NFL Coaching Based on Lots of Data

Andrew C. Gallagher
andrew.c.gallagher@gmail.com

Abstract

A dataset of 430,022 football plays[1] from 2,654 games of the National Football League are analyzed. The main result is shown in Figure 6, and shows optimal fourth down decisions (punt vs. field goal vs. go for first) as a function of field position and yards to go. Coaches should almost never punt in short yardage situations, no matter the field position. Further, even in long yardage situations (fourth and six yards to go), the best decision can be to go for the first down.

1. Introduction

The game of American football is one where teams take turns attempting to advance the football down a 100 yard playing field. A team has a series of four downs to advance the ball from a starting point (the line of scrimmage) by at least 10 total yards. If successful, the team is awarded another first down. The drive continues until a score happens, the possession of the ball changes or the time expires.

On a fourth down, the coach can have a difficult decision: whether to punt the ball, attempt a field goal (3 points), or run a play to attempt to gain the first down in the hope of eventually attaining a touchdown (6, 7, or 8 points).

Each of these options has advantages and disadvantages. WIth a punt, the offensive team loses the ball, but pushes the ball down the field which increases the difficulty for the other team to score. To attempt a field goal, a team must be relatively close to the opponent's goal, but it is only 3 points. In addition, in the event of a miss, the other team gains possession of the ball at the spot of the attempt (about 7 yards behind the line of scrimmage). Finally, if an attempt for the first down is a success, then the drive continues (with a new first down), but if unsuccessful, the ball it turned over to the other team at the spot.

When a coach makes a decision that works out, he is hailed as a genius. When it doesn't, he often get the blame. NFL fans complain on Monday mornings that the coach should have done this, or should have done that.

[1] http://www.armchairanalysis.com/nfl-play-by-play-data.php

While the choices can be hard, there is a lot of historical data (267 games per season, in recent years) to look at. We can look at these decisions, and statistically model their outcomes.

2. Statistical Analysis of Football

Each drive has only a discrete number of outcomes for that drive, including scoring, punting, attempting a field goal, or running out of time.

3. Finding the expected value of a first down

Suppose a team (team A) has the ball with a first down at their own 20 yard line. How much is that field position worth? Here is the approach: Look at all the times from the 2,654 games that a team had a first down at the 20 yard line, and see what happened. Who is the next team to score, and how much is that score worth? It turns out that this situation occurred 7181 times in the dataset. The team with the ball is slightly more likely to score the next points in the game (1628 times the next score was a touchdown by team A, and 1413 times the next score was a touchdown by the other team). Of course, the next score could be a safety or a field goal too. By averaging all of these scores (points by team A are positive and points by team B are negative), we get the *expected point value* of that field position (0.24 points). See Figure 3 for the complete plot.

This same process is used to find the value of punting from a given position.

Also, by looking at all the field goal attempts (9,850) in the dataset, we can show the probability (see Figure 1)of making a field goal from a given distance. Of course, this is averaged over all fields (e.g. domes are usually easier than open stadiums) and all kickers (some are more skilled than others).

Finally, we also need a proxy for the probability that a team would succeed on a fourth down with a given yardage to go. We can estimate this be looking at the performance of teams on third down. Generally, on third down, teams are trying to achieve a first down (to avoid the hard decision), so that success rate should be roughly similar to the success rate that would occur if a team were to go for the fourth

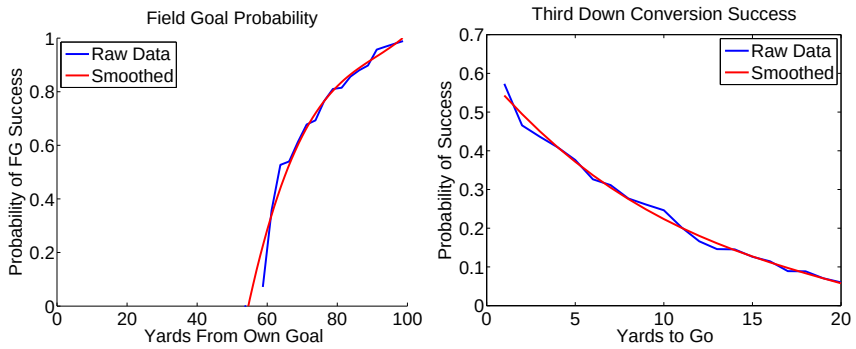

Figure 1. (Left) The probability of a sucessful field goal. (Right) The probability of converting a third down, given yards to go.

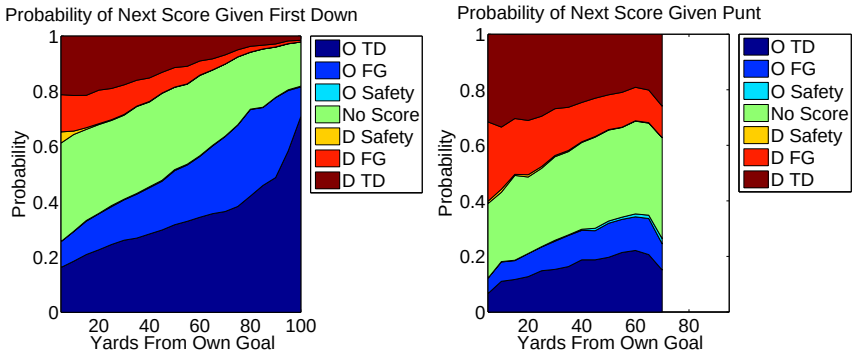

Figure 2. Suppose a given team has a first down at a particular spot on the field. What is the distribution of the next scoring event? The result is shown on the left, where one can make several interesting observations. For example, when the first down in within 5 yards of their own end zone, there is a 4% that the next score (though not necessarily on that drive) will be a saftey. (Right) A similar distribution of the next scoring event is shown for punting from a given location.

down.

At this point, all the pieces of the puzzle are there. In the early part of a game, teams want to score as many points as possible. So, they want to make decisions that maximize the expected points that they will achieve. For each decision, we can compute the expected value of each possible outcome, and we know the probability that each of those outcomes will occur.

For example, suppose it is fourth and four and team A is 38 yards from the opponent's end zone. A field goal could be attempted, but the success rate is 37.1% for a 55 yard field goal. If the field goal misses, then team B will have the ball at their 45 yard line. This field position has an expected value of 1.57 points for team B. So, the action of attempting a field goal 62 yards from team A's goal has an expected outcome of $E(F_{62}) = 0.371 * 3 - (1 - 0.371) * 1.572 = 0.124$.

The other choices can be analysed in a similar way. Punting at the opponent's 38 yard line has an expected value of 0.08 points, an inferior choice to kicking a field goal.

Attempting to get the first down has two possible outcomes. If successful in the 4^{th} and 4 (likelihood of success is about 40.9%) , then we assume that team A has a first down with the ball is at the opponent's 34 yard line (this is the most pessimistic possible ball placement to successfully convert). This field position has an expected payoff of 2.71 points. In the event the attempt fails, we assume team B gains control of the ball at their 38 yard line (an expected value of 1.23 points for team B). Taken together, attempting to get the first has an expected value of $E(G_{62}) = 0.409 * 2.71 - (1 - 0.409) * 1.23 = 0.382$ points. So, the best choice is going for the first down, even with conservative assumptions to ensure that we don't overrecommend going for the first down.

This analysis is repeated for every field position, and for 0 to 10 yards to go on fourth down, to get the optimal coaching decision Figure 4. One amazing thing about the recommended strategy is that for very short yardage-to-go on fourth down (e.g. less than 1 yard), punting is never optimal!

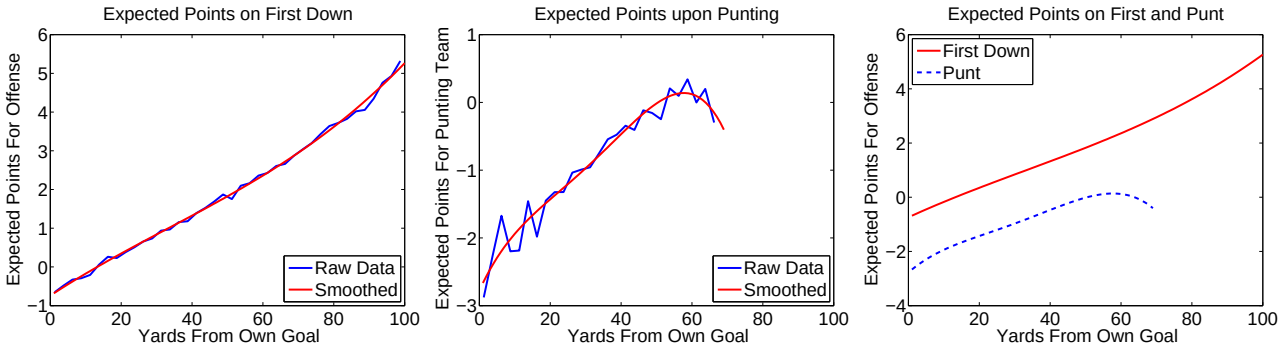

Figure 3. Using the scoring distributions in Figure 2, the expected value of a first down at any location on the field can be found. When a team has a first down at about the 13 yard line, the expected value of a first down is zero, meaning that the offense and defense are equally likely to score in that situation. Also, this plot shows that advancing the ball by 20 yards is worth about 1 point. (Middle) The expected value of punting at a given location. (Right) Punting from a spot on the field is worth about 1.8 (1.766) points less that having a first down at that location. So a 3-and-out essentially costs a team about 1.8 points. Another way of looking at this data is that a team needs to advance the ball by about 37 years to "break even".

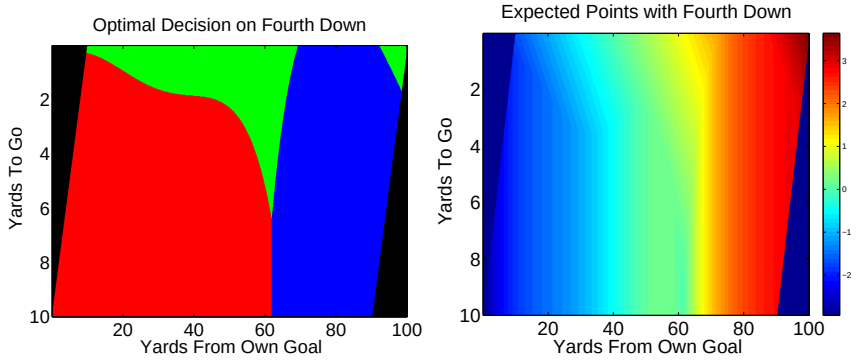

Figure 4. (Left) The optimal coaching decision on fourth down given yards to go and field position. (Red= punt, Green = go for it, Blue = Field goal. (Right) The expected points that the optimal strategy achieves.

4. Punting

If a punt lands into the end zone, a touchback occurs and the other team begins its offensive drive 20 yards from its own goal (80 yards to go). This effectively takes away from the length of the punt, and makes it more difficult to achieve long punts when there is a chance of the ball reaching the end zone.

Aaron D. had the idea to decouple the punting from the end result of the drive by considering the distance that punts travel. Then, the expected value of a punt is:

$$E(p_0) = -\sum_f P(f|p_0)E(f) \qquad (1)$$

where $E(p_0)$ is the expected value of a punt from a particular year line, $P(f|p_0)$ is the probability that a punt from yard line p_0 results in a drive at the opponents f yard line, and $E(f)$ is the expected value of the drive with a first down at the f yard line. The negative sign is there becuase the ball possession changes. Note that 1 assumes that $E(f)$ is independent of the field position where the ball was punted, probably a safe assumption.

Figure 5 shows the mean punt distance given field position, and the resulting expected value of a punt. This is basically a smoother version of Figure 2 and is used to produce a slightly refined version of the decision plot.

5. Comparing with Actual Coaching Decisions

In the beginning of a football game (roughly the first three quarters), coaches should basically have the goal of scoring as many points as possible. As the game nears its ending, sometimes other factors become more important than scoring as many points as possible (e.g. running out the clock).

In the dataset, there are 28,403 fourth down decisions in the first three quarters of the game. For this part, we ignore

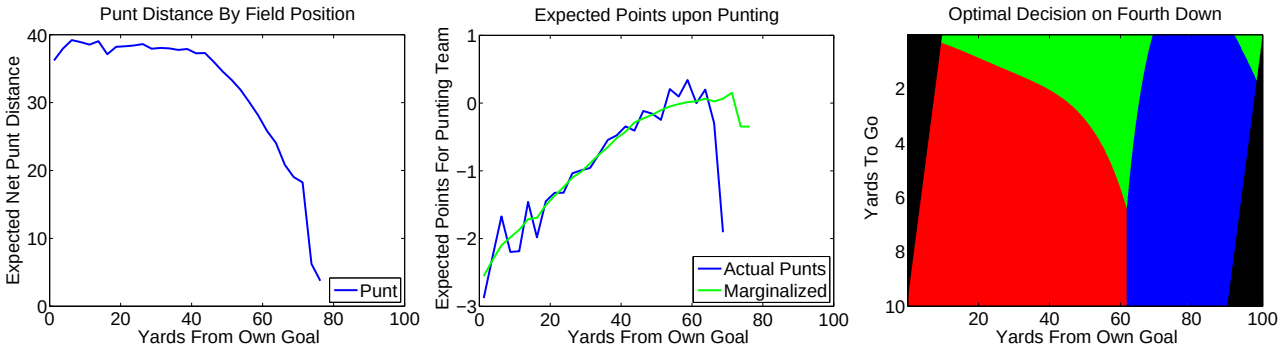

Figure 5. (Left) The mean net yards from punts. (Middle) The expected next score from a punt using the marginalization. (Right) The resulting optimal coaching decisions. (Red = punt, Green = go for first, Blue = field goal attempt.)

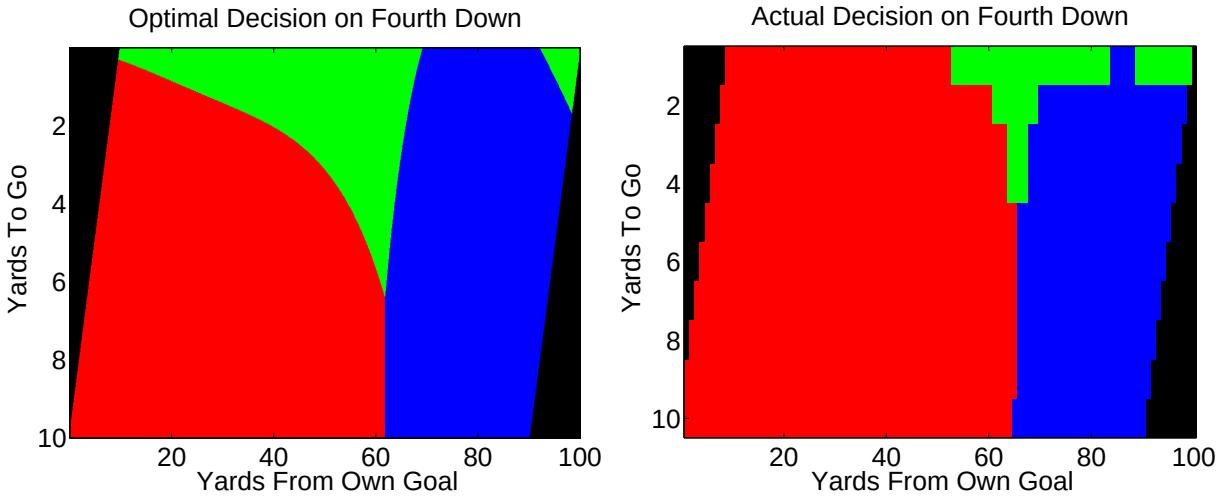

Figure 6. (Left) The optimal decisions. (Right) The most popular decision from actual NFL games given a certain fourth down situation. (Red = punt, Green = go for first, Blue = field goal attempt.) Coaches overuse the punt (instead of going for the first down) on their own side of the field, and perhaps go for the first down too much on the opponents side when they should settle for the field goal in fourth and short situations.

the fourth quarter (and overtime) because coaches might need to consider the other aforementioned factors. So, we can look at what coaches *actually* do when presented with the choice of punting, going for the first down, or attempting the field goal. Figure 6 shows the optimal decisions (left) and what coaches most often do in a given fourth down situation.

There is actually a fairly close similarity in shape, but there are some big differences. The threshold between punting and field goal attempts is almost a match. Coaches are incorrectly reluctant to go for fourth down before crossing the midfield line, when they should be going for the first down whenever the yardage is small. Coaches do get some boldness up to fourth and four on the opponents side of the field. However, coaches should go for the first all the way up to fourth and nearly seven yards to go.

In all, 85.4% of the coaching decisions (in the first three quarters) matched the analysis, but 14.6% are difference (usually by not being agressive enough). The point cost of making these suboptimal decisions is about 0.71 points per game, not a lot on average. However, in a particular game, a better decision strategy could lead to say, an extra touchdown and a win instead of a loss.

6. Some Other Analysis

In this section, there are a few other graphs from the data. Figure 7 shows the outcomes of a particular drive, starting at a particular location on the field with a first down.

Figure 8 shows, for a first drive, the result of that series of plays (i.e. that first, second, third, and forth down). In many cases, that series simply results in another first down further down the field.

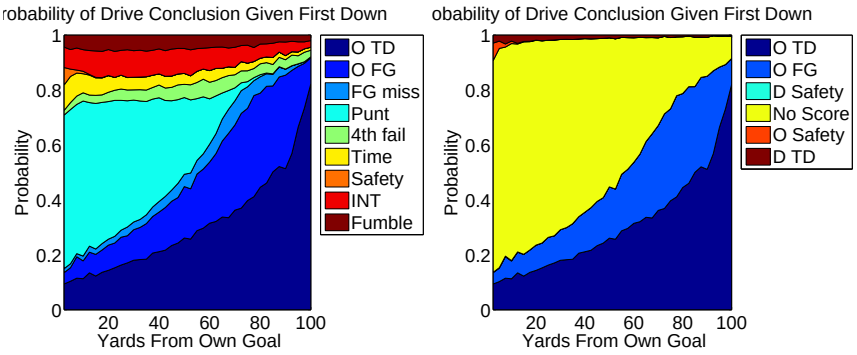

Figure 7. (Left) The result of a particular drive. (Right) The expected scores that the drive achieves.

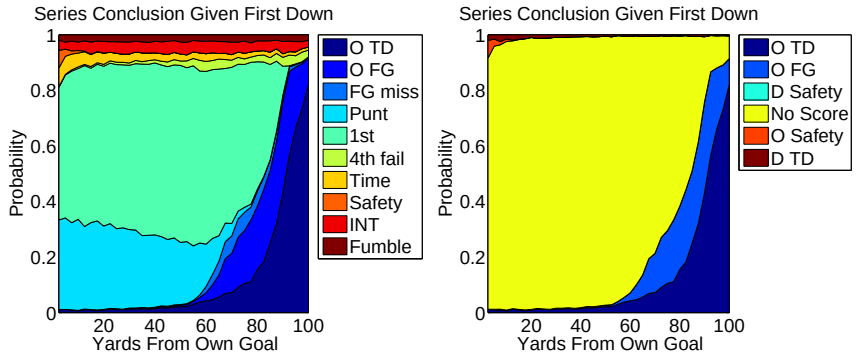

Figure 8. (Left) The result of a particular series. (Right) The expected scores that the series achieves.

7. Discussion

The basic conclusion is that the punt is overused. However, there are some weaknesses to the analysis. First, in computing the expected values of given field positions for punts and first downs, the data is based on outcomes that occurred with "conventional" coaching decisions. For example, a good number of the times that Team A fails to score and punts, Team B then punts back. So I suspect that the estimated expected points from a given field position is underestimated because of the overuse of the punt. We really want our expected values of field positions to be based on the optimal coaching strategy, and then we could use those expected values to find *even more optimal* coaching strategy.

I have a few ideas on ways to compensate for that bias, and I think that there will be even more situations where teams should go for the first down.

Finally, I want to give credit to Aaron Deever and Majid Rabbani for lots of great discussions on the probabilistic aspects of football!

8. Conclusions

NFL Football teams should punt less often!