2019

Win Some, Lose Some: Regression Toward the Mean in College Football Wagers

CONNOR LOVELAND

Introduction

Hall of Fame coach Vince Lombardi once said, "football is a game of inches and inches make the champion." However, these inches aren't easily earned, as luck plays a huge factor in determining a game's outcome. Many of the most memorable moments in football history contain a great deal of luck, such as the famous helmet catch in Super Bowl XLII, or the band running onto the field during the final kickoff in the Stanford/Cal game known as "The Play". These seemingly impossible moments in football history, coupled with unpredictable officiating, play calling, and other factors lead to imperfect measurements in ability. These imperfect predictions tend to exaggerate performance differences in match-ups, resulting in a phenomenon known as *regression towards the mean*.

This paper will investigate whether sports bettors and bookmakers consider this common phenomenon in spread and over/under betting on NCAA football games.

Regression toward the Mean

Regression toward the mean refers to the theory that an extreme performance will revert closer to the mean over time. The first description of the common phenomenon came from Sir Francis Galton[1] in 1886 after studying parents' heights compared to their children. Since then, studies of regression toward the mean have come from a wide range of disciplines.

For example, Linden[2] found that subjects chosen to participate in health care interventions based on an initially high "risk" score tend to have lower actual risk. Pritchett and Summers[3] examined regression toward the mean in national income growth of "middle-income" countries and showed that slowdowns in growth are the consequence of initial growth levels being extreme. Smith and Smith[4] found regression toward the mean in group test scores used to judge students, teachers, and schools.

Outlier performances tend to exaggerate high and low ability, thus leading to follow-up performances much closer to the mean. This tendency, if not properly considered, will overestimate extreme performances as a reflection of extreme ability. However, these spectacular performances tend to come from those with a much lower true ability than performance shows, and vice versa for extremely underwhelming performances.

Regression toward the mean in sports is common, but often misinterpreted. Some of the most popular application in sports are the "Sophomore Slump" referring to a player performing worse after their great rookie season, or the "Madden Curse", where players on the cover of the popular sports game perform worse after appearing on the cover. These players had extraordinary performance throughout that year, leading them to regress back towards their true ability and seem worse than previously.

There are many papers that consider NFL regression toward the mean. Sapra[5] found that point spread lines are efficient in that they accurately reflect victory probability. Furthermore, he also found significant regression toward the mean season-over-season as reflected through team *alpha*, a metric used to show better or worse than expected records for a season. Vergin[6] found that bettors tend to overreact to the recent positive performance, but not to recent negative performance. Lee and Smith[7] find that regression toward the mean is evident in the NFL and NBA using efficient spread and totals to achieve consistently positive win rates above the assumed 50% average.

The Smith and Capron[8] model, where observed performance *Y* fluctuates randomly about ability μ, is:

$$Y = \mu + \varepsilon$$

where ε is a random error term with an expected value of zero. A person's ability is the expected value of their performance, and thus an unbiased predictor of performance. Since the variance of performance is:

$$\sigma_Y^2 = \sigma_\mu^2 + \sigma_\varepsilon^2$$

we see that performance varies by a greater amount that ability, thus extraordinary performance usually reflects more ordinary ability.

However, aside from Lee and Smith[7] and Smith and Capron[8], few papers consider the potential profits of betting with regression toward the mean in mind. Smith[9] finds that the average point spreads in the NFL range from -7 to +7. In contrast, NCAA football tends to range from -12 to +12, indicating a wider performance margin in college football. The main goal of this paper will be to examine whether the wider margin in college football yields similar results to the aforementioned papers regarding NFL wagers and regression toward the mean.

Wagers in Football: Against the Spread and the Over/Under

This paper will focus on two major types of football bets. Betting *against the spread (ATS)* means betting on the predicted margin of victory. In the 2018 Bedlam Battle, where the Oklahoma Sooners played the Oklahoma State Cowboys, bookmakers set the spread, also referred to as the *line*, such that the Sooners were (-21.5), signifying a twenty-one-and-a-half-point favorite to win the game. Betting on Oklahoma paid out if the Sooners won the game by more than 21.5 points, whereas betting on Oklahoma State paid out if the Cowboys lost by less than 21.5, or won outright. Winning the bet is known as *covering*. The 0.5 in the line, known as the *hook*, is added by bookmakers for many games in order to avoid a break-even situation and better balance money on both sides. Otherwise, should the line have been Sooners (-21), and the games ends with the Sooners winning by 21 points exactly, the game is a *push* and the bettor neither wins nor loses.

The other bet this paper considers is the *over/under (O/U)*, also known as the *total*. This bet is based on the total number of points in the game. In the previous example, the total was 80, meaning that bettors would take the line based on whether they believed the total points would be over or under

80. If the two teams combined for more than 80 points, the over bet would pay out; otherwise, if the teams combined for under 80 points, the under would pay out.

The Bedlam Battle ended in a 48-47 victory for the Oklahoma Sooners over the Oklahoma State Cowboys. Using the above examples, the Cowboys covered the spread as they did not lose by more than 21 points. Furthermore, the over covered, as the total for the game was 95 points.

__Betting Payouts and Theory__

For both against the spread and over/under bets a bettor risks more than they will be paid should their bet cover. Typically, payouts are given in terms of $100. Given we are only focusing on the two aforementioned bet types, the only payout relevant to our analysis is (-110). This payout says that, in order to win *$100*, a bettor must risk *$110*. Therefore, in order to have a positive expected value given these bets, a bettor must win at least 52.38% of the time, as shown by:

$$(\$100)P + (-\$110)(1 - P) > 0 \; if \; P > 0.5238$$

Thus, if gambling the same amount each time, a bettor would have to win more than 52.38% of their bets to win money. Considering how much money is wagered and how few (if any) of the uber wealthy amassed their fortune through gambling, it can be inferred that winning this often is much harder than gamblers believe.

Bookmakers try to set the lines for both the spread and the total so there is approximately an even amount of money on both sides. If this is true, bookmakers can make a riskless profit regardless of the outcome of the game. The losers of the bet would give $100 to the winners, and the extra $10 risked in the bet goes to the bookmakers. A $10 profit on a $220 total results in a 4.55% *vigorish* for the bookmaker. The vigorish is compensation for bookmakers making the market and taking on the risk of poorly set lines. While this seems like free money, there have been many cases of uneven sets of money costing bookmakers millions. For example, English Premier League longshot Leicester City was set at a 5000-to-1 odds of winning the league before the 2015-2016 season. After a historic run the team pulled off the impossible, forcing bookies to pay out extraordinary sums. The three largest bookmakers in England lost a combined $11.4 million[10].

As Sapra[5] showed, betting lines, on average, reflect the victory probability for the favorite team. Bookmakers consider every factor from of the game: injuries, location, coaching, player circumstance, rivalries, referees. This results in lines set such that the average bettor should have an equal chance of covering or losing the bet. Therefore, if a regression strategy wins more than 50% of the time at a significant level, that will be an indication of bettors not considering regression toward the mean.

Betting Strategy

This paper will use the betting strategy devised by Lee and Smith[7] for the NFL and will examine weekly matchups for the 64 teams in the Power 5 conferences. The strategy considers each team previous games with regard to covering their bets, and uses that knowledge to bet in favor of regression toward the mean.

With point spread bets, the success differential *D* is measured as the difference in points covered between Team 1 and Team 2 over the previous *n* weeks:

$$D = \sum_{t=1}^{n} S_{1,t} - \sum_{t=1}^{n} S_{2,t}$$

A positive success differential *D* indicates Team 2 is doing worse in total points covered. Betting with regression toward the mean in mind suggests putting $*D* on Team 2. If *D* is negative, the opposite is true and $*D* is placed on Team 1. Each week $1,000 worth of bets are placed, where *D* serves as the multiplier for every game. Therefore, a matchup where the success differential *D* = 50, would have twice as much money on it as a matchup with a D = 25. The total dollars bet are scaled to match $1000 exactly for every week.

Cumulative games will also be considered. In this strategy only the number of times the teams covered the spread or not are considered without care for the number of points. The total dollars are again scaled to $1000.

A similar strategy pertains to total line bets. The success differential for total lines between Team 1 and Team 2 over the previous n weeks is given by the total points over or under the line by the teams combined:

$$D = \sum_{t=1}^{n} S_{1,t} + \sum_{t=1}^{n} S_{2,t}$$

A positive D indicates the teams have combined to be over the total line in the previous games, thus the strategy bets *$D* on the under. The opposite is true for a negative D, where the over would be bet. As with point spread bets, $1000 will be placed every week, and scaled between matchups to ensure bets of that exact amount.

A cumulative games strategy is considered here as well, where only the number of times the teams were over or under the total points spread is considered.

Bookmakers play a large part in allowing these two methods to succeed. The previous matchups of Teams 1 and 2 are built into the spread and over/under already, allowing us to bet without worry for who the actual teams are playing, their coaching styles, or any other factors regarding the two teams. By virtue of the information already being built in, this allows for an easy way to test whether regression toward the mean is taken into account.

In this paper, the bets are recalculated looking back from only the previous game, all the way to the last 10 games. Doing so considers different assumptions for how many previous games a gambler looks back when deciding who to bet. A horizon of one would be only considering what happened in the previous game, where a horizon of 10 considers the summation of the previous 10 games.

Different cutoff levels for weekly bets are considered as well, with the backwards horizon being the entirety of the season. In theory, a higher success differential should result in higher winning percentages if regression toward the mean is occurring, as there is a higher level of abnormal performance. Therefore, for point spread bets, cutoff levels of 50, 100, 150, and 200 are considered while the cutoff levels for cumulative game bets are 2, 4, 6, 8, and 10.

Data

Matchup betting lines and results for NCAA football games for the Power 5 conferences from 2001-2018 came from The Gold Sheet[11]. Betting lines have some variation between sites, but generally lines are identical. Otherwise, large differences in spreads would result in bettors utilizing arbitrage, resulting in unequal money on each side. Game locations and matchups were cross-referenced with Sports Reference[12] to ensure accuracy between scores. Each season follows the current 64 Power 5 conference teams in NCAA football. These teams make up the most prevalent and competitive conferences: *Pac 12, Big 10, Big 12, ACC, SEC*. However, there are a considerable number of non-conference games played by these teams every year. While those games are used to determine the success differential for matchups, the only games bet on are games consisting of two Power 5 teams playing each other, regardless of conference. Furthermore, games that end is a push are not considered in our bets, as the money would be paid back without win or loss.

Table 1 displays summary statistics for point spreads. Whether the favorite covers or does not cover omits any pushes, where neither the team neither covers or does not. Averages spread and actual margins are from the standpoint of the favored team, while standard deviations and correlation coefficients are from the standpoint of the home team. The favored team covered the spread only 47.6% of the time, meaning they did not the other 52.4%. Average point spreads are significantly lower than the average actual margin of victory, showing a large discrepancy between the line set and the actual outcome. Further evidence comes from a low correlation between point spreads and victory margin at 0.31. The average point spread and actual margin over the six-year period are significantly larger than what Smith and Capron[8] found for the NFL. This is in line with Smith[9] findings, where point spreads are larger for college football compared to the NFL.

Table 2 provides similar summary statistics for total lines. The over hits significantly less than the under (46.8% compared to 53.2%), but the average total score and average total line are almost identical. The correlation between the two is even lower than point spreads at 0.27. Interestingly, the correlation is the exact same as that found in the NFL by Smith and Capron[8], indicating the low number is not a singularity within college football, but more likely a product of the bet type.

The low correlation levels between the two imply large amounts of randomness in college football. Unpredictable events drive differences between lines and actual outcomes. One team can have all the right things go their way, while the other can have nothing go right. A 109-yard field goal return to win the game. A ball thrown out of bounds that is batted back in for an interception. A one-handed catch pinned to the defender's back. There are hundreds of famous examples in college football history of events that are impossible to predict beforehand. Due to these, the best team doesn't always win, just as the favorite doesn't always win. These uncertainties cause performance to regress toward the mean, as a team isn't always going to have things go their way.

Results

Betting with cutoffs give a better picture of whether regression to the mean is built into betting lines or not, as seen in *Table 3*. This shows the cutoff bets for cumulative points in against the spread bets. Using a binomial distribution to test the hypothesis that the probability of making a winning bet is 0.50, the probability of picking this many or more winners in 1673 games (without cutoff) is 0.366, a number that is neither high nor low, but better than the null hypothesis of 50%. Interestingly, the cutoff at 50 points performed far and away the best, with a p-value of .0375. Afterwards, the numbers rise significantly, which is the opposite of expectations with cutoff bets. Higher cutoff levels should raise the overall win percentage if regression toward the mean is not considered within the lines. Here, the higher cutoff levels have very poor results, evident by the win percentage of 0 for the 200-point cutoff.

Table 4 provides similar data for various cutoff levels with total line bets. With no cutoff, the binomial distribution states that the probability of picking this many or more winners in 1675 games is 0.5. This provides no evidence for regression toward the mean betting tactics as an effective means in over/under betting. The cutoff levels for over/under bets perform much worse compared to the point spread bets. Each cutoff boasts a negative return, some extremely so. Furthermore, the p-values, with the exclusion of the 50+ cutoff, are much higher than the null hypothesis of 0.50.

Table 5 shows the point spread bets when considering cumulative games as opposed to points. Here, there is clearly higher evidence of regression toward the mean betting working. Overall, the first four cutoff levels show promise with a low p-value. Although the 0 and 2 cutoff levels do not have positive returns, they still win more than 50% of the time. The higher cutoff levels perform similarly when considering cumulative games or points, both having very low win rates.

Table 6 shows the cumulative games cutoffs with regards to total points bets. Here, we see similar results to the cumulative points bets for the lower cutoffs. However, the high cutoffs preform extraordinarily well, with extremely high returns for the 10-game cutoff. This is completely opposite from the cumulative points cutoffs in *Table 4*. However, the poor performance at lower levels suggest this is more of coincidence with a low number of bets at this level instead of evidence of regression toward the mean.

Figure 1 shows the different horizons for cumulative points bets for both against the spread and over/under betting. While most of the horizons boast win rates higher than 50%, the levels are not enough to cover the vigorish. Alternatively, *Figure 2* shows the same horizons for cumulative games bets. Here we see much higher win rates for the bets, but still very few positive returns over the period.

<u>Summary</u>

Smith and Capron[8] found compelling evidence that gamblers do not take regression toward the mean into account fully with NFL wagers on point spreads and over/under lines. This paper aimed at finding whether the same holds within college football, particularly given the wider margin of difference between college teams when compared to NFL teams. However, the general results did not find conclusive evidence for this theory.

There was little rhyme or reason to the varying win rates over years, especially when looking over different game horizons. While many of these ended up with positive win rates defined as greater than the null hypothesis of 50%, none won enough to cover the vigorish.

Overall, there is little evidence to suggest that regression toward the mean is not being taken into account by bettors in college football. There are many reasons as to why college football wagering does not follow the NFL, but there is no doubt that college football wagers do not see the same levels of regression toward the mean as seen in NFL wagers.

Table 1:

Point Spreads and Outcomes								
		Favorite		Point Spreads		Actual Margin		
	# of Games	Cover	No Cover	Mean	SD	Mean	SD	Correlation
2013	275	143	130	11.89	9.06	17.74	13.68	0.51
2014	282	117	162	10.54	9.98	15.16	11.78	0.20
2015	281	143	134	10.04	8.10	15.15	12.48	0.30
2016	287	139	144	11.61	8.65	17.39	14.48	0.36
2017	292	137	143	11.31	8.52	17.12	13.10	0.42
2018	289	134	153	10.60	8.25	16.52	13.15	0.43
Total	1706	813	866	11.00	8.79	16.52	13.16	0.31

Table 2:

Total Lines and Outcomes								
		Score		Total Line		Total Score		
	# of Games	Over	Under	Mean	SD	Mean	SD	Correlation
2013	275	138	133	58.03	26.80	57.00	19.22	0.20
2014	282	115	163	55.79	11.46	55.08	19.97	0.36
2015	281	131	147	57.12	30.64	56.02	21.08	0.22
2016	287	138	144	57.38	10.32	58.01	19.51	0.52
2017	292	134	156	55.39	8.89	54.90	17.87	0.49
2018	289	142	144	55.73	8.08	56.20	20.37	0.49
Total	1706	798	887	56.56	18.29	56.20	19.69	0.27

Table 3:

Cumulative Points Cutoff ATS					
Cutoff	Bets	Won	% Won	P-Value	% Return
0	1673	844	50.45%	0.366076	-0.94%
50	611	328	53.68%	0.037491	0.80%
100	137	75	54.74%	0.152621	7.38%
150	28	13	46.43%	0.714206	-16.43%
200	3	0	0.00%	-	-100.00%

Table 4:

Cumulative Points Cutoff O/U					
Cutoff	Bets	Won	% Won	P-Value	% Return
0	1675	838	50.03%	0.5	-5.58%
50	626	314	50.16%	0.484061	-3.95%
100	129	57	44.19%	0.920673	-14.63%
150	16	7	43.75%	0.772751	-5.18%
200	3	1	33.33%	0.875	-36.67%

Table 5:

Cumulative Games Cutoff ATS					
Cutoff	Bets	Won	% Won	P-Value	% Return
0	1404	726	51.71%	0.104852	-1.21%
2	1173	603	51.41%	0.175069	-0.89%
4	537	284	52.89%	0.097707	1.82%
6	204	113	55.39%	0.070647	1.71%
8	55	28	50.91%	0.5	-2.77%
10	13	5	38.46%	0.866577	-22.27%

Table 6:

Cumulative Games Cutoff O/U					
Cutoff	Bets	Won	% Won	P-Value	% Return
0	1395	714	51.18%	0.195791	-4.46%
2	1143	587	51.36%	0.187447	-4.74%
4	556	273	49.10%	0.679553	-5.48%
6	214	102	47.66%	0.773919	-6.37%
8	70	38	54.29%	0.275209	6.93%
10	22	14	63.64%	0.143139	21.67%

Figure 1:

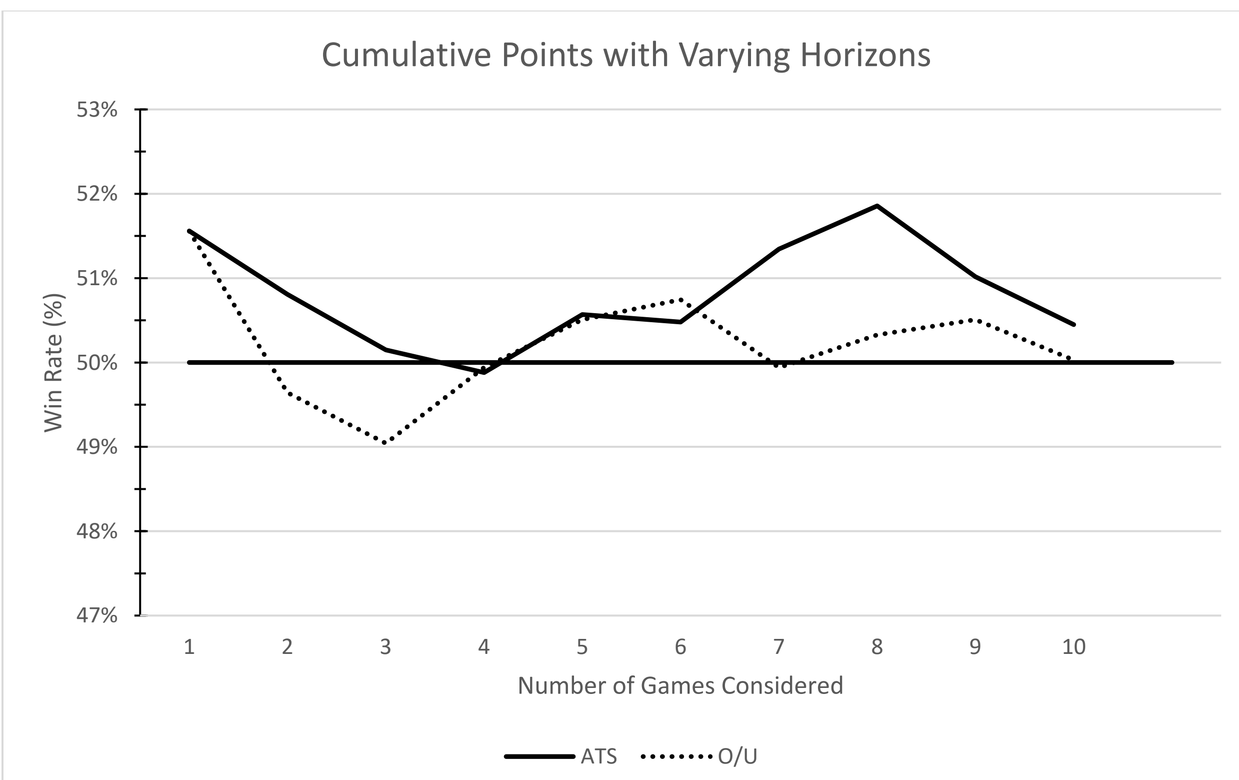

Figure 2:

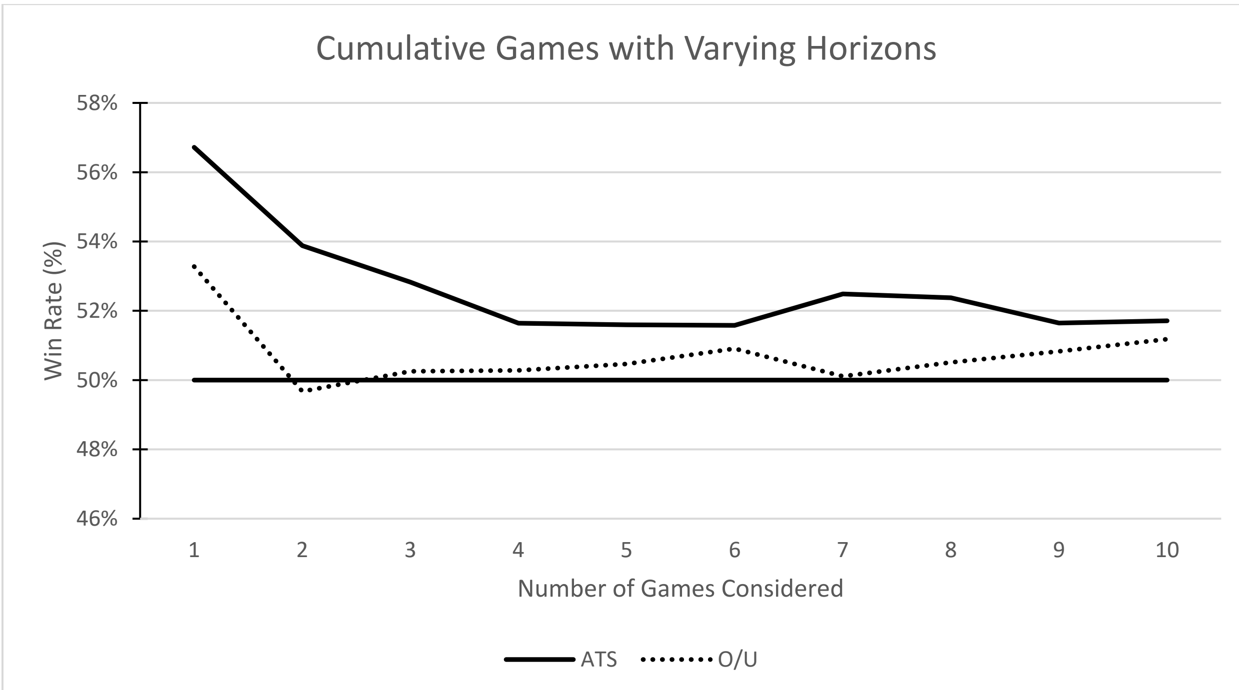
Cumulative Games with Varying Horizons
58%
56%
54%
52%
50%
48%
46%
Win Rate (%)
1
2
3
4
5
6
7
8
9
10
Number of Games Considered
ATS
O/U

References

1. Galton, Francis. "Regression towards mediocrity in hereditary stature." *The Journal of the Anthropological Institute of Great Britain and Ireland* 15 (1886): 246-263.
2. Linden, Ariel. "Assessing regression to the mean effects in health care initiatives." *BMC medical research methodology* 13.1 (2013): 119.
3. Pritchett, Lant, and Lawrence H. Summers. *Asiaphoria meets regression to the mean*. No. w20573. National Bureau of Economic Research, 2014.
4. Smith, Gary, and Joanna Smith. "Regression to the mean in average test scores." *Educational Assessment* 10.4 (2005): 377-399.
5. Sapra, Steven G. "Evidence of betting market intraseason efficiency and interseason overreaction to unexpected NFL team performance 1988-2006." *Journal of Sports Economics* 9.5 (2008): 488-503.
6. Vergin, Roger C. "Overreaction in the NFL point spread market." *Applied Financial Economics* 11.5 (2001): 497-509.
7. Lee, Marcus, and Gary Smith. "Regression to the mean and football wagers." *Journal of Behavioral Decision Making* 15.4 (2002): 329-342.
8. Smith, Gary, and Andrew Capron. "Overreaction in football wagers." *Big data* 6.4 (2018): 262-270.
9. Smith, Robert James. *How to Beat the pro Football Pointspread: A Comprehensive, No-Nonsense Guide to Picking NFL Winners*. Skyhorse Publishing, 2015.
10. Gonzalez, Roger. “Bookmakers Lose $11.4 Million Thanks to 5000-to-1 Leicester City.” *CBSSports.com*, CBS Sports, 3 May 2016, www.cbssports.com/soccer/news/bookmakers-lose-114-million-thanks-to-5000-to-1-leicester-city/.
11. The Gold Sheet. (n.d.). Retrieved February 4, 2019, from www.goldsheet.com
12. Sports Reference. (n.d.). Retrieved February 7, 2019, from www.sports-reference.com/cfb