL Salary Database; ESPN; Fantasy Football Today)

Variable	2006 Coefficient	Variable	2007 Coefficient
Age in 2005	0.061 (.027)	Games Started in 2006	-0.043 (.014)
Receptions in 2005	0.021 (.008)	Fantasy Points per Game in 2006	-.248 (.072)
Points in 2005	-0.009 (.004)		
R-squared=.137	N=92	R-Squared=.106	N=112
Variable	2008 Coefficient	Variable	2009 Coefficient
Square Root of Fantasy Points 2007	0.112 (.042)	Age in 2008	-0.092 (.045)
Games Started in 2007	-0.065 (.024)	Yards in 2008	0.003 (.001)
		Games Started in 2008	-0.099 (.048)
R-squared=.096	N=78	R-squared= .172	N=86

Table 4: All Explanatory Variables of Future Salary Growth OLS Regression (Author's Calculations Using: USA Today NFL Salary Database; ESPN; Fantasy Football Today)

Variable	2006 Coeff	2007 Coeff	2008 Coeff	2009 Coeff
Age	0.052 (.028)	-0.031 (.0267)	-0.001 (.025)	-.101* (.048)
Games Played	-0.044 (.042)	-0.005 (.047)	0.019 (.057)	0.033 (.068)
Games Started	0.006 (.025)	-.068** (.023)	-0.053 (.033)	-0.086 (.053)
Receptions	.023* (.011)	0.023 (.012)	0.003 (.014)	0.002 (.025)
Yards	-0.0005 (.001)	-0.001 (.001)	0.0003 (.001)	0.002 (.003)
Fumbles	-0.068 (.019)	0.030 (.094)	0.035 (.069)	-0.129 (.139)
Fantasy Points	-0.011 (.009)	0.0006 (.013)	-0.011 (.009)	0.002 (.020)
Targets			-0.009 (.009)	-0.009 (.017)
Square Root of Points	0.146 (.126)	.253* (.125)	.323* (.149)	-0.015 (.017)
Fantasy Points per Game	-0.064 (.092)	-0.227 (.172)	0.003 (.138)	0.060 (.141)
N	92	112	78	86
R- Squared	.190	.194	.168	.194

* $p < 0.05$; ** $p < 0.01$; *** $p < 0.001$

Table 5: Summary of the Significant ($p < 0.05$) Explanatory Variables of Future Performance OLS Regression (Author's Calculations Using: ESPN; Fantasy Football Today)

Variable	2006 Coefficient	2007 Coefficient	2008 Coefficient	2009 Coefficient	Average Coefficient
Fantasy Points per Game	0.637 (.075)		0.512 (.082)	0.712 (.088)	0.465
Receptions		0.066 (.010)			0.016
N	117	114	125	117	
R- Squared	.388	.281	.336	.444	

Table 6: All Explanatory Variables of Future Performance OLS Regression (Author's Calculations Using: ESPN; Fantasy Football Today)

Variable	2006 Coeff	2007 Coeff	2008 Coeff	2009 Coeff
Age	-0.131 (.081)	-0.080 (.087)	-0.165 (.100)	-0.117 (.108)
Games Played	-0.086 (.103)	-0.143 (.113)	-0.244 (.162)	0.078 (.131)
Games Started	-0.0002 (.081)	-0.094 (.079)	0.141 (.126)	-0.021 (.118)
Receptions	-0.016 (.039)	.0889* (.041)	0.076 (.057)	-0.046 (.053)
Yards	0.002 (.003)	-0.001 (.004)	0.004 (.0005)	0.003 (.004)
Fumbles	0.300 (.198)	0.355 (.321)	-0.106 (.277)	0.094 (.289)
Targets			-0.068 (.038)	0.019 (.035)
Fantasy Points per Game	0.587** (.222)	0.265 (.302)	0.242 (.270)	.495* (.241)
N	117	114	78	86
R- Squared	.424	.325	.401	.489

* $p < 0.05$; ** $p < 0.01$; *** $p < 0.001$

Table 7: Summary of the Significant ($p < 0.05$) Explanatory Variables of Future Performance over Two Seasons OLS Regression (Author's Calculations Using: ESPN; Fantasy Football Today)

Variable	2007 Coefficient	2008 Coefficient	2009 Coefficient	Average Coefficient
Yards(t-1)		0.004 (.001)	0.005 (.001)	0.003
Yards(t-2)	0.003 (.001)			0.001
Receptions(t-1)	0.069 (.018)			0.023
Receptions(t-2)		0.036 (.011)		0.012
Games Started (t-1)	-0.209 (.082)			-0.069
N	93	94	90	
R- Squared	.319	.465	.554	

NB: "(t-1)" refers to the previous season, while "(t-2)" refers to two seasons previous

Table 8: Summary of the Significant ($p < 0.05$) Explanatory Variables of Performance Growth OLS Regression (Author's Calculations Using: ESPN; Fantasy Football Today)

Variable	2007 Coefficient	2008 Coefficient	2009 Coefficient	Average Coefficient
Games	-0.153 (.033)	-0.145 (.045)	-0.075 (.035)	-0.121
Fantasy Points per game	-0.134 (.031)			-0.045
N	114	79	84	
R- Squared	.333	.119	.054	

Table 9: All Explanatory Variables of Future Performance Growth OLS Regression (Author's Calculations Using: ESPN; Fantasy Football Today)

Variable	2007 Coeff	2008 Coeff	2009 Coeff
Age	-0.025 (.033)	-0.030 (.037)	-0.041 (.043)
Games Played	-.1479*** (.039)	-.119* (.051)	-0.071 (.042)
Games Started	-0.050 (.030)	0.002 (.046)	-0.004 (.048)
Receptions	0.023 (.015)	0.023 (.021)	-0.004 (.022)
Yards	-0.02 (.001)	0.0003 (.001)	0.001 (.001)
Fumbles	0.0800 (.117)	-0.037 (.099)	0.067 (.117)
Targets		-0.019 (.013)	-0.002 (.014)
N	114	78	86
R- Squared	.362	.168	.108

* $p < 0.05$; ** $p < 0.01$; *** $p < 0.001$

Table 10: Lagged Performance on Future Performance OLS Regression (Author's Calculations Using: ESPN; Fantasy Football Today)

Variable	2008 Coefficient	2009 Coefficient	Average Coefficient
Square root of Points (t-1)	8.373*** (1.482)	8.796*** (1.568)	8.585
Square root of Points (t-2)	3.817** (1.289)	4.116** (1.493)	3.967
Ratio of Coefficients	2.194	2.137	
N	94	90	
R- Squared	.462	.494	

* $p < 0.05$; ** $p < 0.01$; *** $p < 0.001$

Table 11: Lagged Performance on Future Salary OLS Regression (Author's Calculations Using: USA Today NFL Salary Database; ESPN; Fantasy Football Today)

Variable	2008 Coefficient	2009 Coefficient	Average Coefficient
Square root of Points (t-1)	293794*** (73174)	268983*** (67610)	281388.5
Square root of Points (t-2)	289550*** (63667)	328162*** (64346)	308856
Ratio of Coefficients	1.0147	.8197	
N	94	90	
R- Squared	.4517	.528	

* $p < 0.05$; ** $p < 0.01$; *** $p < 0.001$

Table 12: Lagged Performance on Future Fantasy Price OLS Regression (Author's Calculations Using: ESPN; Fantasy Football Today) Extension from Chakravarthy (2010)

Variable	2009 Coefficient
Square root of Points (t-1)	0.068 (.019)
Square root of Points (t-2)	0.045 (.014)
Ratio of Coefficients	1.51
N	40
R- Squared	.530

Figures

Figure 1: Fantasy Percent Started in 2011 vs. Average Fantasy Points per game in 2011 (Source: ESPN)

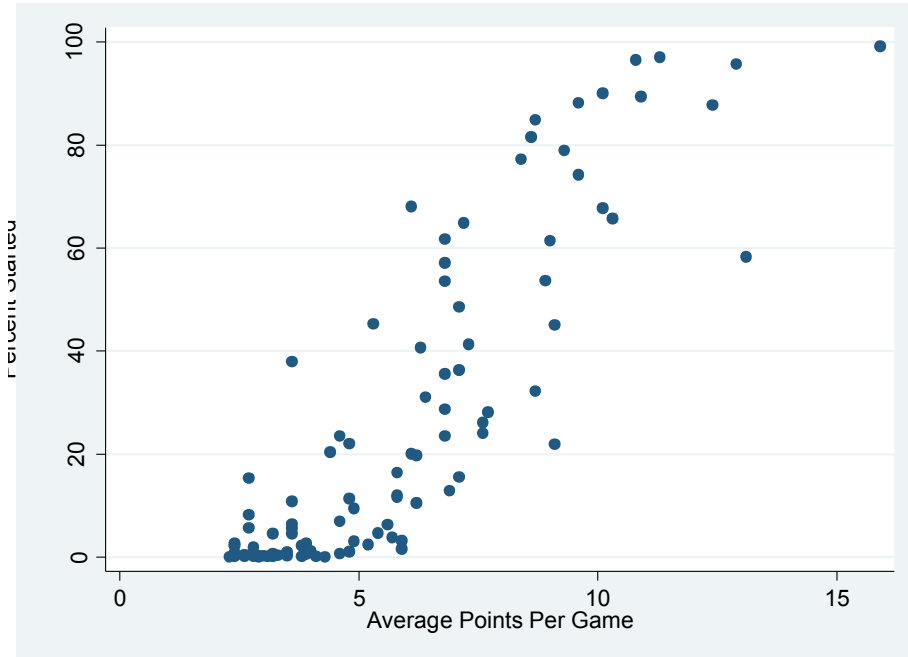


Figure 2: Fantasy Rank in 2011 vs. Fantasy Percent Started in 2011 (Source: ESPN)




Figure 3: Fantasy Percent Started in 2010 vs. Total Fantasy Points Scored in 2010 (Source: ESPN)




Figure 4: Effective Fantasy Points in 2009 vs. Fantasy Points in 2009 (Author's Calculations Using: ESPN), Extension from Chakravarthy (2010)

- The red line is the OLS regression for bench players.
- The green line is the OLS regression for starting players. $p < 0.001$
- Slopes of red and green line are significantly different.





Figure 5: Fantasy Price in 2009 vs. Fantasy Points in 2009 (Author's Calculations Using: ESPN),²⁵ Extension from Chakravarthy (2010)

- The red line is the OLS regression for bench players. $p < 0.001$
- The green line is the OLS regression for starting players. $p < 0.001$
- Slopes of red and green line are not significantly different.



²⁵ The reason Figure 5 and Figure 6 are limited to N=60 is the lack of fantasy pricing data further than that point.

Figure 6: Fantasy Price in 2008 vs. Fantasy Points in 2008 (Author's Calculations Using: ESPN), Extension from Chakravarthy (2010)

- The red line is the OLS regression for bench players. $p < 0.001$
- The green line is the OLS regression for starting players. $p < 0.001$
- Slopes of red and green line are not significantly different.

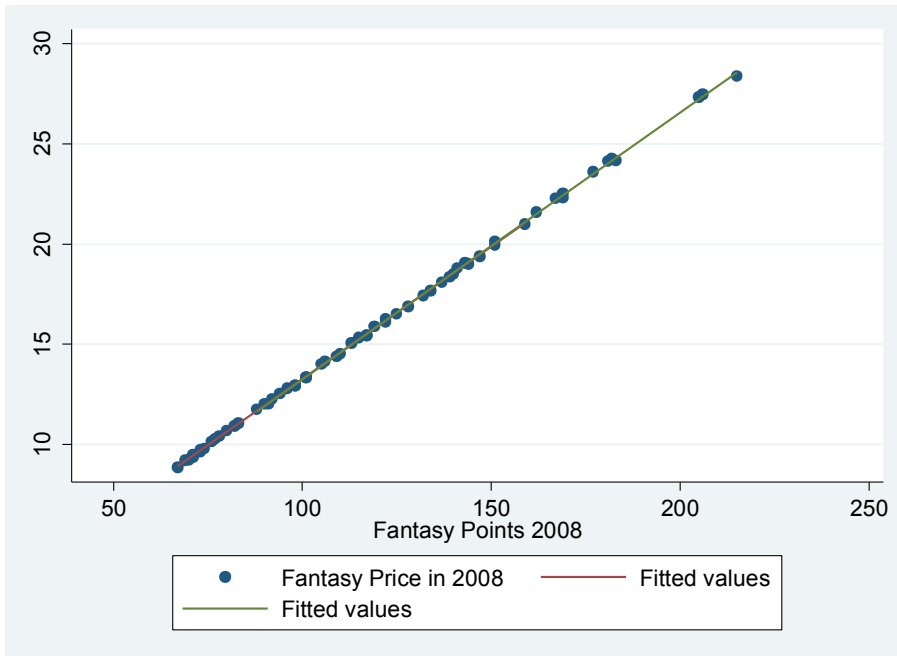


Figure 7: Fantasy Points in 2009 vs. Targets in 2009, $p < 0.001$ (Author's Calculations Using: ESPN; Fantasy Football Today)




Figure 8: Fantasy Points in 2009 vs. Targets in 2008, $p < 0.01$ (Author's Calculations Using: ESPN; Fantasy Football Today)




Figure 9: Fantasy Points per Target in 2009 vs. Fantasy Points in 2009 (Author's Calculations Using: ESPN; Fantasy Football Today)

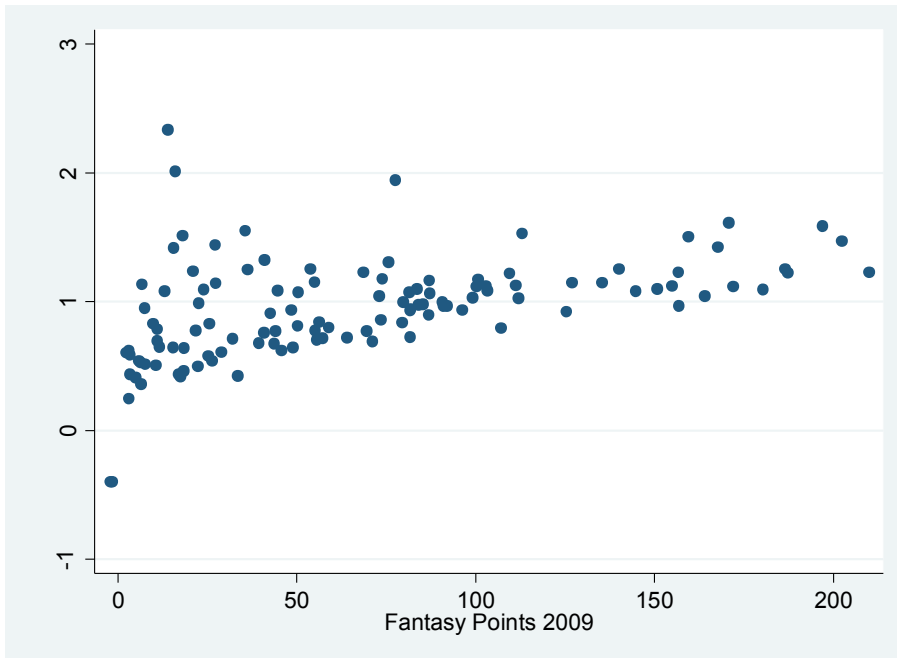


Figure 10: Fantasy Points in 2009 vs. Fantasy Points per Target in 2008, $p < 0.001$
(Author's Calculations Using: ESPN; Fantasy Football Today)

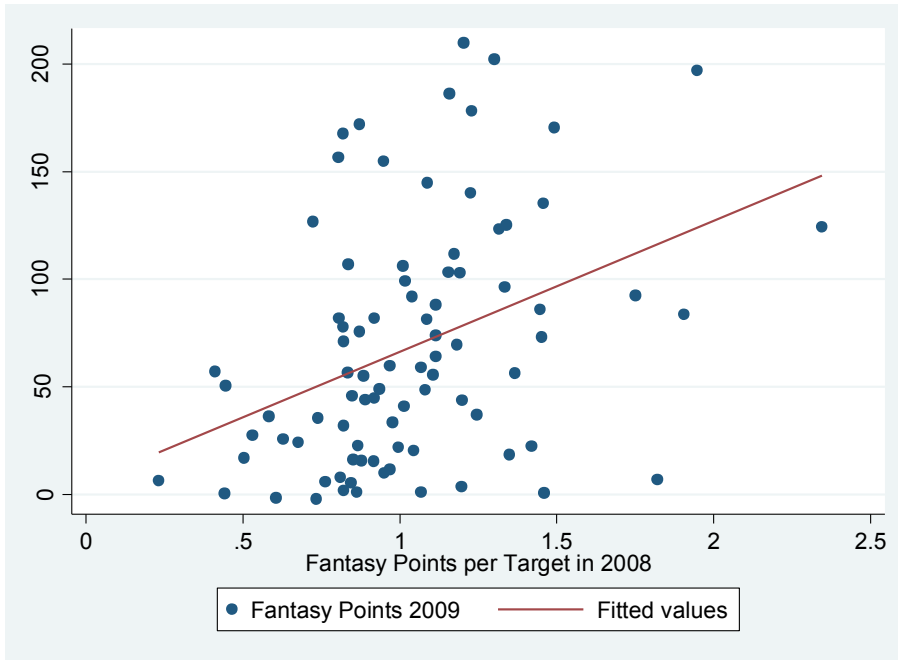


Figure 11: Cap Hit in 2009 vs. Fantasy Points per target in 2008, $p < 0.05$ (Author's Calculations Using: ESPN; Fantasy Football Today; USA Today NFL Salary Database)




Figure 12: Cap Hit in 2009 vs. Fantasy Points in 2009 (Author's Calculations Using: ESPN; Fantasy Football Today; USA Today NFL Salary Database)

- The red line is the OLS regression for bench players. $p < 0.001$
- The green line is the OLS regression for starting players. $p < 0.001$
- Slopes of red and green line are significantly different.



Figure 13: Cap Hit in 2008 vs. Fantasy Points in 2008 (Author's Calculations Using: ESPN; Fantasy Football Today; USA Today NFL Salary Database)

- The red line is the quadratic regression for bench players of Cap Hit in 2008 vs. Fantasy Points in 2008. $p < 0.001$
- The green line is the OLS regression for starting players. $p < 0.001$
- Slopes of red and green line are significantly different.




Figure 13a: Inset of bench players from chart above, $p < 0.001$




Figure 14: Cap Value in 2008 vs. Age in 2008 among Bench Players, $p < 0.001$ (Author's Calculations Using: ESPN; Fantasy Football Today; USA Today NFL Salary Database)²⁶


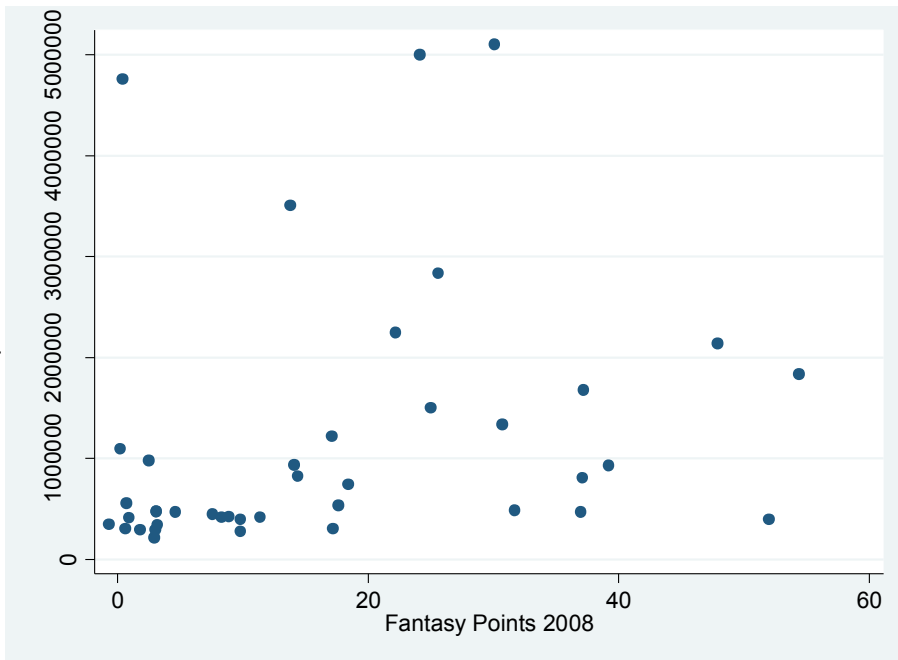


Figure 15: Cap Hit in 2008 vs. Fantasy Points in 2008 among Players Cut in 2009 (Author's Calculations Using: ESPN; USA Today NFL Salary Database)



²⁶ N=64, despite appearances it may be smaller.

Figure 16: Targets in 2008 among Players Cut in 2009 (Source: ESPN; Fantasy Football Today; USA Today NFL Salary Database)

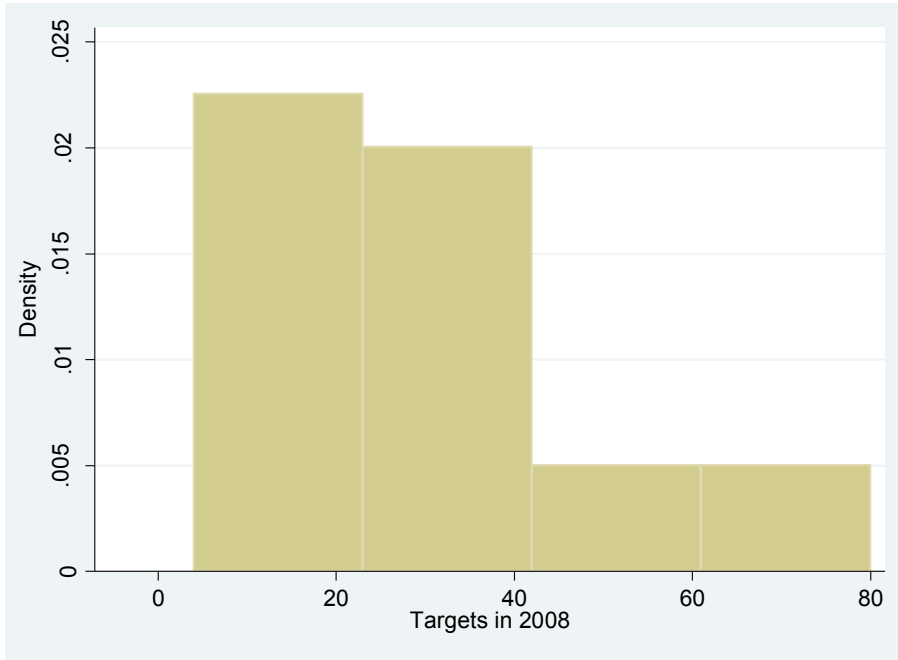


Figure 17: Fantasy Points per Dollar in 2008 among Players Cut in 2009 (Author's Calculations Using: ESPN; USA Today NFL Salary Database)

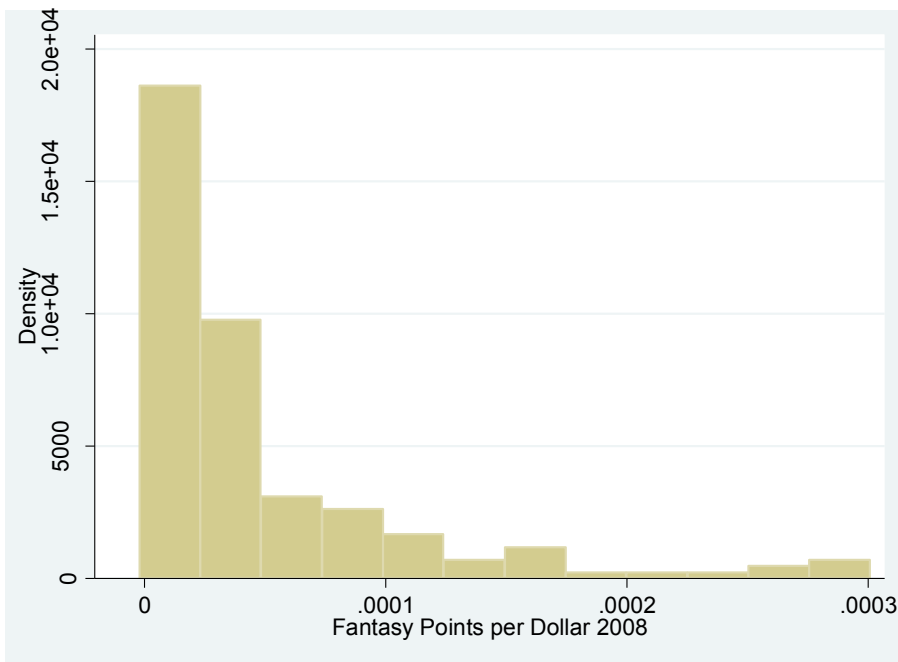
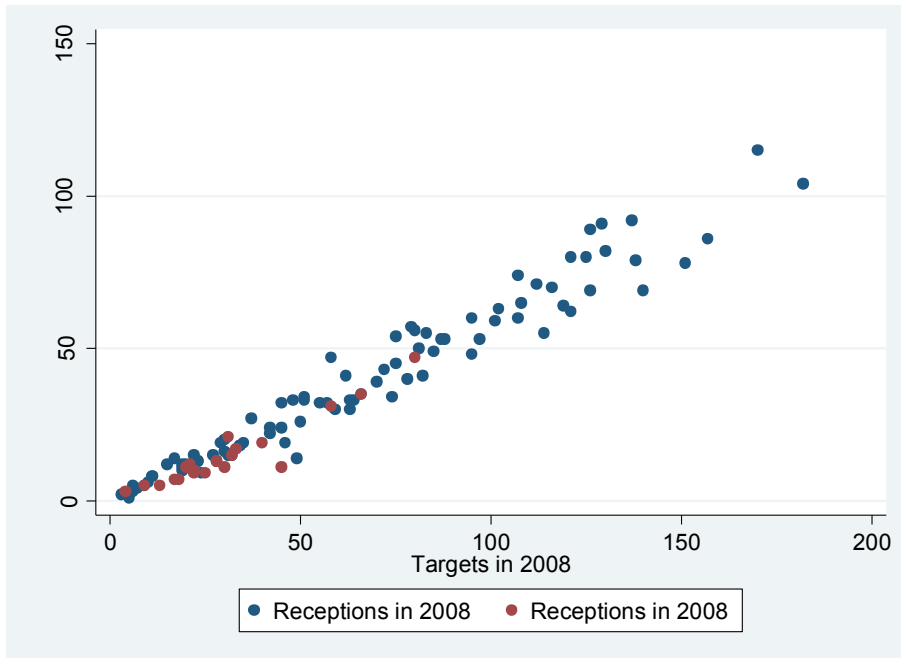


Figure 18: Receptions in 2008 vs. Targets in 2008 (Source: ESPN; Fantasy Football Today; USA Today NFL Salary Database)

- Red points denote players who were cut.
- Blue points denote players who were not cut.



References

- Armstrong, J.** 2011. "2011 Drive Stats," J. Armstron, *Drive Stats*. Football Outsiders, NFL statistics on a per drive basis.
- Backman, S.E.** 2002. "Nfl Players Fight for Their Freedom: The History of Free Agency in the Nfl." *Sports Law. J.*, 9, 1.
- Chakravarthy, P.** 2010. "Optimizing Draft Strategies in Fantasy Football."
- Enterprises, NFL.** 2010. "Year-by-Year Salary Cap," Staff, *NFL Communications*. New York: NFL, Yearly listing of maximum team salary.
- Garvey, E.** 1989. "Foreword to "the Scope of the Labor Exemption in Professional Sports": A Perspective on Collective Bargaining in the Nfl." *Duke Law Journal*, 1989(2), 328-38.
- Goodell, R.** 2011. "Official Playing Rules and Casebook of the National Football League," N. F. League, New York: National Football League, 113.
- Heller, A.W.** 2000. "Creating a Win-Win Situation through Collective Bargaining: The Nfl Salary Cap." *Sports Law. J.*, 7, 375.
- Kahn, L.M.** 2000. "The Sports Business as a Labor Market Laboratory." *The Journal of Economic Perspectives*, 14(3), 75-94.
- Larsen, A.; A.J. Fenn and E.L. Spenner.** 2006. "The Impact of Free Agency and the Salary Cap on Competitive Balance in the National Football League." *Journal of Sports Economics*, 7(4), 374.
- Leeds, M.A. and S. Kowalewski.** 2001. "Winner Take All in the Nfl." *Journal of Sports Economics*, 2(3), 244-56.

Massey, Cade and Richard Thaler. 2005. "Overconfidence Vs. Market Efficiency in the National Football League." *National Bureau of Economic Research Working*, No. 11270.

Nissim, A. 2004. "Trading Game: Nfl Free Agency, the Salary Cap, and a Proposal for Greater Trading Flexibility, The." *Sports Law. J.*, 11, 257.

Redding, M.J. and D.R. Peterson. 2009. "Third and Long: The Issues Facing the Nfl Collective Bargaining Agreement Negotiations and the Effects of an Uncapped Year." *Marq. Sports L. Rev.*, 20, 95, 102.

Richard, B. 2008. "Allocation of Scarce Resources: Insight from the Nfl Salary Cap." *Journal of Economics and Business*, 60(6), 536-50.

Rösch, D.; R. Hodgson; L. Peterson; T. Graf-Baumann; A. Junge; J. Chomiak and J. Dvorak. 2000. "Assessment and Evaluation of Football Performance." *The American Journal of Sports Medicine*, 28(suppl 5), S-29.

Rosen, S. 1981. "The Economics of Superstars." *The American Economic Review*, 71(5), 845-58.

Scully, G.W. 2004. "Player Salary Share and the Distribution of Player Earnings." *Managerial and Decision Economics*, 25(2), 77-86.

Truelock, M. 1993. "Free Agency in the Nfl: Evolution or Revolution." *SMUL Rev.*, 47, 1917.

Walder, S. 2011. "Evaluating fantasy football advanced statistics." *NY Daily News Fantasy Huddle*, September 29, 2011.

Zimbalist, A. 2010. "Reflections on Salary Shares and Salary Caps." *Journal of Sports Economics*, 11(1), 17.