

Going for Three: Predicting the Likelihood of Field Goal Success with Logistic Regression

Torin K. Clark, Aaron W. Johnson, Alexander J. Stimpson
Massachusetts Institute of Technology,
Cambridge, MA, USA, 12139
Email: torin@mit.edu

Abstract

The field goal is a critical scoring play in the National Football League. Coaches and fans alike are interested in the probability that a field goal attempt will be made or missed. Traditional analyses assume that the attempt distance is the primary factor determining success; however, we believe that other environmental and situational factors cannot be ignored. We constructed a binary logistic regression model based on data from the 2000-2011 NFL seasons to identify factors that have a significant effect on the likelihood of field goal success. Distance and most environmental factors were significant. Altitude and artificial turf improved the likelihood of a make, while cold temperatures, wind, and precipitation reduced it. Contrary to popular belief, not one situational factor (regular season vs. postseason, home vs. away, whether a timeout was called before the attempt, and situational pressure) was significant. We used our comprehensive model to evaluate kicker careers, seasons, and stadiums between 2000-2011. This evaluation is superior to pure make percentage, which is ignorant of the difficulty of a kicker's field goal attempts. By more accurately predicting the outcome of field goal attempts, coaches can make better in-game decisions and fans can gain a greater understanding of kicker ability.

1 Introduction

On January 19, 2002, New England Patriots kicker Adam Vinatieri lined up for a 45-yard field goal attempt that no NFL fan would have considered normal. It was a cold, windy, and snowy divisional playoff game, and the Patriots were down 3 points to the Oakland Raiders with 27 seconds left in regulation. Many fans were understandably nervous, wondering how the weather and pressure would affect Vinatieri. Perhaps they fretted about the distance, wishing that Patriots coach Bill Belichick had pushed farther downfield, giving Vinatieri a shorter kick. But, how much farther would have made a difference? We aim to quantitatively address these questions for *any* field goal attempt by developing a comprehensive model that assesses the probability that the attempt will be successful.

Field goals are a critical element of the NFL game; on average nearly four field goals are attempted per game with wins and losses often being determined on their outcomes. During every game, NFL coaches must decide whether or not their kicker can successfully make a field goal based on the current game conditions. But what game conditions actually have an influence on the success of a field goal? Distance is an obvious one, but are the environmental factors – wind speed, precipitation, field surface, etc. – also important? What about the situational, or psychological, factors – the score differential, the time left in the game, added pressure in the playoffs, etc.? By estimating the effect of the game conditions with respect to different distances, coaches can determine the range of yard lines where they are willing to attempt a kick.

In this paper, we aim to accomplish four goals:

- 1) Determine what factors influence the likelihood that a field goal attempt will be successful.
- 2) Develop a comprehensive logistic regression model to quantify this likelihood and the difficulty of field goal attempts.
- 3) Use this model to more accurately evaluate and compare individual kickers, seasons, and stadiums.
- 4) Apply techniques to classify field goal attempts as either makes or misses based upon relevant factors.

2 Background

There is a great deal of anecdotal evidence as to what factors influence the outcome of a field goal, but thorough, comprehensive statistical analysis is scarce. Many analyses have focused on environmental factors. Brian Burke investigated field goal success in different temperatures and wind speeds [1], and in games at high altitude (Denver) [2]. He found that cold temperatures (21-30 °F) did appear to have an effect, as the overall make percentages for field goals ≥ 25 yards were lower than the make percentage for attempts of the corresponding distance in mild (51-

60 °F) or hot temperatures (81-90 °F). Altitude also appeared to have an effect, with higher make percentages for field goals ≥ 36 yards at Denver, as compared to other stadiums in the NFL. Burke was unable to draw conclusions about the effects of wind speed due to noise in the data. In *The Complete Guide to Kickology, 3rd Edition*, Mike Harman found that the make percentage of kickers in domes was slightly higher than outdoors, and the make percentage of kickers on artificial turf was slightly higher than on grass [3]. Lastly, in their proprietary analytics FootballOutsiders.com adjusts field goal difficulty for altitude (Denver vs. non-Denver), temperature, and if the game was indoors [4].

While most analyses have examined these environmental factors, Football Freakonomics looked at one particular situational and psychological factor – “icing the kicker.” This is when the opposing coach calls a timeout just prior to the field goal in an effort to make the kicker overthink the attempt and miss [5]. This analysis mainly relies on anecdotal evidence and personal opinions, but it does compare the make percentage of kickers who were “iced” and those who were not in the last two minutes of a game. The data shows that icing the kicker actually has the opposite effect from what is desired: kickers who were iced have a higher make percentage than those who were not.

There are a number of serious limitations with each of these studies. Firstly, none of them test for statistical significance. Their analysis may suggest trends, but it is unknown if this is an actual effect or simply the result of natural random variability. Secondly, and more importantly, these studies average their data over multiple potentially important factors. For example, an analysis that only investigates temperature would consider a windy, rainy playoff game to be the same as a calm, sunny regular season game. Building a *comprehensive* model that considers all factors together is the only way to conclude with certainty that an effect actually exists. All of the above studies leave out potentially significant factors, which simplifies the analysis but may miss important effects.

We aim to build a comprehensive binary logistic regression model that includes all statistically significant explanatory variables and estimates the probability that a field goal attempt will be successful. The model will then be able to accurately predict the likelihood the outcome of future kicks.

3 Construction of the Binomial Logistic Regression Model

All 11,896 field goal attempts from the 2000-2011 NFL seasons (excluding preseason games) were parsed out of a complete play-by-play dataset obtained from ArmchairAnalysis.com. For every field goal attempt, the distance was identified along with the values of environmental (temperature, field surface, altitude, precipitation, wind speed, and humidity) and situational (regular season vs. postseason, situational pressure, home vs. away, and “icing”) explanatory variables. Prior to statistical analysis most raw continuous explanatory variables (e.g., temperature in °F) were converted into reasonable categorical variables (e.g., cold: $< 50^\circ\text{F}$; warm: $\geq 50^\circ\text{F}$). This reduces model complexity and improves ease of interpretation. Details and justification for these categorizations can be found in the Appendix. A combination of SPSS and SYSTAT were used for model construction and analysis.

Table 1 displays all tested explanatory variables, their definitions, and p-values. Variables with a p-value < 0.05 were considered significant (meaning that their effect was unlikely to have been created by random chance) and were included in the final model. The p-value for each variable found to have a non-significant effect was determined from a model that included that particular non-significant variable and all significant variables. Variables with significant effects have coefficients, which explain the magnitude and direction of the effect, listed in Table 1. A positive coefficient indicates the presence of this condition improves the likelihood of a make relative to the baseline condition (which is a kick on grass at low altitude with mild temperatures, low wind, and no precipitation), while a negative coefficient reduces the likelihood. These effects are graphically displayed for each individual variable, with respect to distance, in the Appendix.

As seen in Table 1, the distance and most of the environmental factors significantly impact the probability of making a field goal. The model coefficients show that longer kicks, cold temperatures, precipitation, and high winds reduce conversion rates, while kicking on turf and at altitude improve the likelihood of a made field goal. Humidity is the only environmental factor tested that was not significant, although it was highly correlated with precipitation, which was significant. *None* of the situational or psychological factors have a significant impact on the kick outcome, despite what fans, coaches, and the media may think.

These coefficients are used in the following model equation, which calculates the probability that a particular field goal will be successful:

$$P(\text{make} = 1) = \frac{1}{1 + e^{-(\beta_0 + \beta_{\text{dist}}x_1 + \beta_{\text{turf}}x_2 + \beta_{\text{alt}}x_3 + \beta_{\text{cold}}x_4 + \beta_{\text{wind}}x_5 + \beta_{\text{preip}}x_6)}} \quad (1)$$

In the model, X_1 is a continuous variable that represents the distance of the kick in yards and X_2 - X_7 are binary (0 or 1) variables dependent upon whether or not that particular condition applies. If the variable's condition is different than the baseline (wind ≥ 10 mph or turf field), then the respective binary X variable is equal to 1.

Table 1: Comprehensive Logistic Regression Model

	Variables	Coefficient	Standard Error	Wald Statistic	DoF	Significance
General	Constant	$\beta_0 = 5.953$	0.220	25.2	1	$p < 0.0005$
	Distance (yards)	$\beta_{\text{dist}} = -0.106$	0.003	35.5	1	$p < 0.0005$
Environmental¹	Cold temperature ($< 50^\circ\text{F}$)	$\beta_{\text{cold}} = -0.341$	0.061	5.63	1	$p < 0.0005$
	Field surface (artificial turf)	$\beta_{\text{turf}} = 0.299$	0.053	5.62	1	$p < 0.0005$
	Altitude ($\geq 4000\text{ft}$) ²	$\beta_{\text{alt}} = 0.694$	0.157	4.43	1	$p < 0.0005$
	Precipitation (rain, snow, etc.) ³	$\beta_{\text{precip}} = -0.280$	0.099	2.84	1	$p = 0.005$
	Windy ($\geq 10\text{mph}$)	$\beta_{\text{wind}} = -0.140$	0.055	2.55	1	$p = 0.011$
	Humid ($\geq 60\%$)			0.20	1	$p = 0.844$
Situational/ Psychological	Postseason			1.29	1	$p = 0.196$
	High situational pressure ⁴			0.61	1	$p = 0.539$
	Away game ⁴			0.67	1	$p = 0.501$
	"Icing the kicker" (TO before) ⁶			1.56	1	$p = 0.118$

- ¹ All environmental conditions are at kickoff and not specific to the time of each individual kick.
- ² The only games played at altitudes greater than 4,000 ft were those in Denver or Mexico City (Oct. 2, 2005).
- ³ A "chance of rain" is categorized as no precipitation. The model was tested with the alternative categorization, and this had a negligible impact on the value and significance of the coefficient.
- ⁴ Several alternate methods of categorizing pressure were tested, none of which were significant. See Appendix.
- ⁵ Neutral site games were categorized based upon official distinctions. Treating neutral site games as a third category was still not significant.
- ⁶ A timeout called by either head coach was considered "icing the kicker." Categorizing "icing" as either no timeout, timeout by opposing coach, or timeout by own coach was still not significant.

While the individual effect of each variable may not appear large on its own, they can quickly add up – just as they did for the Patriots and Adam Vinatieri in 2002. Our model indicates that the pressure of the field goal (being a critical make late in a playoff game) has no significant influence on the outcome. However, the cold, the snow, and the wind did. Figure 1 shows our model's prediction for these conditions as a function of distance. The model estimates that Vinatieri's attempt is converted only 53% of the time. But what if the game took place in Oakland? The raucous Raider Nation wouldn't have a significant influence, but what about the weather that day – temperatures just above 50 degrees, no rain, and low winds? In that case, the model estimates that the likelihood of Vinatieri's field goal would increase by 0.18! The effect of environmental variables is certainly not trivial.

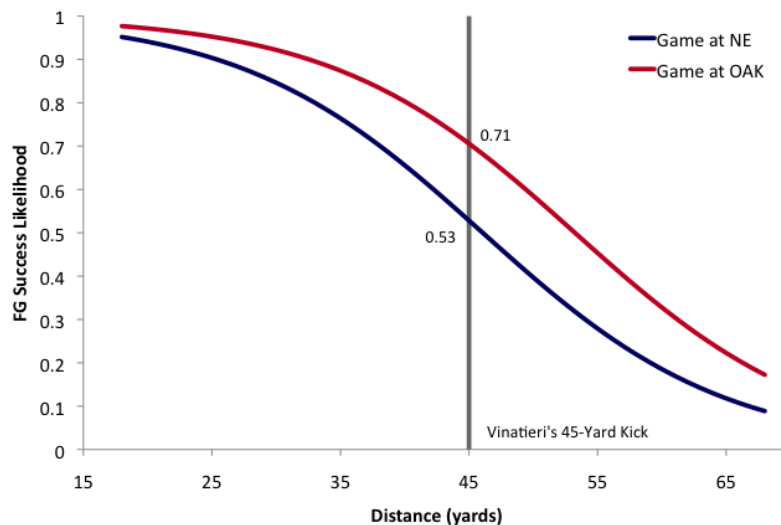


Figure 1: Likelihood of Success for Vinatieri's Field Goal in the 2001-2002 New England-Oakland Divisional Playoff Game

4. Considering Field Goal Difficulty when Ranking Kickers, Seasons, and Stadiums

One of the main metrics used to evaluate and compare kickers is their overall field goal make percentage. This statistic is based on the assumption that better kickers make a higher percentage of their field goal attempts. What this statistic misses, however, is that not all field goal attempts are created equal. Some analyses will account for make percentage at different distances, but as our model shows, not even all field goal attempts of the same distance are created equal. The likelihood that a field goal attempt will be successful, as indicated by our model, is related to the attempt difficulty. A higher likelihood of success means a lower difficulty, and vice versa. By evaluating kickers accounting for the difficulty of their attempts, we can get an unbiased perspective on the kicker's skill.

This kicker evaluation was performed using a metric called "added points". The metric subtracts the model's predicted likelihood from the actual outcome of a field goal (1 for a make, 0 for a miss), and multiplies this difference by the 3 points awarded for a make. Therefore, the units are standard points. As an example, a successful field goal that had a 77% probability of being made would provide $3*(1-0.77) = +0.69$ added points. If the same kick was missed, the added points would be $3*(0-0.77) = -2.31$. Kickers with positive added points are providing additional points beyond what would be expected of an average kicker given the same opportunities. Similarly, kickers with negative added points are taking points away from their team, compared to an average kicker in the same situation.

We have used our difficulty-sensitive added points metric to perform a number of kicker and stadium evaluations and rankings over the 2000-2011 seasons.

Who had the best and worst careers?

The added points per attempt were calculated for each kicker with more than 50 attempts between 2000-2011. Table 2 shows the top five, led by current Tennessee Titans kicker Rob Bironas with 0.262 added points per attempt, and the bottom five, with Wade Richey losing the most added points per attempt (-0.467). Table A.2 in the Appendix gives a complete listing of all kickers. For comparison, the top five best- and worst-ranked kickers based upon raw make percentage (which is ignorant to the difficulty of the kicks attempted) are given in Table 3, with the full data given in Table A.3 in the Appendix.

Evaluating kickers on pure make percentage underrates kickers who attempt proportionately more difficult kicks, and overrates kickers who make proportionately easier kicks. Table 4 shows the ranking differential between kickers' rankings in Tables 2 and 3. Kickers with a positive number are underrated based on make percentage, and kickers with a negative number are overrated. The complete ranking for all kickers can be found in the Appendix, Table A.4.

Table 2: Kicker Careers Ranked by Added Points

Rank	Kicker	Added Points per Attempt
1	Rob Bironas	0.262
2	Robbie Gould	0.204
3	Connor Barth	0.195
4	John Kasay	0.160
5	Dan Carpenter	0.134
...
51	Steve Christie	-0.200
52	Nick Novak	-0.201
53	Tim Seder	-0.342
54	Jose Cortez	-0.405
55	Wade Richey	-0.467

Table 3: Kicker Careers Ranked by Make Percentage

Rank	Kicker	Make %	Number of Kicks
1	Garrett Hartley	87.7	57
2	Matt Stover	86.8	335
3	Robbie Gould	86.2	224
4	Rob Bironas	86.1	223
5	Shayne Graham	85.4	254
...
51	Dave Rayner	72.2	90
52	Nick Novak	71.9	64
53	Tim Seder	71.0	62
54	Jose Cortez	70.7	75
55	Wade Richey	66.1	56

Table 4: Kickers Underrated and Overrated by Make Percentage

Rank	Kicker	Degree Underrated / Overrated by Make %
1	Sebastian Janikowski	25
2	Dan Carpenter	19
3	Ryan Succop	16
4	Josh Scobee	15
5	Mason Crosby	14
...
51	Lawrence Tynes	-11
52	Shayne Graham	-12
53	Gary Anderson	-13
54	Mike Vanderjagt	-17
55	Stephen Gostkowski	-20

Who had the best and worst seasons?

To evaluate which kickers had the best and worst single seasons, we looked at total added points in a season. This tells us the direct effect of a kicker on his team – how many points he gained/lost compared to an average kicker. Table 5 shows the five best and worst seasons (in which the kicker had at least 20 attempts). Players selected to the Pro Bowl that year are marked in green.

It is interesting to note that three players appear on the list twice – Sebastian Janikowski, with two of the best seasons, Kris Brown, with two of the worst (8 seasons apart), and Neil Rackers, who was able to make the jump from fourth-worst season to second-best in only five seasons!

Are kickers getting better over time?

There is an interesting pattern in Table 5 – four of the five worst seasons are from 2001 or 2003, and three of the five best seasons are from 2009 or 2011. But is this truly a trend? Have kickers, as a whole, really been improving over time? We find that while the average kick difficulty and the number of attempts (Figures A.6 and A.7, Appendix) are not significantly different over the 2000-2011 seasons, kickers are indeed getting better! Figure 2 displays the average added points per attempt by season. A linear regression shows added points per attempt has been significantly improving over the time period (+0.017 added points per attempt/season, $t(11)=5.34$, $p<0.0005$).

So kickers as a group are improving. But why? One possibility is that kickers have been staying in the NFL throughout our dataset, gaining more experience. To test this hypothesis, we re-ran our logistic regression with the addition of experience (the kicker’s number of seasons at the time of a field goal attempt) as an explanatory variable. To avoid the confounding variable of kickers who do not make it in the NFL, and therefore only attempt field goals when experience, we removed any kick made by a kicker with < 50 attempts in our dataset. We found that experience significantly improved the prediction of field goal success likelihood (coefficient=+0.017/season of experience, Wald Statistic=3.21, $p=0.001$). Furthermore, we found that the number of years of experience for the average field goal attempt (Figure A.8, Appendix) significantly increased over each season from 2000-2011 (linear regression, $t(11)=6.24$, $p<0.0005$). These pieces of evidence taken together do support the notion that more kicks are being taken by more experienced kickers. However, this may not be the only factor. We also tested our logistic regression with season year (2000, 2001, etc.) as an explanatory variable. Even with the experience variable included, season year also had a significant effect (coefficient=+0.037, Wald Statistic=4.86, $p<0.0005$). Kicker experience is not the only factor; something else is at work. It may be that better kickers are entering the league, or that training techniques have been improving and making kickers better.

In which stadiums are field goal attempts easiest and most difficult?

Beyond evaluating individual kickers, our model can be used to determine which stadiums, on average, have the most difficult kicking conditions. We used our model to calculate the predicted difficulty of a 45-yard attempt (the distance where the effect of environmental factors is approximately greatest) in every game between 2000-2011. This encompassed 3,410 games played across 51 stadiums. For the purposes of this analysis, when a stadium switched playing surfaces from grass to artificial turf it was treated as two different stadiums (e.g. Paul Brown Stadium with grass and Paul Brown Stadium with turf). We found the average kick difficulty to be significantly affected by different stadiums (one-way ANOVA, $F(50,3089)=131.8$, $p<0.0005$). Table 6 shows the *current* stadiums in which field goal attempts are most and least difficult, on average. Negative numbers indicate a lower likelihood of success,

Table 5: Single Seasons Ranked by Added Points

	Kicker	Season	Added Points
1	Sebastian Janikowski	2009	19.4
2	Neil Rackers	2005	18.7
3	Sebastian Janikowski	2011	18.4
4	Rob Bironas	2011	17.8
5	Mike Vanderjagt	2003	16.6
...
341	Kris Brown	2001	-15.1
342	Neil Rackers	2001	-15.2
343	Kris Brown	2009	-15.7
344	Wade Richey	2001	-16.4
345	Seth Marler	2003	-20.0




Figure 2: Improvement in Kickers’ Added Points per Attempt

and therefore a higher difficulty. Positive numbers indicate a higher likelihood of success, and therefore a lower difficulty. The complete listing (Table A.5, Appendix) includes all stadiums used for at least one season between 2000-2011.

5 Discussion

The results of our model contradict some of the convention wisdom about what factors influence the success of a field goal attempt. We found that psychological factors such as high situational pressure or "icing" the kicker have no statistically significant effects on the likelihood that a field goal attempt will be successful. This is completely contrary to the popular belief held by most fans and commentators. The

kicker position appears to be one with moments of incredible stress, where winning or losing the game may essentially depend on your performance alone for one single play. However, our analysis suggests that in the NFL, these situational factors play far less of a role in the success of field goals than environmental factors. We hypothesize this may be because NFL kickers have a high level of experience in dealing with high-stress situations. Whether non-professional leagues such as the NCAA or high school kicking would similarly lack psychological effects is yet to be analyzed.

An additional goal of field goal models is classification; that is, the prediction of whether any particular kick will be a success or a failure. Using the logistic regression model, we might use the rule of thumb that any kick with a probability of success of greater than 50% should be classified as a make, with the rest classified as a miss. Using this sort of strategy, we can try to maximize the number of correct classifications by comparing the predicted classifications to the actual outcomes. Due to the high proportion of made kicks in our data set, using a cutoff of 50% probability does not maximize the number of correct classifications (the optimal likelihood threshold for this dataset is around 37%). Alternatively, we can apply methods that are intended to directly provide a classification, rather than a probability. One commonly used method for this type of classification are Support Vector Machines (SVMs). These models construct a maximum-margin hyperplane to classify new data. Applied to the field goal dataset, both of these strategies failed to produce correct classification rates significantly higher than naive classification approaches (such as predicting all kicks as makes). The failure of these methods to generate a useful classifier likely lies within the data. First of all, there is inherent randomness contained within human performance data such as field goal kicking. If a kicker attempted 100 field goal kicks under the same conditions, it is unlikely that all would be makes or all would be misses. By assigning a set of conditions to either a classification of make or miss, we are inherently creating an error based on the fact that the kicker will not always succeed on the same kick under the same conditions every time. Secondly, the availability of make and miss data is not uniform; coaches fundamentally tend to attempt kicks that are likely to succeed. By having few datapoints on kicks that have a high chance of failure, we are forced into forming classifiers that are heavily weighted towards predicting success.

While there are numerous applications of the presented model and analysis, there are also some limitations. This study is observational, and not experimental. Therefore, the sampling is not random, and there is a high degree of correlation between some of the variables used in the model (particularly environmental variables). For example, altitude and precipitation are correlated, as it often snows in Denver during the NFL season. Because of these correlations, it is impossible to unequivocally separate the effect of altitude and precipitation. However, knowing the absolute, individual effect of one particular variable is not necessary to make confident predictions of field goal success likelihood or difficulty. And as such, our model still has great value. Additional limitations of the model actually lie in the limitations of our data set. All of the environmental variables are only provided for the beginning of the game, and not at the time of each individual field goal attempt. Turning temperature and wind speed into categorical variables helps to deal with this limitation, and misclassifications will only occur if these variables cross a categorical "boundary" between the beginning of the game and the time of the field goal (e.g., the temperature at the beginning of the game is greater than 50 °F, but it drops below this threshold by the end). Secondly, wind speed was included in the model, but wind direction was not. The wind direction, whether at one's back, in one's face, or a crosswind, is an important factor that should be considered. With good data on wind direction and the end zone towards which the offense is driving, wind direction could be included in future iterations of our model. Lastly, the

Table 6: Average Kick Difficulty by Stadium

Rank	Stadium	Team	Ave. Likelihood Relative to Mean
2	Lambeau Field	GB	-0.063
4	Heinz Field	PIT	-0.049
5	Cleveland Browns Stadium	CLE	-0.049
6	Soldier Field	CHI	-0.047
7	Arrowhead Stadium	KC	-0.042
...
39	Georgia Dome	ATL	0.0461
41	Mall of America Field at H.H.H. Metrodome	MIN	0.0461
42	Mercedes-Benz Superdome	NO	0.0461
44	Ford Field	DET	0.0461
45	Sports Authority Field at Mile High	DEN	0.0957

data does not contain spatial information about field goal attempts – where on the field the ball was placed and how the field goal may have been missed (wide left, wide right, short, etc.) In the future, we plan to seek out other databases and sources of information to make these improvements to our model.

6 Conclusions

Field goals, as relatively isolated plays, represent a major opportunity to develop models that assess the contributing factors to their success. Through the development of a binomial logistic regression model, we found that situational (psychological) factors have no statistical effect on the outcome of field goal attempts, despite what fans, coaches, and the media believe. On the other hand, environmental factors – temperature, precipitation, wind, field surface, and altitude – *do* have major impacts upon the difficulty of a field goal attempt. This model allows for a more comprehensive analysis of kickers, in that we can judge ability based on the difficulty of the field goals attempted. Our added points metric helped identify the best and worst kickers and single seasons from 2000-2011, and shown us that kickers, as a group, are improving. We have also used our model to identify the stadiums that are more and less difficult than average for kickers. Our analyses have helped to better quantify the likelihood of field goal success and the difficulty of each attempt, putting into numbers all the thoughts and worries of Patriots fans on January 19, 2002, as they watched Adam Vinatieri jog out onto the snowy field.

Acknowledgements

We would like to acknowledge ArmchairAnalysis.com for providing the detailed dataset for this analysis. We would also like to thank our reviewers – Amanda Jason, Gretchen Johnson, and Dustin Kendrick – for their comments.

References

- [1] Burke, Brian. “Temperature and Field Goals.” AdvancedNFLStats.com. Advanced NFL Stats, 17 Jan 2012. Web. 16 May 2012.
- [2] Burke, Brian. “Altitude and Field Goals.” AdvancedNFLStats.com. Advanced NFL Stats, 9 Jan 2013. Web. 12 Jan 2012.
- [3] Herman, Mike. *The Complete Guide to Kickology*. 3rd edition. Footballguys, 2009. Web. 16 May 2012.
- [4] Dubner, Stephen J. “Why Even Ice a Kicker?” Football Freakonomics Episode 2. NFL, 13 Nov 2011. Web. 16 May 2012.
- [5] Schatz, Aaron. “Methods to Our Madness: Special Teams.” FootballOutsiders.com. Web. 16 May 2012.

Appendices

Categorization Definitions and Justification

Continuous and complex categorical environmental variables were reduced to simpler categorical variables for analysis. There are several advantages of this. First, the interpretation of the model is often simpler. Second, the environmental data are all at the beginning of the game and not at the specific time of the kick. It is likely that some variation in conditions may occur, especially for kicks near the end of the game. Gross categorization reduces the impact of this uncertainty. For example, while the temperature may fluctuate during the course of a game, for the majority of kicks if the temperature was “cold” or “warm” at kickoff, this will remain in this category for each kick throughout the game. Finally, having many levels of a complex categorical variable will reduce the statistical power.

During the process of categorization, we aimed to reduce the model complexity with maintain as much critical information as possible. For example, when determining the number and division of categories for temperature, the temperature was first broken into 10 bins to gain a big-picture understanding. The make probability nearly monotonically increased with increasing temperature, however it did so in a non-linear fashion. Temperatures below about 50 °F were all relatively similar, while above 50 °F there was a large jump in make probability. Thus the temperature was binned into two categories: cold (<50 °F) and warm (≥50 °F). Similar techniques were used to categorize the other variables.

Situational pressure was categorized by combining score differential and time remaining in the game in meaningful ways. There are obviously many ways in which this could be done. Two of the categorizations we tested are shown in Table . There was no strong evidence that situational pressure has any effect using any of the categorizations we tested. Thus it is quite unlikely that a different method for categorization would yield a significant effect.

Graphical Effects of Significant Explanatory Variables




Figure A.1: Effect of Field Surface on Likelihood of Field Goal Success




Figure A.2: Effect of Altitude on Likelihood of Field Goal Success

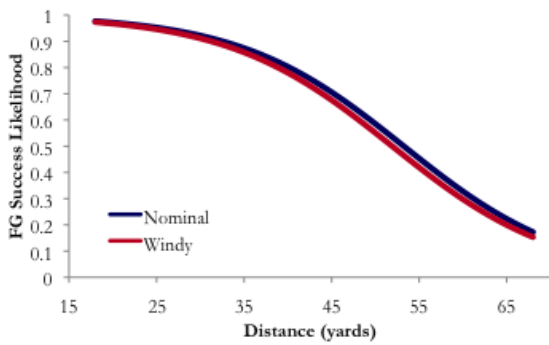


Figure A.3: Effect of Wind on Likelihood of Field Goal Success




Figure A.4: Effect of Precipitation on Likelihood of Field Goal Success

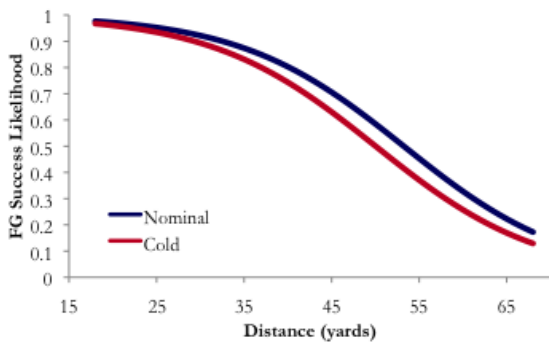


Figure A.5: Effect of Temperature on Likelihood of Field Goal Success

Categorization of Situational Pressure

Table A.1: Categorization of Situational Pressure

Time Remaining	Score Differential	Effect of Kick	6-category Pressure	20category Pressure
4 th quarter	$> 21 $	No effect	No	Low
4 th quarter, < 2 minutes	$> 8 $	No effect	No	Low
4 th quarter, < 2 minutes	< -7	No effect	No	Low
1 st -3 rd quarters	Any	Regular effect	Low	Low
4 th quarter, > 2 minutes	$< 21 $	"Close" 4 th quarter	Medium	Low
4 th quarter, < 2 minutes	5, 6, 7, 8	Helps seal game	Medium	Low
4 th quarter, < 2 minutes	-4, -5, -6	Come within 3	Medium	Low
4 th quarter, < 2 minutes	4	Opponent needs TD	Medium-high	High
4 th quarter, < 2 minutes	3	Opponent needs TD	Medium-high	High
4 th quarter, < 2 minutes	1, 2	Opponent needs TD	High	High
4 th quarter, < 2 minutes	0	Win. If miss, OT	Higher	High
Overtime	0 (Any)	Win. If miss, more OT	Higher	High
4 th quarter, < 2 minutes	-3	OT. If miss, lose	Highest	High
4 th quarter, < 2 minutes	-1, -2	Win. If miss, lose	Highest	High

Kicker Rankings Over Career

Table A.2: Kicker Careers Ranked by Added Points

Rank	Kicker	Added Points per Attempt	Rank	Kicker	Added Points per Attempt
1	Rob Bironas	0.262	28	Stephen Gostkowski	-0.0040
2	Robbie Gould	0.204	29	John Carney	-0.0175
3	Connor Barth	0.195	30	Steven Hauschka	-0.0263
4	John Kasay	0.160	31	Jason Elam	-0.0268
5	Dan Carpenter	0.134	32	Olindo Mare	-0.0269
6	Joe Nedney	0.130	33	Rian Lindell	-0.0271
7	Sebastian Janikowski	0.117	34	Jay Feely	-0.0280
8	Garrett Hartley	0.108	35	Mike Hollis	-0.0559
9	Matt Stover	0.101	36	Mike Nugent	-0.0568
10	Ryan Succop	0.0911	37	Paul Edinger	-0.0587
11	Phil Dawson	0.0889	38	Shaun Suisham	-0.0635
12	David Akers	0.0865	39	Matt Prater	-0.0695
13	Jason Hanson	0.0816	40	Nick Folk	-0.0721
14	Ryan Longwell	0.0740	41	Brett Conway	-0.0742
15	Jeff Wilkins	0.0737	42	Lawrence Tynes	-0.0823
16	Adam Vinatieri	0.0732	43	Kris Brown	-0.0838
17	Shayne Graham	0.0705	44	Doug Brien	-0.0846
18	Nate Kaeding	0.0692	45	Martin Gramatica	-0.102
19	Mason Crosby	0.0594	46	John Hall	-0.115
20	Jeff Reed	0.0558	47	Todd Peterson	-0.138
21	Josh Brown	0.0536	48	Billy Cundiff	-0.141
22	Gary Anderson	0.0490	49	Graham Gano	-0.147
23	Mike Vanderjagt	0.0481	50	Dave Rayner	-0.170
24	Josh Scobee	0.0377	51	Steve Christie	-0.200
25	Matt Bryant	0.0273	52	Nick Novak	-0.201
26	Morten Andersen	0.0078	53	Tim Seder	-0.342
27	Neil Rackers	0.0021	54	Jose Cortez	-0.405
			55	Wade Richey	-0.467



Table A.3: Kicker Careers Ranked by Make Percentage

Rank	Kicker	Career Make %	# of Kicks	Rank	Kicker	Career Make %	# of Kicks
1	Garrett Hartley	0.8772	57	28	Olindo Mare	0.8110	328
2	Matt Stover	0.8687	335	29	Rian Lindell	0.8086	324
3	Robbie Gould	0.8616	224	30	Mike Nugent	0.8086	162
4	Rob Bironas	0.8610	223	31	Lawrence Tynes	0.8019	212
5	Shayne Graham	0.8543	254	32	Sebastian Janikowski	0.7995	384
6	Mike Vanderjagt	0.8510	208	33	Mason Crosby	0.7989	174
7	Joe Nedney	0.8480	250	34	Brett Conway	0.7966	59
8	Stephen Gostkowski	0.8470	183	35	Neil Rackers	0.7959	343
9	Gary Anderson	0.8443	122	36	Matt Prater	0.7934	121
10	Adam Vinatieri	0.8431	376	37	Doug Brien	0.7895	114
11	Jeff Wilkins	0.8419	253	38	Steven Hauschka	0.7885	52
12	Nate Kaeding	0.8419	215	39	Josh Scobee	0.7844	218
13	Connor Barth	0.8391	87	40	Shaun Suisham	0.7824	170
14	John Kasay	0.8385	322	41	Nick Folk	0.7785	149
15	Phil Dawson	0.8385	322	42	Todd Peterson	0.7754	138
16	Matt Bryant	0.8382	241	43	Billy Cundiff	0.7742	186
17	David Akers	0.8315	445	44	Mike Hollis	0.7701	87
18	Jason Hanson	0.8302	324	45	Kris Brown	0.7622	307
19	John Carney	0.8299	288	46	Martin Gramatica	0.7596	183
20	Morten Andersen	0.8297	182	47	Paul Edinger	0.7514	181
21	Jeff Reed	0.8267	277	48	Steve Christie	0.7500	120
22	Jason Elam	0.8264	311	49	John Hall	0.7485	171
23	Ryan Longwell	0.8249	354	50	Graham Gano	0.7375	80
24	Dan Carpenter	0.8217	129	51	Dave Rayner	0.7222	90
25	Jay Feely	0.8143	350	52	Nick Novak	0.7188	64
26	Ryan Succop	0.8118	85	53	Tim Seder	0.7097	62
27	Josh Brown	0.8110	291	54	Jose Cortez	0.7067	75
				55	Wade Richey	0.6607	56



Table A.4: Kickers Underrated and Overrated by Make Percentage

Rank	Kicker	Degree Underrated / Overrated by Make %	Rank	Kicker	Degree Underrated / Overrated by Make %
1	Sebastian Janikowski	25	28	Wade Richey	0
2	Dan Carpenter	19	29	Tim Seder	0
3	Ryan Succop	16	30	Nick Novak	0
4	Josh Scobee	15	31	Jose Cortez	0
5	Mason Crosby	14	32	Steve Christie	-3
6	Paul Edinger	10	33	Matt Prater	-3
7	John Kasay	10	34	Rian Lindell	-4
8	Connor Barth	10	35	Olindo Mare	-4
9	Ryan Longwell	9	36	Jeff Wilkins	-4
10	Mike Hollis	9	37	Todd Peterson	-5
11	Steven Hauschka	8	38	Billy Cundiff	-5
12	Neil Rackers	8	39	Nate Kaeding	-6
13	Josh Brown	6	40	Mike Nugent	-6
14	Jason Hanson	5	41	Morten Andersen	-6
15	David Akers	5	42	Adam Vinatieri	-6
16	Phil Dawson	4	43	Matt Stover	-7
17	Rob Bironas	3	44	Garrett Hartley	-7
18	John Hall	3	45	Doug Brien	-7
19	Shaun Suisham	2	46	Brett Conway	-7
20	Kris Brown	2	47	Matt Bryant	-9
21	Robbie Gould	1	48	Jay Feely	-9
22	Nick Folk	1	49	Jason Elam	-9
23	Martin Gramatica	1	50	John Carney	-10
24	Jeff Reed	1	51	Lawrence Tynes	-11
25	Joe Nedney	1	52	Shayne Graham	-12
26	Graham Gano	1	53	Gary Anderson	-13
27	Dave Rayner	1	54	Mike Vanderjagt	-17
			55	Stephen Gostkowski	-20

Kicker Improvement Over Time

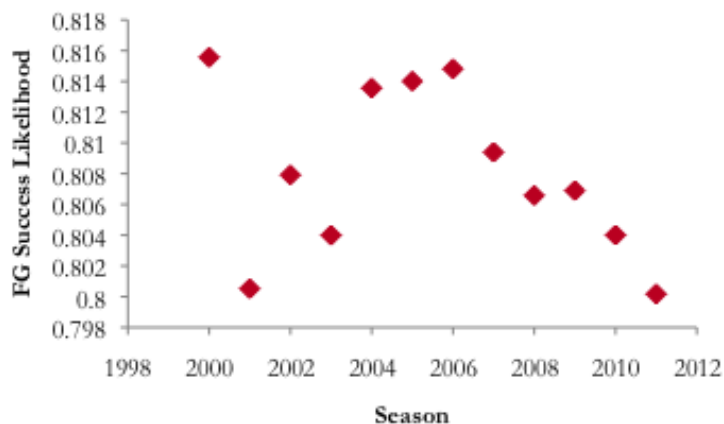


Figure A.6: Likelihood of Field Goal Success (Attempt Difficulty) from 2000-2011

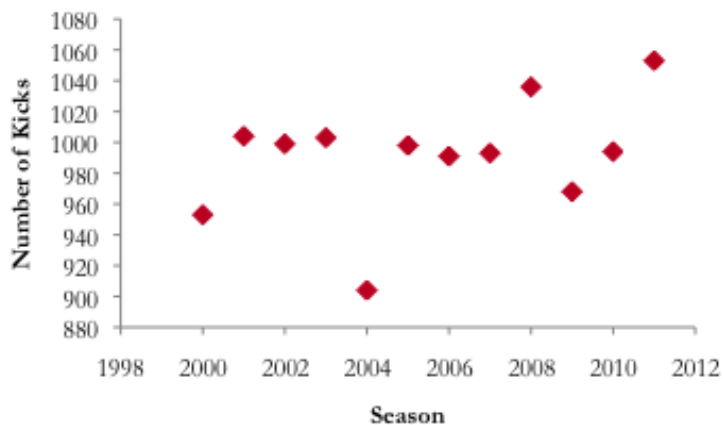


Figure A.7: Number of Kicks from 2000-2011

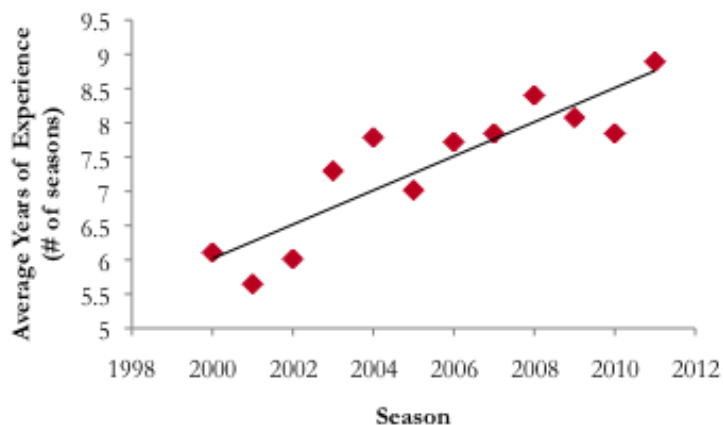


Figure A.8: Years of Experience for the Average Field Goal Attempt from 2000-2011

Difficulty of the Average Field Goal Attempt by Stadium

Table A.5: Difficulty of the Average Field Goal Attempt by Stadium (with a minimum of 5 games)

Rank	Stadium	Team	Average Difficulty Relative to Mean	Seasons	Games
1	Gillette Stadium, Grass	NE	-0.074	4.55	39
2	Lambeau Field	GB	-0.063	12	103
3	Foxboro Stadium	NE	-0.059	2	17
4	Heinz Field	PIT	-0.049	11	95
5	Cleveland Browns Stadium	CLE	-0.049	12	95
6	Soldier Field	CHI	-0.047	11	92
7	Arrowhead Stadium	KC	-0.042	12	97
8	Paul Brown Stadium, Grass	CIN	-0.041	4	32
9	Giants Stadium, Grass	NY	-0.04	3	51
10	M&T Bank Stadium, Grass	BAL	-0.038	3	25
11	Lincoln Financial Field	PHI	-0.037	9	78
12	Sun Life Stadium	MIA	-0.032	12	98
13	LP Field	TEN	-0.028	12	99
14	EverBank Field	JAC	-0.028	12	95
15	O.Co Coliseum	OAK	-0.027	12	99
16	Candlestick Park	SF	-0.027	12	97
17	Raymond James Stadium	TB	-0.026	12	99
18	Bank of America Stadium	MIN	-0.025	12	98
19	FedExField	WAS	-0.024	12	96
20	Sun Devil Stadium (ASU)	ARI	-0.019	6	47
21	Qualcomm Stadium	SD	-0.018	12	98
22	Reliant Stadium	HOU	-0.012	10	78
23	University of Phoenix Stadium	ARI	-0.012	6	52
24	Veterans Stadium	PHI	0.0109	3	28
25	Ralph Wilson Stadium	BUF	0.011	12	91
26	Gillette Stadium, Turf	NE	0.013	5.45	48
27	Husky Stadium (U of Washington)	SEA	0.0153	2	16
28	Paul Brown Stadium, Turf	CIN	0.0225	8	65
29	Texas Stadium	DAL	0.0231	9	71
30	MetLife Stadium	NY	0.026	2	33
31	Giants Stadium, Turf	NY	0.0269	7	113
32	CenturyLink Field	SEA	0.0287	10	85
33	M&T Bank Stadium, Turf	BAL	0.0302	9	73
34	Memorial Stadium (U of Illinois)	CHI	0.036	1	8
35	Three Rivers Stadium	PIT	0.0428	1	8
36	Lucas Oil Stadium	IND	0.0439	4	36
37	Cowboys Stadium	DAL	0.045	3	26
38	Edward Jones Dome	STL	0.0461	12	98
39	Georgia Dome	ATL	0.0461	12	97
40	Pontiac Silverdome	DET	0.0461	2	15
41	Mall of America Field at Metrodome	MIN	0.0461	12	95
42	Mercedes-Benz Superdome	NO	0.0461	11	90
43	RCA Dome	IND	0.0461	8	69
44	Ford Field	DET	0.0461	10	80
45	Sports Authority Field at Mile High	DEN	0.0957	12	98

Current stadiums marked in green.