The Determinants of Scoring in NFL Games and Beating the Over/Under Line

C. Barry Pfitzner*, Steven D. Lang*, and Tracy D. Rishel**

Abstract

In this paper we attempt to predict the total points scored in National Football League (NFL) games for the 2005-2006 season. Separate regression equations are identified for predicting points for the home and away teams in individual games based on information known prior to the games. The predictions from the regression equations (updated weekly) are then compared to the over/under line on individual NFL games in a wagering experiment to determine if a successful betting strategy can be identified. All predictions in this paper are out-of-sample. Using this methodology, we find that several successful wagering procedures could have been applied to the 2005-2006 NFL season. We also estimate a single equation to predict the over/under line for individual games. That is, we test to see if the variables we have collected and formulated are important in predicting the line for NFL games.

I. Introduction

Bookmakers set over/under lines for virtually all NFL games. Suppose the over/under line for total points in a particular game is 40. Suppose further that a gambler wagers with the bookmaker that the actual points scored in the game will exceed 40, that is, he bets the "over." If the teams then score more than 40 points, the gambler wins the wager. If the teams score under 40 points, the gambler loses the bet. If the teams score exactly 40 points, the wager is tied and no money changes hands. The process works symmetrically for bets that the teams will score fewer than 40 points, or betting the "under." The over/under line differs, of course, on individual games. Since losing bets pay a premium (often called the "vigorish," "vig," or "juice" and typically equal to 10 percent of the wagered amount), the bookmakers will profit as long the money bet on the "over" is approximately equal to the amount of money bet on the "under" (bookmakers also sometimes "take a position," that is, will welcome unbalanced bets from the public if the bookmaker has strong feelings regarding the outcome of the wager. See Levitt (2002). It is widely known a gambler must win 52.4 percent of wagers to be successful. That particular calculation can be established simply. Let P_w be the proportion of winning bets and $(1 – P_w)$ the proportion of losing

*Department of Economics/Business, Randolph-Macon College, Ashland, VA, bpfitzne@rmc.edu, 804-752-7307, **slang@rmc.edu**
**Belmont Abbey College, Belmont, NC, tracyrishel@bac.edu

bets. The equation for breaking even where every winning wager nets \$10 and each losing wager represents a loss of \$11 is:

$P_w(\$10) = (1 – P_w)(\$11)$,

and solving for P_w,

$P_w = 11/21 = .5238$, or approximately 52.4 percent.

This research attempts to identify methods for predicting the total points scored in a particular game based on information available prior to that game. The primary research question is whether or not these methods can then be utilized to formulate a successful gambling strategy for the over/under wager.

The remainder of this paper is organized as follows: in section II we describe the efficient markets hypothesis as it applies to the NFL wagering market; section III offers a brief review of the literature; in section IV we describe the data and method; section V provides descriptive statistics and the main regression results; section VI presents the wagering simulations; section VII contains our investigation of the determinants of the over/under line; and the final section offers conclusions.

II. NFL Betting as a Test of the Efficient Markets Hypothesis

A number of important papers have treated wagering on NFL games as a test of the Efficient Market Hypothesis (EMH). This hypothesis has been widely studied in economics and finance, often with focus on either stock prices or foreign exchange markets. Because of the difficulties of capturing EMH conclusions given the complexities of those markets, some researchers have turned to the simpler betting markets, including sports (and the NFL) as a vehicle for such tests.

If the EMH holds, asset prices are formed on the basis of all information. If true, then the historical time series of such asset prices would not provide information that would allow investors to outperform the naïve strategy of buy-and-hold. See, for example, Vergin (2001). As applied to NFL betting, if the use of past performance information on NFL teams cannot generate a betting strategy that would exceed the 52.4 percent win criterion, the EMH hypothesis holds for this market. Thus, the thrust of much of the research on the NFL has taken the form of attempts to find winning betting strategies, that is, strategies that violate the weak form of the EMH.

III. A Brief Review of the Recent Literature

Nearly all of the extant literature on NFL betting uses the point "spread" as the wager of interest. The spread is the number of points by which one team (the favorite) is favored over the opponent (the underdog). Suppose team A is favored over team B by 7 points. A wager on team A is successful only if team A wins by more than 7 points (also known as "covering" the spread). Symmetrically, a wager on team B is successful if team B loses by fewer than 7 points or, of course, team B wins or ties the game—

in any of these cases, team B “covers”. Vergin (2001) and Gray and Gray (1997) are examples of research that focuses on the spread.

Based on NFL games from 1976 to 1994, Gray and Gray (1997) find some evidence that the betting spread is not an unbiased predictor of the actual point spread on NFL games. They argue that the spread underestimates home team advantage, and overstates the favorite’s advantage. They further find that teams that have performed well against the spread in recent games are less likely to cover in the current game, and teams that have performed poorly in recent games against the spread are more likely to cover in the current game. Further Gray and Gray find that teams with better season-long win percentages versus the spread (at a given point in the season) are more likely to beat the spread in the current game. In general, they conclude that bettors value current information too highly, and conversely place too little value on longer term performance. That conclusion is congruent with some stock market momentum/contrarian views on stock performance. Gray and Gray then use the information to generate probit regression models to predict the probability that a team will cover the spread. Gray and Gray find several strategies that would beat the 52.4 percent win percentage in out-of-sample experiments (along with some inconsistencies). They also point out that some of the advantages in wagering strategies tend to dissipate over time.

Vergin (2001), using data from the 1981-1995 seasons, considers 11 different betting strategies based on presumed bettor overreaction to the most recent performance and outstanding positive performance. He finds that bettors do indeed overreact to outstanding positive performance and recent information, but that bettors do not overreact to outstanding negative performance. Vergin suggests that bettors can use such information to their advantage in making wagers, but warns that the market may adjust and therefore this pattern may not hold for the future.

A paper by Paul and Weinbach (2002) is a departure from the analysis of the spread in NFL games. They (as do we in this paper) target the over/under wager, constructing simple betting rules in a search for profitable betting strategies. These authors posit that rooting for high scores is more attractive than rooting for low scores. Ceteris paribus, then, bettors would be more likely to choose “over” bets. Paul and Weinbach show that from 1979-2000, the under bet won 51 percent of all games. When the over/under line was high (exceeded the mean), the under bet won with increasing frequency. For example, when the line exceeded 47.5 points, the under bet was successful in 58.7 percent of the games. This result can be interpreted as a violation of the EMH at least with respect to the over/under line.

Levitt (2002) approaches the efficiency question from a different perspective. It is clear that if NFL bets are balanced, the bookmaker will profit by collecting $11 for each $10 paid out. As we suggested earlier, bookmakers at times take a “position” on the assumption that they know more (or think they do) about a particular wager than the bettors. Levitt presents evidence that the spread on games is not set according to market efficiency. For example, using data from the 2001-2002 seasons, he shows that home underdogs beat the spread in 58 percent of the games, and twice as much was bet on the visiting favorites. Bookmakers did not "move the line" to balance these bets, thus increasing their profits as the visiting favorite failed to cover in 58 percent of the cases.

Dare and Holland (2004) re-specify work by Dare and MacDonald (1996) and Gray and Gray (1997) and find no evidence of the momentum effect suggested by Gray and Gray, and some, but less, evidence of the home underdog bias that has been consistently pointed out as a violation of the EMH. Dare and Holland ultimately conclude that the bias they find is too small to reject a null hypothesis of efficient markets, and also that the bias may be too small to exploit in a gambling framework.

Still more recently, Borghesi (2007) analyzes NFL spreads in terms of game day weather conditions. He finds that game day temperatures affect performance, especially for home teams playing in the coldest temperatures. These teams outperform expectations in part because the opponents were adversely acclimatized (for example, a warm weather team visiting a cold weather team). Borghesi shows this bias persists even after controlling for the home underdog advantage.

IV. The Current Project: Data and Method

This project differs from the extant literature in several ways. First, we focus on the over/under wager. The vast majority of previous work relates to the spread on NFL games. Second, virtually all prior work employs betting rules that are commonly referred to in the literature as "naïve" strategies. For example, betting the home team against the spread when they are the underdog is a simple rule, or naïve strategy. Third, our method is not backward looking—for example, looking back through prior NFL seasons to test the efficiency of the over/under wager. Our method generates "out-of-sample" predictions for each week of the season and tests those predictions against the outcomes. Finally, we also offer evidence on the determinants of the line on NFL games.

The paper by Paul and Weinbach (2002) is a departure from the norm in that they focus on the over/under bet. They, however, like much of this literature also employ a simple rule, suggesting that betting the "under" can be a successful strategy in some circumstances, because bettors prefer to wager that teams will score many points. In this paper, in contrast to prior work, we attempt to formulate regression equations to model points for the home and away teams for specific team match-ups, and to use the sum of predictions from the regression equations for comparison to the over/under. Our modeling method produces out-of-sample symmetrical wagering opportunities, since our predictions for total points suggest wagering the "over" in some cases and the "under" in others. To our knowledge, no prior work attempts to use this type of econometric modeling as a guide to wagering strategies.

With the objective of estimating regression equations for home and away team scoring, data were gathered for the 2005-06 season. The data were collected from the NFL website (www.NFL.com), the Super NFL website (www.supernfl.com), and *The Richmond Times-Dispatch* newspaper. The variables include:

TP = total points scored for the home and visiting teams for each game played
PO = passing offense in yards per game
RO = rushing offense in yards per game
PD = passing defense in yards per game
RD = rushing defense in yards per game
D = a dummy variable equal to 1 if the game is played in a dome, 0 otherwise
PP = points scored by a given team in their prior game
L = the over/under betting line on the game.

1. Match-ups Matter

The general regression format is based on the assumption that "match ups" are important in determining points scored in individual games. For example, if the team "A" with the best passing offense is playing the team "B" with the worst passing defense, ceteris paribus, team "A" would be expected to score many points. Similarly, a team with a very good rushing defense would be expected to allow relatively few points to a team with a poor rushing offense. In accord with this rationale, we formed the following variables:

PY = PO + PD = passing yards
RY = RO + RD = rushing yards.

For example, suppose team "A" is averaging 325 yards (that's high) per game in passing offense and is playing team "B" which is giving up 330 yards (also, of course, high) per game in passing defense. The total of 655 yards, comprised of the excellent passing offense for one team and relatively poor passing defense for the opponent, should indicate an advantage for team "A". Such an advantage is likely to lead to more scoring, at least in that phase of the game.

The dome variable will be a check to see if teams score more (or fewer) points if the game is played indoors, which eliminates weather as a factor.

The variable for points scored in the prior game (PP) is intended to check for streakiness in scoring. That is, if a team scores many (or few) points in a given game, are they likely to have a similar performance in the ensuing game?

We also test to ascertain whether or not scoring is contagious. That is, if a given team scores many (or few) points, is the other team likely to score many (or few) points as well? We test for this by two-stage least squares regressions in which the predicted points scored by each team serve as explanatory variables in the companion equation.

2. <u>General Regression Equations</u>

The general sets of regressions are of the form:

$$TP_{hi} = \beta_0 + \beta_1(PY_{hi}) + \beta_2(RY_{hi}) + \beta_3(D_i) + \beta_4(PP_{hi}) + \varepsilon_{hi} \qquad (1)$$

and

$$TP_{vi} = \alpha_0 + \alpha_1(PY_{vi}) + \alpha_2(RY_{vi}) + \alpha_3(D_i) + \alpha_4(PP_{vi}) + \varepsilon_{vi} \;, \qquad (2)$$

where the subscripts h and v refer to the home and visiting teams respectively, and the i subscript indicates a particular game.

Equations such as 1 and 2 are estimated using data for weeks 5 through 17 of the 2005-06 season. We choose to wait until week five to begin the estimations so that statistics on offense and defense are more reliable than would be the case for earlier weeks. For example, a team might have two exceptional games against poor opponents that would bias early season statistics in a positive way. Such aberrations will average out as more games are played.

V. Descriptive Statistics and Regression Results

1. <u>Descriptive Statistics</u>

Table I contains some summary statistics for the data set. Teams averaged approximately 211 yards passing per game (offense or defense, of course) for the season, and they averaged approximately 109 yards rushing. The statistics reported on the rushing and passing standard deviations without parentheses are for the offenses and the defensive standard deviations are in parentheses. Interestingly, both passing defense and rushing defense are less variable across teams than are the offensive counterparts. We hypothesize that teams must be more balanced on defense to keep other teams from exploiting an obvious defensive weakness, but teams may be relatively unbalanced offensively and still be successful. (The 2007 Patriots would be an example.) Home teams scored approximately 23 points on average for the season and outscored the visitors by about three points. Total points averaged 41.8 in 2005-2006 and the over/under line averaged 40.8. The difference between these means is not statistically significant; the calculated value for the t-test of paired samples is approximately 1. Not surprisingly, the standard deviation was much smaller for the line than for total points.

Table I: Summary Statistics

Variable	Mean	Standard Deviation
Passing Yards	211.4	53.9 (44.3)
Rushing Yards	109.3	29.3 (23.2)
Visitor Points	19.4	10.1
Home Points	22.5	9.7
Total Points	41.8	13.1
Line	40.8	5.0

2. Regression Results

In the estimations of equations 1 and 2, we find no role played by points scored in the prior week and thus we do not report regressions with that variable included. These estimations are available from the authors upon request. The estimated equations (at the end of the 16th week) are given in Table II.

Table II: Regression Results for Total Points Scored

Explanatory Variable	Dependent Variable = TP_{hi}	Dependent Variable = TP_{vi}
Intercept	-9.13 (-1.20)	-22.87 (-2.93)
PY_{hi}	0.0248** (2.00)	
RY_{hi}	0.0948* (4.55)	
PY_{vi}		0.038* (3.18)
RY_{vi}		0.115* (5.17)
D_i	1.66 (0.88)	3.76** (1.98)
$\overline{R}^2$	0.102	0.115
SEE	9.22	9.34
Observations	178	178
F-statistic	7.74*	11.85*

(The numbers in parentheses are t-statistics)
** represents significance at the 95 percent level of confidence or better and * represents significance at the 99 percent level of confidence or better for one-tailed tests

For the home points equation, the dome effect is not statistically significant, but the passing yardage and the rushing yardage are significant at the 95 and 99 percent confidence levels, respectively. The equation explains a modest 10.2 percent of the variance in home points scored. On the other hand, the F-statistic indicates that the overall equation meets the test of significance at the 99 percent level of confidence. The estimated coefficients for all of the variables have the anticipated signs. To interpret those coefficients, an additional 100 yards passing (recall that this is the sum of the home team's passing offense and the visitor's passing defense) implies approximately 2.5 additional points for the home team, whereas an additional 100 yards rushing implies approximately 9.5 additional points.

The visiting team estimation yields a very slightly better fitting equation. All explanatory variables are statistically significant—the yardage variables are each significant at the 99 percent level of confidence, and the dome dummy variable is significant at the 95 percent level. The equation explains almost 12 percent of the variance in visiting team points, and the F-statistic implies overall significance far greater than the 99 percent level of confidence. The coefficients for passing and rushing suggest a greater effect for the visiting team than the home team. The coefficients imply that an additional 100 yards passing yields approximately 3.8 points for the visiting team, and an additional 100 yards rushing is worth 11.5 points. The dome effect implies that the visiting team scores almost 4 additional points in indoor games. The combined dome effect suggests that (ceteris paribus) approximately 5.4 additional points are scored per game in domed stadiums. In fact, for the 32 games in our sample played in domes, the mean number of points was 46.34, whereas the mean was 40.91 for the 162 outdoor games. That difference, 5.43 points, is nearly identical to that predicted by our two equations. A t-test for different mean points scored between domed stadiums and outdoor stadiums is statistically significant at greater than the 95 percent level of confidence.

3. Other Hypotheses

Another hypothesis we wished to test is whether scoring is contagious. A priori, we surmised that points scored in given games for visiting and home teams would be positively related. This conjecture does not look promising. The estimated simple correlation coefficient between home team and visiting team points is -0.106, which is not statistically different from zero and has the “wrong” sign according to our intuition. Our initial thinking was that if team “A” scores and perhaps takes a lead, team “B” has a greater incentive to score. An obviating factor to this line of reasoning is that a given team may dominate time of possession, thus denying the opposing team opportunities to score. We also experimented with two-stage least squares to test the hypothesis that scoring was contagious. In this formulation we developed a “predicted points” variable for the home team, entered that variable as an independent variable in the visiting team equation, and reversed the procedure for the home team equation. Neither of the predicted points variables were statistically significant. The variable was positively signed for the home team equation, and negatively signed for the away team equation.

As indicated above, we find no evidence that teams are “streaky” with respect to points scored. In short, we find that points scored in the preceding week do not contribute to the explanation of points scored in the current week. That conclusion holds up for the regressions in section VI as well.

VI. Wagering on the Over/Under Line

In this simulated wagering project we use the estimated equations to predict scores of the home and away teams for all of the games played over weeks 8 through week 17 (end of the regular season). The points predicted in this manner are then compared to the over/under line for each game. We then simulate betting strategies on those games.

1. Out-of-Sample Method

Since it is widely known that betting strategies that yield profitable results "in sample," are often failures in "out-of-sample" simulations, we use a sequentially updating regression technique for each week of games. That is, we estimate equations TP_{hi} and TP_{vi} with the data from weeks 5, 6, and 7, then "feed" those equations with the known data for each game through the end of week 7, generating predicted points for the visiting and home team for all individual games in week 8. The predicted points are then totaled and compared to the over/under line for each game. Next we add the data from week 8, re-estimate equations TP_{hi} and TP_{vi}, and make predictions for week 9. The same updating procedure is then used to generate predictions for weeks 10 through 17. This method ensures that our results are not tainted with in-sample bias.

2. Betting Strategies

We entertain four betting strategies for the predicted points versus the over/under line on the games. These strategies are:

1. Bet only games for which our predicted total points differ from the line by more than 10 points. (For example, if the line is 40 and the model predicts more than 50, a simulated bet would be placed on the over. With the same line, if the model predicted less than 30, the bet would be on the under.)
2. Bet only games for which our predicted total points differ from the line by more than 7 points.
3. Bet only games for which our predicted total points differ from the line by more than 5 points.
4. Bet all games for which our predicted total points differ from the line by any amount—in our case, all games.

As stated previously, a betting strategy on such games must predict correctly at least 52.4 percent of the time to be successful.

Table III contains a summary of the results for the four betting strategies. The first betting strategy yields only nine "plays" over weeks 8 to 17. That betting strategy would have produced five wins, three losses, and one tie. For this small sample this strategy is profitable, with a 62.5 winning percentage. The second strategy (a differential greater than 7 points) yields 21 plays and a record of 11-9-1—a winning percentage of 55. The 5 point strategy yields more action, 39 bets and a 60.5 percent success rate. Finally for every game played, the method produces a still profitable record of 80-68-4 (notice that two games were eliminated from the full sample, because there were no lines for those games), with the winning percentage at 54.

While we do not find the winning results surprising for the 10 and 7 point criteria, we were surprised to find a winning percentage at lesser differentials, especially the strong performance of the model at the 5 point differential.

Table III: Results of Different Betting Strategies

Betting Strategy (Differential)	*Games "Played"*	*W-L-T Record*	*Win Percentage*
> 10 points	9	5-3-1	62.5
> 7 points	21	11-9-1	55.0
> 5 points	39	23-15-1	60.5
All Games	152	80-68-4	54.0

An important question is whether results for a single season will be robust over future seasons. Indeed, certain biases can disappear over time as bettors and bookmakers incorporate the knowledge of historical biases into their behavior (e.g., bettors with knowledge of the home team underdog bias may choose to wager more often on those teams). We suggest that the information used in this paper is less easily incorporated into market behavior, making the likelihood greater that the model will perform well in future seasons. Simulating the results of this paper requires weekly data collection as well as statistical modeling. Most bettors and bookmakers are unlikely to engage in such work.

Note also that we make no adjustment for injuries, weather, and the like that would be considered by those who make other than simulated wagers. We offer these methods only as a guide, not as a final strategy.

VII. Another Method of Predicting the Line and Total Points

Since we have collected and created variables that may be relevant to determining the betting line (and total points), in this section we investigate the relevance of our variables in this context. For purposes of comparison, we estimate an equation for the over/under line and, separately, for the actual points scored. These equations may be useful in confirming (or contradicting) the results of the previous sections, and may provide useful information applicable to wagering strategies.

The results of those regressions are contained in Table IV. We estimated a regression equation with the line as the dependent variable and all of the right-hand side variables specified in equations 1 and 2. Every coefficient estimate is correctly signed, statistically significant, and $\overline{R}^2 = .671$. As a comparison, we also estimated (far less successfully) an equation for total points with the same set of explanatory variables.

Perhaps the most striking result of these regressions is that the regression for the line explains fully two-thirds of the variance in that dependent variable and the equation for the actual points explains only 5.2 percent of the variance in total points, with only four of the seven explanatory variables meeting the test for statistical significance at traditional levels. The F-test for overall significance of the equation for total points does indicate, however, that a significant portion of the variance in the dependent variable is explained by the regression equation.

Table IV: Regression Results for the Line and Total Points

Explanatory Variable	Dependent Variable = Line	Dependent Variable = Total Points
Intercept	-21.03 (-5.29)	-10.58 (-0.59)
PY_h	0.0476* (12.10)	0.0166 (0.94)
RY_h	0.0507* (6.87)	0.0559** (1.69)
PY_v	0.0442* (11.52)	0.0376** (2.18)
RY_v	0.0450* (5.93)	0.0576** (1.69)
PP_v	0.0669* (2.86)	0.0766 (0.73)
PP_h	0.0343 (1.53)	0.1100 (1.09)
D	2.21* (4.01)	5.17** (2.09)
$\overline{R}^2$	0.671	0.052
SEE	2.84	12.76
Observations	194	194
F-statistic	57.2*	2.51**

(The numbers in parentheses are t-statistics)
** represents significance at the 95 percent level of confidence or better and * represents significance at the 99 percent level of confidence or better for one-tailed tests

In short, the line is, as expected, much easier to predict than actual points scored. That is, the outcomes of the games and points scored are not easily predicted, which is "why they play the games." At least two further observations are in order. First, consider the coefficients for points scored in the previous game. Those variables matter in determining the line for the game. However, they seem to play an insignificant (statistical or practical) role in determining the actual points scored. This particular result may mean that bettors place too much emphasis on recent information, as other authors have suggested. Finally note that the dome effect (based on the magnitude of the coefficient estimates) is weaker for the line equation than in the equation for total points. In fact, the effect for total points is approximately equal to the sum of those effects for equations TP_{hi} and TP_{vi} in section IV. Based on these results, a tentative conclusion might be that bettors underestimate the dome effect.

VIII. Summary and Conclusions

The regression results in this paper identify promising estimating equations for points scored by the home and away teams in individual games based on information known prior to the games. In a regression framework, we apply the model to four simulated betting procedures for NFL games during weeks 8 through 17 of the 2005-2006 season. Betting strategies based on four differentials between our predictions and the over/under line each produced winning results for the season. The relevant question is, of course, whether these results will hold up in future seasons.

REFERENCES

Borghesi, Richard, “The Home Team Weather Advantage and Biases in the NFL Betting Market,” *Journal of Economics and Business*, 2007, 59, 340-354.

Dare, William H. and S.S. MacDonald, “A Generalized Model for Testing Home and Favourite Team Advantage in Point Spread Markets,” *Journal of Financial Economics*, 40, 1996, 295-318.

Dare, William H. and A. Steven Holland, “Efficiency in the NFL Betting Market: Modifying and Consolidating Research Methods,” *Applied Economics*, 2004, 36, 9-15.

Gray, Philip K. and Stephen F. Gray, “Testing Market Efficiency: Evidence from the NFL Sports Betting Market,” *The Journal of Finance*, 1997, LII(4), 1725-1737.

Levitt, Steven D., “How Do Markets Function? An Empirical Analysis of Gambling on the National Football League,” *National Bureau of Economic Research working paper* 9422, 2002.

NFL Enterprises, LLC. “2005 Regular Season.” <http://www.nfl.com/stats/2005/ regular> Cited 4 August 2006.

Paul, R.J. & Weinbach, A.P. “Market efficiency and a Profitable Betting Rule: Evidence from Totals on Professional Football,” *Journal of Sports Economics*, 2002, 3, 256-263.

The Richmond Times-Dispatch, “NFL Statistics”, Wednesday editions, October 5 through December 28, 2005.

The Super NFL. “2005 NFL Daily Odds.” <http://www.supernfl.com/ NFLodds2005.html> Cited 1 August 2006.

Vergin, Roger C., “Overreaction in the NFL Point Spread Market,” *Applied Financial Economics*, 2001, 11, 497-509.