Department of Statistics

THE PROBABILITY OF WINNING A FOOTBALL GAME AS A FUNCTION OF THE POINTSPREAD

by

Hal Stern
Stanford University

TECHNICAL REPORT NO. 59
MARCH 1986

DEPARTMENT OF STATISTICS
STANFORD UNIVERSITY
STANFORD, CALIFORNIA

THE PROBABILITY OF WINNING A FOOTBALL GAME AS A FUNCTION OF THE POINTSPREAD

by

Hal Stern

TECHNICAL REPORT NO. 59

March 1986

DEPARTMENT OF STATISTICS

STANFORD UNIVERSITY

STANFORD, CALIFORNIA

The Probability of Winning a Football Game as a Function of the Pointspread

Hal Stern
Department of Statistics
Stanford University

Abstract

Based on the results of the 1981, 1983, and 1984 National Football League seasons, the margin of victory over the pointspread (number of points scored by the favorite minus the number of points scored by the underdog minus the pointspread) is not significantly different from the normal distribution $N(0,\sigma^2)$, $\sigma = 13.861$. The probability that a team favored by p points wins the game is then $\Phi(p/13.861)$. For $|p| \leq 6$, $.5 + .03p$ is a good linear approximation to the probability of winning.

The Probability of Winning a Football Game as a Function of the Pointspread

Hal Stern

1. Introduction.

The perceived difference between two football teams is measured by the pointspread. For example, Team A may be a three point favorite to defeat Team B. Bets can be placed at fair odds (there is a small fee to the person handling the bet) on the event that the favorite defeats the underdog by more than the pointspread. In our example, if Team A defeats Team B by more than three points, then those who bet on Team A would win the bet. If Team A wins by less than three points (or loses the game), then those who bet on Team A would lose the bet. If Team A wins by exactly three points, no money is won or lost. The pointspread is set so that the amount bet on the favorite (actually the bet is on the event that the favorite beats the underdog by more than the pointspread) is the same as the amount bet against the favorite.

Alternatively, bets may be placed on the event that the favorite wins the game, at odds determined by the probability of that outcome. Bassett (1981) develops a model which explains the widespread use of pointspreads in betting. The pointspread as a predictor of the outcome of a game is discussed by Pankoff (1968), Vergin and Scriabin (1978), and Zuber, Gandar, and Bowers (1985).

In this paper, the results of National Football League games are used to develop a formula for the probability that the favorite wins a football game as a function of the pointspread.

2. Data.

For the 1981,1983, and 1984 National Football League seasons, the pointspread and score of each game are recorded. The 1982 season is excluded due to a players' strike that occurred that year. The total number of games is 672. For each game the number of points scored by the favorite (F), the number of points score by the underdog (U), and the pointspread (P) are recorded.

The margin of victory over the pointspread, which we will call MARGIN is defined by

$$MARGIN = F - U - P$$

for each game. The distribution of MARGIN is concentrated on multiples of one-half since F and U are integers, while P is a multiple of one-half.

The natural estimate of the probability that a p-point favorite wins a game is the proportion of p-point favorites in the sample that won their game. This procedure leads to estimates with large standard errors due to the small number of games with any particular pointspread. Instead, all of the data is used to find a model which then provides a formula for the probability that a p-point favorite wins a game.

3. Distribution of MARGIN.

A histogram of the margin of victory over the pointspread appears in Figure 1. Each bin of the histogram has a width of 4.5 points. Two statistical tests indicate that the distribution of MARGIN is not significantly different from a Gaussian distribution with mean 0 and standard deviation 13.861. The density of this Gaussian distribution is included in Figure 1.

The chi-square goodness of fit test statistic is computed as:

$$\sum_{i=-8}^{8} \left(\frac{(o_i - e_i)^2}{e_i} \right)$$

where o_i = number of observations in the i^{th} bin, $(4.5i - 2.25, 4.5i + 2.25)$,
and e_i = expected number in the i^{th} bin under the Gaussian distribution
$= 672 \left[\Phi\left(\frac{(4.5i+2.25)}{13.861}\right) - \Phi\left(\frac{(4.5i-2.25)}{13.861}\right) \right].$

The rightmost bin (corresponding to i=8) covers the interval $(33.75, \infty)$ in order that the expected count in each bin is greater than five. A similar adjustment is made to the leftmost bin.

This statistic is compared to the chi-square distribution with 14 degrees of freedom. The number degrees of freedom is computed as the number of bins minus one, less one degree for each parameter estimated in computing the expected counts. In this case there are 17 bins and 2 parameters (the mean and standard deviation of the Gaussian random variable). The observed chi-square statistic is 15.27. Since this is less than the .75 quantile of the chi-square distribution with 14 degrees of freedom, we do not reject the hypothesis of normality.

The histogram in Figure 1 represents a smoothed estimate of the density of MARGIN. The hypothesis of normality can be tested for histograms which are less smooth (histograms with smaller bins). The smallest possible bin size is one-half point since all values of MARGIN are multiples of one-half. A histogram with this bin size shows that almost two-thirds of the data is integer-valued. The quantization of a normal random variable into bins of this size does not have this property. This discrepancy is a result of the procedure used in setting the pointspread.

Figure 1 Histogram of MARGIN with scaled N(0,13.861**2) density

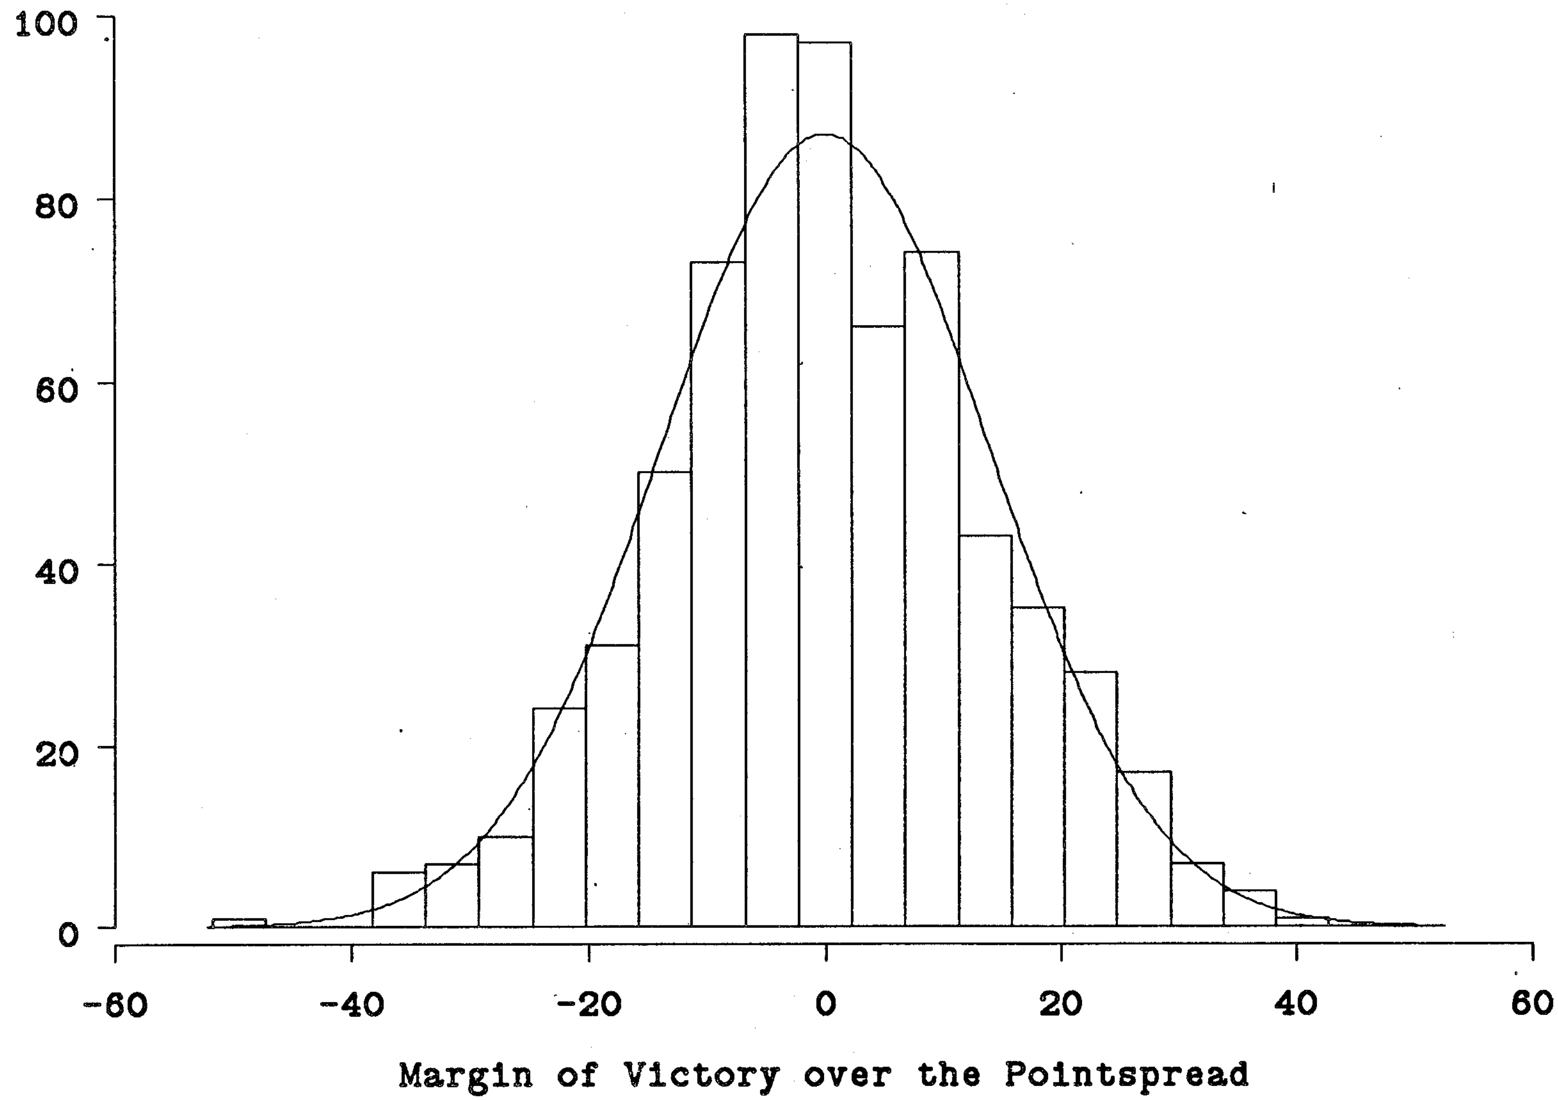

The chi-square test can be applied separately to the integer values of MARGIN, and the non-integer values. In both cases, the normal approximation can not be rejected. The Gaussian distribution is a good approximation even at this level of quantization. The fit of the Gaussian distribution is of course not exact, in fact, the mode of the distribution of MARGIN seems to be less than zero.

The largest absolute difference between the empirical cumulative distribution function of MARGIN and the cumulative distribution function of a $N(0,13.861^2)$ random variable is 0.036. The Kolmogorov-Smirnov test statistic is calculated by multiplying this difference by the square root of the sample size; the resulting value is 0.933. This is less than the .80 quantile of this statistic's limiting distribution, so the normal approximation can not be rejected using this test.

4. Probability of Winning as a Function of the Pointspread.

The favorite wins the football game if it's score F, is greater than the underdog's score U. Using the approximate normality of MARGIN, the probability that a team favored by p points wins the game is

$$\Pr(F > U) = \Pr(F - U - p > -p) = 1 - \Phi\left(\frac{-p}{13.861}\right) = \Phi\left(\frac{p}{13.861}\right)$$

where $\Phi(\cdot)$ is the cumulative distribution function of the standard normal.

For convenience, the quadratic approximation to the normal cumulative distribution function of Shah (1985) may be used. The approximation in this case is

$$\Pr(F > U) \approx 0.50 + 0.10\left(\frac{p}{13.861}\right)\left(4.4 - \left(\frac{p}{13.861}\right)\right).$$

This approximation is accurate to 0.0052 for $0 \leq p \leq 30$. Pointspreads of more than 30 points are extremely rare. In such cases, the probability that the favorite wins the game is approximately 0.99.

As a linear approximation to the above quadratic we get

$$\Pr(F > U) \approx 0.50 + 0.03p.$$

This formula is accurate to within 0.0126 for $p \leq 6$. For the remainder of the paper, we will use the quadratic expression as our formula.

For some sample pointspreads, the probability of victory is calculated using the quadratic expression. The observed proportion of p-point favorites that won their game, $\hat{P}$, and an estimated standard error are also computed. The results are:

Pointspread	$\Pr(F > U)$	$\hat{P}$ (S.E.)
1	0.531	0.571 (.071)
3	0.591	0.582 (.055)
5	0.646	0.615 (.095)
7	0.697	0.750 (.065)
9	0.744	0.650 (.107)

The estimates from the formula are consistent with the estimates made directly from the observations. In addition, the estimates from the formula have the property that they are monotone increasing in the pointspread. This is consistent with the interpretation of the pointspread as a measure of the difference between two teams.

5. Applications.

The function developed in the previous section can be used to analyze the results of a series of football games. Conditioned on the value of the pointspread, the outcome of each game (measured by $F - U$) can be thought of as the sum of the pointspread and a zero mean Gaussian random variable. This is a consequence of the normal distribution of MARGIN. We assume that the zero mean Gaussian random variables associated with different games are independent. Then the probability of a sequence of events is computed as the product of the individual event probabilities.

As an example, the pointspreads of the 1984 New York Giants football games can be used to compute the distribution of the number of games won by the Giants during the sixteen game season. The pointspread for each game and the corresponding probability of a Giants victory were:

Game #	Pointspread	Pr(Giants win)
1	2.0	0.5614
2	-5.0	0.3543
3	-6.0	0.3283
4	6.0	0.6717
5	-3.0	0.4095
6	-3.5	0.3953
7	-5.0	0.3543
8	0.0	0.5000
9	-6.0	0.3283
10	-7.0	0.3033
11	3.0	0.5905
12	-1.0	0.4688
13	7.0	0.6967
14	3.5	0.6047
15	-4.0	0.3814
16	9.0	0.7435

A negative pointspread indicates that the Giants were underdogs in that game.

Adding the probabilities for all sequences of sixteen games that have a particular number of wins leads to the following distribution:

# of wins	probability
0	0.0000
1	0.0002
2	0.0020
3	0.0100
4	0.0332
5	0.1426
7	0.1938
8	0.2025
9	0.1635
10	0.1017
11	0.0482
12	0.0171
13	0.0044
14	0.0008
15	0.0001
16	0.0000

The Giants actual total of 9 wins was slightly higher than expected.

6. Summary.

What is the probability that a p-point favorite wins a football game? It turns out that the margin of victory for the favorite over the pointspread is approximated by a Gaussian random variable. This approximation satisfies the chi-square goodness of fit test. The normal cumulative distribution function can then be used to compute the probability that the favored team wins a football game. Quadratic and linear approximations are given to allow easier calculations.

Acknowledgements

The author would like to thank Tom Cover for many helpful discussions.

References

Bassett,G.W. (1981) "Point Spreads versus Odds," *Journal of Political Economy*, 89, 752-768.

Pankoff,L.D. (1968) "Market Efficiency and Football Betting," *Journal of Business*, 41, 203-214.

Shah,A. (1985) "A Simple Approximation for Areas Under the Standard Normal Curve," *The American Statistician*, 39, 80.

Vergin,R.C. and Scriabin,M. (1978) "Winning Strategies for Wagering on Football Games," *Management Science*, 24, 809-818.

Zuber,R.A. ,Gandar,J.M. ,and Bowers, B.D. (1985) "Beating the Spread: Testing the Efficiency of the Gambling Market for National Football League Games," *Journal of Political Economy*, 93, 800-806.