

Beating the Book: Are There Patterns in NFL Betting Lines?

Michael R. Summers

Abstract

Las Vegas sports books provide two even-money bets (not counting commission, or “vigorish”) regarding National Football League games – the point-spread between each pair of teams and the total number of points scored by each pair of teams (the “over-under”). Odds makers set the numbers (“lines”) for these two bets almost a week in advance of the games, and they may change these numbers throughout the week as new information becomes available regarding the teams or as bettors bet unequally on either side of the line. This study examines whether there are any predictable patterns in the betting lines over the course of the preceding week that bettors could exploit to improve their expected returns, including whether it would have been profitable to try for “middles” as the betting lines changed over the course of the week. We will also note whether during the years studied any simple betting strategy, such as betting on favorites, underdogs, home teams, or visiting teams, would have been profitable.

Key words: NFL betting, gambling, point spread, sports book, line

Introduction

Sports betting, a legalized gambling option in the State of Nevada, presents a betting market that has spawned several scientific studies over the past few decades. Most studies center on ways players attempt to engage in sports betting profitably. The National Football League (NFL) is a popular field of wagering among sports gamblers.

Terms, such as odds and point spreads loom large in the literature. Odds makers set a point spread between two teams, and a bettor can bet on either side of the point spread. For example, Team A may be favored by 3 points over Team B. A bettor choosing Team A will win if Team A wins by more than 3 points. A bettor choosing Team B will win if Team B loses by fewer than 3 points or wins the game. A winning bettor pays a commission of 1/11 of the bet (9.09%); thus, for example, a bettor would need to bet \$110 to potentially win \$100. Similarly, odds makers set a total number of points for the two teams in “over-under” betting. A bettor can bet that the total points in the game will be over or under that number. If the point spread or point total of the game is exactly equal to the line, the bet is refunded.

Sports books try to set their lines so that an equal amount of money will be bet on either side of the line. If the initial line accomplishes this, the book will make a profit of 9.09% on half of the total amount bet without any risk. However, if new information is received that changes the game’s outlook, or if bettors tend to bet more on one side of the line, the odds maker may change the line in the time leading up to the game in order to regain the desired balance of bets. Of course, it may not be possible always to balance the bets on either side of the line, and bets made earlier in the week with a different line remain in effect (unlike pari-mutuel betting, where all bettors are subject to the final odds). Therefore, sports books will generally have some degree of risk and will make more or less than the amount of vigorish (but a good one will come out ahead at least as often as behind!).

Professor of Business
Administration
Pepperdine University
Malibu, CA 90263
Phone: 310-506-4536
E-mail: msummers@
pepperdine.edu

Previous Studies

Studies on the NFL betting market have concentrated on whether or not this market is “efficient” or “rational.” In other words, are there betting rules that can take advantage of inefficiencies in the market to make a profit? For example, maybe it is best always to bet on the favorite, or on the underdog – or the home team, or the visitor – or home teams that are underdogs – or the team that has recently beaten the spread the most often – or the team that was in the playoffs last year – or the team that bettors are moving away from. The list of possible betting strategies is obviously long and includes some strategies that are quite complicated.

Evidence of exploitable inefficiencies in the NFL betting market has been spotty at best. Pankoff (1968) found the market to be efficient in the aggregate. Canes (1976) found no profitable betting strategies. Vergin and Scriabin (1978) found winning strategies for the 1969 – 1974 seasons, but Tryfos et al (1984) found the same strategies to be ineffective for other years. Zuber, Gandar, and Bowers (1985) could not reject market rationality but found a profitable regression model for the 1983 season. Gandar, Zuber, O’Brien, and Russo (1988) could find no statistical evidence of inefficiency in the market. However, they did find three betting rules out of seven to be profitable in the 1980 – 1985 period. Lacey (1990) found several profitable betting rules for the 1984 – 1986 seasons, but Vergin (1998) found none of them to be profitable in the 1987 – 1995 period. While Golec and Temarkin (1991) and Gray and Gray (1997) found it profitable to bet on home teams that were underdogs, Dare and McDonald (1996) found no significant effect. Dare and Holland (2004) found home underdogs to perform better than expected, but the profit was not significant, and 9 other betting rules were worse. Borghesi (2007) found that weather could help predict winners, and Osbourne (2001) successfully used previous scores to make a profit. However, Boulier, Stekler, and Amundson (2006) could not find significant gains using power rankings or stadium characteristics. Overall, it would seem that any inefficiencies in the NFL betting market are small and rather transitory.

Searching for Patterns in the Line

While the aforementioned studies used a single number for each game’s point-spread (usually the closing line), the current study examines whether there are any predictable patterns in the point-spread over the course of the week preceding the games. For example, might odds makers set an opening line that is on average overly generous to home teams and then adjust throughout the week? Perhaps there is a burst of betting near the end of the week that tends to pull the line in a particular direction. If there is any significant pattern in the betting lines, there is an opportunity for the bettor to exploit it, either by getting a line that is more favorable to a bettor’s position or by betting on both teams in a game and trying for a “middle.” The study will also examine whether different sports books present the bettor with any significant differences in the betting lines or have different patterns throughout the week. Finally, any obvious inefficiencies that can be exploited by simple betting rules will be noted.

Data

Data were collected for all Sunday games during the 17 regular-season weeks of the 2003 – 2007 seasons. The few games played on other days of the week were not included since they did not have the same time span for betting lines to change. Using five seasons creates a very large sample size (3399 data points based on 1133 games at 3 sports books) while focusing only on more recent years rather than those long past. The current results are based on the lines of 3 sports books recorded at the same time each day (2:45 AM PST) from the VegasInsider.Com Odds Archive. This archive maintains records for betting lines at these three sports books for every day of every NFL season beginning in 2000. While there are many other sports books, including some online, they do not provide such archival data.

The 1133 Sunday games of the last five years represent the entire population for those years. However, in the statistical tests conducted they were considered as a sample of the population of all NFL games. Even those tests relating to an individual year included a large sample of more than 200 games.

Figure 1 shows an example of the data collected regarding the opening and closing point-spreads for the Las Vegas Line, Hilton, and Mirage sports books for all games played on September 9, 2007, the first week of the 2007 season. Each week the Las Vegas Line (formerly the Stardust) is the first book to post odds, with the others following shortly. Most opening lines are set on Monday; however, if a team is involved in a Monday game or if there is some uncertainty regarding the availability of key players, the opening line may not come out until later in the week. Closing lines, of course, are those on Sunday immediately before the games. The change from the opening line to the closing line is calculated, and the actual point-spread is listed.

For simplicity point-spreads are listed from the point of view of the home team and are listed as positive if the home team is the favorite. For example, if a home team is favored by 3 points, the betting line would list that team at -3, meaning that if the home team would still win after subtracting 3 points from its total, then bettors picking the home team will win. Here, however, the same line is listed as +3, meaning that the home team is favored by 3 points. If the visiting team is favored by 3 points, the spread listed here would be -3 from the home team's point of view. Recording data from the home team's point of view avoids the problem of the favored team changing during the week when recording point-spreads from the favorite's point of view. The actual point spreads are listed similarly. A bold "0" in the opening or closing line indicates a point-spread of 0, meaning that there is no favorite. A bold "0" in the deviations from the lines to the actual point-spread indicates a tie, where the bet is refunded. A bold "X" next to a particular game indicates a case where a bettor could have made money by betting a "middle," as discussed later.

Figure 1: Example of Data Sheet

Spreads										
Week 1 -- September 9										
	LV Line			Hilton			Mirage			Actual
	Open	Close	Chg.	Open	Close	Chg.	Open	Close	Chg.	
K. C. -- Houston	1	3	2	1	3	2	3	3	0	17
Denver -- Buffalo	-3	-3	0	-4	-3	1	-3.5	-3	0.5	-1
Pitts. -- Cleveland	-3	-6	-3	-4.5	-6	-1.5	-4.5	-5.5	-1	-27
Tenn. -- Jacks.	5.5	7.5	2	6	7	1	6.5	7.5	1	-3
Carolina -- St. Louis	1	1.5	0.5	0	1.5	1.5	2.5	2.5	0	-14
Phil. -- Green Bay	-2.5	-3.5	-1	-3	-3	0	-3	-3	0	3
Atlanta -- Minn.	2.5	3	0.5	2.5	3	0.5	3	3	0	21
Miami -- Wash.	3.5	3	-0.5	3	3	0	3	3	0	3
New Eng. -- NY Jets	-5.5	-6.5	-1	-6.5	-6.5	0	-6.5	-6.5	0	-24
Tam. Bay -- Seattle	6	6	0	6	6.5	0.5	6	6	0	14
Chic. -- San Diego	5.5	7	1.5	6	7	1	6	7	1	11
Detroit -- Oakland	2.5	3	0.5	2.5	3	0.5	1.5	3	1.5	-15
NY Giants -- Dallas	4	7	3	4	7	3	4	7	3	10
Home team wins:	7	7		7	7		7	7		
Favorite wins:	7	7		7	7		7	7		
Sum of changes (home team):			4.5			9.5			6	
Sum of absolute changes:			15.5			12.5			8	

LV Line	Hilton	Mirage	LV Line	Hilton	Mirage		Open	Close
Act.-Open	Act.-Open	Act.-Open	Act.-Close	Act.-Close	Act.-Close		Ranges	Ranges
16	16	14	14	14	14		2	0
2	3	2.5	2	2	2		1	0
-24	-22.5	-22.5	-21	-21	-21.5		1.5	0.5
-8.5	-9	-9.5	-10.5	-10	-10.5		1	0.5
-15	-14	-16.5	-15.5	-15.5	-16.5		2.5	1
5.5	6	6	6.5	6	6		0.5	0.5
18.5	18.5	18	18	18	18		0.5	0
-0.5	0	0	0	0	0		0.5	0
-18.5	-17.5	-17.5	-17.5	-17.5	-17.5		1	0
8	8	8	8	7.5	8		0	0.5
5.5	5	5	4	4	4		0.5	0
-17.5	-17.5	-16.5	-18	-18	-18		1	0
6	6	6	3	3	3		0	0
						Sum:	12	3

Results

Are there any obvious market inefficiencies that can be exploited by simple betting rules?

The proportion of time that the home team beat the spread and that the favorite beat the spread (“Home team wins” and “Favorite wins”) was calculated for each of the three books both for the opening line and the closing line and summarized in Table 1. For example, a bettor always betting on the home team at the Mirage with the opening line during the 2003 season would have won 54.5% of the time. Because a winning bettor must pay 1/11 of his winnings to the book, however, a bettor needs to win 52.4% of his bets to break even. While a 54.5% winning percentage would more than break even, with a sample size of 220 games, the winning percentage can not be proven to be significantly greater than 52.4%. (All tests are 1-tailed, with a p-value no higher than .05 to be considered significant. Tests that produce lower p-values are noted.) Betting on the home team at any of the three sports books, and betting on either the opening or closing line, would be above the break-even point in 2003, but none significantly.

2004 and 2006, however, were not so kind to home teams. In fact, betting on visiting teams in 2004 using the opening lines at the Hilton or at the Mirage would produce wins significantly above the break-even point. Results at the other lines that year were not quite significant. In 2006 all three opening lines produced significant profits in betting for visiting teams.

In 2005 and 2007 results slightly favored betting on the home teams, but none of these results were above the break-even point. There is obviously some variability in home team results from year to year, but the overall totals were very close to 50% for all lines, showing no consistent bias in favor of home teams or visitors that could be exploited by bettors.

Betting always on the favorites also showed a great deal of variability from year to year. 2006 was a disastrous year for favorites; all lines would have produced winning percentages greater than the break-even point at well beyond the .01 level of significance for betting on underdogs. None of the other opening or closing lines at any of the three books showed any significant differences from the break-even point either for favorites or underdogs. Overall, favorites won just under 50% of the time at all three books and for both opening and closing lines, again showing no consistent bias that could be exploited during these years. Because of the home-team advantage, home teams are also the favorites more often than not, so we see fairly similar results from year to year for both home teams and favorites.

Table 1: Results of Home Teams and Favorites

% Home Team Beats the Spread									
	LV Line		Hilton		Mirage		Total		
	<u>Open</u>	<u>Close</u>	<u>Open</u>	<u>Close</u>	<u>Open</u>	<u>Close</u>	<u>Open</u>	<u>Close</u>	
2003	0.551	0.545	0.546	0.540	0.545	0.543	0.548	0.543	
2004	0.473	0.476	0.469	0.471	0.469	0.473	0.470	0.473	
2005	0.507	0.519	0.507	0.524	0.505	0.524	0.506	0.522	
2006	0.450	0.489	0.462	0.482	0.466	0.487	0.459	0.486	
2007	0.518	0.511	0.514	0.507	0.509	0.507	0.514	0.508	
Total	0.500	0.508	0.500	0.505	0.499	0.506	0.499	0.506	
% Favored Team Beats the Spread									
	LV Line		Hilton		Mirage		Total		
	<u>Open</u>	<u>Close</u>	<u>Open</u>	<u>Close</u>	<u>Open</u>	<u>Close</u>	<u>Open</u>	<u>Close</u>	
2003	0.528	0.505	0.502	0.500	0.509	0.498	0.513	0.501	
2004	0.470	0.475	0.471	0.477	0.473	0.482	0.471	0.478	
2005	0.543	0.563	0.564	0.563	0.553	0.557	0.554	0.561	
2006	0.379	0.416	0.400	0.400	0.430	0.402	0.402	0.406	
2007	0.502	0.511	0.521	0.509	0.509	0.509	0.511	0.510	
Total	0.484	0.493	0.491	0.489	0.495	0.489	0.490	0.490	

Are there significant changes between opening and closing lines that could lead to profitable betting strategies?

At two of the three books (LV Line and Mirage) the average point-spread moved slightly more in favor of home teams throughout the week (Table 2). The overall average change was positive for home teams in all years except 2006. The only change that was significant at the .05 level for a particular book was the .175 average increase at the LV Line in 2005. Bettors wanting to bet on home teams would have been better off betting at the start of the week, and those betting on visitors at the end of the week, in order to obtain the most advantageous point-spread at the LV Line that year. The total average change for all three books in all five years was just a .021 increase, however. Therefore, changes throughout the week showed no significant pattern up or down that could be exploited by bettors.

Looking at the magnitude of the average change throughout the week without regard to its direction, we see that the lines changed an average of just less than 1 point in most cases, with an overall average change of .874. All of these absolute changes for each book and for each year are significantly different from 0 at well beyond the .001 level of significance. Therefore, the lines do change significantly from the opening line to the closing line each week, but not always in the same direction. The fact that such changes exist suggests that there may well be more complicated strategies that can take advantage of these changes.

Table 2: Changes in Betting Lines

Average Change During the Week				
	LV Line	Hilton	Mirage	Total
2003	0.035	-0.030	0.100	0.035
2004	0.017	0.015	0.024	0.019
2005	0.175	0.055	0.041	0.091
2006	-0.132	-0.128	0.055	-0.068
2007	-0.037	0.037	0.096	0.032
Total	0.010	-0.011	0.064	0.021

Average Absolute Change During the Week				
	LV Line	Hilton	Mirage	Total
2003	0.775	0.710	0.827	0.771
2004	1.013	0.841	0.880	0.912
2005	0.885	0.779	0.793	0.819
2006	1.075	0.938	0.976	0.996
2007	1.029	0.809	0.772	0.870
Total	0.955	0.816	0.850	0.874

Does it matter which sports book a bettor chooses?

Among the three sports books examined, ranges (high – low) were calculated both for opening point spreads and for closing point spreads to see if there were any significant differences between books. As shown in Table 3, the overall average range among the three books was .877 points for opening spreads. Interestingly, 2006 and 2007 showed much more variability in opening lines among the books than did the previous three years. We see this trend also in the percentage of time that there was no difference among the three books’ opening lines. In the first three years they all agreed (0 range) almost 30% of the time, while this happened less than 20% of the time during the last two years. The books tended to get even closer in point-spreads over the course of the week. At the close the average range was only .338 points overall, and a full 46.2% of the games had no differences among the three books. Sports books necessarily need to keep an eye on each other in order to balance their bets on both sides of the line, as bettors can easily go elsewhere if they perceive better odds. This seemed especially true near closing. Therefore, any shopping among sports books should probably be done early in the week.

Table 3: Differences Among the Books

Range Between the Three Lines				
	Average Range		% of 0 Ranges	
	Open	Close	Open	Close
2003	0.697	0.364	0.316	0.433
2004	0.685	0.378	0.248	0.404
2005	0.700	0.286	0.327	0.521
2006	1.174	0.346	0.172	0.445
2007	1.125	0.316	0.197	0.509
Total	0.877	0.338	0.252	0.462

Can a bettor profit by seeking “middles” as the line changes throughout the week?

A bettor can try to find a “middle,” wherein he bets both for and against the home team and hopes to win both bets. For example, if the bettor bets for the home team when the point-spread is 3 in favor of the home team, he wins if the home team wins by more than 3 points. If later in the week the point spread increases to 6 in favor of the home team, he can bet against the home team and win whenever the home team fails to win

by at least 6 (Figure 2). If the home team wins by 4 or 5 points, though, the bettor wins both bets. He also comes out ahead if one of his bets is a tie. At worst he wins one bet and loses the other. The latter case would result in a small loss because of the vigorish he pays on his winning bet. Therefore, the question arises as

to whether there are enough middles to offset these small losses. If so, how much of a change in point-spread is desirable before a bettor should try for a middle?

Figure 2: Trying For Middles

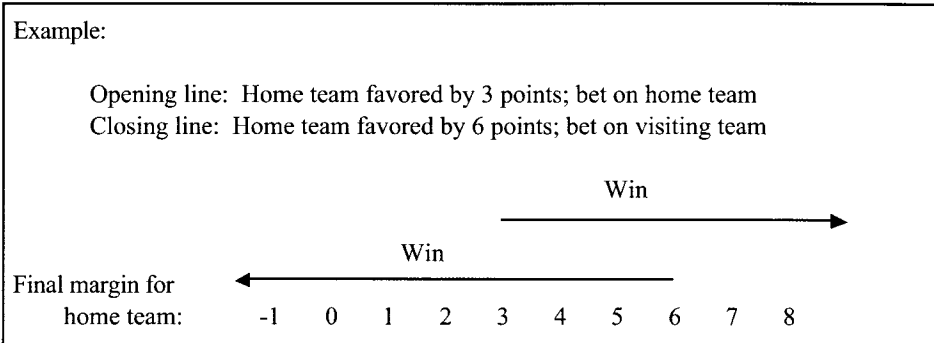


Table 4 shows the proportion of times that the point-spread changed by a certain amount at all three of the books during the week. Overall, the spread didn't change at all 30% of the time, making it impossible to try for a middle. 40% of the time the change was at least 1 point, but larger changes were much less frequent. Changes of at least 2 points occurred about 14% of the time, and changes of 3 or more points less than 6% of the time. Of course, these percentages varied somewhat from year to year, but the pattern was fairly consistent.

Table 4: Proportion of Changes in Point-Spread During the Week

Year	Change in Point-Spread				
	0	.5+	1+	2+	3+
2003	0.335	0.665	0.367	0.121	0.038
2004	0.283	0.717	0.452	0.142	0.062
2005	0.329	0.671	0.381	0.137	0.055
2006	0.257	0.743	0.423	0.170	0.069
2007	0.307	0.693	0.399	0.137	0.051
Total	0.302	0.698	0.405	0.142	0.055

Table 5 shows the number of times that a bettor would have won both bets, won one bet and tied the other, or won one bet and lost the other if he had tried for a middle after point-spread changes of at least either .5 point, 1 point, 2 points, or 3 points. If a bettor bet an amount X on both sides of a game, the payoffs after vigorish would be 20X/11 for winning both bets, 10X/11 for winning one bet and tying the other, and -1X/11 for winning one bet and losing the other. Using the sample frequencies for each result, expected payoffs are calculated.

Overall, the highest positive expected payoff (.070X) occurs when a bettor tries for a middle whenever point-spreads change by at least 3 points (in the right direction). As expected, it turned out that the larger the change in point-spread, the more likely a profitable middle. However, in our sample we have seen that changes of at least 3 points occurred only 5.5% of the time. Also, the change would presumably be in the right direction to try for a middle only half of those times, or 2.75% of the time. A bettor could get a lot more action trying for a middle whenever the point-spread changed by 2 or more points, which happened 14% of the time (7% of the time in the right direction) and still

had a positive expected profit of .022X. Also, we have looked only at bets placed at the opening and closing lines *at the same book*; there should be many more opportunities to try for middles by betting at two different books.

These expected profits are small, and a strategy of betting on both sides of a game might not be terribly exciting to a typical nonprofessional bettor. However, the .070X expected profit for changes of at least 3 points is above 0 at the .025 level of significance. The .022X profit for changes of at least 2 points has a p-value of just over .10, not significant by our standards.

The results of trying for middles varied quite a bit from year to year. In particular, 2004 would have been a terrible year to try for middles, with only 1 double win and 6 cases of 1 win, 1 tie in the entire season at these three books. On the other hand, the two most recent years, 2006 and 2007, would have produced a double win almost once per week, with many more cases of 1 win, 1 tie. These were also the two years that showed a much greater variability in opening lines among the three books, as discussed earlier.

Table 5: Results of Trying For Middles By Betting X Dollars on Each Side

2003						
Change	Two Wins	One Win, One Tie	One Win, One Loss	Total	Expected Profit	
.5+	4	19	438	461	-.033X	
1+	4	9	241	254	-.025X	
2+	2	6	76	84	.026X	
3+	2	3	21	26	.171X	
Payoff	20X/11	10X/11	-1X/11			
2004						
Change	Two Wins	One Win, One Tie	One Win, One Loss	Total	Expected Profit	
.5+	1	6	488	495	-.075X	
1+	1	2	309	312	-.078X	
2+	0	0	98	98	-.091X	
3+	0	0	43	43	-.091X	
Payoff	20X/11	10X/11	-1X/11			
2005						
Change	Two Wins	One Win, One Tie	One Win, One Loss	Total	Expected Profit	
.5+	6	31	400	437	.006X	
1+	6	11	231	248	-.000X	
2+	2	7	80	89	.031X	
3+	1	4	31	36	.073X	
Payoff	20X/11	10X/11	-1X/11			

Conclusions

While not the main purpose of this study, we noted whether there were any obvious inefficiencies in the NFL betting market that could be exploited. For the seasons of 2003 – 2007 neither home teams nor visitors performed significantly better than the other, and neither did either favorites or underdogs. The simple strategies of always betting on one of these categories, which have been addressed in other studies covering other time periods, did not prove profitable during these years.

Although there were very significant changes in point-spreads from the opening line to the closing line at all three books, the positive and negative changes balanced out,

leaving no significant bias one way or the other to be exploited. However, the fact that point-spreads tended to change by almost 1 point on the average throughout the week leaves open the possibility of finding a way to exploit these changes. A bettor who favors a particular team when it is a 5-point underdog might well be inclined to bet even more if it changes to a 7-point underdog, for example.

One way to exploit the changes in point-spreads is in the strategy of trying for middles by betting on both sides of a game when the point-spread changes. When the point-spread changed by at least 3 points during the week, there was a significant profit made by playing for middles. In such situations a bettor averaged a profit of 7% of the amount placed on each team. Even when the point-spread changed by just 2 or more points, the expected profit was positive (2.2% of the amount placed on each team), although not statistically significant. When the point-spread was less than 2 points, though, there were not enough double wins to outweigh the vigorish paid on the winning side of the two bets, causing the expected result to be negative.

The three sports books examined showed little variability in betting lines at the close of betting. Opening lines did show some differences, though, especially in the two most recent years. Therefore, bettors should do any book shopping at the start of the week to get the best line in the direction that they want. Placing a bet on the opening line at one book and on the closing line at another book should dramatically increase the number of occasions when a bettor can obtain larger changes between the two, providing more opportunities to get an attractive spread or to try for middles. The expected winnings for trying for middles when the line changes by at least 2 points or by at least 3 points should apply regardless of which sports book a bettor uses at the opening or at the close.

In performing statistical tests we necessarily assume that our results are representative of all NFL games. In reality, though, our results suffer from the same problems exhibited in the other studies mentioned – they actually apply only to a particular sport in a particular time period and to particular sports books. In order to test the persistence of our findings, more seasons could be examined as they occur, as well as more sports books. A similar analysis could be applied to other sports involving point-spreads, such as college football and basketball.

Further research could include a closer examination of point-spread changes from day to day throughout the week. There may be other profitable strategies that exploit point-spread changes besides trying for middles. It will be interesting to see whether the recent variability in opening lines among the books continues and whether the more frequent opportunities for middles in the last two years continues. Also, a similar analysis can be made regarding over-under betting as well as point-spreads.

References

- Borghesi, R. (2007). The home team weather advantage and biases in the NFL betting market. *Journal of Economics & Business*, 59(4), 340-354.
- Boulier, B. L., Stekler, H. O., & Amundson, S. (2006). Testing the efficiency of the National Football League betting market. *Applied Economics*, 38, 279-284.
- Canes, M. E. (1976). The market for pro football betting. In W. R. Eadington (Ed.), *Gambling and Society: Interdisciplinary Studies on the Subject of Gambling*. Springfield, IL: Charles C. Thomas, 108-137.
- Dare, W. H., & Holland, A. S. (2004). Efficiency in the NFL betting market: Modifying and consolidating research methods. *Applied Economics*, 36, 9-15.
- Dare, W. H., & McDonald, S. S. (1996). A generalized model for testing the home and favorite team advantage in point spread markets. *Journal of Financial Economics*, 40, 295-318.
- Even, W. E., & Noble, N. R. (1992). Testing efficiency in gambling markets. *Applied Economics*, 24, 85-88.

- Gandar, J. M., Zuber, R. A., O'Brien, T., & Russo, B. (1988). Testing rationality in the point spread betting market. *Journal of Finance*, 43, 995-1008.
- Golec, J., & Tamarkin, M. (1991). The degree of inefficiency in the football betting market: Statistical tests. *Journal of Financial Economics*, 30, 311-323.
- Gray, P. K., & Gray, S. G. (1997). Testing market efficiency: Evidence from the NFL sports betting market. *Journal of Finance*, 52, 1725-1737.
- Lacey, N. (1990). An estimation of market efficiency in the NFL point spread betting market. *Applied Economics*, 22, 117-129.
- Osborne, E. (2001). Efficient markets? Don't bet on it. *Journal of Sports Economics*, 2(1), 50-61.
- Pankoff, L. D. (1968). Market efficiency and football betting. *Journal of Business*, 41, 203-214.
- Tryfos, P., Casey, S., Cook, S., Leger, G., & Pylypiak, B. (1984). The profitability of wagering on NFL games. *Management Science*, 30, 123-132.
- VegasInsider.Com. *Odds Archive*. Available from <http://www.dbcsports.com/pastscores>
- Vergin, R. C. (1998). The NFL point spread market revisited: Anomaly or statistical aberration? *Applied Economics Letters* 5, 175-179.
- Vergin, R. C. (2001). Overreaction in the NFL point spread market. *Applied Financial Economics*, 11(5), 497-509.
- Vergin, R. C., & Scriabin, M. (1978). Winning strategies for wagering on National Football League games. *Management Science* 24, 809-817.
- Zuber, R. A., Gandar, J. M., & Bowers, B. D. (1985). Beating the spread: Testing the efficiency of the gambling market for NFL games. *Journal of Political Economy*, 92, 800-806.

Article submitted: 8/10/08

Sent to peer review: 8/26/08

Reviewer comments sent to author: 10/06/08

Author's revised version received: 10/13/08

Accepted for publication: 10/13/08