Paper 2101-2016

Using the Kaplan-Meier Product-Limit Estimator to Adjust NFL Yardage Averages

Keith Curtis, USAA

ABSTRACT

Average yards per reception, as well as number of touchdowns, are commonly used to rank National Football League (NFL) players. However, scoring touchdowns lowers the player's average since it stops the play and therefore prevents the player from gaining more yardage. If yardages are tabulated in a life table, then yardage from a touchdown play is denoted as a right-censored observation. This paper discusses the application of the SAS/STAT® Kaplan-Meier product-limit estimator to adjust these averages. Using 15 seasons of NFL receiving data, the relationship between touchdown rates and average yards per reception is compared, before and after adjustment. The modification of adjustments when a player incurred a 2-point safety during the season is also discussed.

INTRODUCTION

It is intuitive that a player who tends to gain more yards from a forward pass reception is also likely to score more touchdowns than an average player. So 2 metrics commonly used to compare players are their average yards-per-reception and number of touchdowns. However, scoring a touchdown is detrimental to the yardage average, since the touchdown stops the play, and therefore prevents the player from gaining more yardage. If we assume the player would have continued gaining yards on an analogous play further from their opponent's goal line, then the yardage value is a right-censored observation. With a right-censored observation, the "true" value is higher than the recorded observation. This contrasts from a left-censored observation, where the recorded observation is higher than the "true" value.

Consider a passing play where a team has possession at their own 20-yard line, and a player gains 8 yards to the 28-yard line. However, if the identical play is executed at their opponent's 5-yard line, then the player can only gain 5 yards before scoring a touchdown. Since we assume this play would yield 8 yards when the line of scrimmage is 8 or more yards from their opponent's goal line, this 5-yard touchdown is considered a right-censored observation.

The touchdown rate, or the ratio of the number of touchdowns scored and number of passes caught, is studied in this analysis. Rates are needed because of the large variation in reception counts among players. The hypothesis proposed is that 2 conditions hold when player yardage averages are displayed against touchdown rates on a scatterplot. These 2 conditions are shown below:

1. A positive correlation exists between yardage averages and touchdown rates.
2. The strength of this positive correlation is diminished with high touchdown rates.

If both conditions hold true, then correcting this downward bias through survival analysis would be attempted. The chosen technique for performing the correction is the Kaplan-Meier product-limit estimator. It is ideal because of the properties listed below:

1. It is designed to handle right-censored observations.
2. It is a non-parametric technique, and so no assumptions about the shape of the distribution are necessary. The shape of the survival curve is not of interest for this analysis. The sole purpose for calculating the curve is to derive an adjusted yardage average from it. One property of a survival distribution is that the area under the curve is an estimate of the mean of the random variable.
3. When a Kaplan-Meier is performed on data with no censored observations, the empirical survival distribution calculated by traditional means is obtained. So the mean calculated from this survival distribution is identical to the sample average among the observations.

KAPLAN-MEIER EXAMPLE

Consider a player who caught 5 passes and scored 1 touchdown. The touchdown reception was 3 Yards, while the remaining 4 catches gained 0, 2, 3 and 5 yards. If the touchdown reception is not treated as a censored observation, then the average yards-per-reception is estimated to be 2.6 yards, and the estimate of the probability of a reception longer than 3 yards is 0.2 or 1/5. Both values are underestimated without applying the Kaplan-Meier technique.

Kaplan-Meier will estimate the probability of a reception longer than 3 Yards using the formula below:

$$P[\text{Reception} > 3 \text{ Yards}] = P[\text{Reception} \geq 3 \text{ Yards}] * P[\text{Reception} > 3 \text{ Yards} \mid \text{Reception} \geq 3 \text{ Yards}]$$

The rationale for this formula is that a reception cannot be longer than 3 yards unless 3 yards is first reached. So the probability of a reception of 3 yards or more is first obtained. The conditional probability of a reception longer than 3 yards, given that the reception is at least 3 yards, is then determined. The product of these values is then taken to yield the final estimate. The application of products is why the full name of the technique is the Kaplan-Meier product-limit estimator.

In our example above, 3 of the 5 receptions are 3 Yards or longer, including the touchdown reception. So P[Reception ≥ 3 Yards] = 3/5. Among these 3 receptions, the touchdown value is treated as being above 3 Yards. When grouped with the 5-Yard reception, the P[Reception > 3 Yards | Reception ≥ 3 Yards] = 2/3. The product of these fractions is 2/5 or 0.4, which is higher than the 0.2 estimate that does not consider censoring.

The non-adjusted and Kaplan-Meier adjusted survival curves are shown in Figure 1:

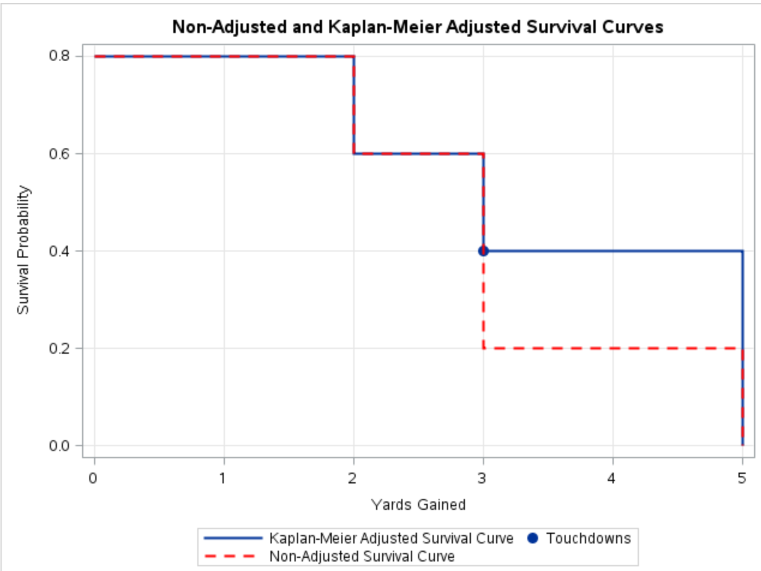

Figure 1. Example Survival Curves

Each yardage average can be determined by calculating the area under each curve. For the non-adjusted survival curve, which is represented by the red dotted line, the area is calculated by the following formula:

$$(0.8 * 2) + (0.6 * 1) + (0.2 * 2) = 2.6 \text{ Yards}$$

This 2.6 yards value matches the mean of the 5 yardage values.

The area under the blue curve, which represents the Kaplan-Meier survival curve, is 3.0 yards. It includes the 2.6 yards under the non-adjusted curve, plus the area of 0.4 yards between the curves.

ANALYSIS OBJECTIVES

The objectives of this analysis include the following:

1. Demonstrating the detrimental impact that a high rate of scoring touchdowns has on a player's yards-per-reception average.
2. Elimination of this impact by treating a touchdown as a right-censored observation, and then calculating an adjusted yards-per-reception average using the Kaplan-Meier product-limit estimator.
3. Highlight player/season combinations that yield large changes in their yards-per-reception rankings after their averages are adjusted.
4. Discuss why negative yardage from a play that yields a 2-point safety is left-censored, and present an approach for adjusting a player's average to accommodate both touchdowns and 2-point safeties.

This analysis primarily utilized 5 SAS® procedures, and use of each procedure is described below. SAS® University Edition is where all these procedures were run. IMPORT, SQL and RANK are Base SAS® procedures. LIFETEST and SGPLOT are SAS/STAT® and SAS/GRAPH® procedures, respectively:

1. <u>IMPORT</u>: Read relational databases found in Comma Separated Values (.csv) files into SAS.
2. <u>SQL</u>: Merge and aggregate the relational databases, and identify player/season combinations that caught 48 or more receptions.
3. <u>LIFETEST</u>: Calculate the Kaplan-Meier product-limit estimator for each identified player/season combination.

4. SGPLOT: Produce all graphs included in this within this paper, including scatterplots, bubble plots, a stacked bar graph, and step graphs used to represent survival distributions.
5. RANK: Rank player/season combinations in descending order by their yardage averages. Rankings were performed by primary position, and separate non-adjusted and Kaplan-Meier adjusted rankings were determined.

This analysis will emphasize passing plays over rushing plays, since touchdown rates tend to be higher on passing plays. For example, running back LaDainian Tomlinson scored a record 28 rushing touchdowns during the 2006 Season. But this season still only produced a touchdown rate of 8.0%, since Tomlinson had 348 rushing attempts. Receivers have sometimes achieved touchdown rates over 20% by contrast. However, this approach can still be used to adjust rushing averages, as well as yards-per-completion averages for quarterbacks.

The methodology for accommodating 2-point safeties will be demonstrated with rushing plays, since only 1 play with a completed forward pass produced a safety in this dataset. However, during the same fifteen seasons analyzed, 61 rushing plays yielded 2-point safeties for the defensive team.

DATA ANALYZED

Fifteen seasons of National Football League (NFL) data from 2000 – 2014 were purchased from Armchair Analysis (www.armchairanalysis.com). Within each season, players that caught 48 or more receptions were then identified. A minimum reception threshold was chosen so that more stable yards-per-reception averages and touchdown rates would be obtained. A threshold of 48 ensures that a player averages at least 3 catches per game during the 16-game NFL season. The primary offensive position was also obtained for each player, namely wide receiver (WR), tight end (TE) or running back (RB). A separate analysis will be performed for each primary position, since their yards-per-reception average and touchdown rate distributions vary greatly. The average values and relative frequencies by position are summarized in Table 1:

Primary Position	Player/Season Count	Relative Frequency	Non-Adjusted Average	Touchdown Rate
RB	150	13.90%	8.23	3.05%
TE	179	16.59%	11.49	8.43%
WR	750	69.51%	13.53	8.34%
Total	**1,079**	**100.00%**	**12.45**	**7.62%**

Table 1. Summary Statistics by Primary Position

OVERALL IMPACT OF TOUCHDOWN RATE ON RECEIVING AVERAGES

Figure 2 shows the relationship between touchdown rate and average yards-per-reception. Each symbol represents 1 of the 1,079 player/season combinations with 48 or more receptions between 2000 – 2014. A 2nd Order Locally Weighted Scatterplot Smoothing (LOESS) regression was also performed, with touchdown rate and average yards-per-reception representing the independent and dependent variable, respectively. LOESS regression is a non-parametric technique, and is preferred over parametric methods such as Ordinary Least Squares (OLS) regression for the following reasons:

1. No assumptions are required regarding the shape of the curve, such as a linear or parabolic relationship. This is ideal for this analysis, since the goal is to simply observe the shape, and not to perform statistical tests.
2. The estimated curve is dynamic, and therefore changes as touchdown rate increases. This occurs because a nearest neighbor algorithm is applied at each touchdown rate to determine which player/season combinations will contribute to the model at that touchdown rate. This differs from OLS regression, since OLS performs a single, composite fit based on all the observations.

The LOESS regression results are conveyed by the black curve. Based on this curve, it is apparent there is a positive relationship between touchdown rate and average yards. As the touchdown rate increases, however, the rate of change in average yards-per-reception decreases. This decrease in the line's slope is apparent until near the 75th Percentile of average yards (14.31 yards). At this point, the relationship shifts from a quadratic to linear one, and the line is shallower than the preceding curve.

The above player/season combinations are also color-coded by the player's primary position. It is apparent that running backs typically have the lowest touchdown rates and yards-per-reception averages, while wide receivers typically have the highest yardage averages. Tight end yardage averages fall between the other 2 primary positions, and their touchdown rates are similar to wide receivers. However, as the NFL has increasingly emphasized passing plays over running, many tight ends have begun to yield results more similar to wide receivers. One contribution is the increasing impact of "joker" tight ends. Tight ends traditionally lined up beside either the left or right tackle, and were often valued more for their blocking ability than for receiving. Joker tight ends often line up with space in-between them and the interior linemen (tackles, guards and center). They can be found in the slot between a tackle and wide receiver, or even in the wide receiver position:

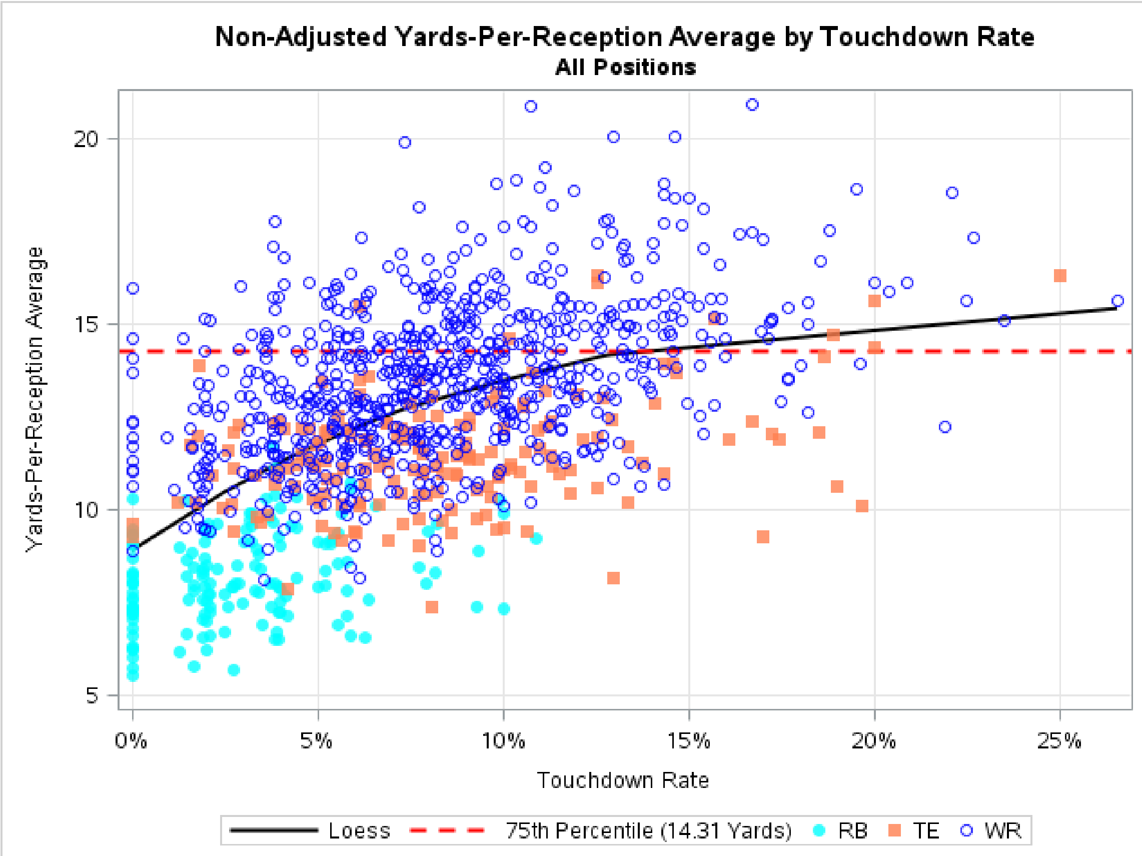

Figure 2. Scatterplot Between Non-Adjusted Average Yards and Touchdown Rate

Whether to compensate a joker as a tight end or wide receiver was a discussion during Jimmy Graham's contract negotiation in 2014. Although Jimmy Graham's primary position is listed as tight end, he only lined up as a tight end on 33% of his plays during the 2013 Season. Graham lined up in the slot and as a wide receiver on 45% and 22% of his plays, respectively. Each NFL team can apply a franchise tag to 1 player per year, and Jimmy Graham received that designation from his team at that time, the New Orleans Saints. Because wide receivers are guaranteed a higher salary than tight ends under the franchise tag, Graham wanted to be considered a wide receiver. NFL arbitrator Stephen Burbank ultimately ruled that he was a tight end.

As the role of tight ends has changed, the frequency of tight ends catching 48 or more passes in a season has also increased. Since 2011, at least 16 tight ends per season have achieved this feat, as compared to 7 at most from 2000 – 2004. The opposite trend has occurred with running backs. Between 2000 – 2003, at least 12 running backs caught 48 or more passes annually. Except for the 2013 Season, no more than 10 running have reached this threshold since 2004. These results are shown as a stacked bar graph in Figure 3:

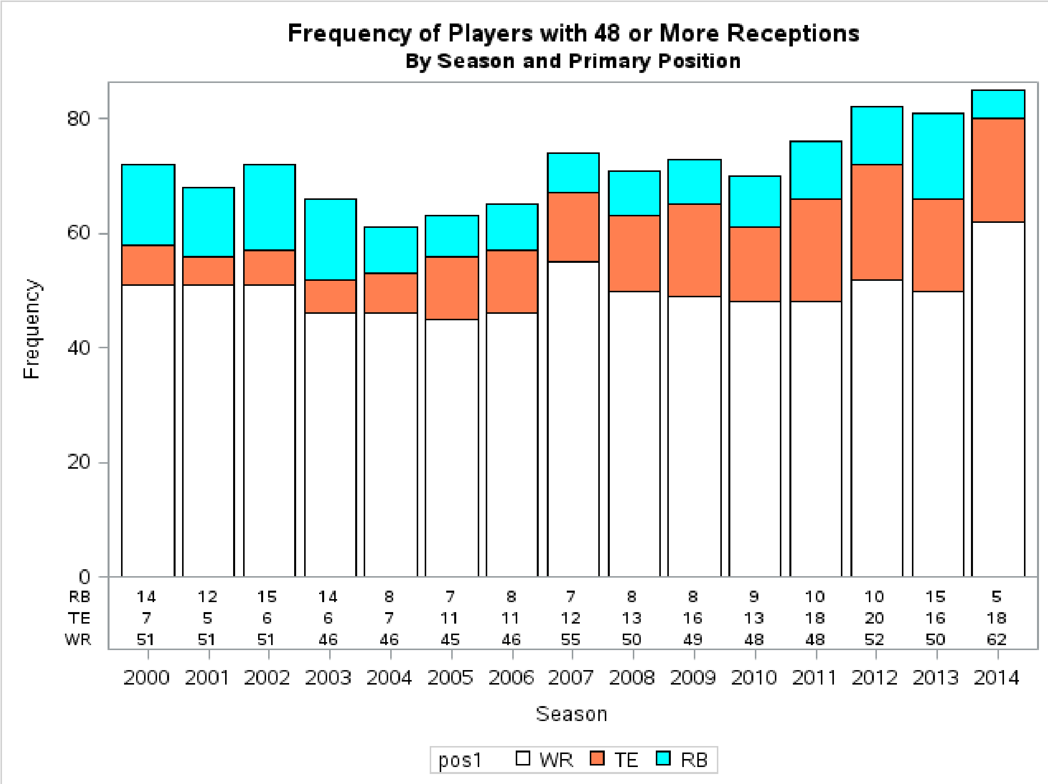

Figure 3. Stacked Bar Graph of Player Frequency by Season and Primary Position

Because of differences by position, as well as how their contributions to the passing game have changed, all subsequent analysis will be performed separately by position.

IMPACT OF TOUCHDOWN RATE ON RECEIVING AVERAGES BY POSITION

Figure 4 shows the relationship between touchdown rate and average yards-per-reception for wide receivers. The 2nd Order LOESS regression shows similar behavior to the regression in the overall graph. The linear relationship is shallower above the 75th Percentile of average yards (14.89 Yards) than below it. The relationship for running backs is shown in Figure 5. It also shows a steeper line followed by a shallower line. However, the inflection occurs before the 75th Percentile of 9.16 Yards is reached. So scoring a higher rate of touchdowns is detrimental to the yardage average for both wide receivers and running backs.

Tight ends exhibit a different relationship. This can be attributed to differences in the distributions of touchdown rates and yardage averages between traditional and joker tight ends. Three examples of joker tight ends are Vernon Davis, Antonio Gates and Rob Gronkowski. Although they account for only 19 of the 179 player/season combinations (10.6%) among tight ends, all 3 yield 2 seasons among the highest 10 non-adjusted yards-per-reception averages. Also, only 3 of the 179 tight end player/season combinations (1.7%) produced touchdown rates of 20% or better. Davis, Gates and Gronkowski each yielded 1 of these 3 seasons.

By calculating hypergeometric distribution probabilities from the PROBHYPR Function in SAS, the probability of these exceptional performance rankings under random chance can be determined. The 4 parameters for the PROBHYPR Function are explained below:

PROBHYPR = (N, K, n, x),

where N = Overall population size

K = Number of observations within a subset of the population

n = Number of observations sampled without replacement from the population

x = A random variable that indicates the number of observations from the sample of size n belonging to the population subset of size K.

Note that the returned probability is a cumulative probability, or the probability of observing a value between zero and x. Also note that all 4 parameter values are integers:

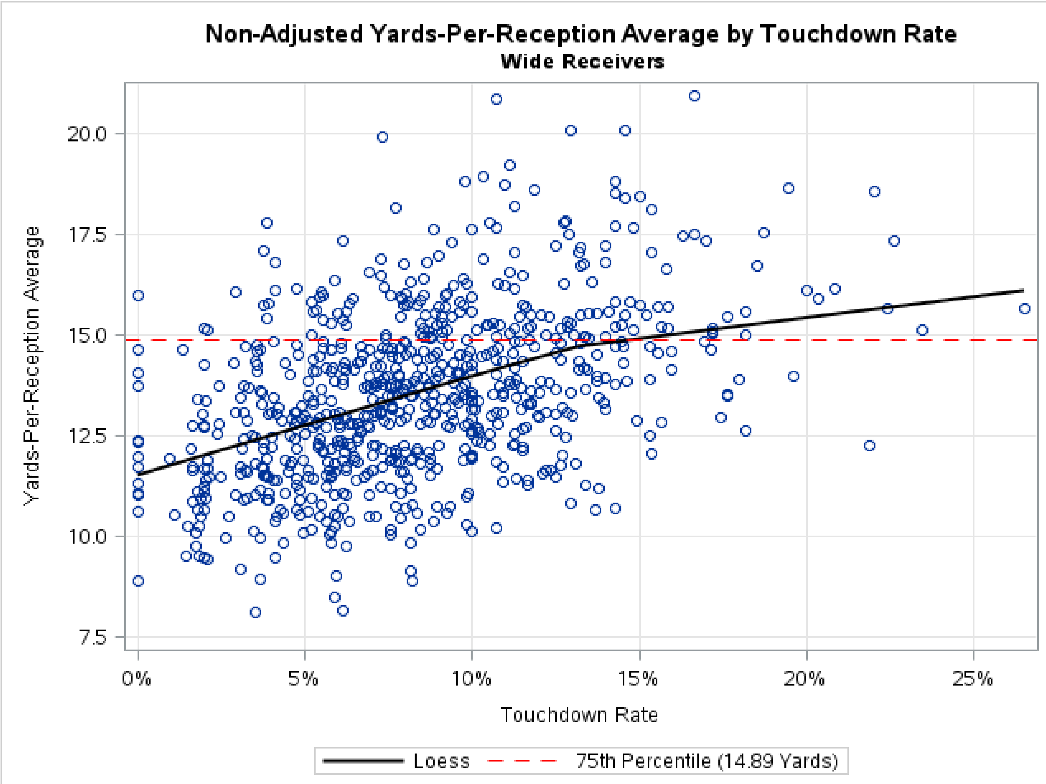

Figure 4. Scatterplot Between Non-Adjusted Average Yards and Touchdown Rate (Wide Receivers)

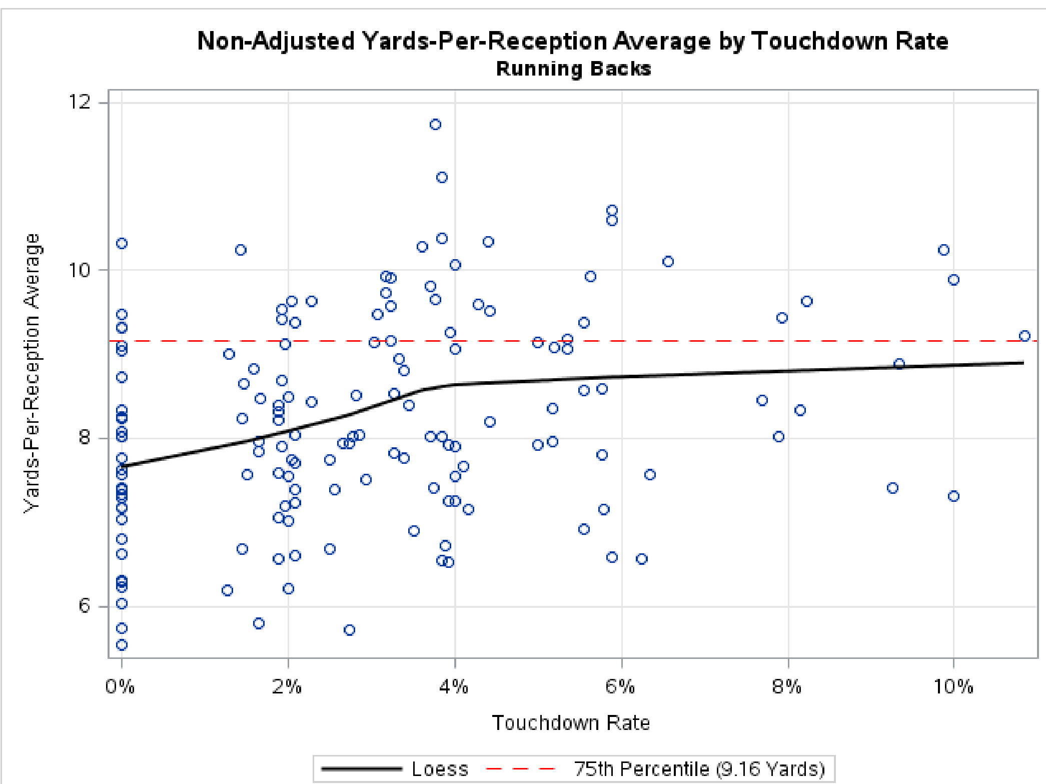

Figure 5. Scatterplot Between Non-Adjusted Average Yards and Touchdown Rate (Running Backs)

For the tight end data, N = 179, which corresponds to the number of player/season combinations. Since Vernon Davis, Antonio Gates and Rob Gronkowski account for 19 combinations, K = 19. To determine the probability of these 3 players yielding at least 6 of the highest 10 receiving averages, the following line is executed in SAS:

```
top10_probability = 1 - probhypr(179,19,10,5);
```

This command yields a probability of only 0.0001.

For the probability of these 3 tight ends yielding the 3 best touchdown rates, the following line is executed:

```
top3_probability = 1 - probhypr(179,19,3,2);
```

This command yields a probability of only 0.001.

Because the accomplishments of Davis, Gates and Gronkowski are so exceptional, 2 separate scatterplots were created. Figure 6 contains all 179 player/season combinations for tight ends, and the 19 player/season combinations of Davis, Gates and Gronkowski are highlighted. The LOESS regression in the first scatterplot shows an inflection near the 75th Percentile, just like the overall and wide receiver results. However, the linear relationship becomes steeper above this percentile, instead of shallower. This is because the higher average yardage performances for joker tight ends outweigh the detrimental impact of the high touchdown rates.

These 19 seasons of Davis, Gates and Gronkowski are excluded from Figure 7, and a modest increase in average yards is yielded across all touchdown rates:

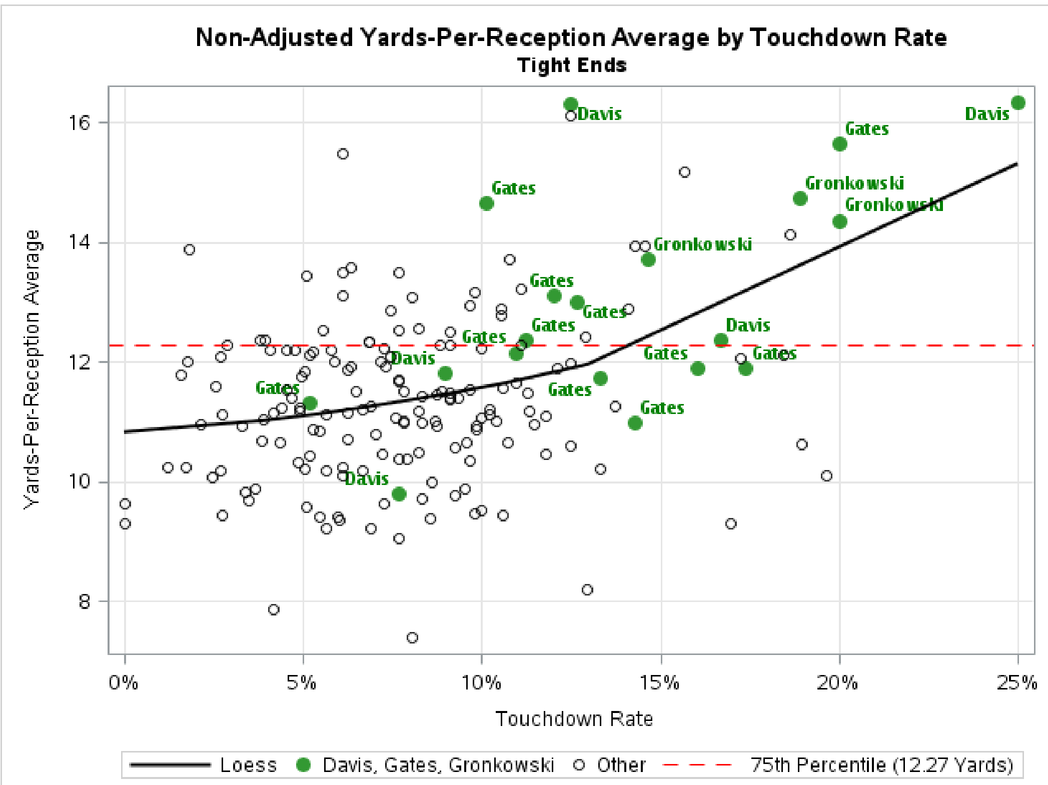

Figure 6. Scatterplot Between Non-Adjusted Average Yards and Touchdown Rate (All Tight Ends)

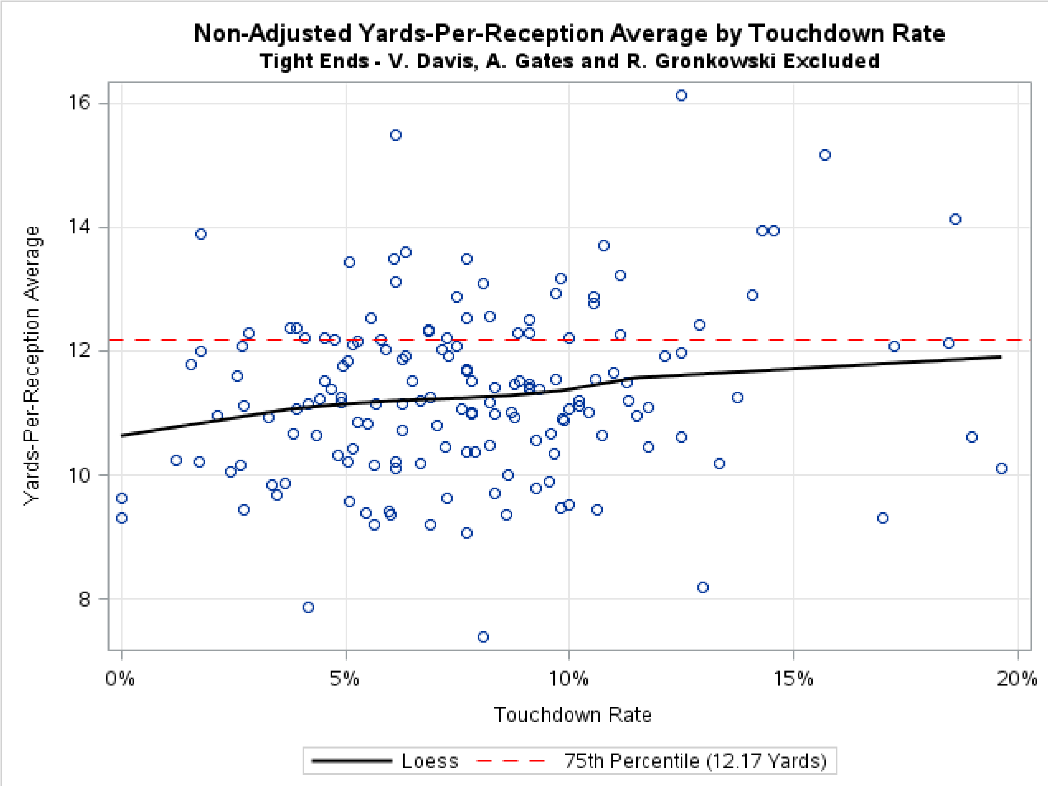

Figure 7. Scatterplot Between Non-Adjusted Average Yards and Touchdown Rate (3 Tight Ends Excluded)

CALCULATION OF THE KAPLAN-MEIER ADJUSTED AVERAGES

PROC LIFETEST code was executed separately on the 1,079 player/season combinations, and a generic version is shown below. The (1) following td_censoring specifies that yardage values associated with a value of 1 for td_censoring are censored observations. The weight statement is used because the data was aggregated to each yardage and censoring status (0 or 1) combination, and play_count indicates the number of receptions that yielded the corresponding yardage/censoring combination. The mean value resulting from the Kaplan-Meier survival curve is output to the nonneg_avg dataset:

```
proc lifetest data = player;
  time yds * td_censoring(1);
  weight play_count;
  title "Kaplan-Meier Product-Limit Estimator for a Single Player/Season";
  ods output means = nonneg_avg;
run;
```

Note that PROC LIFETEST does not accommodate negative observations. So the mean stored in the nonneg_avg dataset only reflects receptions of 0 or more yards. So a player's final Kaplan-Meier adjusted average is a weighted average of the raw negative reception average and the non-negative average estimate obtained from PROC LIFETEST. The number of receptions within each group provide the weights. The formula for the Kaplan-Meier adjusted average is shown below:

$$\frac{[\text{ (\# of Negative Rec.) * (Observed Negative Avg.) + (\# of Non-Neg. Rec.) * (Kaplan-Meier Non-Neg. Avg. Estimate) }]}{\text{Total \# of Receptions}}$$

The longest reception for each player/season was always coded as a non-censored observation, even if they scored a touchdown on the play. This was done because the mean estimate will be biased downward when the largest observation is censored. This estimate can even become lower than the non-adjusted average. There are alternative techniques to hard-coding the observation as non-censored, and they extend the survival curve beyond the largest observation. These techniques are beyond the scope of this paper. PROC LIFETEST provides the following warning when the largest observation is censored:

> **Note**: The mean survival time and its standard error were underestimated because the largest observation was censored and the estimation was restricted to the largest event time.

Figure 8 compares the non-adjusted and Kaplan-Meier adjusted curves for 1 player/season. Jordy Nelson's 2011 Season was chosen because it yielded the largest average yardage increase from the censoring of touchdowns, and therefore the greatest amount of area between the survival curves:

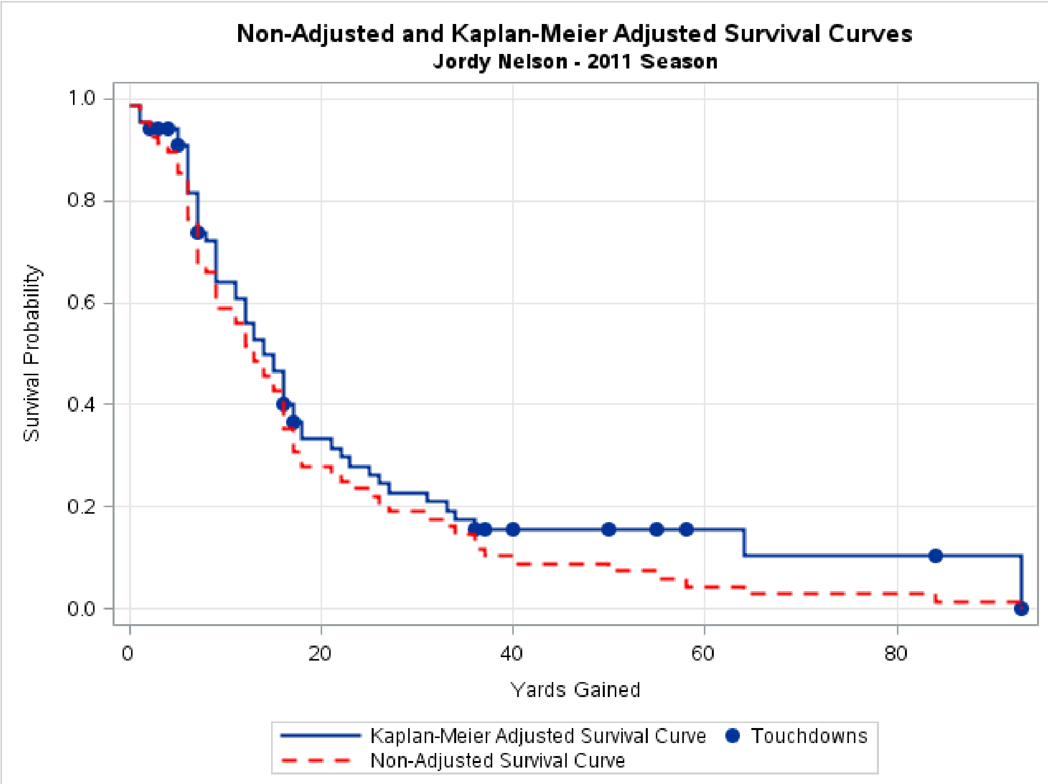

Figure 8. Survival Curves for Jordy Nelson's 2011 Season

Nelson's 2011 Kaplan-Meier adjusted average is 24.67 Yards, which is 6.10 Yards higher than his non-adjusted average of 18.57 Yards. The red dashed line represents the non-adjusted curve, and it is never greater than the solid blue line that represents the adjusted curve. Jordy Nelson's 15 touchdowns are shown as circles on the blue curve. At these 15 yardage values, the survival probability will decrease more for the non-adjusted curve than the adjusted curve, and therefore the curves diverge more. This divergence becomes more pronounced over 35 Yards, as 8 of Jordy Nelson's 10 receptions of 36 Yards or longer yielded touchdowns. Nelson's 93 Yard touchdown was not censored, however, as it was also his longest reception during the 2011 Season. If the area between the 2 curves is calculated, then the 6.10 Yard difference will be obtained.

Table 2 contains the PROC LIFETEST table for Jordy Nelson's 2011 Season. It is found on the last page. Note that the survival times marked with an (*) are censored observations. Since Nelson had no receptions that lost yardage during that season, all 68 of his receptions from 2011 are included.

BEFORE/AFTER COMPARISONS OF THE RELATIONSHIP BETWEEN TOUCHDOWN RATES AND RECEIVING AVERAGES

This section contains 2 scatterplots for each primary position. The first scatterplot displays non-adjusted averages by touchdown rate, while the second shows the Kaplan-Meier adjusted average on the y-axis. Each scatterplot contains a red dashed line that represents an Ordinary Least Squares (OLS) regression line. If the LOESS and OLS regression results align more closely in the second graph, then the LOESS line is becoming more linear. This alignment shows that censoring touchdowns through the Kaplan-Meier technique successfully reduced the detrimental effect that touchdowns have on average yardage average.

WIDE RECEIVERS

Wide receiver before/after scatterplots are shown in Figures 9 and 10, respectively. In Figure 9, the OLS and LOESS regression values diverge by more than a yard as touchdown rates increase. But the 2 regressions align closely in Figure 10. The LOESS regression still exhibits some shallowing, however. Further research could include choosing settings other than the default parameters for the LOESS command in PROC SGPLOT, and assessing how these changes impact the shape of the curve. Another consideration is the impact of other factors on receiving averages, such as the original line-of-scrimmage distribution. For example, if a player catches a high frequency of passes while their team is in the red zone, or within 20 yards of their opponent's goal line, this will increase their frequency of short pass receptions. So these receptions could be detrimental to their non-adjusted average. However, unless the play resulted in a touchdown, no adjustment is currently applied to these red zone receptions:

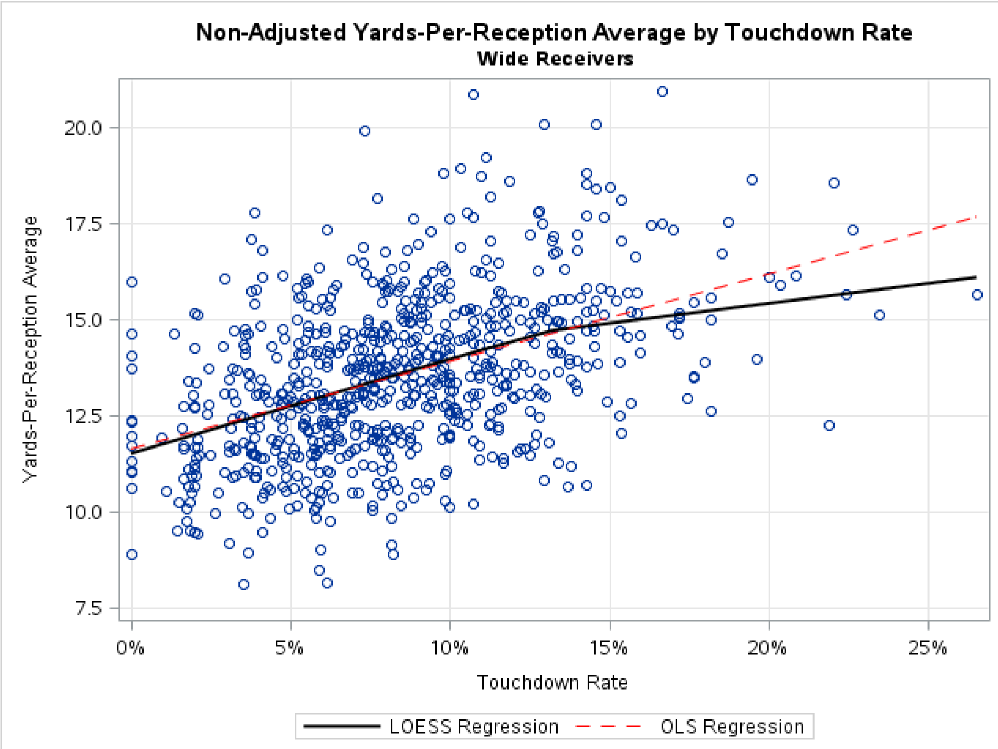

Figure 9. Scatterplot Between Non-Adjusted Average Yards and Touchdown Rate (Wide Receivers)

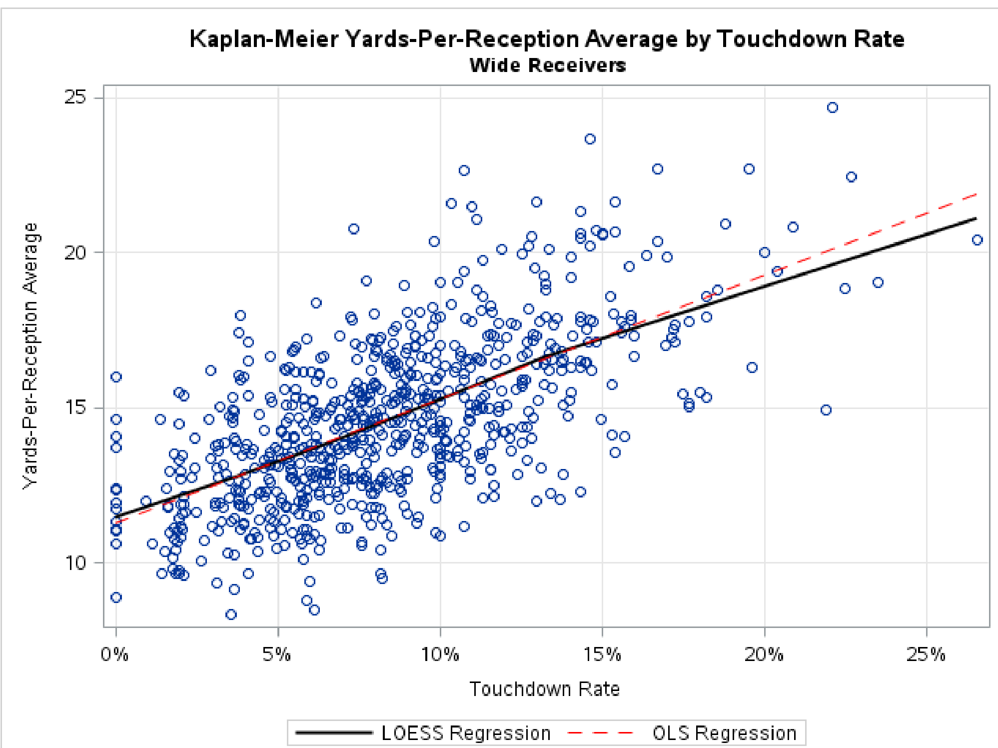

Figure 10. Scatterplot Between Kaplan-Meier Adjusted Average Yards and Touchdown Rate (Wide Receivers)

RUNNING BACKS

The running back before/after scatterplots are shown in Figures 11 and 12, respectively. Just like wide receivers, the LOESS and OLS regressions align more closely under the Kaplan-Meier average. The disparity between the before/after graphs is larger for wide receivers, however, since their highest touchdown rates are double the best rates for running backs:

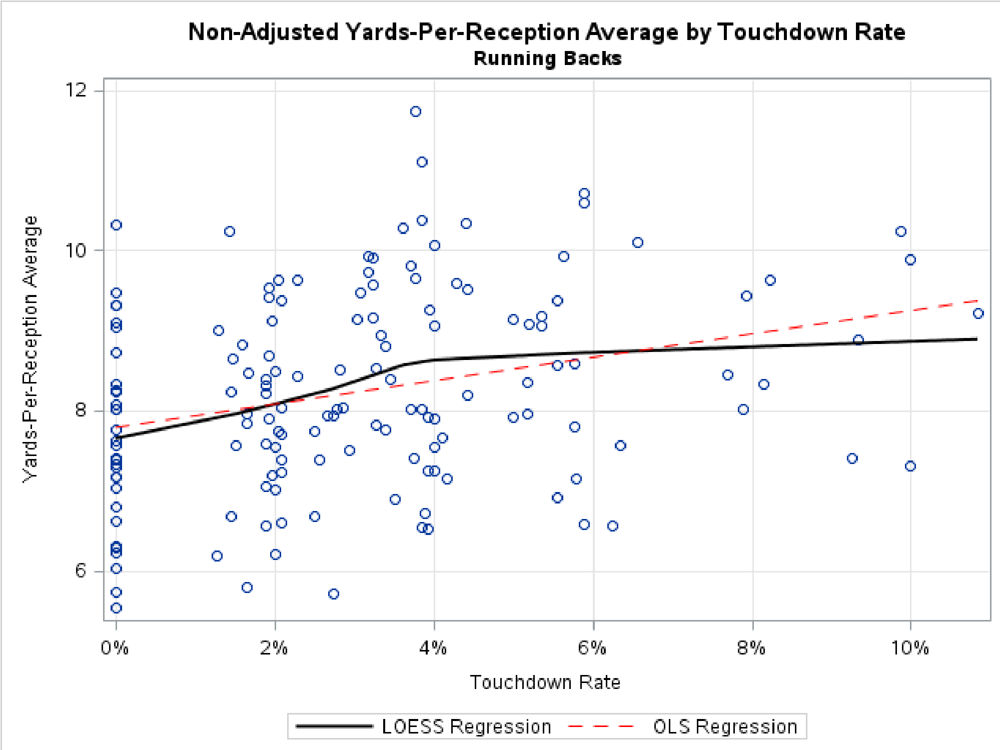

Figure 11. Scatterplot Between Non-Adjusted Average Yards and Touchdown Rate (Running Backs)

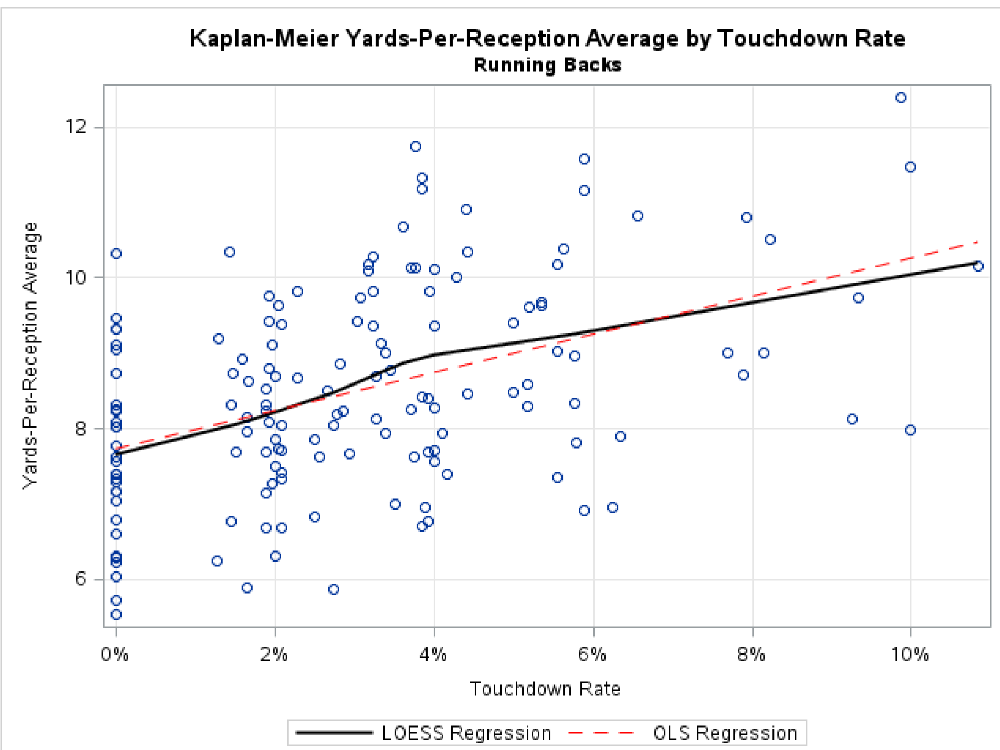

Figure 12. Scatterplot Between Kaplan-Meier Adjusted Average Yards and Touchdown Rate (Running Backs)

ALL TIGHT ENDS

The before/after scatterplots for all tight ends are shown in Figures 13 and 14, respectively. Unlike the other primary positions, the LOESS and OLS regressions do not align better under the Kaplan-Meier adjusted averages. In fact, the

LOESS regression is more volatile in the Kaplan-Meier graph, since the adjustments for touchdown rate accentuate the exceptional performances of joker tight ends:

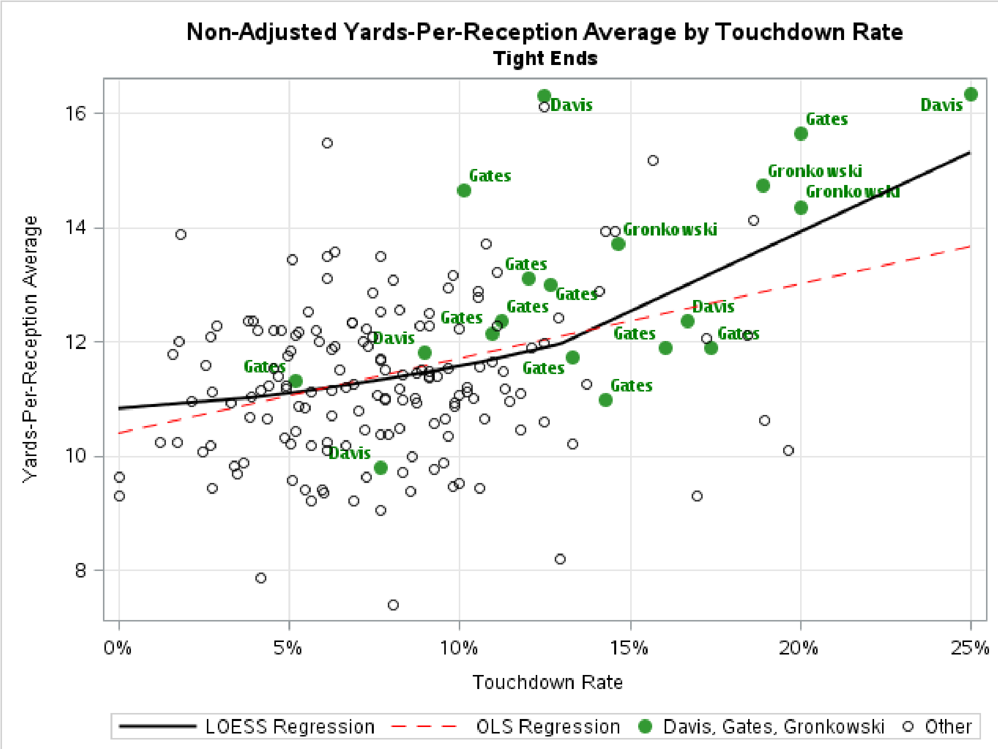

Figure 13. Scatterplot Between Non-Adjusted Average Yards and Touchdown Rate (All Tight Ends)

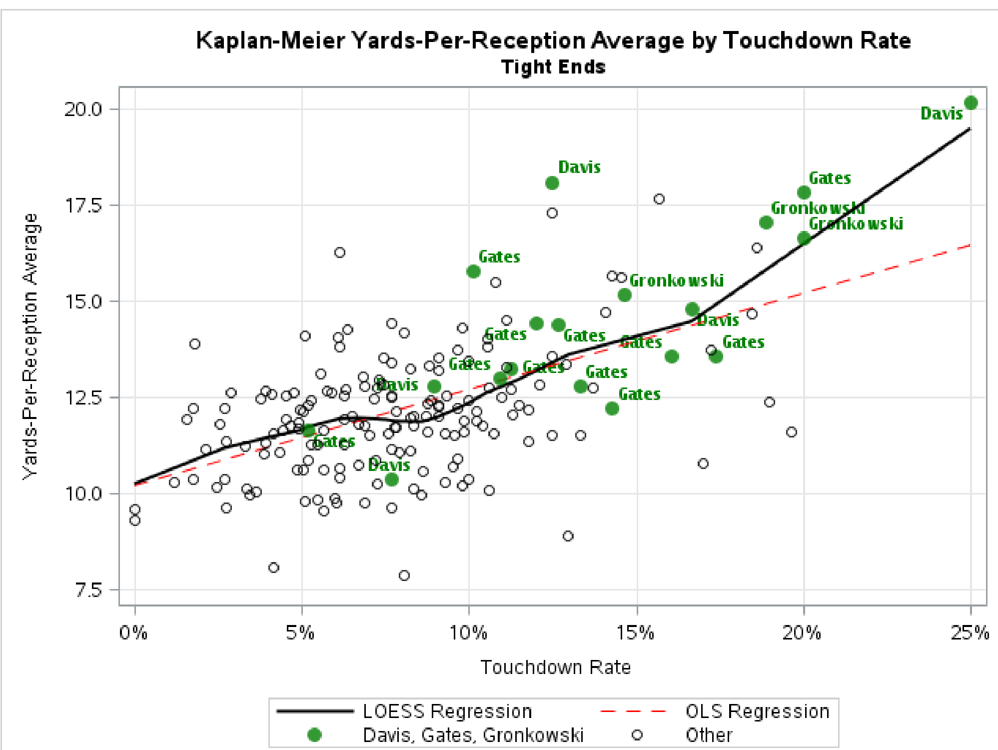

Figure 14. Scatterplot Between Kaplan-Meier Adjusted Average Yards and Touchdown Rate (All Tight Ends)

TIGHT ENDS – DAVIS, GATES AND GRONKOWSKI EXCLUDED

The tight end before/after scatterplots with Vernon Davis, Antonio Gates and Rob Gronkowski's 19 combined seasons excluded are shown in Figures 15 and 16, respectively. The LOESS and OLS regressions are closely aligned in both graphs. However, the significance tests of the slope from the OLS regressions still indicate that touchdown rates are detrimental to yardage averages. In Figure 15, the slope of the OLS regression line is marginally significant, as the p-value is only 0.0189. However, the OLS slope in Figure 16 yields a highly significant p-value (< 0.0001):

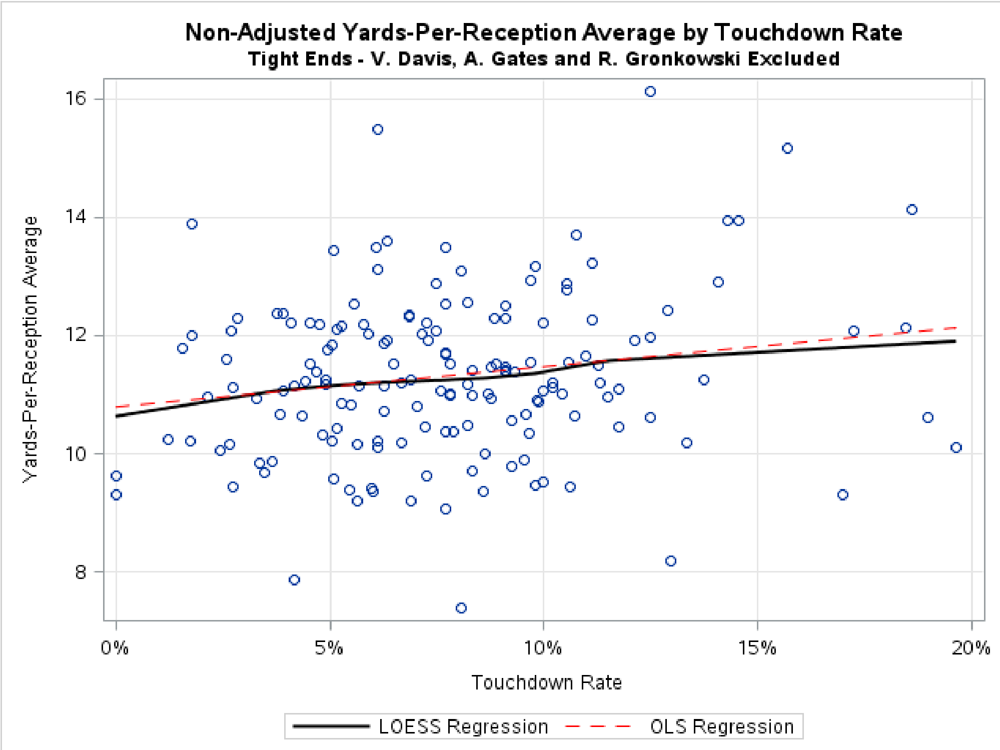

Figure 15. Scatterplot Between Non-Adjusted Average Yards and Touchdown Rate (3 Tight Ends Excluded)

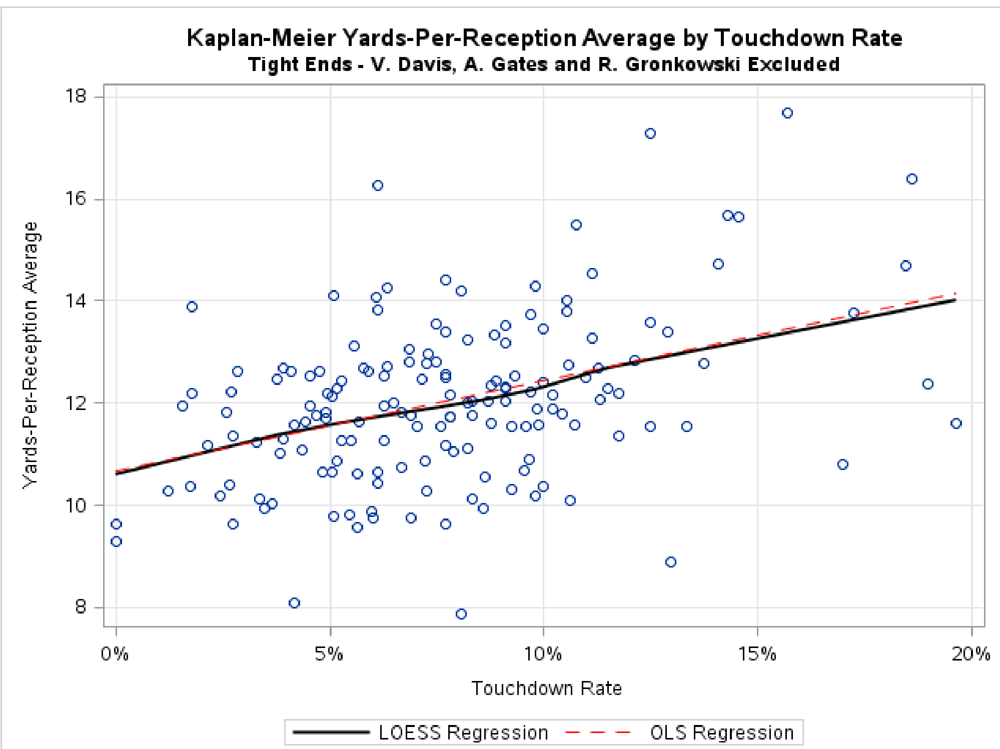

Figure 16. Scatterplot Between Kaplan-Meier Adjusted Average Yards and Touchdown Rate (3 Tight Ends Excluded)

BEFORE/AFTER COMPARISONS OF INDIVIDUAL PLAYER RANKINGS

This section contains bubble plots that display how player/season rankings compare between the non-adjusted and Kaplan-Meier adjusted receiving averages. There are 3 bubble plots, or 1 for each primary position, and larger bubbles correspond to higher touchdown rates. Each plot contains 50 player/season bubbles, or 1 for each of the 50 highest Kaplan-Meier averages. The x-axis contains non-adjusted averages, while Kaplan-Meier averages are represented by the y-axis. Each plot also contains a 45-degree line that represents where the 2 average values are equal, and therefore players who scored no touchdowns during a season will reside on that line. Higher touchdown rates will increase the positive difference between the Kaplan-Meier and non-adjusted averages, and therefore move the player/season bubble further above the 45-degree line.

Top-ranked player/seasons for either average are color-coded and labelled on each graph. Each label contains the player's abbreviated name, namely their first initial and last name, the corresponding season, and 2 numbers separated by a slash. The first and last numbers are their non-adjusted and Kaplan-Meier average rankings, respectively.

WIDE RECEIVERS

Figure 17 contains the wide receiver bubble plot, and twelve players who achieved a ranking of 8th or better for either average are highlighted. The 4 players highlighted in red, namely Jordy Nelson, Randy Moss, Greg Jennings and Santonio Holmes, are ranked in the Top 8 for their Kaplan-Meier average, but not for their non-adjusted average. Jordy Nelson's 2011 Season moved from 13th to 1st, due to a touchdown rate of 22.1%, which is the fifth highest rate among wide receivers. The 2007 Season for Greg Jennings gained 27 positions under the Kaplan-Meier Average, as his touchdown rate of 22.6% is the third highest. Both Nelson and Jennings played for the Green Bay Packers during their respective seasons. The 2 remaining receivers in this group, namely Randy Moss and Santonio Holmes, produced touchdown rates in the top decile during their 2000 and 2007 Seasons, respectively.

The 4 players highlighted in blue achieved a Top 8 ranking for their non-adjusted average only. This is due to them yielding 4 of the 5 lowest touchdown rates among the twelve players highlighted. Johnny Knox lost 17 positions, due to yielding a touchdown rate of only 9.8%. Torry Holt's touchdown rate was only 7.3%, and his ranking dropped from 5th to 15th. One factor in Torry Holt scoring only 6 touchdowns in 2000 was the scoring ability of a teammate on the St. Louis Rams, running back Marshall Faulk. Faulk scored 8 receiving touchdowns in 2000, and his exceptional season is reflected in running back bubble plot.

The 4 players highlighted in yellow maintained a Top 8 ranking under both averages. Three players, namely Bernard Berrian, Mike Wallace and Ashley Lelie, maintained their ranking by yielding touchdown rates in the top quintile. The exception is the 2014 Season for DeSean Jackson. His touchdown rate of 10.7% is in the 73rd Percentile, and is more comparable to the rates for the players highlighted in blue. However, Jackson only lost 3 positions in his Kaplan-Meier ranking. One potential factor is that the yardage benefit provided by the Kaplan-Meier technique not only depends on the touchdown rate, but also where the touchdowns yardage values are distributed among all the yardages. So analysis of the impact of where the censoring occurs, as opposed to just the volume of censoring, is a desired future research topic:

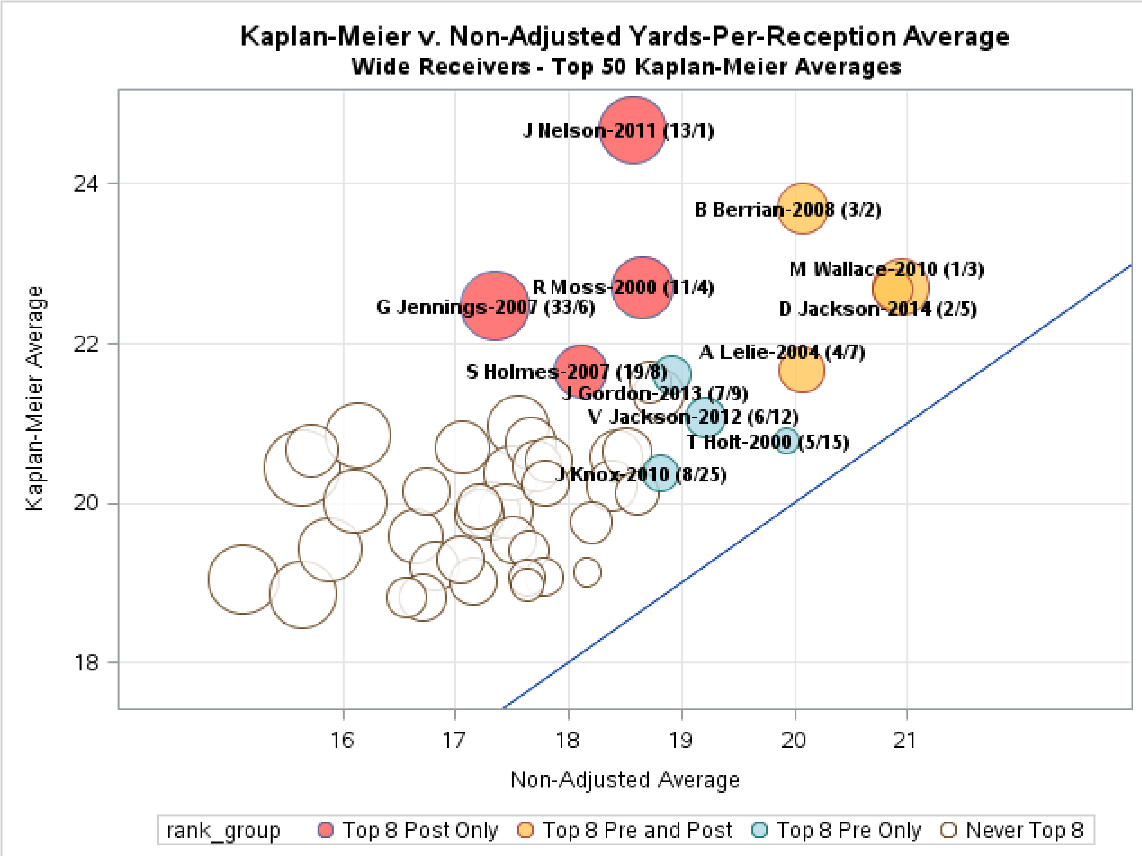

Figure 17. Bubble Plot of Top 50 Wide Receiver Rankings by Average

RUNNING BACKS

Figure 18 contains the running back bubble plot. Ten players who achieved a ranking of eighth or better for either average are highlighted. The 2 players highlighted in red, namely Marshall Faulk and Jamaal Charles, are ranked in the Top 8 for their Kaplan-Meier average, but not for their non-adjusted average. Faulk and Charles achieved touchdown rates of 9.9% and 10.0%, respectively, which rank among the top 4 rates within the 150 running back player/season combinations. Marshall Faulk's average ranking moved from 9th to 1st, and his Kaplan-Meier average of 12.39 Yards is 0.65 Yards more than his nearest competitor, namely the 2011 Season for Arian Foster. Jamaal Charles moved from 16th to 4th.

The 2 players highlighted in blue, namely Joique Bell and Le'Veon Bell, rank in the Top 8 for only the non-adjusted average. Joique Bell caught no touchdown passes in 2013, so his bubble resides on the 45-degree line, and his non-adjusted and Kaplan-Meier averages are both 10.32 Yards. Scoring no touchdowns resulted in an average ranking decrease from 7th to 16th. Le'Veon Bell's ranking only dropped 3 positions. Although is touchdown rate of 3.6% is higher than the 3.1% average rate among running backs, it is still lower than all 8 players in the graph shaded in either red or yellow.

The 6 running backs highlighted in yellow maintained a position in the Top 8 under both averages. Their touchdown rates range from 3.8% - 5.9%, and this range represents the 66th and 90th Percentiles of rates produced by running backs:

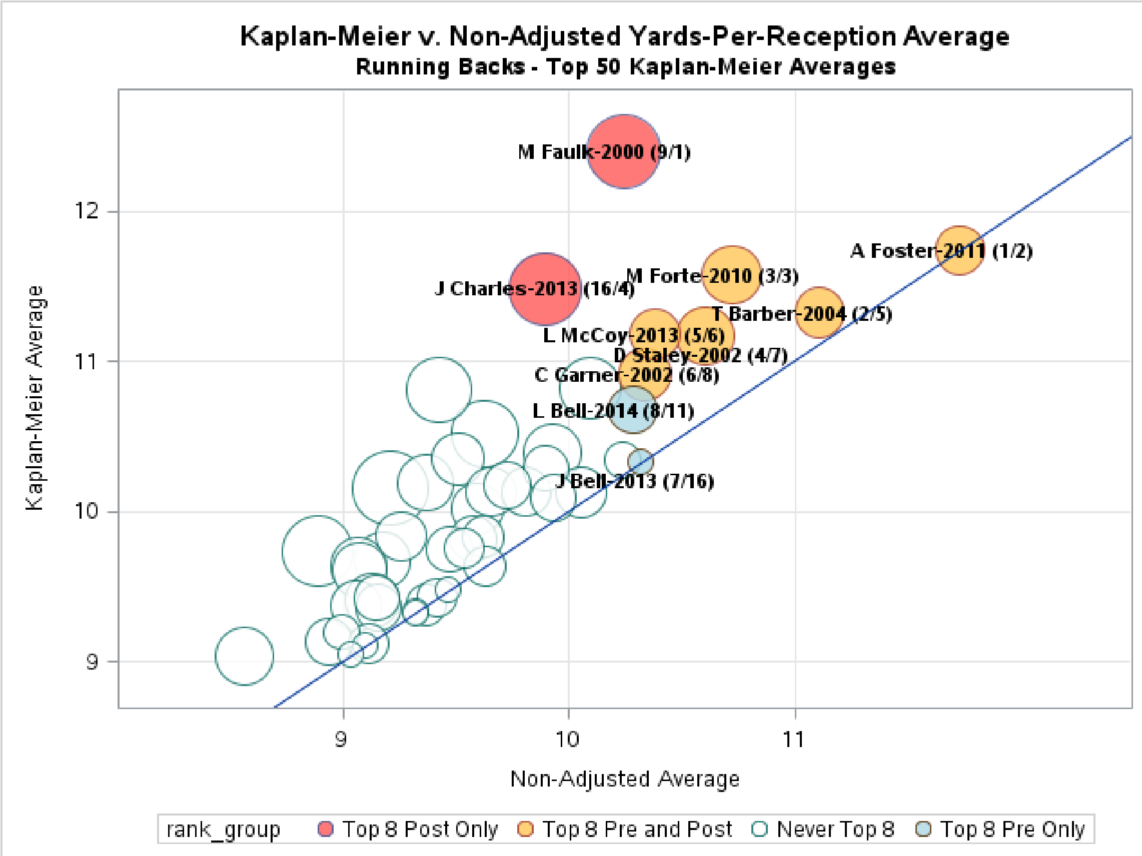

Figure 18. Bubble Plot of Top 50 Running Back Rankings by Average

TIGHT ENDS

Figure 19 contains the tight end bubble plot, and the same player/season combinations are represented in the Top 10 for both averages. As previously mentioned, Vernon Davis, Antonio Gates and Rob Gronkowski each yield 2 of the Top 10 performances. The remaining 4 were produced by Coby Fleener, Alge Crumpler, Jimmy Graham and Jared Cook. The largest ranking change among them belongs to Jared Cook, who dropped 4 positions from 5th to 9th. This is due to achieving a touchdown rate of only 6.1%, while all other tight ends in the Top 10 produced rates above 10%.

Applying the Kaplan-Meier adjustment makes the strong 2013 Season for Vernon Davis even more exceptional. Among the 179 player/season combinatons for tight ends, Davis's 2013 Season is the best for both the non-adjusted average (16.35 yards) and touchdown rate (25.0%). This touchdown rate exceeds the second best rate by 5%, as Antonio Gates and Rob Gronkowski have each achieved 20% rates once. Vernon Davis achieved the best non-adjusted average by only 0.03 Yards, and this second-best average of 16.32 yards also belongs to Davis. He accomplished the latter average during the 2010 Season, and also produced a touchdown rate of 12.5%. Although this touchdown rate is strong enough to maintain the second best Kaplan-Meier average, the disparity between Davis's 2010 and 2013 Seasons is much wider under the Kaplan-Meier average, since the 12.5% rate is only half the 25% rate that Davis yielded in 2013. Davis's Kaplan-Meier average of 20.18 Yards in 2013 is more than 2 yards better than his 2010 value of 18.10 Yards.

The stability of the same player/season combinations being the best for both averages extends through the Top 15, with the exception of 1 player. During the 2008 Season, Zach Miller achieved a non-adjusted average of 13.89 Yards. However, he only scored 1 touchdown in 56 receptions, which resulted in a touchdown rate of 1.8%. Since Miller's lone touchdown of 63 Yards was also his longest reception in 2008, it was coded as a non-censored observation, so that a Kaplan-Meier Average that is lower than his non-adjusted average would not result. So Miller's Kaplan-Meier average is also 13.89 yards, and his average ranking drops from 13th to 28th. Zach Miller's 2008 Season is shown in blue on the bubble plot, and it falls on the 45-degree line.

So while the best player/season combinations for wide receivers and running backs vary with the average used, the best tight end performances are consistent across them. This is due to the shift in recent years towards joker tight ends. As tight ends have lined up more frequently in the slot or as wide receivers, they have produced statistics more similar to wide receivers. For example, Vernon Davis's Kaplan-Meier average in 2013 would rank 28th among the 750 wide receiver seasons. His 25% touchdown rate in 2013 is second only to the 26.5% rate produced by Randy Moss in 2004. This shift is also apparent in the seasons represented among the Top 10 tight ends. With the exception of Alge Crumpler's 2004

Season, all the Top 10 performances occurred after 2008. Three recent seasons, namely 2010, 2011 and 2013, are represented twice in the Top 10. Because of this shift in the use and performance of tight ends, a future research project would be to segment players on attributes other than their primary position:

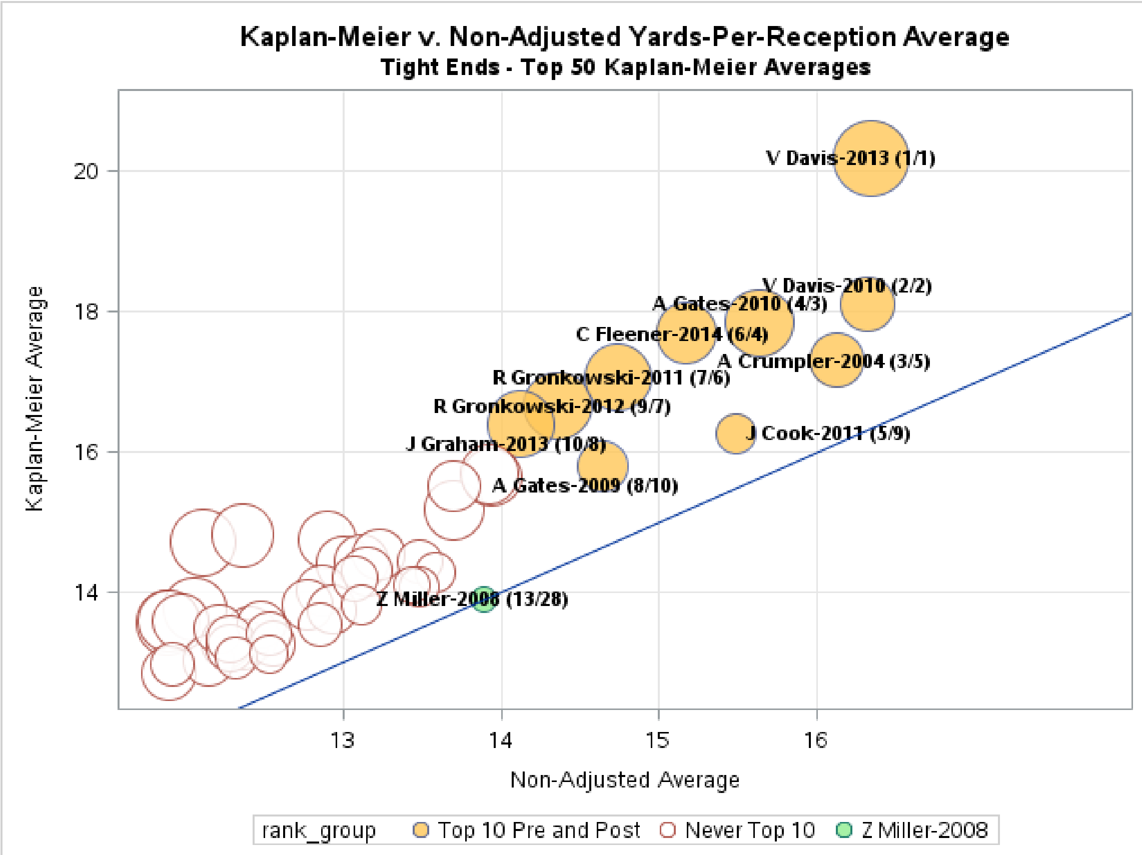

Figure 19. Bubble Plot of Top 50 Tight End Rankings by Average

ADJUSTING FOR 2-POINT SAFETIES

The Kaplan-Meier product-limit estimator is designed to accommodate right-censored observations. An observation is right-censored when the "true" value is above the observed value. This censoring is applicable when a touchdown is scored, if we assume the player would have continued gaining yardage on an analogous play where no touchdown resulted. Consider a play where a team has possession at their own 20-yard line, and a player runs for 8 yards to the 28-yard line. However, if the identical play is executed at their opponent's 5-yard line, then the player can only rush for 5 yards before scoring a touchdown. Since we assume this play would yield 8 yards when the line of scrimmage is 8 or more yards from their opponent's goal line, this 5-yard touchdown is considered a right-censored observation.

A safety results from a negative yardage play where a player is tackled in their own end zone, and 2 points are awarded to the defense. If we assume the player would have lost more yardage on an analogous play where no safety resulted, then we consider the yardage loss as a left-censored observation. A left-censored observation occurs when the "true" value is below the observed value. For example, if a team has possession at their opponent's 45-yard line, and a runner loses 7 yards, then the new line of scrimmage is their own 48-yard line. However, if the identical play is executed from their own 4-yard line, the play will be recorded as a 4-yard loss and a safety. Since we assume the play would have yielded a 7-yard loss when executed 7 or more yards their own goal line, this 4-yard loss for a safety is considered a left-censored observation.

Because the Kaplan-Meier technique only accommodates right-censored observations, and therefore cannot account for touchdowns and safeties simultaneously, an alternate methodology was developed. This approach first obtains separate negative and non-negative averages that are adjusted for safeties and touchdowns, respectively. These 2 averages are then weighted to obtain an overall average, with the number of negative and non-negative plays determining the weights. The steps to this approach are listed below:

1. Count the number of non-negative and negative yardage values.
2. Determine which non-negative yardage values correspond to a touchdown, and designate these values as censored observations. However, the longest gain should never be coded as a censored observation, even if a touchdown

occurred. When the largest observation is censored, the mean estimate from the survival function produced by Kaplan-Meier is biased downward. PROC LIFETEST also gives the following warning:

> **Note**: The mean survival time and its standard error were underestimated because the largest observation was censored and the estimation was restricted to the largest event time.

3. Execute PROC LIFETEST on the non-negative yardages, as shown by the SAS Code below. The time statement designates the yardage variable (yds) and the touchdown censoring variable (td_censoring). The (1) following td_censoring specifies that yardage values associated with a value of 1 for td_censoring are censored observations. The weight statement is used because the Marshawn_Lynch_2014_nonneg dataset is aggregated to each yardage and censoring status combination, and play_count indicates the number of carries that yielded the corresponding yardage/censoring combination:

```
proc lifetest data = Marshawn_Lynch_2014_nonneg;
  time yds * td_censoring(1);
  weight play_count;
  title "Marshawn Lynch – 2014 Season (Non-Negative Carries)";
  ods output means = nonneg_avg;
run;
```

4. Obtain the adjusted non-negative yardage average. It is the mean variable within the nonneg_avg dataset. This dataset was created from the ods command line within the above PROC LIFETEST code.

5. Determine which negative yardage values correspond to a 2-point safety, and designate these values as censored observations. However, the largest loss should never be coded as a censored observation, even if a safety occurred. This is due to the downward bias discussed in Step 2.

6. Convert each negative value into a positive one by taking its absolute value. This step turns the left-censored observations (safeties) into right-censored values, which is necessary when calculating the Kaplan-Meier estimates. This conversion is also a requirement of PROC LIFETEST, since this procedure excludes negative values.

7. Execute PROC LIFETEST on the absolute value of the negative yardages, as shown by the SAS code below. The time statement designates the variable containing the absolute value of negative yardage (abs_yds) and the safety censoring variable (safety_censoring). The (1) following safety_censoring specifies that yardage values associated with a value of 1 for safety_censoring are censored observations. The weight statement is used because the Marshawn_Lynch_2014_neg dataset is aggregated to each yardage and censoring status combination, and play_count indicates the number of carries that yielded the corresponding yardage/censoring combination:

```
proc lifetest data = Marshawn_Lynch_2014_neg;
  time abs_yds * safety_censoring(1);
  weight play_count;
  title "Marshawn Lynch – 2014 Season (Negative Carries)";
  ods output means = neg_avg;
run;
```

8. Obtain the adjusted negative yardage average. It is the mean variable within the neg_avg dataset. This dataset was created from the ods command line within the above PROC LIFETEST code.

9. Calculate an overall weighted average of the adjusted non-negative and negative averages obtained during Steps 4 and 8, respectively. Weights are provided by the counts determined during Step 1. The calculation is performed using the formula below:

$$\frac{[\,-\,(\text{\# of Neg. Plays} * \text{Neg. Yards Avg.}) + (\text{\# of Non-Neg. Plays} * \text{Non-Neg. Yards Avg.})\,]}{[\,\text{\# of Neg. Plays} + \text{\# of Non-Neg. Plays}\,]}$$

This methodology is illustrated with rushing plays from the 2014 Season of Marshawn Lynch, a running back for the Seattle Seahawks. Lynch gained 1,306 yards on 280 carries in 2014, which resulted in a 4.66 yards-per-carry average. He also scored 13 touchdowns, and once was tackled for a safety. 260 of Lynch's carries gained zero or positive yardage, while the remaining 20 carries yielded losses. Note that his longest carry in 2014 was a 79-yard touchdown run, so only 12 touchdowns were designated as censored observations.

By censoring the 1-yard loss that resulted in 2-point safety, Lynch's average negative carry increased in absolute value from 2.15 yards to 2.238. His non-negative average increased from 5.19 to 5.526 by censoring 12 of his 13 touchdowns. By applying the weighted average shown below, Marshawn Lynch's yards-per-carry average increased from 4.66 to 4.97. This is an increase of 0.31 yards, or 6.7% on a relative basis:

$$\begin{aligned}\text{Kaplan-Meier Adjusted Average} &= [\,-\,(20 * 2.238) + (260 * 5.526)\,] / 280 \\ &= 4.97\end{aligned}$$

CONCLUSION

High touchdown rates have a detrimental impact on yards-per-reception averages. This is more apparent for wide receivers than running backs, since on average they yield more than twice the touchdown rate of running backs. For both positions, Locally Weighted Scatterplot Smoothing (LOESS) regression curves show a stronger linear relationship when the Kaplan-Meier adjusted yardage average replaces the non-adjusted average. The best player/season rankings also vary with the average chosen for both positions.

For tight ends, however, the adverse effect is not apparent from LOESS regression curves. This is due to the advent of joker tight ends, and their ability to produce results more similar to wide receivers than traditional tight ends. When 3 exceptional tight ends are removed, the detrimental impact of touchdown rates is apparent for the remaining tight ends. This is based on statistical tests of the slope coefficient from Ordinary Least Squares (OLS) regression. The outlier behavior of top tight end performances is also apparent in the stability of the Top 15 rankings between both averages.

Potential future research projects include the following:

1. Consider alternative survival analysis techniques to Kaplan-Meier, such as the Nelson-Åalen estimator.
2. Consider alternative techniques for eliminating the downward bias in the Kaplan-Meier average when the longest reception produced a touchdown, and would therefore normally be censored.
3. Assess how the distribution of touchdown receptions, as well as the touchdown rate, impacts the average yardage gained from the Kaplan-Meier estimate.
4. Determine an alternative to the primary position for segmenting players.
5. Control for other factors that impact receiving averages, such as offensive and defensive formations, down and distance, and the frequency of receptions inside the red zone.

REFERENCES

Armchair Analysis is the data source. Available at www.armchairanalysis.com

Jones, Lindsay H. "Jimmy Graham ruled tight end for franchise tag by arbitrator" *USA Today* July 2, 2014. Available at http://www.usatoday.com/story/sports/nfl/saints/2014/07/02/jimmy-graham-ruled-tight-end-for-franchise-tag-by-arbitrator/12046925/

Lee, Elisa T. 1992. *Statistical Methods for Survival Data Analysis*. 2nd ed. New York, NY: John Wiley & Sons, Inc.

Pro-Football-Reference was used to check the accuracy of the aggregated results. Available at http://www.pro-football-reference.com

Triplett, Mike. "Jimmy Graham: Tight End or Receiver" *ESPN* January 15, 2014. Available at

http://espn.go.com/blog/new-orleans-saints/post/_/id/4598/jimmy-graham-tight-end-or-receiver

Weishaus, Abraham. 2013. *ASM Study Manual for Exam C/Exam 4*. 16th ed. Greenland, NH: a/s/m Actuarial Study Materials

Wikipedia 2016. "Franchise tag" Accessed February 24, 2016. Available at https://en.wikipedia.org/wiki/Franchise_tag

CONTACT INFORMATION

Your comments and questions are valued and encouraged. Contact the author at:

Keith Curtis
keithc71@yahoo.com

Product-Limit Survival Estimates							
					Weighted		
yds		Survival	Failure	Survival Standard Error	Number Failed	Number Left	Weight
0.0000		1.0000	0	0	0	68.0000	0
0.0000		0.9853	0.0147	0.0232	1.0000	67.0000	1.0000
1.0000		0.9559	0.0441	0.0397	3.0000	65.0000	2.0000
2.0000		0.9412	0.0588	0.0456	4.0000	64.0000	1.0000
2.0000	*	.	.	.	4.0000	63.0000	1.0000
3.0000	*	.	.	.	4.0000	62.0000	1.0000
4.0000	*	.	.	.	4.0000	61.0000	1.0000
5.0000		0.9103	0.0897	0.0563	6.0000	59.0000	2.0000
5.0000	*	.	.	.	6.0000	58.0000	1.0000
6.0000		0.8161	0.1839	0.0785	12.0000	52.0000	6.0000
7.0000		0.7377	0.2623	0.0874	17.0000	47.0000	5.0000
7.0000	*	.	.	.	17.0000	46.0000	1.0000
8.0000		0.7216	0.2784	0.0886	18.0000	45.0000	1.0000
9.0000		0.6415	0.3585	0.0928	23.0000	40.0000	5.0000
11.0000		0.6094	0.3906	0.0929	25.0000	38.0000	2.0000
12.0000		0.5613	0.4387	0.0925	28.0000	35.0000	3.0000
13.0000		0.5292	0.4708	0.0916	30.0000	33.0000	2.0000
14.0000		0.4971	0.5029	0.0904	32.0000	31.0000	2.0000
15.0000		0.4651	0.5349	0.0890	34.0000	29.0000	2.0000
16.0000		0.4009	0.5991	0.0852	38.0000	25.0000	4.0000
16.0000	*	.	.	.	38.0000	24.0000	1.0000
17.0000		0.3675	0.6325	0.0818	40.0000	22.0000	2.0000
17.0000	*	.	.	.	40.0000	21.0000	1.0000
18.0000		0.3325	0.6675	0.0780	42.0000	19.0000	2.0000
21.0000		0.3150	0.6850	0.0759	43.0000	18.0000	1.0000
22.0000		0.2975	0.7025	0.0736	44.0000	17.0000	1.0000
23.0000		0.2800	0.7200	0.0714	45.0000	16.0000	1.0000
25.0000		0.2625	0.7375	0.0690	46.0000	15.0000	1.0000
26.0000		0.2450	0.7550	0.0666	47.0000	14.0000	1.0000
27.0000		0.2275	0.7725	0.0641	48.0000	13.0000	1.0000
31.0000		0.2100	0.7900	0.0615	49.0000	12.0000	1.0000
33.0000		0.1925	0.8075	0.0588	50.0000	11.0000	1.0000
34.0000		0.1750	0.8250	0.0560	51.0000	10.0000	1.0000
36.0000		0.1575	0.8425	0.0531	52.0000	9.0000	1.0000
36.0000	*	.	.	.	52.0000	8.0000	1.0000
37.0000	*	.	.	.	52.0000	7.0000	1.0000
40.0000	*	.	.	.	52.0000	6.0000	1.0000
50.0000	*	.	.	.	52.0000	5.0000	1.0000
55.0000	*	.	.	.	52.0000	4.0000	1.0000
58.0000	*	.	.	.	52.0000	3.0000	1.0000
64.0000		0.1050	0.8950	0.0556	53.0000	2.0000	1.0000
84.0000	*	.	.	.	53.0000	1.0000	1.0000
93.0000		0	1.0000	.	54.0000	0	1.0000

Table 2. PROC LIFETEST Life Table for Jordy Nelson's 2011 Season