

The Price of a Pick

**Examining the Efficiency of the NFL Draft Pick Market by
Measuring the Elasticity of the Relationship Between Draft Pick
Value and Player Performance**

Joseph Grisso, B.A. in Economics (University of Colorado Boulder)

*Advised by Donald Waldman, Professor in the Economics Department at the
University of Colorado Boulder*

*Honors Council included Waldman, Martin Boileau (Professor in the Economics
Department at the University of Colorado Boulder), and Fred Pampel (Professor
in the Sociology Department at the University of Colorado Boulder)*

Defended on April 11th, 2017

Paper Overview

This paper is divided into six sections. Except for Sections I and VI, each section is further divided into subsections, including a general overview at the beginning. A rudimentary knowledge of the National Football League (NFL) is required to understand the concepts discussed in this paper, as well as the terminology used.

Section I: ABSTRACT. I provide a synopsis of my question, research, and findings. P. 2

Section II: INTRODUCTION. I introduce my topic and explain what my contribution will be to existing research. P. 3-8

Section III: METHODOLOGY. I describe my data and research method. P. 9-22

Section IV: RESULTS. I summarize the results of my research. P. 23-31

Section V: CONCLUSIONS. I draw conclusions based on the results of my research. P. 32-38

Section VI: BIBLIOGRAPHY. I cite the authors whose work I relied on to write my paper. P. 39-40

I. Abstract

In the NFL Draft, teams take turns (known as “picks”) selecting (or “drafting”) college football players with the intention of drafting players that will perform well in the NFL and thus cause them to win more games. Since the pool of available college football talent is limited, players who are projected to have greater NFL performance are selected with earlier picks, causing the value of a draft pick to decrease as its number increases. With my research, I seek to determine a.) the relationship between the value of a particular pick in the draft (DPV) and the performance of the player selected with that pick (PP) and b.) whether or not this relationship differs between quarterbacks and non-quarterbacks. First, I devise appropriate measures for DPV and PP. Second, I regress the latter on the former using a log-log regression, including a dummy QB variable and an interaction term to determine if this relationship differs between quarterbacks and non-quarterbacks. I also include a control for team quality. Third, I constrain the model to include only drafted players. I find PP to be relatively inelastic to changes in DPV for all players, but I do not find a different relationship for quarterbacks and non-quarterbacks, nor do I find team quality to impact PP.

II. Introduction

i.) General Overview

Each year, NFL teams attempt to draft the best players possible in order to maximize their regular season win totals and have a shot at winning the Super Bowl. The impact of draft pick value on player performance in the NFL has been extensively analyzed in the economic literature, primarily due to its implications regarding market efficiency. Does the market value of NFL draft picks accurately reflect the performance of players selected with those picks? Existing literature seems to say that it does not, but fails to distinguish between quarterbacks and non-quarterbacks. This is an important distinction because of how valuable quarterbacks are to NFL teams in relation to other positions.

ii.) Topic Introduction

Every spring, the 32 teams of the National Football League (NFL) convene for several days to participate in the NFL Draft. This process allows teams to select (or “draft”) talented college football players to fill the holes on their roster through taking turns: each turn is known as a “pick”, and there are seven total rounds of picks with each team being assigned one pick per round (additional picks, known as compensatory picks, may be assigned to certain teams as well, causing there to be more than 224 picks in most drafts).

It can be presumed that the short term goal of all 32 teams in a given season is to maximize the number of regular season games they win. This is because the long term goal of all 32 teams each season is to win the Super Bowl, which a team can only accomplish if it qualifies for the postseason through winning a sufficiently high percentage of its regular season games. Thus, while a team’s drafted players will depend on what positions it needs to fill, teams

generally select players they believe will perform at a high level in the NFL. This causes players who have been highly rated by scouts and the media to be selected earlier than players who haven't been as highly rated as they are predicted to generate more wins for the team that drafts them through a higher level of individual performance.

In fact, the rationale behind the Draft's order is to allow bad teams to improve by selecting better players than the good teams that draft after them. In each of the Draft's seven rounds, teams select in inverse order of the previous year's standings, with the number one pick of the round going to the team that owned the worst record in the league the previous season and the final pick of the round going to the team that won the previous season's Super Bowl.

Yet a team is not required to simply keep the picks it is assigned. Instead, draft picks can be traded for players, other draft picks, or a combination of both. Since higher (i.e. earlier) picks are valued more than lower (i.e. later) picks due to the expected performance of players taken earlier in the Draft being greater than the expected performance of those taken later, a team has two options when trading only draft picks with another team. It can either "trade up" by giving up a greater number of low picks for a smaller number of high picks, or "trade down" by giving up a smaller number of high picks for a greater number of low picks.

iii.) Market Efficiency

Much like the price of a stock, the concept of market efficiency dictates that the price of a draft pick (expressed as the draft picks that must be dealt to acquire the pick) should reflect all available information, making it impossible to "beat the system" by engaging in systematic behavior (Malkiel, 2003). This means that the expected performance of players taken with picks acquired by the team trading up should equal the expected performance of players taken with

picks acquired by the team trading down. If one side has a systematic advantage, then the market for draft picks is inefficient as it's possible for a team to increase the expected performance of players taken with its draft picks through systematically trading either up or down.

iv.) My Question

In order to determine whether or not the market for NFL draft picks is efficient, I will calculate the elasticity of the relationship between draft pick value (DPV), a measure of how much teams value particular picks in the Draft, and player performance (PP), the individual performance of players drafted with those picks. The market for draft picks is efficient if PP is unitarily elastic to changes in DPV. If PP is relatively inelastic to changes in DPV, there is a systematic advantage to trading down: the market overvalues high draft picks, causing DPV to decline more rapidly than PP as the Draft progresses. Likewise, if PP is relatively elastic to changes in DPV, there is a systematic advantage to trading up: the market undervalues high draft picks, causing DPV to decline at a slower rate than PP as the Draft progresses. Based on the existing literature, I anticipate that PP will be relatively inelastic to changes in DPV, making the market for draft picks inefficient as there is a systematic advantage to trading down.

I also seek to determine if this relationship differs between quarterbacks and non-quarterbacks. Anecdotal evidence suggests that the PP of quarterbacks is less elastic to changes in DPV than the PP of non-quarterbacks. For example, four time Super Bowl MVP Tom Brady is widely considered to be the greatest quarterback of all time and was drafted 199th overall by the New England Patriots. JaMarcus Russell, however, lasted a mere three seasons in the NFL after the Oakland Raiders drafted him first overall. Since a quarterback taken with the 199th pick has enjoyed a much more successful career than one taken first overall despite the

DPV of the first pick likely being much greater than the DPV of the 199th, I suspect PP will be almost entirely inelastic to changes in DPV for quarterbacks.

v.) Literature Review

Similar economic studies have indirectly examined the relationship between PP and DPV by examining the relationship between factors used by teams to predict NFL performance and the actual performance of these players. (Since these factors influence which players get drafted first, they serve as a proxy for DPV.) Most measurements used to predict which college players will enjoy successful NFL careers have been found to be unreliable: a quarterback's Wonderlic score (the Wonderlic is an intelligence test, similar to an IQ test), forty yard dash time, and height all have a statistically significant impact on his draft position, yet none have a statistically significant impact on his NFL performance (Berri and Simmons, 2009). Teams also rely on collegiate team success as a predictor of NFL individual success, causing players from highly ranked college teams to be drafted earlier than equally skilled players from unranked teams even though a player's collegiate affiliation has no long term effect on his career success (Kitchens, 2014).

Despite the unreliability of their measurements, teams have been shown to overestimate their predictive abilities, causing DPV to decline much more rapidly during the draft than PP (Massey and Thaler, 2005). Overconfidence in one's predictive abilities is a well documented phenomenon among human beings. For example, an overconfidence in predictive ability was a key contributor to the volatility of financial markets during the 2008 global economic collapse (Abbes, 2013). Human beings have also been shown to become more confident in their predictions with the accumulation of more information, even when this information adds little to

no value to their predictive models. In a classic psychological study from the 1960's, individuals became more confident in their judgements regarding a hypothetical patient's psychological state as the patient grew older, even though their judgement accuracy remained constant across the patient's lifetime (Oskamp, 1965). Much like financial markets and psychological analysis, drafting college football players is a process that involves a significant amount of information collecting, so if NFL executives (who are human beings) have an overconfidence in their predictive abilities that is made worse with more information, one would expect high draft picks to be overvalued by the market for draft picks, causing it to be inefficient.

vi.) My Contribution

My research builds on a study by Massey and Thaler published in 2005 that directly examines the relationship between PP and DPV. Although it found both PP and DPV to decline from the draft's beginning to its end, the decline in DPV is found to be much more rapid, indicating a market inefficiency that causes high draft picks to be overvalued. However, this study fails to distinguish between quarterbacks and non-quarterbacks. This distinction is necessary because of how important quarterbacks are to NFL teams. Quarterbacks must be capable of running the offense by knowing every single play in the playbook (Tang 2015) and have the highest positional wins above replacement (WAR) by a wide margin, indicating that their individual performance is crucial in determining whether or not their team wins (Hughes, et al., 2015). Because quarterbacks are so important, they are often drafted much earlier than non-quarterbacks: in 14 of the 19 drafts between 1998 and 2016, a quarterback was selected first overall, and in five of these years (1998, 1999, 2012, 2015, and 2016) they were followed by

another quarterback at number two. But do teams overvalue drafting quarterbacks early, thus causing PP to be less elastic to changes in DPV for quarterbacks than non-quarterbacks?

III.) Methodology

i.) General Overview

Comparing the DPV of NFL players to their on-field performance first requires devising measurements for both variables. Once these measurements have been appropriately devised, I regress PP on DPV using a sample of both quarterbacks and non-quarterbacks to determine a.) what the general relationship between PP and DPV is for all players and b.) how this relationship differs between quarterbacks and non-quarterbacks. To separate PP from team related factors, I also control for team quality, a variable that is positively correlated with team success and thus likely positively correlated with individual success, or PP. Furthermore, I constrain my model to include only drafted players: undrafted players should have a DPV of 0, which would prevent me from running a log-log regression and thus calculating how elastic PP is to changes in DPV.

ii.) Measuring Draft Pick Value

The value of a particular draft pick may vary considerably from year to year depending on the the presence of particularly desirable prospects. For example, former Stanford quarterback Andrew Luck was considered such a sure bet in the NFL that fans of a number of teams wanted their franchises to deliberately lose games during the 2011 season in order to earn the number one pick in the 2012 Draft and be able to select him, a phenomenon known as “Suck for Luck” (Politi, 2011). By contrast, not many fans were begging their teams to “Fail for (Eric) Fisher,” the offensive tackle selected first overall a year later. Thus, the value for a particular pick (e.g. the number ten pick of the 2010 NFL Draft) is difficult to determine. However, it’s considerably easier to determine the value for a particular pick number (e.g. the number ten pick of any given

NFL Draft) by using historical data to determine how teams have valued different picks throughout the years.

DPV can be approximated by examining draft day trades involving only draft picks and assuming equal value on each side of the trade to estimate the relative value of any given pick number. For instance, if a team trades the 20th overall pick for the 33rd and 100th pick, it can be inferred that the 20th pick has as much value as the 33rd and 100th picks combined. This is done by Massey and Thaler, who assume a Weibull distribution in DPV from the draft's beginning to its end. Such a distribution allows the DPV of each pick number (relative to the DPV of the first overall pick) to decline in either a decreasing, constant, or increasing manner, depending on the distribution's parameters. According to this method, the relative DPV of any given pick number can be calculated with the following equation:

$$DPV(PickNum) = e^{-\lambda(PickNum-1)^\beta}$$


Where λ and β are the Weibull distribution parameters to be estimated using data on draft pick trades. (Plugging 1 into the equation for pick number will yield a DPV of 1, regardless of the values of these parameters.) Massey and Thaler used 213 same year draft pick trades from 1988 to 2004 as their data set, which fit the Weibull model incredibly well as R-squared was equal to 0.999. Their estimated values for λ and β were .148 and .7 respectively, resulting in the following equation for DPV:

$$DPV(PickNum) = e^{-.148(PickNum-1)^{0.7}}$$

The estimated DPV of any given pick number from 1 to 250 using this equation is shown in Figure 1. Massey and Thaler also incorporate trades involving future draft picks into their model to derive alternative estimates for λ and β , but I did not feel this was appropriate since

teams trading for future draft picks are unaware of the exact pick number they will receive in future drafts. For example, a team trading picks in a given draft for another team's first round pick in the following draft could receive anywhere from the first overall pick of the following year's draft to the 32nd, depending on how well the team trading the future pick performs during the upcoming season. Thus, I ignored these alternative estimates and only used their λ and β estimates for same year draft pick trades.

Figure 1




iii.) Measuring Player Performance

To measure PP, I use data from Pro Football Focus (PFF). PFF is a firm that watches film of every preseason, regular season, and postseason game and assigns each player a grade of -2 to 2 for each play. Players then receive a game grade (an accumulation of his grades for each play during a given game), which are further accumulated into season grades. These season grades, in turn, can be accumulated into team grades. Since all teams seek to win the greatest amount of games possible, examining the relationship between team grades and team winning percentage

will reveal whether or not individual grades are a good way of measuring PP: an increase in an individual player's season grade will obviously increase his team's grade, but will it increase his team's winning percentage? The short answer is yes: using data from all thirty two teams over a ten year period from 2007-2016 (320 total observations), the relationship between team grade and team winning percentage in the regular season is displayed in Figure 2. R-squared is greater than 0.5, indicating a strong relationship between the two variables, which is justification for using individual grades as a measurement of PP.

Figure 2



Although PFF keeps track of postseason and preseason data, I only use regular season data for my analysis. Using postseason performance data would result in an uneven data set since only 12 of the league's 32 teams participate in the postseason during any given year, while I did not feel the preseason would yield appropriate performance data since teams tend to rest their projected starters for the majority of preseason games due to fear of losing them to injury before the regular season begins. I also restrict my sample to players who played at least 25% of their

team's snaps on their side of the ball (either offense or defense, depending on which position they play) as players who play under 25% of their team's snaps on their side of the ball are unlikely to play enough for their performance to impact whether or not their team wins.

In addition to single season performance grades, PFF also produces data on snap count, which is simply the number of snaps a player played during a given year. Dividing a player's season grade by his snap count will yield his points per snap, or PPS. I decided to use this as the basis for my performance measurement as existing studies have shown early round draft picks to receive significantly more snaps than late round draft picks (Berri and Simmons, 2009), so I wanted to negate the effect snap count may have on PP. My initial sample consisted of 10,063 player grades and snap counts over a ten year period from 2007 to 2016 across thirteen different position groups, including quarterbacks, my main focus.

iv.) Restricting My Sample

Before regressing PP on DPV, I needed to further restrict my sample by focusing only on players whose performance was roughly similar in distribution to quarterbacks. The reason I had to do this was simple: each position group had a different mean and standard deviation in grade, causing a particular grade to be vastly different depending on the position of the player. For example, halfbacks had an average grade of -0.31 with a standard deviation of 7.52, meaning a halfback with a grade of 7.21 would be one whole standard deviation above the average for his position. However, he would only be 0.06 standard deviations above the average for quarterbacks, who had an average grade of 6.13 with a standard deviation of 17.88. A halfback who was well above average would, therefore, be considered average if compared alongside quarterbacks, making performance grades an unreliable way to directly compare the performance

of quarterbacks to halfbacks. Thus, I wanted to restrict my sample only to position groups whose distribution in grade was similar to that of quarterbacks: these positions were found to be tight ends, offensive guards, centers, and edge rushers, bringing my sample down to 3,171. (See Section IV.ii for the specifics on how I decided on this restriction.) 387 of these players were quarterbacks, and the remaining 2,784 were one of the four aforementioned non-quarterback positions.

v.) Controlling for Team Quality

Individual players play their position to the best of their ability with the intention of contributing to team success, or wins. But this relationship likely works in both directions: being on a good team might also cause a player to perform better. Not only does having teammates who play their position well make it easier for a player to play his own position effectively, but teams can contribute to effective play from their players through coaching, motivation, game strategy, and play design.

Even though a team's coaching and ownership may change from year to year, some teams are consistently better than others at getting the most out of their players. For instance, the Cleveland Browns have languished in obscurity since rejoining the NFL as an expansion franchise in 1999, after the original franchise was relocated to Baltimore and renamed the Ravens. They've only qualified for the postseason once in that span, and have finished last in their division in 12 of the last 14 NFL seasons. Part of this is certainly because they lacked talented players, but it's also likely true that the quality of team factors external to the level of on-field talent has been subpar throughout this period. Conversely, the New England Patriots have been the league's gold standard for the last decade and a half, winning five Super Bowls

since 2002 and qualifying for the postseason in 13 of the last 14 NFL seasons. This success is partially due to the Patriots having a high level of on-field talent, but it's also likely a byproduct of particularly competent ownership and a coaching staff that is great at maximizing what it gets out of its players.

I expect team quality to be positively correlated with PP and negatively correlated with DPV. Good teams win more games and thus have lower draft picks, leading to a lower DPV. But good teams also generate a higher level of performance out of the players they draft, so team quality should be positively correlated with PP. Therefore, controlling for team quality should increase the effect DPV has on PP.

vi.) Regression Setup

Once my sample was appropriately restricted, I found data on draft pick number for each of the 3,171 players by using both Wikipedia and Pro-Football-Reference. While single seasons are used as individual data points, a player with multiple seasons in the sample has multiple data points with the same draft pick number and thus DPV, making a player's DPV static. Once a player is drafted with a particular draft pick number, they are assigned the DPV of that pick number for each of their seasons in the sample. After a player's draft pick number was found, DPV was easily calculated by applying the formula described in Section III.ii.


PP, meanwhile, was calculated by adding 0.08 to each player's PPS. This adjustment allows every single PP to be positive (the lowest PPS found in the sample has a negative value with an absolute value lower than 0.08), enabling me to run a log-log regression. Since I am testing the efficiency of the market for NFL draft picks through measuring the elasticity of the relationship between PP and DPV, I decided that the market would be efficient if the coefficient

on $\ln DPV$ equaled 1 (i.e. was not statistically different from 1). Since I believed the market to be inefficient based on prior research, I expected the value of this coefficient to be less than 1, indicating an overvaluation of high draft picks. However, I expected it to be positive: if it's negative, either a.) teams value late draft picks more than early ones or b.) players taken with late draft picks perform better than those taken with early ones, and neither scenario seems plausible without the other also being true (if both were true, the coefficient would still be positive).

I am also interested in whether the relationship between DPV and PP differs between quarterbacks and non-quarterbacks, so I included a dummy variable for quarterbacks equal to 1 if a player was a quarterback and 0 if he wasn't, along with an interaction term between the dummy variable and $\ln DPV$. Since I expected PP to be less elastic to changes in DPV for quarterbacks, I expected the coefficients on both the dummy variable and the interaction term to be negative.

Finally, controlling for team quality required me to take ten year averages of each franchise's winning percentage and team grade over the ten year period from 2007-2016. The 320 data points displayed in Figure 2 were grouped into 32 separate teams and averaged out, giving me 32 data points for team winning percentage and grade. (Each team now had a single data point, which was their ten year average for both variables, rather than data points for each season.) The relationship between team grade and team winning percentage over a ten year period is shown in Figure 3. Averaging out extreme data points reduced outliers, causing R -squared to increase to .636.

Figure 3



Since I am defining team quality to mean “factors other than on-field talent that contribute to whether or not a team wins games,” I calculated team quality by taking the difference between a team’s actual ten year winning percentage and their predicted ten year winning percentage given their ten year average team grade, calculated using the formula for the trendline in Figure 3. Teams above the line had a positive team quality, while teams below the line had a negative team quality, with most teams having a team quality close to zero. Each team’s ten year winning percentage, average grade, predicted winning percentage, and team quality is shown in Table 1. (LA refers to the Los Angeles Rams, known as the St. Louis Rams for nine of the ten seasons in the sample, while SD refers to the team now known as the Los Angeles Chargers.)

Table 1

TeamName	10 yr. Grade	10 yr. Win %	Expected Win %	TeamQuality	TeamName	10 yr. Grade	10 yr. Win %	Expected Win %	Team Quality
AZ	-15.24	0.5409	0.40121136	0.13968864	LA	-3.29	0.2909	0.41359156	-0.12269156
ATL	149.76	0.5565	0.57215136	-0.01565136	MIA	87.76	0.444	0.50791936	-0.06391936
BAL	237.45	0.5627	0.6629982	-0.1002982	MIN	123.98	0.5033	0.54544328	-0.04214328
BUF	1.52	0.4127	0.41857472	-0.00587472	NE	290.83	0.7876	0.71829988	0.06930012
CAR	102.81	0.5158	0.52351116	-0.00771116	NO	139.81	0.5692	0.56184316	0.00735684
CHI	-3.36	0.4628	0.41351904	0.04928096	NYG	80.73	0.5502	0.50063628	0.04956372
CIN	106.92	0.5282	0.52776912	0.00043088	NYJ	76.72	0.4627	0.49648192	-0.03378192
CLE	-4.89	0.3003	0.41193396	-0.11163396	OAK	-9.23	0.3752	0.40743772	-0.03223772
DAL	167.18	0.5752	0.59019848	-0.01499848	PHI	125.11	0.5283	0.54661396	-0.01831396
DEN	162.02	0.5877	0.58485272	0.00284728	PIT	131.17	0.6439	0.55289212	0.09100788
DET	-10.27	0.394	0.40636028	-0.01236028	SD	49.54	0.5191	0.46832344	0.05077656
GB	154.8	0.6658	0.5773728	0.0884272	SF	108.24	0.4846	0.52913664	-0.04453664
HOU	120.09	0.5127	0.54141324	-0.02871324	SEA	155.14	0.5596	0.57772504	-0.01812504
IND	38.48	0.6252	0.45686528	0.16833472	TB	-33.48	0.394	0.38231472	0.01168528
JAX	-54.87	0.3316	0.36015468	-0.02855468	TEN	71.84	0.4565	0.49142624	-0.03492624
KC	63.49	0.4502	0.48277564	-0.03257564	WAS	-26.95	0.4158	0.3890798	0.0267202

The relationship between team quality and team winning percentage is shown in Figure 4. As expected, R-squared is equal to .365, meaning the R-squared values for the relationship between a.) team quality and team winning percentage and b.) team grade and team winning percentage add up to equal roughly one, so that all of the variation in team winning percentage is explained by variation in either team quality or team grade.

Figure 4



Since the reason NFL teams seek to win as many games as possible is so they can a.) qualify for the postseason, also known as the playoffs and b.) have a shot at winning the Super

Bowl, I also examined the relationship between team quality and whether or not a team makes the playoffs, as well as the relationship between team quality and whether or not a team wins the Super Bowl. I did this by creating two dummy variables for each season in my sample equal to 1 if the team made the playoffs/ won the Super Bowl in that season and 0 if they didn't. Once I took the ten year averages, I was left with percentages as my playoff and Super Bowl variables. A team with a playoff percentage of .5 made the playoffs in five of the sample's ten seasons, for instance, while a team with a Super Bowl percentage of .2 won the Super Bowl in two of the sample's ten seasons. Figure 5 shows the relationship between team quality and playoff percentage, while Figure 6 shows the relationship between team quality and Super Bowl percentage.

Figure 5





Figure 6



These relationships, while positive (the direction I anticipated), were not as strong as the relationships team grade had with both playoff and Super Bowl percentage, as shown in Figures 7 and 8 respectively.

Figure 7





Figure 8



While higher R-squared values indicate that more variation in team winning, playoff, and Super Bowl percentage is explained by variation in talent (i.e. team grade), team quality also contributes to each and thus must be controlled for when measuring DPV's impact on PP. I suspected there to be a positive relationship between team quality and PP, given the positive relationship between both a.) team quality and winning percentage and b.) team grade and winning percentage. Like DPV, team quality is static in my model: for each year a player played for a particular team, they were assigned that team's value for team quality. Although I suspect that teams vary in quality from year to year, making the team qualities players were assigned inexact, I have a large enough sample size of players to expect the disturbances on either side of this variation (i.e. a team with a high quality having a particularly bad year the season a particular player played for them or a team with a low quality having a particularly good year the season a particular player played for them) to cancel each other out.

Once I have measured PP and regressed it on DPV, the QB dummy, the interaction term, and my control for team quality, my resulting regression is as follows:

$$\ln PP = \beta_0 + \beta_1 \ln DPV + \beta_2 QB + \beta_3 QB * \ln DPV + \beta_4 TeamQuality + U$$

Where the β values are parameters to be estimated and U is a random disturbance assumed to a.) be uncorrelated with $\ln DPV$, QB , and $TeamQuality$ and b.) have zero mean and constant variance.

vi.) Constraining the Model

Additionally, I further constrained my model to include drafted players only (by constrain, I mean limit). When a player does not get drafted in a particular draft, his DPV can be thought of as zero, since DPV declines throughout the draft and ultimately approaches zero at the draft's end. But because a DPV of zero would not enable me to run a log-log regression between PP and DPV, I had to artificially assign a draft pick number to each undrafted player, which was determined by adding one to the final pick of the NFL Draft the year they went undrafted. This produced a positive (albeit incredibly small) DPV for each undrafted player. Since DPV should be zero for undrafted players, I constrained my model by dropping them altogether. However, I do not ignore the unconstrained model: instead, I run separate regressions and report results for both the unconstrained and constrained models in Section IV.

IV.) Results


i.) General Overview

First, I narrowed my focus to compare quarterbacks only to the four position groups most similar to them in grade distribution. These were found to be tight ends, offensive guards, centers, and edge rushers. On average, they had a higher PPS (and thus PP) than quarterbacks but a much lower DPV. Second, I created scatter plots measuring the impact of DPV on PPS for both quarterbacks and non-quarterbacks, both with and without constraining the model to include drafted players only. Constraining the model increased the value for R-squared on the non-quarterback scatter but not the quarterback scatter. Finally, I ran a log-log regression for each model with a dummy QB variable and interaction term, plus a control for team quality. In each model, the coefficient on $\ln DPV$ was statistically smaller than 1, indicating that PP is relatively inelastic to changes in DPV. The coefficients on the rest of the variables were not statistically significant in either model, but all coefficients had their expected sign.

ii.) Narrowing My Focus

Before regressing PP on DPV, I needed to determine which of the 10,063 observations of the total sample to use. Since performance grades vary widely among positions, I only wanted to consider players which played positions with a similar performance distribution to quarterbacks. The sample is broken down by position in Figure 9. 13 position groups were included: quarterbacks, wide receivers, tight ends, halfbacks, fullbacks, offensive tackles, offensive guards, centers, edge rushers, defensive interiors, linebackers, cornerbacks, and safeties. (Kickers and punters were excluded from the sample because snap count data was unavailable.)

Figure 9



To determine which position groups I would compare alongside quarterbacks, I found the mean and standard deviation of grade, snap count, and PPS for each of the 13 position groups. The results are shown in Table 2. Since quarterbacks had the highest mean grade, I decided to consider only those positions whose mean grade was at least fifty percent of the mean grade of quarterbacks. This distinction applied to four non-quarterback positions, which are highlighted along with quarterbacks: tight ends, offensive guards, centers, and edge rushers. I initially planned to apply the same rule to the standard deviation of grade, along with the mean and standard deviation of both snap count and PPS. Thus, in order to be considered, a position group would need to have a mean and standard deviation of grade, snap count, and PPS equal to at least

fifty percent of the mean and standard deviation of grade, snap count, and PPS of quarterbacks. However, these four position groups all met each of the other five criteria, so further restriction was unnecessary.

Table 2

Position	N	MGrade	StGrade	MSnap	StSnap	MPPS	StPPS
QB	387	6.127906977	17.87546721	816.0361757	273.6097193	0.004117576132	0.02128253164
WR	1119	2.290884718	8.187136329	664.3923146	227.5480385	0.002126042879	0.01189091925
TE	632	3.11028481	9.870149265	612.7531646	231.6729329	0.003706636749	0.01580413473
HB	606	-0.3102310231	7.516663205	470.1039604	172.9762724	-0.00138048253	0.01601100244
FB	264	1.631818182	6.461713132	307.1439394	114.60374	0.003454867764	0.02145339272
T	775	2.656903226	17.32034698	838.123871	261.06912	0.001001858394	0.02145557903
G	789	5.55095057	15.30132338	811.5703422	264.116229	0.004905967548	0.01905257042
C	367	4.411989101	14.28254956	884.8283379	253.828974	0.003595617785	0.01682282554
ED	996	4.173493976	15.33177409	639.7821285	223.1197432	0.004653189745	0.02272521782
DI	1197	-0.01779448622	14.63081205	534.663325	192.6567149	-0.00237689010	0.02382381029
LB	955	-0.297591623	10.86295922	723.1623037	258.8586581	-0.00100165440	0.01519242995
CB	1105	2.175475113	8.242847436	709.1855204	254.1730356	0.001995582599	0.01185181518
S	871	1.243053961	7.538718105	766.9908152	262.4684588	0.0009770029270	0.009944364932
ALL	10063	2.239829077	12.4551716	680.3767266	267.6764745	0.001614472753	0.01790001804

iii.) Comparing Quarterbacks to Non-Quarterbacks

Once I had my four non-quarterback position groups chosen, I pooled them together as a single non-quarterback group. I then found the mean and standard deviation of grade, snap count, PPS, pick number, DPV, and team quality for each group, and created a dummy variable which equals 1 if the player was undrafted and 0 if they were drafted. (The average of this variable is the proportion of each group that was undrafted: this variable would prove useful once I constrained the model to include drafted players only.) Each group's mean PP can be calculated by adding .08 to the group's mean PPS, and the standard deviation of each group's PP simply equals the standard deviation of the group's PPS. The results are summarized in Table 3.

Table 3

QB (N=387)	Mean	St. Deviation
Grade	6.127906977	17.87546721
SnapCount	816.0361757	273.6097193
PPS	0.004117576132	0.02128253164
PickNum	74.08010336	90.08893874
DPV	0.3822760507	0.3857578607
TeamQual	-0.004142321034	0.064731857
Undrafted	0.09302325581	0.2908410298
Non-QB (N=2784)	Mean	St. Deviation
Grade	4.353951149	14.14286223
SnapCount	714.6350575	262.7527329
PPS	0.004370535752	0.01953362482
PickNum	115.0193966	85.04708373
DPV	0.1102621573	0.1738179649
TeamQual	-0.000948917955	0.06470066724
Undrafted	0.1505028736	0.3576278777

The four most important numbers in Table 3 are the mean PPS of quarterbacks and non-quarterbacks and the mean DPV of quarterbacks and non-quarterbacks. Mean PPS is slightly higher for non-quarterbacks, indicating non-quarterbacks have, on average, higher performance than quarterbacks. But mean DPV is significantly lower for non-quarterbacks. In fact, the average non-quarterback has less than a third of the DPV of the average quarterback, despite performing slightly better. This further led me to believe that PP is less elastic to changes in DPV for quarterbacks than non-quarterbacks.

iv.) Constraining the Model

My next step was to create scatter plots for both quarterbacks and non-quarterbacks to see if DPV has an obviously different effect on PPS (and thus PP) for the two groups. Figures 10 and 11 are scatter plots comparing DPV and PPS for the two groups when the model is unconstrained, i.e. includes undrafted players.

Figure 10

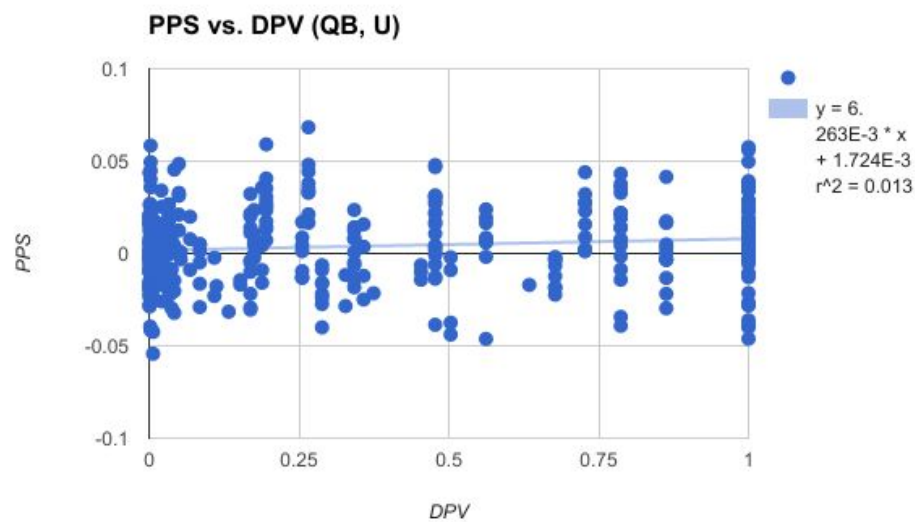



Figure 11



DPV doesn't seem to have a strong effect on PPS for either group, although there appears to be a slightly stronger relationship between DPV and PPS for non-quarterbacks as evidenced

by the larger value of R-squared. I then constrained the model to include drafted players only and created constrained scatter plots of the two groups, shown in Figures 12 and 13.

Figure 12





Figure 13



Although constraining the model increases R-squared for non-quarterbacks, it does not change R-squared for quarterbacks. I took this as more evidence that PP was less elastic to changes in DPV for quarterbacks than non-quarterbacks, particularly in the constrained model.

v.) Running my Regression

My final step was to run a log-log linear regression for each model (constrained and unconstrained). Table 4 shows the regression table for the unconstrained model, which yields the following equation:

$$\ln PP = -2.456 + 0.013 \ln DPV - 0.028 QB - 0.002 QB * \ln DPV + 0.074 TeamQuality + U$$

Table 4

SUMMARY OUTPUT								
Regression Statistics								
Multiple R	0.10542321							
R Square	0.01111405							
Adjusted R Square	0.00986467							
Standard Error	0.26286438							
Observations	3171							
ANOVA								
	df	SS	MS	F	Significance F			
Regression	4	2.45867244	0.61466811	8.89564024	3.853E-07			
Residual	3166	218.763257	0.06909768					
Total	3170	221.22193						
	Coefficients	Standard Error	t Stat	P-value	Lower 95%	Upper 95%	Lower 95.0%	Upper 95.0%
Intercept	-2.4549783	0.01006583	-243.89234	0	-2.4747145	-2.4352421	-2.4747145	-2.4352421
LnDPV	0.01272217	0.00230386	5.52210931	3.6201E-08	0.00820496	0.01723938	0.00820496	0.01723938
QB	-0.0283008	0.021224	-1.333432	0.182486	-0.0699149	0.01331342	-0.0699149	0.01331342
QB*LnDPV	-0.0017491	0.00579859	-0.3016416	0.7629451	-0.0131185	0.00962027	-0.0131185	0.00962027
TeamQual	0.07431839	0.07237578	1.02684053	0.30457401	-0.0675898	0.21622656	-0.0675898	0.21622656

I also created a constrained model, dropping players who went undrafted by NFL teams and thus had a DPV of zero. Table 5 shows the regression table for the constrained model, which follows the exact same linear regression as the unconstrained model, just with a smaller sample size. It yields the following equation:

$$\ln PP = -2.439 + 0.019 \ln DPV - 0.039 QB - 0.003 QB * \ln DPV + 0.058 TeamQuality + U$$

Table 5

SUMMARY OUTPUT								
Regression Statistics								
Multiple R	0.1326479							
R Square	0.01759547							
Adjusted R Square	0.01614596							
Standard Error	0.26097638							
Observations	2716							
ANOVA								
	df	SS	MS	F	Significance F			
Regression	4	3.30706198	0.82676549	12.1389171	8.8088E-10			
Residual	2711	184.642603	0.06810867					
Total	2715	187.949665						
	Coefficients	Standard Error	t Stat	P-value	Lower 95%	Upper 95%	Lower 95.0%	Upper 95.0%
Intercept	-2.439018	0.01101618	-221.40321	0	-2.460619	-2.4174171	-2.460619	-2.4174171
LnDPV	0.01941875	0.0029999	6.47313239	1.1354E-10	0.01353643	0.02530108	0.01353643	0.02530108
QB	-0.038957	0.02197354	-1.7729057	0.07635659	-0.0820436	0.00412957	-0.0820436	0.00412957
QB*LnDPV	-0.0033811	0.00718055	-0.4708755	0.63776752	-0.0174611	0.01069877	-0.0174611	0.01069877
TeamQual	0.0584192	0.07843854	0.74477671	0.45647127	-0.0953862	0.21222458	-0.0953862	0.21222458

vi.) Interpretation of the Coefficients

In the unconstrained model, a one percentage point increase in DPV increases a player's expected PP by .013-.002QB percentage points, meaning expected PP increases by 0.013 percentage points if the player is not a quarterback and 0.011 if he is. Being a quarterback, meanwhile, decreases a player's expected PP by 0.028+0.002LnDPV percentage points. Increasing team quality by 0.1 (about the difference between the worst team and an average team) will increase expected PP by 0.743.

In the constrained model, a one percentage point increase in DPV increases a player's expected PP by 0.019-0.003QB percentage points, meaning expected PP increases by 0.019 percentage points if the player is not a quarterback and 0.016 if he is. Being a quarterback, meanwhile, decreases a player's expected PP by 0.039+0.003LnDPV percentage points. Increasing team quality by 0.1 (about the difference between the worst team and an average team) will increase expected PP by 0.584.

All coefficients have their expected sign. In both models, the coefficients on $\ln DPV$ and $TeamQuality$ are positive, while the coefficients on QB and $QB*\ln DPV$ are negative. However, of the eight coefficients total ($\beta_1 - \beta_4$ in the unconstrained model and $\beta_1 - \beta_4$ in the constrained model), only two are statistically significant at the 5% level (P value equal to or less than 0.05). These are the coefficients on $\ln DPV$ in both the constrained and unconstrained model. At the 10% level (P value equal to or less than 0.10), the coefficient on QB in the constrained model is also statistically significant.

Since β_1 is statistically different from zero in both models, I can conclude that it is also statistically different from (and smaller than) 1 in both models. The reason I am able to do this is because the estimated values for β_1 are much closer to zero than they are to one, meaning that if β_1 is statistically different from zero it must also be statistically different from 1.

V.) Conclusions

i.) General Overview

My findings are consistent with previous research on the relationship between DPV and PP in the NFL. PP is relatively inelastic to changes in DPV, indicating a market inefficiency allowing teams to increase the expected PP of the players they select by continuing to trade down in the draft. There is insufficient evidence, however, to conclude that this relationship differs between quarterbacks and non-quarterbacks, or that team quality impacts PP. These findings have implications both inside and outside the NFL. My research is limited, however, by the lack of a cost control, the lack of other potentially relevant controls, and the potential unreliability of my measurements for DPV, PP, and team quality. Because I suspect the performance of quarterbacks to be less elastic to changes in DPV, performing my regression with more controls and better measurements for DPV and PP might allow me to conclude this by making the QB coefficients statistically significant. I could also devise a better measurement for team quality, making this coefficient statistically significant as well. Alternatively, if it's true that team and individual performance are not strongly related, I could extend my research by studying this relationship in greater detail.

ii.) Key Findings

Since the coefficient on $\ln DPV$ is statistically less than 1 in both of my models, I conclude that PP is relatively inelastic to changes in DPV. This indicates that the market for NFL draft picks is inefficient: high draft picks are overvalued, and there is a systematic advantage to teams trading down in draft pick trades. There is insufficient evidence to suggest that this relationship differs between quarterbacks and non-quarterbacks. While β_2 is statistically

significant at the ten percent level in the constrained model, it is not statistically significant at the five percent level, nor is it statistically significant at either the five or ten percent level in the unconstrained model. There is also insufficient evidence to suggest that team quality impacts PP. The magnitude of β_4 is economically significant in both models: a 0.1 increase in team quality causes expected PP to increase by 0.743 in the unconstrained model and 0.584 in the constrained model, which are both much greater than the difference between the highest and lowest PP in the sample (roughly 0.168). However, the coefficient itself is not statistically significant in either model at either the five or ten percent level.

iii.) Implications Within the NFL

First, PP being relatively inelastic to changes in DPV confirms the findings of previous research on the relationship between PP and DPV. Late draft picks have significantly less value in the market than early draft picks despite the expected performance of players taken with those picks being only slightly lower, which means that the team “trading down” should be expected to gain more PP with the picks it acquires than the team “trading up.” Since draft pick trades are assumed to have equal value on each side, a team that trades down deals a few highly valued picks for a greater number of picks that are less valued. But since the PP of the later picks is only slightly lower than the PP of the earlier ones, the team trading down gains more PP through acquiring more draft picks. This has its limitation, of course, in that teams must trim their roster to 53 players before the regular season begins, but if a team is well below that limit it can, theoretically, increase the expected PP of its draft picks simply by continuing to move down in the draft.

Second, the lack of statistical significance on most of the QB related coefficients suggests that perhaps the PP of quarterbacks is no less elastic to changes in DPV than the PP of non-quarterbacks, contrary to anecdotal evidence which points to a less elastic relationship for quarterbacks. This may be due to the fact that quarterbacks receive more media attention than non-quarterbacks do, causing cases of late drafted quarterbacks succeeding (such as Tom Brady) and early drafted quarterbacks failing to succeed (such as JaMarcus Russell) to be more well known than analogous cases involving non-quarterbacks.

Third, the lack of statistical significance of the team quality coefficient in each model suggests that perhaps the quality of the team a player is on does not impact PP. This implies that teams may be ineffective at increasing the PP they get from their players through effective coaching. It also suggests that perhaps team and individual performance aren't as highly correlated as I had previously thought, since team quality has an impact on the former but apparently not the latter. Evidence to suggest that this may indeed be the case is found by studying the pick value chart, or PVC, a chart designed by NFL head coach Jimmy Johnson in the 1990's to help NFL executives create fair draft pick trades. An alternative method of measuring DPV, the PVC simply assigns each pick number a value ranging from 3,000 to 2. Using this method produces a moderately strong positive correlation between a team's DPV in a particular Draft (the summation of the values of all its draft picks) and its change in winning percentage from the year prior to the Draft to the year after (Bonds, et al., 2015). But despite its continued use in draft day trades, the PVC has been found to be an inaccurate predictor of individual success (Barney, et al., 2013), indicating a weaker than expected relationship between team and individual performance.

iv.) Implications Outside the NFL

In addition to confirming the findings of the economic literature on the relationship between PP and DPV, my research also confirmed the findings of the psychological literature on human overconfidence. Because teams overvalue early draft picks, they believe their predictive abilities are greater than they actually are, since they overvalue the opportunity to have more players available to choose from. It's easy to see how this could create inefficiencies outside the NFL. Suppose a firm is considering hiring two different workers for the same position, worker A and worker B. If the firm believes there is a sufficiently high chance that worker A will be more valuable to the company than worker B, it will offer worker A a much greater salary. However, if the firm is overconfident in its predictive abilities, it might have severely overestimated this probability, meaning that salary it should offer worker A is much lower than the salary it actually offers worker A.

The idea that team and individual performance are not strongly related also has implications outside the NFL. Much like NFL teams draft players with the intention of increasing their performance, firms hire workers with the intention of increasing their profits. But if worker productivity and profits are also weakly related, then perhaps firms are not effectively maximizing their profits through their hiring behavior. It might also be the case that firms are unable to impact worker productivity through effective motivation, much like teams appear to be unable to impact PP through effective coaching. There are obviously differences between firms in most industries and NFL teams, but each is attempting to generate group success through generating individual success, a strategy that is flawed if the two are not strongly related to each other.

iv.) Possible Limitations

The biggest limitation of my research was my lack of a cost variable. NFL players are paid a yearly salary, and this salary contributes to the team's salary cap: each year, each NFL team receives an equal amount of revenue through the league's revenue sharing program to spend on player salaries. Thus, controlling for player salary would enable me to draw conclusions about the cost effectiveness of particular draft picks: a player taken later in the draft may perform slightly worse but be paid a lot less in salary, thus making a late draft pick more cost effective than an early draft pick. Since salary declines throughout the draft, almost without exception, it is likely positively correlated with DPV. I also expect it to be positively correlated with PP, indicating that better players are being paid more on average. Thus, including a salary variable in my regression would decrease the impact that DPV has on PP. If early round draft picks are already overvalued (as my research indicates) then including a cost variable would make them even more overvalued due to the fact that these players must be paid a higher salary, perhaps even to the point of making the coefficient on $\ln DPV$ negative.

Additionally, I could have included controls for a player's age and the quality of the team that drafted him, as opposed to the team he played for in a given season. PP might increase as a player ages initially and gets more NFL experience but decrease once he reaches a certain age and begins losing his speed and athletic ability, making the overall effect of age on PP ambiguous. The team quality variable I included only applied to the team the player was on during the sampled season, but a player could have been drafted by a team with a much different team quality than the one he played for. Since high quality teams likely do a much better job of developing their draft picks than low quality ones, I would expect this variable to be positively

correlated with PP and negatively correlated with DPV (high quality teams win a lot of games and thus have lower draft picks), causing it to increase the impact DPV has on PP if I include it in my regression. However, since the control for team quality I did include did not have a statistically significant impact on PP, it could be the case that another type of team quality would also not have a statistically significant impact on PP.

Additionally, the measurements I used for DPV and PP may have been inaccurate themselves. I mentioned in Section III.ii that the exact value of a particular pick in the Draft varies depending on the quality of players available. Instead of simply ignoring this variance, I could have included a measure for the quality of players available and included it in my DPV measurement. The DPV formula I use is also based on old data: team behavior might have changed significantly in recent years, causing the values of λ and β to change as well. As for PP, team quality's lack of an impact on PP indicates that perhaps individual performance data from PFF is not the best way to measure it. Rather than concluding that team and individual performance are unrelated, I could have used an individual performance measurement more closely related to team performance, which would have helped me determine the extent to which teams are succeeding at maximizing their performance through the draft. I could also improve my measure for team quality by measuring it each season rather than basing it on ten year averages for each team. Finally, standardizing performance grades within each of the 13 position groups would have allowed me to measure my entire initial player sample rather than just five position groups by resulting in a uniform distribution in grade among all positions. I could have also done this by including a dummy position variable for all positions, not just quarterbacks.

v.) Future Research Opportunities

Despite my findings, I still suspect that there may be a different relationship between PP and DPV for quarterbacks and non-quarterbacks. Because all the quarterback coefficients are negative and the coefficient on β_2 in the constrained model is statistically significant at the ten percent level, I have some evidence to *suggest* that PP is less elastic to changes in DPV for quarterbacks than non-quarterbacks. There just isn't enough evidence to *conclude* that it's less elastic. Running the same regression with the controls and/ or improvements mentioned in Section V.iv might enable me to conclude that it's less elastic by making the QB coefficients statistically significant. The same is true for team quality: I suspect that there may indeed be a relationship between team quality and PP, so running a regression with an improved measure for team quality might make the coefficient on β_4 statistically significant.

On the other hand, my findings on the relationship between team quality and PP might indeed be accurate: it could be the case that a team's quality has little to no impact on the performance of its players. More research can be done on the specific factors that contribute to team quality, rather than simply lumping them all together into a single measurement. Studying the precise relationship between team and individual performance, meanwhile, could yield valuable implications regarding the methods currently used by teams to increase their performance through increasing the performance of their players. Are they effective at increasing team performance? If so, which ones are the most effective? Such questions can also be applied to scenarios outside the NFL. Examining the relationship between group success and individual success may allow one to determine whether or not the idea of increasing the former through promoting the latter has any merit, and if so, which methods are the most effective at accomplishing this goal.

VI.) Bibliography

Abbes, M.B. "Does Overconfidence Bias Explain Volatility During the Global Financial Crisis?"

Transition Studies Review 19.3 (2013): 291-312. *Springer Link*. Web. 31 March 2017.

Barney, Casey, Anthony Caravella, Michael Cullen, Gary Jackson, and Craig Wills E.

Evaluating Talent Acquisition via the NFL Draft. Thesis. Worcester Polytechnic Institute, 2005.

N.p.: UMI, n.d. *ProQuest Central*. Web. 25 Oct. 2016.

Berri, David J and Rob Simmons. "Catching a Draft: On the Process of Selecting Quarterbacks

in the National Football League Amateur Draft" *Journal of Productivity Analysis* 35.1 (2009):

37-49. *Springer Link*. Web. 25 Oct. 2016.

Bonds, Traci, Carrie LeCrom W, Zachary Reynolds, and Steele Thompson. "Deconstructing the

Draft: An Evaluation of the NFL Draft as a Predictor of Team Success" *Journal of Applied Sport*

Management 7.3 (2015): 74-89. *ProQuest Central*. Web. 25 Oct. 2016.

Hughes, Andrew, Cory Koedel, and Joshua Price A. "Positional WAR in the National Football

League" *Journal of Sports Economics* 16.6 (2015): 597-613. *Sage Journals*. Web. 25 Oct. 2016.

Kitchens, Carl T. "Are Winners Promoted Too Often? Evidence From the NFL Draft

1999–2012." *Economic Inquiry* 53.2 (2015): 1317-330. *Wiley Online Library*. Web. 25 Oct.

2016.

Malkiel, Burton G. "The Efficient Market Hypothesis and Its Critics" *The Journal of Economic*

Perspectives 17.1 (2003): 59-82. *Ingentaconnect*. Web. 5 April 2017.

Massey, Cade and Richard Thaler H. "The Loser's Curse: Decision Making and Market

Efficiency in the National Football League Draft" *Management Science* 59.7 (2013): 1479-1495.

INFORMS PubsOnline. Web. 25 Oct. 2016.

Oskamp, Stuart. "Overconfidence in Case-Study Judgments" *Journal of Consulting Psychology* 29.3 (1965): 261-265. *ProQuest Central*. Web. 31 March 2017.

Politi, Steve. "'Suck for Luck' Could Be Best Hope for NFL's Worst." *CNN*. Cable News Network, 21 Oct. 2011. Web. 05 Apr. 2017.

Tang, Meng-Chi. "Contract Length, Expected Surplus, and Specific Investments: Empirical Evidence From the National Football League" *Journal of Sports Economics* 16.3 (2015): 295-311. *Sage Journals*. Web. 25 Oct. 2016.