Positional WAR in the National Football League

Andrew Hughes
University of Missouri
Ajhvx8@mail.missouri.edu

Cory Koedel
University of Missouri
909 University Ave
318 Professional Building
Columbia, MO 65211
(573) 882-3871
koedelc@missouri.edu

Joshua A. Price
Southern Utah University
351 W. University Blvd
Cedar City, UT 84720
(435) 586-8682
japrice@suu.edu

Positional WAR in the National Football League

Abstract

We empirically estimate positional "wins above replacement" (WAR) in the National Football League (NFL). Positional WAR measures the value of players in the NFL, by position, in terms of generating wins. WAR is a commonly used metric to evaluate individual players in professional baseball and basketball in the United States, but to the best of our knowledge this is the first study to construct WAR measures for American football. A key challenge in constructing these measures is that individual statistics for many football players are not as well-developed as in baseball and basketball. Related to this point, the productivity of individual football players, perhaps more than players in any other major sport, is highly dependent on context. We circumvent issues related to measuring productivity for individual players by constructing WAR measures at the position rather than individual level. The identifying variation that we leverage in our study is generated by arguably exogenous player injuries and suspensions. Using data from three seasons and all 32 NFL teams, we show that the most valuable positions in the NFL are quarterback, wide receiver, tight end/fullback and offensive tackle. Perhaps our most surprising finding is that positional WAR for all positions on the defensive side of the football is zero.

1. Introduction

"Wins Above Replacement" (WAR) and similar, data-driven metrics of player productivity are increasingly used by front-office personnel for professional sports franchises to evaluate players (Barnwell, 2014; Fleming, 2013; London, 2014). These metrics are statistical constructs based on individual and team performance measures, and their incorporation into player assessment represents a fundamental shift in how professional franchises evaluate athletes (DuPaul, 2012). Many attribute the move toward statistics-based analyses of athletes to the publication of *Moneyball: The Art of Winning an Unfair Game* by Michael Lewis in 2004. The book details how the Oakland Athletics used statistical models to evaluate players. The intuition can be explained as follows: Team wins are a function of runs scored and runs allowed, and runs scored and runs allowed are functions of individual player statistics. Thus, each individual statistic can be expressed as a contribution to team wins, or marginal productivity of labor. Though popularized by the book, this type of model was formally presented by Blass (1992) over a decade earlier with an application to player wages.

Building on the approach popularized in *Moneyball*, Gerard (2007) discusses the challenges of applying statistical measures of player productivity to team sports where individual statistics are dependent on teammate performance. He presents a conceptual model and applies it to soccer players in the English Premier League. A similar approach has also been applied to basketball players by Berri, Schmidt and Brook (2007) and Hollinger (2014). Berri, Schmidt and Brook use a regression framework to create a metric called "Wins Produced." Hollinger's "Estimated Wins Added" converts individual player statistics into measures of team wins.[1] A conceptual similarity between both of these basketball metrics is that they compare each player's

[1] Based on Hollinger's popularized "Player Efficiency Rating" (PER) metric.

contribution to a hypothetical replacement player, thus the “Wins Above Replacement” terminology.

All currently estimated WAR measures of which we are aware are based on the idea that individual statistics can be translated into measures of contributions to team wins. However, linking individual player statistics to wins is particularly challenging in American football. As discussed by Gerard (2007), an individual player’s statistics in football depend on the performance of other teammates more so than in other sports. For example, consider a one-yard touchdown run by a running back. The touchdown is officially attributed to the running back, but how much of that touchdown is due to his performance and how much is due to the blocking performance of the offensive line (the latter being particularly difficult to measure with individual player statistics)? And what if the play was set up by an 80 yard punt return? The interdependence of individual player performance in American football makes it difficult to use individual statistics to measure the marginal productivity of labor.

These challenges notwithstanding, some productivity metrics have been developed for players in the National Football League (NFL). These metrics are summarized by Berri and Burke (2012), but they all assume that individual player statistics are independent of the contributions of teammates. Realizing this issue, Oliver (2014) created a player-productivity metric for NFL quarterbacks called Total Quarterback Rating (Total QBR), which is used by ESPN. Oliver (2014) describes Total QBR as a statistical measure that incorporates the context and details of available statistics and what they mean for wins. However, Total QBR is a proprietary measure and the extent to which it truly accounts for context and player interdependence in the NFL is unknown. A different approach to estimate player productivity is taken by Aktinson et al (1988). Instead of using individual player statistics, they use team statistics to estimate the probability of winning. However, due to their aggregating the statistics to the team level, it is not possible to determine the marginal product of labor for individual players or positions.

The contribution of the present study is to add to the quickly growing, empirically driven literature on player productivity by developing a new, fully-contextualized measure of player productivity in the NFL at the *position* level. We call this new measure "Positional WAR." By constructing our WAR measure at the position rather than individual-player level, we can circumvent many of the above-described challenges associated with attributing wins to individual players, most notably because we estimate positional WAR without relying on individual game statistics for players (see below). Positional WAR can be a useful tool for teams looking for quantitative, objective measures of positional value. It can help to inform decisions about how to invest scarce salary-cap dollars, draft choices, and roster spots.[2] Indirectly, positional WAR represents a first step toward the construction of rigorous, summative performance measures for individual NFL players.

The identification strategy that we employ to estimate positional WAR relies on player injuries and, to a lesser extent, suspensions. Games lost by starting players due to injury and suspension are arguably exogenous and inherently unpredictable. By definition, an injured or suspended player must be replaced. We estimate positional WAR by determining how the use of replacement players for injured and suspended starters affects season win totals for all 32 NFL teams across three seasons.

Our estimates of positional WAR measure the marginal contributions of starters compared to replacement players by position. The measures depend on the relative quality of

[2] We assume that teams want to win. For most teams, wins directly affect revenue (Landin, 2012; Burger, 2003).

starters and replacement players within positions and thus reflect the scarcity of talent at the position level (in contrast, positional WAR does not directly compare the absolute value of player performance across positions, which exists conceptually but would be difficult to measure because players at different positions do different things). If, for example, a team has three defensive tackles of equal quality yet only two can start, the difference in productivity between the starters and the backup would be small, hence a case with low positional WAR. Alternatively, if a starting wide receiver is significantly better than his second-string backup, positional WAR would be high.

Our analysis reveals some predictable patterns. For example, we find that quarterbacks are the most valuable position by a wide margin. Other positions for which we find significantly positive WAR values are wide receiver, tight end/fullback and offensive tackle. Perhaps our most interesting finding is that no position-group on the defensive side of the football has a significant WAR value. Put differently, on average, teams do not suffer in terms of wins when defensive starters miss games due to injury and/or suspension. This suggests that relative to their replacements, defensive starters on average are not as valuable as offensive starters at several positions.

2. Data

We construct a data panel for all 32 NFL teams containing performance (wins) and injury/suspension information for three years: 2008, 2010, and 2012.[3] We use even-numbered years instead of all years to reduce the data collection burden; collecting the injury and suspension data is labor intensive (see Appendix A). We spaced out the three years of data

[3] Obtaining game-by-game injury data prior to 2008 is difficult (especially for non-skill positions).

collection to reduce the intra-team correlation in outcomes in the data panel, which improves statistical power.

The independent variables of interest in our analysis measure games lost to injury and suspension, by position, for each NFL team in each year. To build these variables we first identified the projected starting lineup for each team in each year prior to the first game of the season.[4] We then determined the number of games that each projected starter lost to injury and/or suspension using a variety of sources (e.g., news articles, fantasy football updates, and injury reports).[5] Of course, some starters missed games for other reasons, such as a demotion; however, our identification strategy relies entirely on using variation from games missed due to injury and suspension, so we code up only games missed for these reasons.

To construct the dataset we began by determining the number of projected starters at each position for each team. All defenses use two starting cornerbacks and two starting safeties, and either a base 4-3 or 3-4 system.[6] On offense, every team has one starting quarterback, one starting running back, and the traditional five-man offensive line. Starting wide receivers, tight ends, and fullbacks vary. For example, some teams do not use a fullback in the starting lineup; some use two tight ends; some use three wide receivers, etc. No team includes more than one fullback, more than two tight ends, or more than three wide receivers in the starting lineup.[7]

---

[4] We checked projected starters early during training camp to verify that no potential starters had a pre-season injury.

[5] See appendix A for a detailed example of how we compiled the injury/suspension data.

[6] A team that uses a 4-3 system will start four defensive linemen and three linebackers; a team that uses a 3-4 system does the reverse. We did not come across a single instance in any year where a team could not be classified appropriately by using either scheme for their base-defense. We handled this difference in defensive formations through an arithmetic adjustment, averaging the number of starters in each position over all of our observations. Differences in offensive schemes resulting in varying numbers of starters were handled in a similar fashion.

[7] In situations where a team might regularly employ, for example, four-wide-receiver sets, we still concluded that its starting offense included three wide receivers at most. For example, the 2012 Green Bay Packers periodically used four starting wide receivers. The wide receivers were James Jones, Greg Jennings, Jordy Nelson, Donald Driver, and Randall Cobb. We coded the first three as starting wide receivers for the Packers in 2012, along with tight-end

It can be argued that being a starter in football is less meaningful because of the amount of substitutions that occur during a game. For example, in a soccer match each team is allowed only three substitutions per match and as a result almost all minutes are played by starters. Alternatively, in the NFL a team can substitute on each play so a starter may get significantly less playing time. To determine the importance of being a starter in terms of predicting playing time, we collected data from 2012 on the number of snaps played by starters on a game-by-game basis. On the offensive side of the ball, starting quarterbacks and offensive linemen play 96 percent of the snaps. Starting wide receivers and running backs play 78 and 61 percent of snaps, respectively, while tight ends and fullbacks play 68 and 37 percent of snaps. On defense, starting defensive linemen play 69 percent of snaps, starting linebackers play 80 percent, and defensive backs (safeties and cornerbacks) play 91 percent of all snaps[8]. Thus, despite the prevalence of specialized player packages, starters on offense and defense in the NFL play the overwhelming majority of snaps within a game.

With the list of projected starters at each position for each team in hand, along with the injury and suspension data, we constructed measures of games lost due to injury/suspension for each team at the following positions:

> Offense: Quarterback, Running Back, Tight-End/Fullback, Wide Receiver, Interior Offensive Lineman (center, guard), Exterior Offensive Lineman (tackle)
>
> Defense: Cornerback, Safety, Linebacker, Defensive Line

Jermichael Finley. The reason is that Finley's 14 starts and 61 catches for 667 yards were far greater than Cobb's 8 starts and 292 yards or Driver's single start with 77 yards receiving.

[8] Data on the percentage of snaps taken by players comes from http://www.footballoutsiders.com/stats/snapcounts.

We combine some positions to improve statistical power. An example is the tight-end/fullback position. Similarly, we group interior and exterior defensive lineman. In the extensions section we discuss the robustness of our findings to alternative positional groupings, but use the above groupings for our main analysis.[9] Table 1 shows the average shares of games missed for NFL teams at each position due to injury and suspension over the course of our data panel.

3. Empirical Strategy

To estimate positional WAR in the NFL we estimate the following empirical model:

$$\begin{aligned} W_{it} &= \theta_0 + PW_{it}\theta_1 + QB_{it}\delta_1 + RB_{it}\delta_2 + TF_{it}\delta_3 + WR_{it}\delta_4 + IOL_{it}\delta_5 + EOL_{it}\delta_6 \\ &+ CB_{it}\delta_7 + S_{it}\delta_8 + LB_{it}\delta_9 + DL_{it}\delta_{10} + \varepsilon_{it} \end{aligned} \tag{1}$$

In (1), W_{it} is the number of wins for team *i* in year *t*. PW_{it} is the preseason over-under win total for team *i* in year *t*, taken from a Las Vegas sportsbook. We include the "predicted wins" variable for the sole purpose of improving the predictive power of the model, and thus the statistical precision of our estimates.[10] It is moderately valuable in this role – it explains 12 percent of the variation in win totals across the data panel. The variables QB_{it}, RB_{it}, TF_{it}, WR_{it}, IOL_{it}, EOL_{it}, CB_{it}, S_{it}, LB_{it}, and DL_{it} indicate the share of total starter games lost due to injury or suspension for team *i* in year *t* at each position (as listed in the previous section). These are the independent variables of interest. ε_{it} is the error term. We cluster the standard errors from equation (1) at the team level.

[9] Another advantage of grouping tight ends and fullbacks is that it reduces concerns about the influence of endogenous positional groupings. On the offensive side of the football, the most common positional difference across starting lineups (by far) is whether they include two tight ends or one tight end and one fullback.
[10] We also estimate a model with an additional control variable for strength of schedule based on the record of each team's opponents from the previous year. Our positional WAR findings are essentially unaffected by including this control.

The coefficients $\hat{\delta}_1 - \hat{\delta}_{10}$ estimated from equation (1) will represent unbiased, causal estimates of the value of starters over replacements, by position, under the condition that games lost due to injury and suspension are exogenous. The argument for exogeneity is compelling, but not absolute. On the one hand, it is reasonable to assume that future injuries and suspensions will be very difficult to predict *ex ante* for the vast majority of players. However, one might worry that injuries and suspensions can be predicted to some extent. Suppose, for example, that a team has a projected starter with a long and well-documented history of injuries. We might expect the team to invest in a better backup player, at least at the margin. This type of compensating behavior could lead us to underestimate positional WAR in general because it would imply that the replacement players who replace the injured players in our data are better than a replacement for a randomly-injured player. Comparisons of positional WAR across positions will be preserved, however, unless one believes that teams systematically "insure" players with injury histories more or less depending on their position.

4. Results

Table 2 shows results from the estimation of equation (1), with and without the inclusion of the predicted-wins variable. The table shows that the positions where lost starters are most important are quarterback, tight-end/fullback, wide receiver and exterior offensive lineman. For the other positions there are not statistically identifiable consequences associated with using replacement players in the face of an unexpected absence of a starter.

We note two complications with interpreting the estimates in Table 2. First, the variables of interest are coded as the percentage of games missed at each position. Thus, the strict interpretation of the coefficients is that they indicate the effect of going from 0 to 100 percent games missed on total wins. However, interpreting the estimates in this way involves extrapolating well out of the range of variation in the data, as indicated by Table 1. Second, the percentage variables for games missed by position correspond to different numbers of *total games* missed. For example, losing 50 percent of the games at quarterback corresponds to missing 8 games, whereas losing 50 percent of the games at wide receiver corresponds to missing almost 17 games (on average across teams – see Table 1).

To provide a more direct comparison of positional value based on our estimates, Table 3 shows the effect of losing exactly four games due to injury/suspension at each position based on our estimates in Table 2. We use the estimates from the model that includes the predicted-wins variable (the first set of estimates shown in Table 2). Table 3 shows that on a per-game basis, there is no position more valuable than quarterback. Players at wide-receiver, tight-end/fullback, and exterior-offensive-line positions are all similarly valuable; and losing players to injury at other positions does not affect win totals. Note that for the positions where we do not find statistically significant WAR values, statistical power is not the issue. The point estimates for the effects of losing four games to injury for running backs, interior offensive linemen, and all defensive positions are quite small.

To put our estimates into context, consider that the average number of wins in the NFL is eight. When a quarterback is injured for four games, the team is expected to lose more than one additional game, on average, which is a large effect. We also reinforce the point that our estimates of positional WAR should be interpreted to indicate positional value on average, not the value of individual players. They do not preclude situations where an individual player at a position with low positional WAR is more valuable than an individual player at a position with higher positional WAR.

5. Sensitivity Analysis

Table 4 shows estimates from a model similar to equation (1) except that we subdivide tight ends and fullbacks.[11] For brevity we only show estimates from the full model that includes the predicted-wins variable. Recall from above that not all teams use a fullback – in the data just over a third of the teams (37.5 percent) use a fullback in the starting lineup.

The extended model shown in Table 4 suggests that between tight ends and fullbacks, it is the fullback position that is more important. Note that this is only a nominal result – put differently, we cannot statistically reject the null hypothesis that tight ends and fullbacks are equally valuable. Still, the fact that our combined tight-end/fullback estimate is not driven by tight ends runs counter to what one might expect given the rising prominence of the tight-end position in the NFL (Benoit, 2012).

We can only speculate as to why fullbacks are important for wins. One possibility is that the decline of the prominence of the position in the modern NFL is driven in part by the lack of availability of starter-quality fullbacks, which would make the fullbacks we do see more valuable. It may also be that the use of a fullback implies a dearth of talent on the roster, which would impact our positional WAR estimates driven by the difference in talent between starters and reserves – put differently, maybe all teams would rather start an additional player at a different position rather than a fullback, but some teams do not have a strong "next in line" tight end or wide receiver. Finally, some circumstances might add considerable value to fullbacks; for example, it may be that there are quarterbacks who perform better with backfield blocking, which is typically a role filled by the fullback.

---

[11] In results omitted for brevity we also estimate models that identify positional WAR using variation in games missed due to injury only – that is, we ignore games lost due to suspension. Our findings are substantively unaffected by this adjustment, which is consistent with the fact that games lost due to suspension represent only a very small fraction of total games lost (see Table 1).

6. Discussion

While it is beyond the scope of the present study to perform an in-depth analysis of the mechanisms that contribute to our positional WAR findings, in this section we briefly discuss three plausible explanations for why some positions have higher WAR values than others. In particular, we discuss (1) the importance of scheme in compensating for losses in player quality, (2) positional talent scarcity, and (3) constraints that affect franchises' team-building decisions.

With regard to our findings on the defensive side of the football, one possible explanation for the lack of positive positional WAR is that defensive schemes can be adjusted to account for replacement players more easily than offensive schemes. For example, a downgrade at cornerback can be facilitated by more help from safeties and/or a broader adjustment to the coverage scheme. Losing one player on defense may not have a significant impact if the other players are able to continue functioning as a whole. This would not imply that defensive positions are insignificant; rather, it would imply that defenses overall function as a unit in which individual personnel are more easily interchangeable. It is also important to remember that our findings are estimated within the current investment equilibrium in professional football. It would be a mistake to interpret our findings to imply that a defense consisting entirely of backup players would perform no worse than a defense full of starters. Put differently, our finding that teams can fully compensate for the loss of an injured player on defense with their current personnel, on average, does not mean that a defense would not be affected by the simultaneous downgrade of numerous defensive positions at the same time.[12]

[12] The question of the compounding effects of player absences is an interesting one, but our data are ill-suited to test interaction-based hypotheses due to the 32-team clustering structure (each team is a unique source of identifying variation). Data from more teams – perhaps in college – might be used to look for compounding effects of player absences on both sides of the ball.

Still, our findings for defensive positions are not consistent with the popular “defense wins championships” mantra, although our study is not the first to suggest that offensive efficacy is more important. A more thorough look at regular season success shows that teams with better offenses win more frequently than teams with the better defenses (Moskowitz, 2012). Further research has looked at the distribution of success of past offenses and defenses and found that exceptionally talented offenses are more successful than exceptionally talented defenses (Burke, 2008). Our results complement these previous studies by showing that the value of starting players at several offensive positions is greater than that of their defensive counterparts.

In addition to potentially driving our findings for defensive players, scheme may also be important on offense as well. Consider the comparison of exterior and interior offensive linemen as an example. It may be that teams can largely compensate for the loss of a starting interior offensive lineman by changing responsibilities within the line, while for an exterior lineman compensation might require diverting other players to help as blockers, such as a tight ends, fullbacks or running backs.

A second explanation for our findings is differences across positions in the scarcity of talent. Because positional WAR inherently measures starter quality relative to the quality of replacements, if there is a high supply of talented players at a given position, then the drop off between the starter and the backup will be less and positional WAR will be small. Anecdotally, the position where talent is the most scarce is quarterback, and unsurprisingly we find the largest positive WAR for the quarterback position by a wide margin.

A third explanation for our findings relates to the constraints that NFL franchises operate within to build their teams. These constraints include the salary cap, the cap on roster size, and limited draft picks, all of which force teams to make talent tradeoffs throughout their rosters. For example, with regard to the salary cap, in 2008 the average quarterback salary was $3.47 million dollars (median salary of $1.6 million) and the average salary for a running back was $1.67 million (median salary of $755 thousand). Thus, on average, a team could sign two running backs for the same price as one quarterback. Similarly, the mean salary of a defensive cornerback was $2.06 million (median of $933 thousand), or about 41 percent less than a quarterback.

We use the following example to illustrate how these constraints might affect roster decisions. Consider two teams, A and B, where Team A has two quarterbacks that are of higher quality than the starting quarterback for Team B. The value marginal product (VMP) of second-best quarterback on team A, serving as a backup, will be lower than the VMP of that same player as a starter on team B. The divergence is caused by the expected number of snaps played, which per above is significantly higher for starters. Therefore, while Team A would prefer to carry both quarterbacks on their roster, the willingness to pay of teams who would use that player in a starting role will be higher (whether in terms of salary or other scarce resources, like draft picks). The differential value that the same player might provide to two separate teams is discussed by Leeds and Kowalewksi (2001) in the context of understanding how free agency can affect player salaries. With each team facing a budget and roster size constraint, it is not an equilibrium outcome (subject to contract rigidities) to have players serving as non-starters when inferior players are starting for other teams.[13]

[13] Interestingly, this scenario has played out several times in recent years. After the 2006 season, the Atlanta Falcons had two highly-regarded quarterbacks in Michael Vick and Matt Schaub. They traded their backup, Schaub, to the Houston Texans for three draft picks in the first or second round. Schaub became the immediate starter in Houston. Another case occurred in the 2012 draft when the St. Louis Rams had the second overall pick, which could be used to choose another highly-regarded quarterback, Robert Griffin III. They already had a quarterback of starting quality, Sam Bradford, and were able to trade the second pick, which to them would have been the equivalent of a high-quality backup, to the Washington Redskins for six draft picks.

Finally, we conclude with a brief examination of whether NFL teams act as if they have knowledge of positional WAR in making their personnel decisions. We focus on how decision makers use one of the scarce resources at their disposal – first-round draft picks. Specifically, we correlate the share of first-round draft picks devoted to different positions, weighted by the inverse of the positional share in the average NFL starting lineup, with positional WAR (using the 4-game measures from Table 3). We do this for the five NFL drafts spanning our data panel: 2008-2012. The weighting by positional share is necessary because, for example, all else equal a team will draft more wide receivers than quarterbacks owing to the fact that wide receivers constitute a larger share of the starting lineup.[14]

With the important caveat that the correlation between first-round draft share and positional WAR is calculated based on a small number of items (our 10 primary positional categories), and thus is imprecise, we estimate that it is positive at 0.33. Further investigation reveals that the positive correlation is driven entirely by the overrepresentation of quarterbacks as first-round draft picks given their high positional WAR. Our correlational analysis offers suggestive evidence that NFL decision makers understand some but not all aspects of positional value. In football, like with other sports, it will be of interest in future research to monitor the responsiveness of personnel decisions to the rapidly increasing base of empirical information about player productivity.

---

[14] The "usage weighted" draft share (*UWD*) for each position *p* is calculated using the formula $UWD_p = D_p * (\frac{N_{Tot}}{N_p})$ where D_p is the unadjusted draft share over the five NFL drafts for position *p* (i.e., the number of players drafted at that position divided by the number of first round draft picks), N_{Tot} is the total number of games started on offensive and defense during the season for each team (16*22=352) and N_p is the average number of games started at position *p* in the NFL, taken from Table 1. Intuitively, this formula gives more weight to positions that take up a smaller share of the starting lineup to account for variation in team needs driven by positional share alone.

7. Conclusion

The contribution of the present study has been to develop an empirical approach for estimating positional WAR in the National Football League. Our identification strategy leverages the use of replacement players when starters miss games due to injury and suspension. Using data from all 32 NFL teams over three seasons we find that games lost by starting quarterbacks are by far the most important in terms of affecting win totals. We also show that games lost by starting wide receivers, fullbacks/tight-ends and exterior offensive lineman are important. Games lost by starters at all other positions do not affect wins, on average.

Our study represents the first rigorous attempt to quantify positional value in the NFL of which we are aware and moves us toward an improved understanding how players contribute to team wins in football. Future research can advance this line of study in several ways. One direction would be replicate our approach in other sports where individual performance measures can be constructed that are less dependent on the performance of other players. Aggregating the individual performance measures in these sports to the position level should produce positional WAR measures that align with those estimated using our approach based on injuries and suspensions, and discrepancies would be worthy of investigation. Our approach can also be replicated at the college level, where, like in the NFL, football is a major business and generates significant revenue (Isidore, 2013).[15] In addition, one could imagine using more and better data in the future to identify the value of particular player attributes by position (e.g., speed, height, weight, etc.), and the interaction of player attributes on the field. Finally, our findings can help to inform the development of the theoretical literature on constrained team building (as in Leeds

[15] In addition to the substantive interest in extending our method to learn about American football in another context, an additional benefit is that at the college level there are many more teams. Because each team serves as a unique source of identifying variation given the clustering structure of the data, this offers the potential for improvements in the precision of estimates of positional WAR.

and Kawolewksi, 2001), and relatedly, to help NFL executives aiming to maximize wins subject to constraints such as the salary cap and access to draft picks.

References

Atkinson, S.E., Stanley, L.R., & Tschirhart, J. (1988) "Revenue sharing as an incentive in an agency problem: an example from the National Football League" Rand Journal of Economics, 19:1, pp. 27-43.

Barnwell, Bill. 2014. Bridging the Analytics Gap. 2014. *Grantland* (02.28.2014). http://grantland.com/features/bridging-the-analytics-gap/

Benoit, Andy. 2012. 2012 Film Room: The Rise of the Tight End. Football Outsiders (08.30.2012)
http://www.footballoutsiders.com/film-room/2012/film-room-rise-tight-end

Berri, David J., Martin B. Schmidt and Stacey L. Brook. 2007. *The Wages of Wins: Taking Measure of the Many Myths in Modern Sport.* Stanford University Press.

Berri, David J. and Brian Burke. 2012. Measuring Productivity of NFL Players. *The Economics of the National Football League.* Springer New York. (137-158)

Blass, Asher A. 1992. Does the Baseball Labor Market Contradict the Human Capital Model of Investment? *The Review of Economics and Statistics.* 74.2 (1992): 261-268

Burger, John D., and Stephen J. K. Walters. 2003. Market Size, Pay, and Performance: A General Model and Application to Major League Baseball. *Journal of Sports Economics* 4.2 (2003): 108-25.

Burke, Brian. 2008. Does Defense Win Championships? *Advanced NFL Stats* (01.03.2008) http://www.advancednflstats.com/2008/01/does-defense-win-championships.html

DuPaul, Glenn. 2012. The Hardball Times. *The Hardball Times* (08.08.2012)

Fleming, David. 2013. The Geeks Shall Inherit the Turf. *ESPN The Magazine* (08.20.2013)

Gerard, Bill. 2007. Is the Moneyball Approach Transferable to Complex Invasion Team Sports. *International Journal of Sports Economics.* 2(2007): 214-230

Hollinger, John. 2014. Hollinger NBA Player Statistics (05.20.2014) http://insider.espn.go.com/nba/hollinger/statistics

Isidore, Chris. 2013. Super Bowl: NFL Money Still on Top. *CNNMoney* (02.01.2013) http://money.cnn.com/2013/02/01/news/companies/nfl-money-super-bowl/

Kelly, Danny. 2012. Breaking Down the X, Z, and Slot Receiver. *SB Nation* (07.17.2012) http://www.fieldgulls.com/2012/7/17/3165517/breaking-down-the-x-z-and-slot-receiver

Landin, Tulio. 2012. Priced to Win. MIT Sloan Sports Analytics Conference (04.2012) http://www.sloansportsconference.com/wp-content/uploads/2012/04/Dartmouth-Tuck.pdf

Leeds, M.A. & Kowalewski, S. (2001) "Winner Take All in the NFL: The Effect of the Salary Cap and Free Agency on the Compensation of Skill Position Players." *Journal of Sports Economics* 2:3, pp. 244-256.

Lewis, Michael. 2004. *Moneyball: The Art of Winning an Unfair Game.* W.W. Norton & Company Ltd.

London, Jay.2014. Data-Driven: MIT Infuses the Rise of Sports Analytics. *MIT News* (02.26.2014)

Millman, Chad. 2013. Most Valuable ATS players. *ESPN Insider* (12.24.2013) http://insider.espn.go.com/nfl/story/_/id/10077627/nfl-aaron-rodgers-most-valuable-player-point-spread

Moskowitz, Tobias. 2012. Does Defense Really Win Championships? *Freakonomics* (01.20.2012) http://freakonomics.com/2012/01/20/does-defense-really-win-championships/

Oliver, Dean. 2014. Guide to the Total Quarterback Rating. (05.20.2014) http://espn.go.com/nfl/story/_/id/6833215/explaining-statistics-total-quarterback-rating

Peterson, Adrian. 2013. If He Goes Down, We Go Down: NFL Injury Value in Vegas. *Grantland* (08.28.13) http://grantland.com/the-triangle/if-he-goes-down-we-go-down-nfl-injury-value-in-vegas/

Table 1. Average Shares of Projected Starter Games Missed Due to Injury and Suspension, by Position, for NFL teams in 2008, 2010, and 2012.

	Average Number of Games Played	Share of Games Missed due to: Injury	Suspension	Combined
Offense				
Quarterback	16	0.108	0.003	0.111
Running back	16	0.142	0.005	0.147
Tight end/Fullback	30.3	0.097	0.004	0.101
Wide receiver	33.7	0.109	0.008	0.117
Interior offensive line (center, guard)	48	0.150	0.001	0.151
Exterior offensive line (tackle)	32	0.115	0.003	0.118
Defense				
Cornerback	32	0.153	0.004	0.157
Safety	32	0.125	0.006	0.131
Linebacker	54.2	0.128	0.007	0.135
Defensive line	57.8	0.102	0.003	0.105
N (team-years)		96	96	96

Note: For each position we performed tests to determine if the average number of games missed due to suspension and injury is statistically different from the average at all other positions. There are no statistically significant differences for any position at the 5 percent level, although the differences for cornerbacks (more likely to get injured) and tight ends/fullbacks (less likely to get injured) are statistically significant at the 10 percent level.

Table 2. Estimated Effects on Total Wins of Games Missed Due to Injury/Suspension, by Position.

	Full Model	Restricted Model
Predicted Wins (Vegas Line)	0.47	--
	(0.16)**	
Offense		
Quarterback	-5.19	-5.63
	(1.28)**	(1.39)**
Running back	0.14	1.04
	(1.54)	(1.82)
Tight end/Fullback	-4.27	-4.61
	(1.47)**	(1.42)**
Wide receiver	-4.75	-5.97
	(1.89)*	(2.04)**
Interior offensive line	-0.24	-0.026
	(1.64)	(1.65)
Exterior offensive line	-3.58	-4.10
	(1.53)*	(1.75)*
Defense		
Cornerback	-0.02	-0.88
	(2.04)	(1.79)
Safety	-0.22	0.88
	(2.03)	(1.94)
Linebacker	-1.07	-1.35
	(2.20)	(2.45)
Defensive line	-0.58	-1.01
	(3.19)	(3.30)
R-squared	0.399	0.341
N	96	96

** Indicates statistical significance at the 1 percent level

* Indicates statistical significance at the 5 percent level

Notes: Standard errors clustered at the team level are in parenthesis.

Table 3. Estimated Effects on Total Wins of Missing Exactly Four Games due to Injury/Suspension at Each Position.

Offense	
Quarterback	-1.30**
Running back	0.036
Tight end/Fullback	-0.56**
Wide receiver	-0.56*
Interior offensive line	-0.02
Exterior offensive line	-0.45*
Defense	
Cornerback	-0.002
Safety	-0.027
Linebacker	-0.079
Defensive line	-0.040

** Indicates statistical significance at the 1 percent level

* Indicates statistical significance at the 5 percent level

Notes: The effect of a 4-game absence at each position is estimated based on the full model shown in Table 2.

Table 4. Estimated Effects on Total Wins of Games Missed Due to Injury/Suspension, by Position, With Expanded Positional Categories.

	Full Model
Predicted Wins (Vegas Line)	0.47
	(0.15)**
Offense	
Quarterback	-4.91
	(1.34)**
Running back	-0.027
	(1.53)
Fullback	-3.63
	(1.63)*
Tight end	-2.21
	(1.76)
Wide receiver	-4.88
	(2.00)*
Interior offensive line	-0.46
	(1.63)
Exterior offensive line	-3.98
	(1.65)*
Defense	
Cornerback	-0.091
	(2.11)
Safety	-0.88
	(2.09)
Linebacker	-1.17
	(2.19)
Defensive line	-0.29
	(3.25)
R-squared	0.408
N	96

** Indicates statistical significance at the 1 percent level

* Indicates statistical significance at the 5 percent level

Notes: Standard errors clustered at the team level are in parenthesis.

Appendix A
Data Collection Appendix

This appendix describes the data collection process in more detail – in particular, how projected starters were determined and how injury and suspension data were gathered. Pro-Football-Reference (pro-football-reference.com) was used as the base system for identifying starters. We compared players listed online with pre-season depth charts from OurLad’s NFL Scouting Services (ourlads.com). We checked through depth charts as early as June to catch changes to projected starters. Players that suddenly disappeared from a pre-season depth chart were analyzed with Pro Football Weekly’s (profootballweekly.com) injury reports to determine whether or not they had suffered a pre-season injury. If so, statistics from previous years, in addition to fantasy-football updates, were used to determine whether the player would have started for the team in the absence of the injury.

The most challenging aspect of determining starters relates to fluctuations in position groups such as wide receivers, tight ends, fullbacks, linebackers and defensive linemen. Position shares in the starting lineup were determined by examining past statistics, games started during the current year by position, fantasy football updates, and seasonal depth charts. A player that was listed as having started all 16 games was given the designation of “starter” for that team in that year, so long as there were no injuries to other players that might have led to him being given that opportunity to start. For a player who started some but not all games, we checked his fantasy football updates throughout the entire season to determine whether he should be designated as a starter.

When we learned that a projected started had not started all 16 games, we first determined which games were missed. Sometimes this information was available in the player’s ESPN game logs (espn.com). For more obscure players, missed games were again determined through fantasy football updates. After identifying the specific games missed by a projected starter, team injury reports and fantasy football updates were used to determine the cause for the missed game. Players often missed games as "healthy scratches," due to competition from other players, trades, holdouts, and for rest (particularly at the end of the year for teams with a secure playoff position). Games missed for these reasons were not used in our empirical analysis.

To illustrate the full process of filling out the data, we use the somewhat complicated example of the Arizona Cardinals' starting quarterback position in 2012. Unlike many other teams, there was no clear starting quarterback for the Cardinals going into the season. The quarterbacks on the Cardinals' 2012 roster included Brian Hoyer, Kevin Kolb, Ryan Lindley, and John Skelton. We searched all of the players' names on a fantasy football update website (Rotoworld.com) and found an update from August 31, 2012 – just before the start of the season – indicating that "John Skelton will start at quarterback for the Cardinals in the season opener." Thus, we determined that John Skelton was the initial starter. However, we can see that Skelton started just six games in 2012. Using the same website, we found that Skelton suffered a low ankle sprain in the first game of the season, which resulted in him missing the next three games. This is well documented in RotoWorld's fantasy updates. Thus far we have learned that the Cardinals lost at least three games at quarterback due to injury during 2012. Reading forward on the fantasy updates, we determined that Skelton was healthy enough to play in week five but did not start because Kolb was playing well in Skelton's absence. Throughout the rest of the season, we find that Skelton remained healthy but was in and out of the starting lineup, starting only five other games. The reasons for him not staring beyond game-4 were not due to injury. Thus, the missed-games variable for the Cardinals' starting quarterback position in 2012 in our data indicates that 3 out of 16 games were missed due to injury and/or suspension.