Jason Mulholland and Shane T. Jensen*

Predicting the draft and career success of tight ends in the National Football League

Abstract: National Football League teams have complex drafting strategies based on college and combine performance that are intended to predict success in the NFL. In this paper, we focus on the tight end position, which is seeing growing importance as the NFL moves towards a more passing-oriented league. We create separate prediction models for 1. the NFL Draft and 2. NFL career performance based on data available prior to the NFL Draft: college performance, the NFL combine, and physical measures. We use linear regression and recursive partitioning decision trees to predict both NFL draft order and NFL career success based on this pre-draft data. With both modeling approaches, we find that the measures that are most predictive of NFL draft order are not necessarily the most predictive measures of NFL career success. This finding suggests that we can improve upon current drafting strategies for tight ends. After factoring the salary cost of drafted players into our analysis in order to predict tight ends with the highest value, we find that size measures (BMI, weight, height) are over-emphasized in the NFL draft.

Keywords: football; prediction; regression.

DOI 10.1515/jqas-2013-0134

1 Introduction

Tight end is a very unique position in football due to the fact that tight ends have multiple roles that require a varied skill set. Tight ends not only need to be able to block defensive ends and outside linebackers, but also need to be able to run routes and catch passes. As the National Football League (NFL) continues to evolve, tight end has emerged as an extremely important position. Their changing role in NFL offensive schemes and the increasing importance

of effectively drafting tight ends are the motivation for our focus on the tight end position in this paper.

In the past, tight ends were primarily blockers for running plays and only occasionally targets for passing plays. Now, as many NFL teams have turned to a more aggressive, passing-oriented offensive scheme, the tight end has become one of the main playmakers on offense and often a quarterback's primary receiver. One illustration of the larger emphasis on passing is that in 2003, only 3 of the 32 NFL teams called pass plays more than 60% of the time, whereas in 2012, 11 of those 32 teams called pass plays more than 60% of the time.

One of the leading teams in the tight end evolution is the New England Patriots, who have featured tight ends prominently in recent years. In the 2010 NFL Draft, the Patriots selected Rob Gronkowski with the 42nd pick and Aaron Hernandez with the 113th pick. Over the next three seasons, Gronkowski and Hernandez established themselves as two of the premier talents at tight end in the NFL, though both have had off-field struggles that we will discuss later.

This paper will use data available to teams before the NFL draft, specifically size, college and combine performance measures, to address two important questions:

1. What are the best quantitative predictors for NFL draft order of tight ends?
2. What are the best quantitative predictors of NFL career success at the tight end position?

We will address both questions by constructing prediction models with either NFL draft order or NFL career success as the outcome variable. The available predictor variables are any quantitative measures available before the NFL draft, namely, measures from player's college careers and the NFL combine, as well as physical measures.

Our analysis differs in several aspects from past studies that have focused on valuing prospective professional athletes in several sports. Burger and Walters (2003) analyzed the impact of market size and marginal revenue on how teams value players in Major League Baseball. Massey and Thaler (2005) researched when teams find the most value in the NFL draft and found that teams often overvalue the draft's earliest picks. Berri, Brook, and Fenn (2010) analyzed which factors contribute most to the selection of

---

*Corresponding author: **Shane T. Jensen**, Associate Professor of Statistics, Department of Statistics, The Wharton School, University of Pennsylvania, Philadelphia, PA 19104, USA, e-mail: stjensen@wharton.upenn.edu

Jason Mulholland: Undergraduate student, The Wharton School, University of Pennsylvania, Philadelphia, PA 19104, USA

college basketball players in the NBA draft as well as NBA performance. Our own analysis follows a similar process of identifying what factors best predict NFL draft results as well as future NFL performance of college players.

Our analysis indicates that the best predictors of NFL draft order are not the best predictors of NFL career performance, which suggests that tight ends are not currently being drafted in an optimal way with respect to predicted career success. We further explore this result by evaluating NFL performance in the context of the salary cost of each player, so that we can isolate the best predictors of tight end “value” (performance per unit cost). Our overall approach emulates the Dhar (2011) study of NFL wide receivers, though we examine a greater set of NFL performance measures as well as undertaking a cost analysis of performance per salary cost.

Our paper is organized as follows: We first describe the data used in our study in Section 2 and then outline our prediction models in Section 3. We explore the results of our quantitative modeling of the NFL draft order of tight ends in Section 4 and then the results of our quantitative modeling of NFL career performance of tight ends in Section 5. In Section 6, we examine the difference between the selected predictors of NFL draft order and NFL career performance. Finally, we incorporate salary cost into our measures of NFL career performance in Section 7, explore the use of a different measure of NFL performance in Section 8, and then conclude with a summary and discussion in Section 9.

2 Data

For this analysis, we collected data from the NFL Combine, the NFL Draft, and both the college and NFL careers of each tight end that participated in the NFL Combine or was selected in the NFL draft between 1999 and 2013. Incorporating players from earlier than 1999 is difficult due to the lack of NFL combine data. The time period from 1999 to 2013 yielded 315 tight ends, 250 of which participated in the NFL Combine and 65 of which were drafted without participating in the NFL Combine.

The NFL Combine and college performance data are both available to decision-makers prior to the NFL draft, and so we consider any college and combine measures as potential predictors in our modeling of either NFL draft order or NFL career performance.

NFL Combine and size data for each participating player was collected from the public website www.nflcombineresults.com. This data includes the height and weight of each player and the results of six combine drills:

1. Forty Yard Dash, 2. Bench Press, 3. Vertical, 4. Broad Jump, 5. Shuttle, and 6. Three Cone Drill. From the height and weight measures we calculated one additional physical measure, body mass index (BMI), which is often used to measure the muscle mass of athletes.

College football data was collected from two public websites: www.sports-reference.com/cfb and www.ncaa.org. The college measures for tight ends that we collected were each player’s receptions, receiving yards, and receiving touchdowns over their entire college career as well as these same totals specifically in their final season of college football.

From these measures, we created several additional variables to reflect the impact of the player’s final year in college: 1. final year college receptions percentage, 2. final year college yards percentage, and 3. final year college touchdowns percentage. We also created an overall college yards per reception measure for each player as well as an indicator variable, BCS, for whether or not they played at a Bowl Championship Series school. BCS schools tend to receive larger amounts of media and scouting attention and tend to play against more talented competition, which may be predictive of either the NFL Draft or NFL career performance.

NFL Draft data was collected from the public website: www.nfl.com/draft. We collected the draft order of the 223 tight ends (out of 315 total) that were drafted in the 1999–2013 period. In Section 4, we will model the draft order of tight ends that were drafted.

NFL performance data was collected from two public websites: www.pro-football-reference.com and www.nfl.com. The measures of NFL success that we collected for each player were: 1. number of games played, 2. number of games started, 3. total career receptions, 4. total career yards, 5. total career touchdowns. The data used in this study includes players’ NFL statistics through the end of the 2012 NFL season. In Section 5, we will derive several other measures of NFL career success from our collected NFL performance data.

For our NFL performance per cost analysis presented in Section 7, we used data on the rookie wage scale from <http://overthecap.com/nfl-rookie-salary-cap.php>.

3 Statistical methodology

We employ two different statistical approaches for predicting either the NFL draft or NFL career success as a function of pre-draft predictor variables. The first approach is ordinary least squares linear regression. Stepwise variable

selection (Hocking 1976) was used in order to find a subset of the pre-draft variables that are the best linear predictors of the outcome in terms of adjusted R^2 .

Our second approach is a decision tree model fit by recursive partitioning (Breiman et al. 1984). This method creates a decision tree via binary splits of a subset of the predictor variables. Particular predictor variables are chosen as splitting rules in order to maximize the log-worth,

$$\text{log-worth} = -\log_{10}(p\text{-value of } F \text{ statistic})$$

where the F statistic is based on the variance of the outcome variable within versus between nodes of the decision tree. Specifically, the F statistic is large for a good splitting decision that groups observations such that there is small variance in the outcome variable within each node, but large variance in the outcome between the nodes. A large F statistic has a correspondingly small p -value and therefore a large log-worth. In total, a higher value of the log-worth is indicative of greater differences in the outcome variable between nodes of the tree, which intuitively means the tree can produce more refined predictions of the outcome.

Once the number of observations in each terminal node of the tree falls below a pre-determined threshold, the decision tree is pruned. The earlier splits in a decision tree can be interpreted as more important for prediction than the later splits in a decision tree, since the recursive partitioning algorithm proceeds in a greedy fashion: always choosing the next splitting rule based on the best refinement in predictions. Thus, the initial split in the decision tree is always the single predictor variable that, by itself, best predicts the outcome variable. Additional splits are based on predictor variables that improve predictions beyond that initial split. This results in most partition trees having an initial split with the largest log-worth, and then the log-worth decreasing from one level to the next.

Both multiple linear regression and decision trees are designed to accomplish our primary goal, which is the selection of a subset of pre-draft variables that have the best predictive power of the outcome variable (in our case, either NFL draft results or NFL career success). An advantage of linear regression is that the resulting model is relatively easy to interpret in terms of the effects of individual predictor variables. The effects of individual variables in the decision tree approach are somewhat less interpretable, but this method allows for more flexibility in the relationship between the predictors and outcome variable (compared to the linearity assumptions of multiple regression). Both models were implemented using the JMP 10 statistical software.

We are aware that the interpretation of our results will be influenced by multicollinearity, as some of our

predictor variables are correlated with one another. For example, a regression between career college receptions and career college yards provides an R^2 value of 0.925. In highly correlated cases such as this one, these variables act as proxies for each other such that either one entering a model will likely exclude the other from having additional predictive power.

It was somewhat surprising to only observe high correlations within each category of predictor variables but not between categories of predictor variables (e.g., physical attributes vs. college measures). For example, the highest correlation found between inter-category variables is 0.17 between career college yards per reception and Forty Yard Dash time. Other inter-category correlations were surprisingly low, such as a correlation of 0.08 between weight and Bench Press. Thus, we expect that multicollinearity may impact our interpretation of selected variables within a particular category (e.g., college yards vs. college receptions) but not between variable categories (e.g., combine measures vs. college measures).

We elected to avoid interaction terms for use as predictors due to the fact that they would be more difficult to interpret and would increase the potential issue of multicollinearity. Our recursive partitioning tree approach still captures some indirect interaction effects between variables since a split on one variable nested within a split on another variable suggests that an interaction of those two variables leads to different outcome values.

4 NFL draft results

In our first analysis, we model the NFL draft for tight ends as a function of pre-draft variables based on college performance and results from the NFL combine. The specific pre-draft variables that we use as predictors of the NFL draft are discussed in Section 2 above. The college variables consist of seven measures of college receiving performance as well as an indicator of whether the player attended a BCS school. The combine variables consist of six measures of athleticism. We also consider additional size measures of weight and height as well as the measure BMI, which we create from the recorded height and weight of each tight end.

For the remainder of this section (and the next section), we outline our specific results models by model. However, an overall summary of the predictor variables included in each regression models can be seen in Table 5 at the end of the paper, in which a “+” indicates that the predictor was included with a positive coefficient, while a

Figure 1 Distribution of tight end draft order 1999–2013.

“–” indicates that the predictor was included with a negative coefficient.

We first turn our attention to predicting the draft order, using pre-draft college, combine, and size measures, of the 223 tight ends that were drafted between 1999 and 2013. In Figure 1, we see that the draft order of tight ends has been fairly evenly distributed throughout the draft, with the exception of the top 20 picks where tight ends have been drafted less frequently. What factors determine when a tight end is drafted?

We treat draft order as a continuous variable, which runs from 1 to 260, with lower values being “better” in the sense that the player is selected earlier in the draft. We fit a multiple linear regression model on draft order with the pre-draft college and combine measures as predictor variables. Table 1 gives the variables selected by the step-wise procedure as the best subset of predictor variables for draft order. This model had an R^2 of 0.23, which is significant at the 0.01% level, indicating that these pre-draft variables have significant predictive power for draft order.

The selected predictor variables include physical attributes: height, BMI, and two combine measures: Forty Yard Dash and Bench Press. Both of these combine

variables, which measure the speed and strength of the player, are the most significant predictors of draft order, according to this model. The BCS indicator variable is also selected in this model.

The estimated coefficients in Table 1 can be interpreted as the partial effect of that variable holding the other variables constant. Negative coefficients indicate that higher values of that variable predict a lower (i.e., better) draft order. For example, the model suggests that players who have greater muscle mass (as indicated by higher BMI) will be selected earlier. With each additional repetition in the Bench Press, the player is projected to be selected approximately 4 picks earlier in the draft. Each tenth of a second faster in the Forty Yard Dash is associated with being selected approximately 12 picks earlier in the draft.

Attending college at a BCS school is associated with being selected around 33 picks earlier, which is more than an entire round of the draft. However, other than this BCS indicator and the career college yards measure, this model indicates that college performance is not as important in terms of predicting NFL draft order as the combine measures.

To alleviate concerns that our results were driven primarily by our choice of a linear regression model, we also implemented a decision tree model with NFL draft order as the outcome variable and pre-draft variables as predictor variables. Figure 2 gives the decision tree model that was fit by recursive partitioning on this data. This decision tree model of NFL draft order is a little more predictive of NFL draft order, with a root mean square error (RMSE) of 56.52, which is lower than the regression model’s RMSE of 60.08.

Compared with the linear regression model, this decision tree model also has a higher R^2 of 0.35, but this is not surprising since a recursive partitioning decision tree model is more flexible in the sense of not requiring a linear

Table 1 Linear regression of NFL draft order.

	NFL draft order		
	Coefficient estimate	Standard error	p-Value
Intercept	728.97	493.68	0.142
Height	–9.70	5.00	0.055
BMI	–9.32	6.03	0.125
40 yard dash	121.42	45.70	0.009
Bench press	–4.31	1.34	0.002
BCS dummy	–33.09	22.01	0.135
College yards	–0.02	0.01	0.023

Figure 2 Decision tree model of NFL draft order.

relationship between the predictors and the outcome variable. In general, we will focus our comparisons between regression models and decision tree models by examining the root mean square error (RMSE) of the predicted values from each model.

Just as in the linear regression model, Career College Yards was selected as a predictive college performance measure, though the BCS indicator was not selected in the decision tree model. The physical measure of height was selected in both models, whereas the physical measure of weight was selected in the decision tree (while the closely related physical measure of BMI was included in the linear regression). The combine measure of Forty Yard Dash was selected by both models whereas other combine measures were selected by one of the two models: Three Cone Drill and Vertical in the decision tree versus Bench Press in the linear regression model.

Although a somewhat different set of predictor variables were selected by the decision tree model, the results of this model confirm the general theme of the linear regression model: that physical attributes and combine measures dominate the best set of predictor variables for NFL draft order. While the initial split in the decision tree is a college performance measure, every subsequent split in the tree is from either an NFL combine statistic or a size measurement. Overall, our analysis of draft order indicates that for the most part NFL talent evaluators are valuing size and all-around athletic ability though it also seems important to have surpassed a certain level of college production (797 yards according to the partition tree).

We also explored a logistic regression model to predict the binary outcome of whether or not a tight end is drafted

based on pre-draft variables. The selected predictor variables from the logistic regression model suggest that the combine is the most important determinant of whether a player is drafted or not, with four of the six combine measures included in the model. The Forty Yard Dash was the most significant of all the predictor variables. The BCS indicator variable and two college receiving measures were also included in the logistic regression model.

The overall summary of these results is that NFL front offices seem to focus on measures of size, fitness, and all-around athletic ability (based on the combine drills) when determining order for tight ends in the NFL Draft. These models imply that college performance seems generally less influential, but does still have an impact when determining NFL draft order of tight ends.

5 NFL performance results

We now turn our attention to the task of predicting NFL performance of tight ends based on the same set of pre-draft measures from their college performance and the NFL Combine. Just as in our NFL draft analysis, we also consider the additional size measures of BMI, Weight, and Height.

Similar to our analysis of NFL draft order, we will employ both linear regression and decision trees to determine which pre-draft variables are most predictive of NFL performance. Our analysis of NFL performance is complicated by the fact that there is no single definitive outcome variable (unlike the NFL draft where draft order was the

obvious choice of outcome measure). We constructed three different measures of NFL performance that capture potentially different aspects of a successful (or unsuccessful) NFL career for tight ends:

1. NFL Games Started
2. NFL Career Score
3. NFL Career Score per Game

The first measure, NFL Games Started, is the simplest indication of NFL success. Only tight ends that show consistently good performance will be selected as starters over a long time period. When using this outcome variable, we only include the 258 tight ends that were drafted between 1999 and 2010 so that they have at least three years in the NFL to accumulate games started. It is important to note that some NFL teams do sometimes list zero or two tight ends as starters, which may result in some tight ends having more or less starts than they would otherwise have had. Unlike our next two measures of NFL performance that are solely based on receiving statistics, NFL Games Started might also capture more subtle aspects of tight end performance since a tight end could also be chosen to start based on good blocking or being a team leader. We believe that this was the best statistic available to us for capturing these aspects of the position. However, if we had data on percentage of plays on the field or some measure of the players' blocking ability, those variables would have also been appropriate for use.

The second measure, NFL Career Score, is an aggregate measure of receiving performance, which we constructed as

$$\text{NFL career score} = \text{career receiving yards} + 19.3 \times \text{career receiving touchdowns}$$

The coefficient of 19.3 that converts touchdowns to the scale of yards is based on the analysis of Stuart (2008). In that analysis, the coefficient value was computed by finding an estimate of the difference in expected points between possessing the ball at the one-yard line and scoring a touchdown. Stuart (2008) then found that 20.3 was the average number of yards that a team must gain to have that same change in expected points as a touchdown score. One yard was then subtracted to account for the one-yard gain needed to advance from the one-yard line to the end zone, giving an estimated equivalence of 19.3 yards receiving for each receiving touchdown. This aggregated NFL career score measure should be strongly indicative of NFL success across an entire career both in terms of receiving and scoring.

The third measure, NFL Career Score per Game uses the same combination of receiving yardage and touchdown scoring as NFL Career Score but divides by the

number of games played. NFL Career Score per Game is intended to capture a player's average productivity in the NFL instead of the cumulative productivity captured by the previous two measures.

Since this third measure is not cumulative, we can include more recently drafted players. Specifically, when NFL Career Score per Game is the outcome variable, we use the 293 tight ends that entered the NFL between 1999 and 2012, giving us at least one year in the NFL per player for this average measure. In contrast, recall that for NFL Career Score or NFL Games Started we use the 258 tight ends that participated in the combine or were drafted between 1999 and 2010, giving us at least 3 years in the NFL per player for those cumulative measures. We acknowledge that the nature of our sample implies we are over-sampling the early years of players' careers, which will impact the NFL Games Started and NFL Career Score models. However, as our goal is to model future performance using only pre-draft variables, we do not adjust for this impact by including years of experience as a predictor in our model as it is unclear how long a player's career will last prior to them entering the NFL Draft.

For the remainder of this section, we will examine the pre-draft variables (college performance, combine measures, and size measures) that are most predictive of each of our three measures of NFL performance. Ideally, we seek to identify the pre-draft variables that are predictive of success on all three measures of NFL success. Thus, we will give special attention to any pre-draft variables that are selected as predictors across all three NFL outcome measures.

We begin by modeling the NFL Games Started measure as the outcome variable in a linear regression model with all pre-draft measures (size, college, and combine) as predictor variables. Table 2 gives the best subset of predictor variables for NFL Games Started as selected by stepwise linear regression. This linear regression model of NFL Games Started has an adjusted R^2 of 0.28 (which is significant at the 0.01% level).

One overall observation, which will be a running theme throughout this section, is that the set of best predictors of NFL performance contains many variables based on college performance in addition to physical attributes. Recall that college-based measures were less involved in the prediction models of NFL draft order. We will explore this contrast in more detail in Section 6.

The most significant predictors in this model of NFL Games Started are Broad Jump, Career College Receptions, and BMI. Players with explosive power and high muscle mass, along with a large cumulative total of catches in college, are most likely to start many games, which

Table 2 Linear regression models of NFL performance.

	NFL games started			NFL career score			NFL career score per game		
	Coefficient estimate	Standard error	p-Value	Coefficient estimate	Standard error	p-Value	Coefficient estimate	Standard error	p-Value
Intercept	-5189.42	2684.36	0.055	-4205.68	5570.25	0.451	6.99	52.51	0.894
Height	61.08	35.30	0.086	X	X	X	X	X	X
Weight	-8.76	5.29	0.100	11.98	8.53	0.162	X	X	X
BMI	80.79	43.51	0.065	X	X	X	X	X	X
40 yard dash	X	X	X	-1694.69	825.99	0.042	-12.50	7.61	0.102
Bench press	X	X	X	X	X	X	0.44	0.22	0.052
Vertical	X	X	X	X	X	X	X	X	X
Broad jump	2.31	0.43	<0.001	68.38	19.12	<0.001	0.37	0.19	0.054
Shuttle	X	X	X	X	X	X	X	X	X
3 cone drill	X	X	X	X	X	X	X	X	X
BCS dummy	10.48	8.14	0.200	464.05	296.15	0.119	6.31	2.93	0.033
College Yds/Rec	1.87	1.12	0.097	107.55	43.80	0.015	X	X	X
College Rec	0.18	0.07	0.008	14.48	3.66	<0.001	X	X	X
College yards	X	X	X	X	X	X	0.01	<0.01	<0.001
College TDs	X	X	X	-55.30	28.79	0.056	-0.35	0.29	0.232
Final year college Rec, %	X	X	X	X	X	X	X	X	X
Final year college Yds, %	-26.98	17.56	0.127	X	X	X	X	X	X
Final year college TDs, %	8.20	12.43	0.510	X	X	X	X	X	X

suggests these particular variables capture a tight end's durability, consistency, and athleticism.

In addition to BMI, the physical measures of height and weight were also selected as important predictors of NFL Games Started. Height has a positive coefficient and weight has a negative coefficient, but overall, weight has a positive relationship with NFL Games Started due to its positive impact on BMI, which has a large, positive coefficient. This suggests that the larger, taller, and stronger players have longer careers with more games started, likely because their bodies can better handle the stress of the game. Another interesting observation is that Final Year College Yards Percentage has a negative coefficient, whereas Final Year College Touchdowns Percentage has a positive coefficient. We hesitate to over-interpret this particular result due to the potentially destabilizing effect of multicollinearity. However, one plausible explanation is that a player who has shown consistency throughout his college career will have relatively constant yardage totals throughout college. That tight end will likely score more touchdowns in his final year, as he will have more passes thrown his way in the red zone since the quarterback will have greater confidence in him due to his proven consistency.

As an alternative approach, we also estimated a decision tree model with recursive partitioning to NFL Games Started, which is shown in Figure 3. This decision tree model has an R^2 of 0.27 and an RMSE of 30.78, which is slightly lower than the regression model's RMSE of 31.30.

Overall, the set of predictor variables selected in this decision tree model were similar to those in the linear regression model. The side of the decision tree for the more productive tight ends (Career College Yards >797) contains the same variables that were selected by the linear regression model: weight, Broad Jump, and the BCS indicator. For tight ends that were less productive (Career College Yards <797), a couple of additional combine measures, Bench Press and Shuttle, were shown to be predictive of NFL Games Started. It is interesting to note that the initial split in the tree (which had the highest log-worth of 4.72) uses college yards, which is highly correlated with college receptions, the second most significant predictor in the regression model. Additionally, the two second level splits use BMI and Broad Jump, the third and first most significant variables, respectively, in the regression model. Overall, this decision tree shows very similar trends to those shown by the regression model (with some minor differences due to the impact of multicollinearity), which further shows that strong, explosive tight ends that were productive in college will start the most games at the professional level.

We now focus on predicting the second measure of cumulative NFL performance, NFL Career Score. The selected predictor variables from a linear regression model of NFL Career Score can be seen in Table 2, as well. This regression model has an adjusted R^2 value of 0.26 (which is significant at the 0.01% level).

Figure 3 Decision tree model of NFL games started.

Similar to the prediction of NFL Games Started, the most significant predictor variables in this model of NFL Career Score are Broad Jump, Career College Receptions, and Career College Yards per Reception. We see that the set of best predictors is a mix mostly of combine measures and college performance measures. In addition to the other college measures, the BCS indicator variable was also chosen. Unlike the NFL Games Started model, not all three size measures were chosen, but weight does appear in the model. Weight has a positive coefficient in this model, which suggests that larger tight ends are better able to sustain production in the passing game over a long career, in addition to starting more games, as we saw earlier.

The decision tree with NFL Career Score as the outcome variable is shown in Figure 4. This decision tree model has an R^2 value of 0.35. The decision tree has a RMSE of 1190.85, which is higher than 1166.71, the RMSE of the regression model. We see a highly similar set of predictor variables are involved in this decision tree model compared to the regression model (in Table 2), and Forty Yard Dash is the initial splitting variable in the tree and a 5% significant predictor in the regression model. However, in this tree, despite not being the initial split, the college yards second level split has the highest log-worth (of 7.49, while the initial split has a log-worth of 6.35). This indicates that tight ends who run a Forty Yard Dash at slower than 4.69 s can significantly make up for their lack of speed in the NFL if they were able to do so in college, to the extent that

they accumulated at least 1093 college receiving yards. The only additional variable selected in the decision tree and not in the regression model is Shuttle, but it plays a relatively small role, as a fourth level splitting variable.

It is interesting to note that one terminal node of the decision tree stands out as having a substantially larger predicted value than the other nodes. Tight ends with a Forty Yard Dash faster than 4.69, a Broad Jump over 120 inches, and 65 or more Career College Receptions are projected to have far better NFL career scores. Among slower tight ends (the left side of the tree with Forty Yard Dash slower than 4.69), the model indicates that the only path to high NFL career success is to overcome this lack of speed with a large size (Weight >255) and high college production (over one thousand ninety-three Career College Yards).

For success in terms of NFL Career Score, most of the same traits are selected as for success in terms of NFL Games Started. Again, the larger tight end with high explosive athleticism (seen through Broad Jump) and college success will be expected to succeed in the pros. However, for NFL Career Score, speed also is selected as important. To summarize our analysis of NFL performance thus far, a mixture of mostly college and size measures with only a couple of combine measures is the most effective to predict both of the cumulative NFL performance measures (games started and career score).

A combination of college and combine statistics best predicts our third measure of NFL performance,

Figure 4 Decision tree model of NFL career score.

NFL Career Score per Game. Unlike the previous cumulative NFL performance measures, this average NFL performance measure did not have any of the size variables (BMI, weight or height) included in its prediction models.

We present the selected predictors from a regression model with NFL Career Score per Game as the outcome variable also in Table 2. This regression model had an adjusted R^2 of 0.250 (which is significant at a 0.01% level).

As noted above, no size measures were selected as predictive for NFL Career Score per Game. This suggests that size is only important for longevity, as it appeared in the two models predicting cumulative performance models, but not in the average performance model. Thus, the bigger tight ends usually can remain playing at a high level for a longer career, but their size does not necessarily improve their playing ability.

Similar to the other two NFL performance outcomes, the combine measures of Broad Jump and Forty Yard Dash were selected along with the additional Bench Press measure, which was selected only with this particular outcome variable. We again see that college performance is an important indicator of NFL performance: Career College TDs and Yards were both selected as well as the BCS indicator variable.

The decision tree with NFL Career Score per Game as the outcome variable is shown in Figure 5. This decision tree model has an R^2 value of 0.31, while its RMSE of 11.98 is higher than the regression model's RMSE of 11.41.

It is interesting to see the absence of the BCS indicator and Broad Jump from this decision tree model since these two predictors were selected in almost all of the previous models of NFL performance. This may suggest that those

two variables are more important predictors of cumulative NFL performance than average NFL productivity (as measured by NFL Career Score per Game). The combine measure of Forty Yard Dash is selected in this decision tree as well as in several previous models. Several measures of college performance are selected in this decision tree, with College Career Yards appearing to be the most important predictor of average NFL productivity, as the initial split with a log-worth of 11.76. Overall, Career College Yards has been the splitting variable with the highest log-worth in all of the recursive partitioning models, showing that college performance is very indicative of future NFL performance.

As a final summary of our analysis of NFL performance, we compare the set of pre-draft variables that were selected as the best predictors of each of our NFL performance outcomes. Figure 6 provides a Venn diagram of the sets of variables included in the regression model for each of our three NFL performance outcomes. This chart focuses only on the selected predictors from the regression models, though we note that the selected variables from the decision tree models are very similar. A “+” sign in Figure 6 indicates a positive correlation and a “-” sign indicates a negative correlation for the corresponding variable.

A key observation from Figure 6 is that Broad Jump and the BCS indicator variable were selected by all three regression models. This result suggests that these two measures are important predictors of tight end success in the NFL, regardless of how we define that NFL success. The fact that Broad Jump is such an important predictor of NFL performance is especially interesting since it was not an important predictor of NFL draft order in Section 4.

Figure 5 Decision tree model of NFL career score per game.

We will provide a more systematic comparison of predictors of NFL draft order versus those of NFL performance in Section 6.

It is clear from Figure 6 that the best predictors of NFL performance are a mix of size, combine and college variables beyond just the Broad Jump measure and the BCS

indicator. Some form of Career College Yards or Receptions (which are highly correlated with each other) are important predictors of all three NFL performance outcomes, while other combine measures (Bench Press and Forty Yard Dash) and all three size measures also appear in Figure 6. Clearly, one would not want to rely exclusively

Figure 6 Selected predictors in regression models predicting three different NFL performance measures.

on one of college, combine, or size information when evaluating tight ends in terms of their prospective NFL performance, as all include significant predictors.

It is also interesting to note that all predictors selected when NFL Career Score was the outcome variable were also selected when at least one of the other two outcome variables were used. This suggests that there are no additional predictor variables that are important for cumulative productivity (NFL Career Score) beyond those that are predictive of longevity of high-level play (NFL Games Started) and average career productivity (NFL Career Score per Game).

It is also interesting to explore whether there have been changes in the value of particular predictor variables that mirror the changing role of tight ends in the NFL. To examine this issue, we split our dataset of tight ends into two subsets: those who entered the league in 1999–2005 and those who entered the league in 2006–2012. We fit linear regression models for each of our three NFL performance outcomes separately within each of these subsets. Between the two subsets, the selected predictors were very similar when NFL Career Score and NFL Career Score per Game were the outcome variable. However, there was a notable difference in the selected variables between the two subsets when NFL Games Started was the outcome variable. For NFL Games Started, the 1999–2005 subset had BMI and Broad Jump as the most significant indicators, while Vertical was most significant for the 2006–2012 subset, with height also being selected. This result suggests that over the past 15 years, there has been a shift in starting tight ends from those who are large, strong blockers to those who are threats in the passing game with their

height and jumping ability, such as Jimmy Graham of the New Orleans Saints.

6 Comparing predictors of NFL draft order vs. NFL performance

A primary goal of our analysis is an evaluation of the extent to which the best predictors of NFL performance (Section 5) are different from the best predictors of NFL draft order (Section 4), since any differences would suggest that drafting strategies are inefficient.

In Figure 7, we compare the selected predictors from the regression model with NFL draft order as the outcome variable to the selected predictors from the regression model with NFL Career Score as the outcome variable. We get a similar figure if NFL Games Started or NFL Career Score per Game are used as the measure of NFL performance.

Forty yard dash and the BCS indicator variable were the only two variables that were selected as predictors of both draft order and NFL performance. We see, however, that draft order and NFL performance have unique predictors from the both college and combine data. From the college data, college receiving yards is predictive of draft order whereas college receptions is predictive of NFL performance (though these two variables are themselves correlated). From the combine data, Bench Press is predictive of Draft Order whereas Broad Jump is predictive of NFL performance. Therefore, Figure 7 suggests that drafting strategy could be improved (in terms of predictive

Figure 7 Selected predictors from regression models with Draft Order as outcome vs. NFL career score as outcome.

NFL performance) by focusing less on Bench Press (pure strength) and focusing more on Broad Jump (explosive athletic ability).

We further emphasize strategic differences by evaluating the predictive power of the predictors selected when NFL draft order is the outcome variable when these same “draft selected” predictors are used to predict NFL performance. In Table 3, we compare the adjusted R^2 value for predicting each measure of NFL performance when we use the “optimal” predictors as selected by stepwise linear regression (middle column) versus the “draft selected” predictors (right column).

We see in Table 3 that the “draft selected” variables are substantially worse predictors of each NFL performance measure, which confirms that current NFL drafting decisions are not optimally calibrated in terms of subsequent NFL performance.

As an additional analysis, we also fit regression models for our three NFL performance outcomes with draft order included as a predictor variable in addition to our pre-draft college, combine and size measures. Clearly, these models are not useful for prediction of NFL performance prior to the NFL draft (our primary focus) but can still provide insight into variables that are either under- or over-considered when teams are drafting players. In these models, four pre-draft variables had significant non-zero coefficients at the 5% level. College yards per reception and college receptions were significant in the NFL Career Score model, college yards was significant in the NFL Career Score per Game model, and Forty Yard Dash was significant in the NFL Games Started model. All four of these selected predictors had positive coefficients. In the case of the three college receiving statistics, the positive coefficient indicates that these measures were under-considered when drafting, since a higher total in that measure is associated with increased expected NFL performance. This indicates that NFL teams should give college receiving production greater attention when drafting tight ends. In the case of Forty Yard Dash, a positive coefficient indicates that this measure was over-considered when drafting since a slower time (higher measure) will increase

expected NFL performance in this model given the same draft order outcome. This suggests that NFL teams are too focused on speed when drafting tight ends.

In the next section, we will explore our findings further by a cost analysis: we will explore the best predictors of “high value” tight ends that have high NFL performance for a low cost, where cost is the salary implied by their draft order.

7 Predicting NFL performance per salary cost

We will now discuss the predictors of NFL performance per salary cost. To determine salary cost, we use the NFL Rookie Wage Scale that was enacted by the recently negotiated Collective Bargaining Agreement (Overthecap.com, 2013). This rookie wage scale gives the estimated salary cost of a player as a function of their draft order. For each player in our data set, we use the projected draft order (from Section 4) and the rookie wage scale to calculate their projected average salary for a 4-year rookie contract. We then divide our three measures of NFL performance (from Section 5) by the log of their projected average salary to create three measures of NFL performance per salary cost: 1. NFL Career Score per Cost, 2. Games Started per Cost, and 3. NFL Career Score per Game per Cost.

Note that we could have also used actual draft order to determine their rookie salary for our calculations, but we wanted to emulate the setting where we are predicting NFL performance per cost for future players prior to the NFL Draft, so actual draft order and rookie salary will be unknown. However, we did run our analysis with actual draft order and the results were extremely similar to the analysis that follows.

Our primary interest is finding the pre-draft variables that are the best predictors of NFL performance per salary cost. Performance per cost is extremely important in the NFL, as there is a strict salary cap; a team cannot spend more money to compensate for mistaken player selection such as in Major League Baseball. Therefore, it is important to ensure that players are performing well in comparison to the amount of money that they are paid. We fit multiple regression models with each of our three NFL performance per salary cost measures as the outcome variables, and all pre-draft variables (size, college, and combine) as predictors. These regression models are limited to finding linear relationships, and it is unlikely that team utility for performance per salary is linear. However, we believe that

Table 3 Adjusted R^2 for predicting each outcome variable using optimal predictors compared to “draft selected” predictors.

Outcome variable	Adjusted R-squared with optimal predictors	Adjusted R-squared with draft selected predictors
NFL career score	0.258	0.193
NFL career score per game	0.250	0.171
NFL games started	0.245	0.147

at least one of the models. In terms of combine measures, Broad Jump was selected in all three models, while Bench Press was selected in one of the models. The dominance of college performance measures in these three models again indicates that college performance is under-utilized in evaluation of tight end draft prospects, as focus on college performance can help NFL teams find better tight ends for lower cost later in the draft.

We note the absence of any of the size measures (BMI, Weight, Height) from any of the NFL performance per cost models. Size measures were shown in Section 4 to be predictive of NFL draft order and in Section 5 to be predictive of NFL performance. Factoring cost into the analysis, none of these size measures were selected as predictive of NFL performance per cost. This result indicates that while larger tight ends tend to have more successful NFL careers, NFL teams tend to draft them earlier and, thus, must pay a higher salary for their better performance. Although their performance is higher, their value (in terms of performance per cost) is not necessarily higher because of their higher cost.

8 Examination of an alternative outcome variable

In an attempt to take into account blocking and to provide an alternative to our outcome variables in Section 5, we also created a regression model for a statistic from pro-football-reference.com, using all of the same predictors. This statistic is “Approximate Value (AV),” which is intended to assign a number to every player in each season to attempt to measure the player’s value to his team (Paine 2013). The number that we use as our outcome variable is the total AV for each player over the course of his career.

We were only able to obtain the AV for 198 of the 258 tight ends from our dataset of tight ends who entered the league in 2010 or earlier (to allow for at least four seasons to accumulate AV). The AV numbers that we use are accumulated through the 2013 season (while the rest of our data was only through 2012). One issue with our AV analysis, however, is that 160 of the 198 players in our sample had an AV below 20, while the remaining 38 were scattered from 20 to 106, which suggests a skewed distribution that is less desirable for predictive modeling.

Our regression model for AV has an adjusted R^2 of 0.311, which is higher than that of our other models. It also includes a very similar set of predictors, but with slightly more of a focus on the size variables. Career college yards, forty-yard dash, and Broad Jump were significant at the 5%

level; height, weight, and BMI were significant at the 10% level; and, the BCS dummy variable was also included in the model with lower significance.

The results were not extremely surprising for this model, as it would be expected that the predictors for AV would be similar to those of our other models. Players with more receiving production and more games started would likely have a higher AV. Additionally, as AV takes into account overall value, which likely includes blocking, it would be expected that there would be somewhat more of a focus on the size predictor variables.

9 Summary and discussion

In this paper, we have examined the extent to which the NFL draft results and NFL career success for tight ends can be predicted from pre-draft information. Our pre-draft data consisted of college performance, combine measures and physical attributes (BMI, Weight and Height).

We employed both linear regression models and recursive partitioning decision trees to predict NFL draft order and NFL career success. With both modeling approaches, we find that the pre-draft measures that are most predictive of NFL draft order are not necessarily the most predictive measures of NFL career success.

As an example, the combine Bench Press was selected as predictive of NFL draft order, but was not important in predicting NFL career success, whereas the combine Broad Jump was selected as predictive of NFL career success, but not of NFL draft order. The only measures that were selected as predictive of both NFL Draft and NFL Career Score were the combine Forty Yard Dash and an indicator variable for whether the player played at a BCS college. An overall summary of the predictor variables included in the regression models can be seen in Table 5 at the end of the paper, in which a “+” indicates that the predictor was included with a positive coefficient, while a “–” indicates that the predictor was included with a negative coefficient.

These findings suggest that drafting strategy could be improved by focusing more on measures that are predictive of NFL performance. We explore the potential to improve drafting strategy further by also factoring the salary cost of drafted players into our prediction models.

Our cost analysis suggests that a primarily college performance measures (especially college receptions and the BCS college indicator) with a couple combine measures (especially Broad Jump) are the best predictors of high value tight ends. We also find that current NFL

Table 5 Summary of regression results.

	Draft Order (– = earlier and + = later)	NFL games started	NFL career score	NFL career score per game
Size				
Height	–	+		
Weight		–	+	
BMI	–	+		
College				
BCS Dummy	–	+	+	+
Yards/Rec		+	+	
Career Rec		+	+	
Career Yards	–			+
Career TDs			–	–
Final year Rec, %				
Final year yards, %		–		
Final year TDs, %		+		
Combine				
Forty yard dash	+		–	–
Bench press	–			+
Vertical				
Broad jump		+	+	+
Shuttle				
Three cone drill				

draft strategies are efficiently utilize size measures (BMI, Height, Weight), as these measures were predictive of NFL draft order and NFL performance, but were not selected as predictive of NFL performance per cost. This shows that while size is beneficial, NFL teams are aware of this and draft larger tight ends earlier, as is appropriate. We also examined an alternative outcome variable, “Approximate Value (AV)”, and found similar results with this outcome variable, though AV was only available for a smaller set of players.

As a final illustration of a case in which our predictive modeling approach would have been helpful, we consider a case study that compares two tight ends from the 2009 NFL Draft. The Denver Broncos selected Richard Quinn with the 64th pick (second round) of the 2009 NFL Draft, and the Tennessee Titans selected Jared Cook with the 89th pick (third round) of the same draft.

Based on their pre-draft college and combine measures, and using our predictive modeling approach in Section 5, we would have predicted NFL career score of 2308 for Jared Cook versus a predicted NFL career score of 463 for Richard Quinn. Thus, our model suggests in this particular example that Jared Cook should have been drafted earlier than Richard Quinn.

The NFL careers of the two players since the 2009 draft support our model predictions. Jared Cook is currently

the starting tight end for the St. Louis Rams after four productive seasons with the Tennessee Titans, amassing a NFL career score of 2638.9 through the 2013 season. In contrast, Richard Quinn is currently a free agent and has not played a game since the 2011 season, with an amassed NFL career score of only 9.

We should mention that there are several factors that are not taken into account by our prediction models. One result of our analysis is that the indicator of whether or not the player attended a BCS college was predictive of both NFL draft order and NFL performance. Beyond that single variable, there could be predictive power in the specific conference of the college that the player attended, though there are smaller samples within each of these categories. As an example, the Southeastern Conference (SEC) produced the BCS national champion every year from 2006 to 2012 and had 63 players taken in the 2013 NFL Draft, which is more than double the number of players drafted from any other conference (Patterson 2013). This implies that players from the SEC may be superior to those of other conferences, but further analysis would be needed to test this hypothesis.

We have also not accounted for the number of opportunities available to a player while in college. Some players are fortunate enough to come into a situation in which their college football team needs them to start immediately as a freshman, providing them the greatest opportunity to accumulate statistics in college. However, some players do not have that opportunity. As an example, consider the two top tight end prospects in the 2013 NFL Draft: Tyler Eifert is projected by our model to have an NFL Career Score of 2292 whereas Zach Ertz is projected to have an NFL Career Score of only 934. Why the big difference in their projections? Eifert was a starter in college for two full seasons before entering the NFL, whereas Ertz was only a starter for his final season in college, as he was second string to Coby Fleener until Fleener entered the NFL.

The issue of playing time in college is exacerbated by the fact that some tight ends play a different position or sport for a portion (or all) of their college career. A salient example of this is Jimmy Graham, who played college basketball at the University of Miami before joining the football team as a tight end for his final year in college. Graham had low career college receiving statistics, which leads to low predicted NFL performance from our model even though he has been one of the top tight ends in the NFL over the past few years.

The focus of this paper has been the extent to which *quantitative* measures of NFL career success can be predicted based on *quantitative* pre-draft data. There are many non-quantitative factors that are not built into our

models. Injuries are one such factor that can have a large impact on a player's career, and there may be some power to predict injuries from that player's college career. As an example, Rob Gronkowski was drafted 42nd overall in the 2010 NFL Draft by the New England Patriots as he was coming off a missed season at Arizona due to surgery on a herniated disk in his back. Over the past few seasons in the NFL, Gronkowski has been extremely productive when healthy but has missed substantial amounts of time due to multiple injuries including a broken forearm and surgery to repair another herniated disk in his back.

In addition to injuries, NFL career success can also be heavily influenced by off-field personal issues. In the 2010 draft, the New England Patriots also selected Aaron Hernandez with the 113rd overall pick. One reason that Hernandez fell so far in that draft was that many teams were concerned with his psychological makeup. Many NFL teams subscribe to Human Resource Tactics' psychological analyses of NFL prospects. Human Resource Tactics' testing of Hernandez provided a rating of 1 out of 10 in "Social Maturity" (Clegg 2013). Although Hernandez was a productive tight end while on the field in the subsequent three seasons, he was released by the Patriots after being arrested on murder charges and is currently in prison.

Beyond psychological issues, another factor that might have predictive power is the "football intelligence" of a particular player. A tight end's football intelligence is important because they need to be able to learn complicated offensive schemes, memorize the routes they are supposed to run, and recognize patterns in the defense to find weak spots where they can catch passes. NFL scouts can learn about a player's football intelligence through talking with the player in interviews and asking his past coaches how quickly he learned their offensive scheme. There is also the *Wonderlic test*, a quantitative cognitive intelligence test that NFL prospects take at the NFL Combine. This football intelligence data, however, is not made available to the public aside from a few scores that have been leaked to the press. We only had access to Wonderlic scores for 5 out of 315 tight ends, so we were unable to use it in our analysis. However, Berri and Simmons (2009) a study on quarterbacks did have sufficient Wonderlic data to use (likely because quarterback Wonderlic scores are more often leaked to the public). That study did conclude that Wonderlic scores were important for prediction of draft results, but not for future NFL performance.

It would be an interesting future endeavor to model injury potential as a function of available college data, and possibly account for available psychological and "football intelligence" data in our predictions of NFL performance.

Our reason for using tight ends as our position to analyze was this changing role in NFL offensive schemes. Tight ends have recently become a more significant part of NFL passing offenses than they have ever been in the past, as they are now seen as another weapon that can be exploited to pick apart opposing defenses. Though we only analyze tight ends in this study, our general methodology could be applied to other positions in football as well as players in other professional sports.

References

- Berri, D. J., S. L. Brook, and A. J. Fenn. 2010. "From College to the Pros: Predicting the NBA Amateur Player Draft." *Journal of Productivity Analysis*, July 16, 2010.
- Berri, D. J. and R. Simmons. 2009. "Catching A Draft: On the Process of Selecting Quarterbacks in the National Football League Amateur Draft." *Journal of Productivity Analysis*, September 18, 2009.
- Breiman, L., J. Friedman, C. J. Stone, and R. A. Olshen. 1984. *Classification and Regression Trees*. Chapman and Hall/CRC. First edition.
- Burger, J. D. and S. J. K. Walters. 2003. "Market Size, Pay, and Performance: A General Model and Application to Major League Baseball." *Journal of Sports Economics* May 1, 2003.
- Clegg, J. 2013. "Aaron Hernandez: An Early Warning in the 2010 NFL Draft Profile." *Wall Street Journal* July 3, 2013.
- Dhar, A. 2011. "Drafting NFL Wide Receivers: Hit or Miss?" Technical report. Department of Statistics, University of California at Berkeley.
- Hocking, R. R. 1976. "The Analysis and Selection of Variables in Linear Regression." *Biometrics* 32:1–49.
- Massey, T. and R. Thaler. 2005. "Overconfidence vs. Market Efficiency in the National Football League." *National Bureau of Economic Research* April 2005.
- Overthecap.com 2013. "2013 NFL Draft Rookie Contract and Salary Cap Estimates." April 26, 2013. <http://overthecap.com/nfl-rookie-salary-cap.php>.
- Paine, N. 2013. "Approximate Value." *Pro Football Reference Blog* June 14, 2013. <http://www.sports-reference.com/blog/approximate-value/>.
- Patterson, C. (2013). "2013 NFL draft by conference: SEC doubles the competition". *CBSsports.com*, April 29, 2013.
- Stuart, C. 2008. "The Final Value of a Passing Touchdown." *Pro Football Reference Blog* October 3, 2008. <http://www.pro-football-reference.com/blog/?p=633>.