

Hegarty, Tadgh; Whelan, Karl

Working Paper

Comparing two methods for testing the efficiency of sports betting markets

UCD Centre for Economic Research Working Paper Series, No. WP24/03

Provided in Cooperation with:

UCD School of Economics, University College Dublin (UCD)

Suggested Citation: Hegarty, Tadgh; Whelan, Karl (2024) : Comparing two methods for testing the efficiency of sports betting markets, UCD Centre for Economic Research Working Paper Series, No. WP24/03, University College Dublin, UCD School of Economics, Dublin

This Version is available at:

<https://hdl.hanxdl.net/10419/296707>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

UCD CENTRE FOR ECONOMIC RESEARCH

WORKING PAPER SERIES

2024

**Comparing Two Methods for
Testing the Efficiency of Sports Betting Markets**

Tadgh Hegarty

Karl Whelan

School of Economics
University College Dublin

WP24/03

February 2024

**UCD SCHOOL OF ECONOMICS
UNIVERSITY COLLEGE DUBLIN
BELFIELD
DUBLIN 4**

Comparing Two Methods for Testing the Efficiency of Sports Betting Markets

Tadgh Hegarty* Karl Whelan[†]

University College Dublin

January 2024

Abstract

Sports betting markets can be considered strongly efficient if expected returns on all possible bets on an event are equal. If this form of efficiency holds, then there is a direct mapping from betting odds into probabilities of outcomes of sporting events. We compare two regression-based methods for testing this form of efficiency that have been used in previous research: One that uses normalized probabilities as the explanatory variable for event outcomes and one that uses the inverse of the decimal odds. We show that the normalized probability method produces good tests of the null hypothesis of strong market efficiency but that the inverse odds method does not, with results biased against finding favorite-longshot bias. We illustrate this finding using large datasets of bets and outcomes for tennis and soccer and also with realistic simulations.

Keywords: Betting Markets, Market Efficiency, Favorite-Longshot Bias

JEL Classification: G14, L83, Z20, Z21

*tadgh.hegarty@ucdconnect.ie.

[†]karl.whelan@ucd.ie.

1. Introduction

Betting on sports is growing rapidly around the world due to mobile internet technology and legal changes such as the 2018 US Supreme Court abolition of the federal prohibition on sports betting. To understand betting markets, we need a model of how prices in these markets are set. One potential model is that betting odds are similar to prices set in efficient financial markets, defined by Fama (1970) as prices that reflect all available public information. In a betting context, this would imply that bettors cannot use publicly available information to make higher expected returns betting on one outcome in a contest instead of another. If this property holds, expected rates of return on all bets in a contest are equal, which Thaler and Ziemba (1988) termed “strong-form betting market efficiency”.

Whether strong-form betting market efficiency holds matters not just for betting market participants but also for the use of information from betting odds by researchers to assess the probabilities of outcomes from sporting events. If this form of efficiency holds, then there is a direct mapping from betting odds into probabilities. These odds can then be used in quantitative analysis of sporting events to create proxies for which team is more likely to win or for the amount of uncertainty about the outcomes (for example, Forrest and Simmons, 2002 and Paul, Weinbach, Borgesi and Wilson, 2009).

In this paper, we consider two different econometric tests of strong-form betting market efficiency. One of the methods was originally presented by Pope and Peel (1989) and has also been used recently by Angelini and De Angelis (2019) and Elaad, Reade and Singleton (2020). It regresses the outcome of bets (one if the bet is successful and zero otherwise) on the inverse of the decimal odds for the bet, where decimal odds of D imply the bettor gets a payment of D from a one unit bet, inclusive of the original stake. The inverse of the odds is used in these regressions as a proxy for the probability of the bet being successful. However, because bookmakers earn a gross profit margin, the sum of the inverse decimal odds on bets on an event is greater than one, meaning inverse odds cannot be interpreted as probabilities. For this reason, the second method, used for example by Forrest and Simmons (2005, 2008), Strumbelj and Sikonja (2010) and Hegarty and Whelan (2023), regresses outcomes on “normalized probabilities” that re-scale the inverse odds to sum to one.

We show that normalized probabilities are not just a convenient way to estimate probabilities. Under the conditions required for strong-form efficiency (bookmakers having unbiased estimates of the underlying probabilities and using them to set odds so they believe all bets on a contest have equal expected returns), normalized probabilities are unbiased estimates of the true probabilities. This means econometric specifications regressing outcome dummies on normalized probabilities provide a correct test of the null hypothesis of strong-form market efficiency and, conversely, that the inverse odds method does not. We show that the inverse odds method is biased towards accepting the null hypothesis of market efficiency even if the betting market features a favorite-longshot bias inefficiency so that returns on favorites are greater than returns on longshots.

The paper is structured as follows. Section 2 discusses definitions of betting market efficiency. Section 3 describes the link between betting odds under strong efficiency and the expected outcomes of sporting events and derives results on properties of the two econometric tests of market efficiency. Section 4 illustrates our results using data from online fixed-odds betting on soccer and tennis. Section 5 presents simulation-based evidence. We simulate two types of data generating processes. In the first, the market is efficient and the expected return to each bet on an event is the same and in the second, the bookmaker sets odds to have a favorite-longshot bias. We show normalized probability regressions will tend to provide the correct answers while the inverse odds regressions do not.

2. Defining Market Efficiency in Betting

Thaler and Ziemba (1988) defined two forms of betting market efficiency: A strong and weak form.

Weak form: This is the property that no bets should have positive expected value. There has been a large literature attempting to discover pricing anomalies big enough to allow bettors to earn profits. However, even if profitable strategies can be discovered in a particular data sample, it is less clear that these strategies work out of sample, either because the evidence for a strategy's profitability reflected random sampling variability or because bookmakers adapt their business models to stop losing money due to these strategies. A recent comprehensive analysis of potential profit-generating opportunities in betting on European soccer by Winekelmann, Ötting, Deutscher and Makarawicz (2023) summarized its findings as follows: *"most inefficiencies leading to profitable opportunities to bettors are short-lived and do not occur persistently over time or systematically across leagues."* So, on balance, we think the evidence points towards betting markets generally satisfying this weak form of efficiency. Testing this version of efficiency is not the main focus of our paper but we will briefly return later to whether the evidence we present on betting markets suggests violations of weak-form efficiency.

Stronger form: Thaler and Ziemba framed this form of efficiency in terms of pari-mutuel betting as *"All bets should have expected values equal to $(1-t)$ times the amount bet"* where t is the track take from pari-mutuel operators. The vast majority of modern betting now takes the form of fixed-odds betting where bookmakers offer customers odds and, once agreed, that customer's odds do not change after other bettors place subsequent bets. In this context, Thaler and Ziemba's strong form of efficiency can be framed as simply *"all bets on the same event placed at the same time should have the same expected rate of return."* As noted above, this form of efficiency fits well with the traditional Fama-style definition of financial market efficiency because it implies it is not possible to use publicly available information to improve your return by picking one team in a contest over another.

Our focus will be on testing this stronger form of efficiency for online fixed-odds betting markets.

One could define an even stronger form of efficiency which requires the return on all bets *across different events at the same time* to have the same returns. However, it is well known that this property does not hold. It is widely documented, for example, that bookmakers' margins are lower for English Premier League bets than for the top soccer leagues in other countries (Winkelmann et al, 2023) and that margins increase for lower divisions of the soccer club "pyramid." Similarly, Lyócsaa and Fedorko (2014) report that margins are lower for Grand Slam tennis tournaments than for other tournaments and also that margins decline as tournaments progress to their later stages. There are many possible explanations for these patterns, such as lower research costs for bookmakers for higher-profile events or more intense competition among bookmakers for popular events. We will note just that this form of efficiency does not hold and it is not the form we are testing.

3. Regression-Based Tests of Market Efficiency

Consider a series of sporting events, indexed by i , each with K possible mutually-exclusive outcomes. The K outcome dummies Y_{ij} that equal one if event j occurs and zero otherwise have the property that

$$Y_{ij} = \begin{cases} 1 & \text{with probability } P_{ij} \\ 0 & \text{with probability } 1 - P_{ij} \end{cases} \quad (1)$$

where $\sum_{j=1}^K P_{ij} = 1$. The expected values of the outcome variables are given by $E(Y_{ij}) = P_{ij}$. If the probabilities were observable, then the outcomes could be described by a linear probability model of the form

$$Y_{ij} = P_{ij} + \nu_{ij} \quad (2)$$

where $\sum_{j=1}^K \nu_{ij} = 0$ and the ν_{ij} are uncorrelated across the events indexed by i .

We assume that bookmakers formulate unbiased estimates of the underlying probabilities

$$P_{ij}^B = P_{ij} + \psi_{ij} \quad (3)$$

where $\sum_{j=1}^K \psi_{ij} = 0$ and the ψ_{ij} are uncorrelated across the events indexed by i . This means the outcome variable can also be described by the linear probability model

$$Y_{ij} = P_{ij}^B + \epsilon_{ij} \quad (4)$$

where $\epsilon_{ij} = \nu_{ij} + \psi_{ij}$.

Bookmakers are assumed to use their probability estimates to formulate decimal odds O_{ij} on outcome j in event i , meaning a one unit bet returns O_{ij} (inclusive of the original stake) if the bet wins and zero otherwise. In general, bettors and researchers do not observe the bookmaker's assumed

probabilities P_{ij}^B . However, under the assumption of strong-form market efficiency, the bookmaker sets the odds so that they believe all bets on the same event have an equal expected rate of return, and this means their probability estimates can be calculated.

To see this, note that the payout on a one unit bet on outcome j is

$$\pi_{ij} = \begin{cases} O_{ij} & \text{if } Y_{ij} = 1 \\ 0 & \text{if } Y_{ij} = 0 \end{cases} \quad (5)$$

The bookmaker's expected value of this payout is

$$E(\pi_{ij}) = O_{ij}E(Y_{ij}) = O_{ij}P_{ij}^B \quad (6)$$

Market efficiency implies that the bookmaker's expected payouts on all unit bets on event i equal a common value less than one, which we denote as μ_i

$$O_{ij}P_{ij}^B = \mu_i < 1 \quad i = 1, 2, \dots, K \quad (7)$$

This assumption provides a link between the odds and the bookmaker's probability estimates of the form

$$P_{ij}^B = \frac{\mu_i}{O_{ij}} \quad (8)$$

The requirement that the probabilities sum to one implies the following

$$\sum_{j=1}^K \frac{\mu_i}{O_{ij}} = 1 \quad (9)$$

which can be re-expressed as

$$\mu_i = \frac{1}{\sum_{j=1}^K \frac{1}{O_{ij}}} \quad (10)$$

This expected payout is the inverse of the traditional "overround" calculation (the sum of the inverse odds) that bettors use to assess the extent of the bookmaker's margin for a game. With μ_i calculated, the bookmaker's probabilities can now be directly computed as

$$P_{ij}^B = \frac{\mu_i}{O_{ij}} = \frac{1}{\sum_{j=1}^K \frac{1}{O_{ij}}} \frac{1}{O_{ij}} = P_{ij}^N \quad (11)$$

where P_{ij}^N denotes the standard "normalized" probabilities which re-scale the inverses of the decimal odds to sum to one. This normalization approach is widely used as a convenient way to estimate probabilities but what seems to be less commonly understood is that, under the specific assumption of strong-form betting market efficiency, the normalized probabilities equal the bookmaker's proba-

bilities.

From equation 4, this means we can formulate a test of the null hypothesis of strong-form market efficiency using a regression with normalized probabilities as the explanatory variable

$$Y_{ij} = \alpha + \beta P_{ij}^N + \epsilon_{ij} \quad (12)$$

and testing $\alpha = 0$ and $\beta = 1$.¹ This is, of course, an example of the classic Mincer-Zarnowitz regression (1969) for evaluating the accuracy of forecasts. A finding of $\beta > 1$ implies returns for favorites exceed returns for longshots while $\beta < 1$ implies the opposite. Alternatively, one can run the regression

$$Y_{ij} - P_{ij}^N = \alpha + \gamma P_{ij}^N + \epsilon_{ij} \quad (13)$$

and test $\alpha = \gamma = 0$.

In contrast, Pope and Peel (1989) and others have tested efficiency with the following regression

$$Y_{ij} = \alpha + \beta \left(\frac{1}{O_{ij}} \right) + \epsilon_{ij} \quad (14)$$

This approach has been explained as follows. Pope and Peel assumed that bookmakers' margins implied an additive constant term such that

$$\frac{1}{O_{ij}} = P_{ij} + \kappa_{ij} \quad (15)$$

where the term κ_{ij} is due to the bookmaker's margin and is drawn from a random distribution with mean $\bar{\kappa}$. Then the specification in equation 14 can be understood as

$$Y_{ij} = \alpha + \beta \bar{\kappa} + \beta P_{ij} + \beta (\kappa_{ij} - \bar{\kappa}) \quad (16)$$

This specification features the average additive margin as part of the intercept and a zero-mean residual term, $\beta (\kappa_{ij} - \bar{\kappa})$. The test of strong-form market efficiency becomes just a test of $\beta = 1$, i.e. whether the coefficient on inverse odds in the regression given by equation 14 equals one.

A problem with this approach is that the null hypothesis of strong-form market efficiency implies the link between true probabilities and inverse odds is the multiplicative relationship shown in equation 11 rather than additive relationship in equation 15. This means that, under the null, the inverse odds are a multiple (greater than one) of the true probability, so using them in regressions to explain

¹Linear probability models can be objected to because they can predict probabilities outside the $[0, 1]$ interval but, in this case, the null hypothesis is that the expected value of Y_{ij} increases one for one with P_{ij}^N and not the nonlinear relationships implied by Logit/Probit models. Also, any predicted value that lay outside the $[0, 1]$ interval would itself result from a violation of the null hypothesis. However, we find similar results to those reported here in relation to the statistical significance of coefficients when we use these alternatives.

Y_{ij} will produce values of β that are downward biased.

More formally, the regression using inverse odds is related to the true specification under the null hypothesis as follows

$$\begin{aligned} Y_{ij} &= \alpha + \beta P_{ij} + \epsilon_{ij} \\ &= \alpha + \beta \left(\frac{1}{O_{ij}} \right) + \beta \left(P_{ij} - \frac{1}{O_{ij}} \right) + \epsilon_{ij} \\ &= \alpha + \beta \left(\frac{1}{O_{ij}} \right) + \beta \left(1 - \frac{1}{\mu_i} \right) P_{ij} + \epsilon_{ij} \end{aligned} \quad (17)$$

Because $\left(1 - \frac{1}{\mu_i} \right) < 0$, there is a negative correlation between the dependent variable $\frac{1}{O_{ij}}$ and the $\beta P_{ij} \left(1 - \frac{1}{\mu_i} \right)$ element of the error term, resulting in downward-biased estimates of β . Since favorite-longshot bias implies $\beta > 1$, these tests are biased away from finding this bias and towards finding the opposite. Similarly, if efficiency is tested with a regression of the form

$$Y_{ij} - \frac{1}{O_{ij}} = \alpha + \gamma \left(\frac{1}{O_{ij}} \right) + \epsilon_{ij} \quad (18)$$

and a test of $\gamma = 0$, the procedure is biased towards finding a negative γ .

4. Evidence

Here, we describe the data we use for outcomes and odds for professional soccer and tennis matches. We introduce the data and then provide results for the two regression techniques that we discussed.

4.1. Data

We use two datasets made available by gambling expert Joseph Buchdahl. From www.foxtball-data.co.uk, we obtain outcomes and average closing odds on home wins, away wins and draws across a wide range of bookmakers for 84,230 European professional soccer matches, spanning the 2011/12 to 2021/22 seasons for 22 European soccer leagues (listed in the appendix). From www.tenxis-data.co.uk, we have outcomes and average closing odds for 55,988 professional men's and women's tennis matches on the ATP and WTA tours between 2011 and 2022.²

These datasets also report the maximum available odds on contests, which have been used by some researchers to assess whether there are profitable betting opportunities. We use average rather than maximum odds for two reasons. First, our main focus is on testing the strong-form efficiency assumption of equal returns across different bets in the same contest and, for this purpose, average

²We excluded 16 matches from the original dataset because the reported odds had an overround below 1, meaning a bettor could lock in a certain win by taking both bets. These were likely mis-reported.

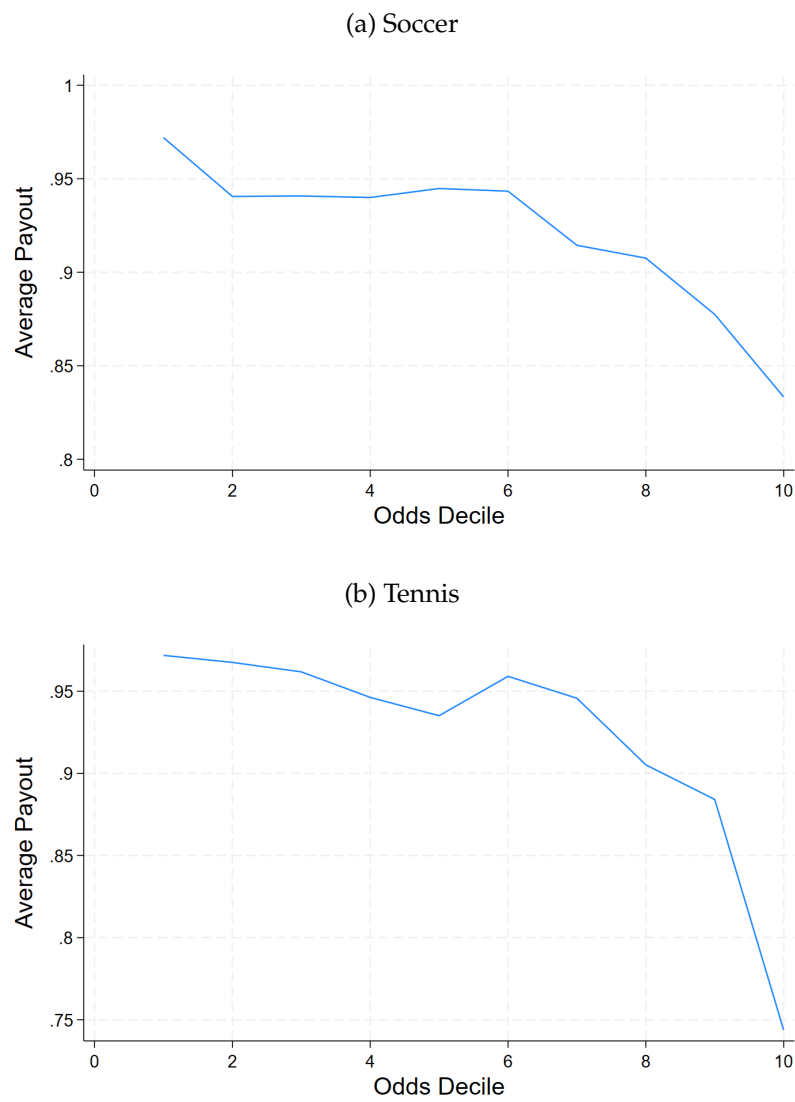
rather than maximum odds are going to be more reflective of the returns that bettors are getting from their stakes. Second, it is not clear that maximum odds data are actually useful for assessing real-world betting strategies. Maximum odds will often tend to be “loss leaders” posted with the intention of attracting new customers. They usually come with restrictions on how much money can be placed. In addition, those who choose to regularly place bets at the best available odds will run the risk of being cut off by retail bookmakers.³

We divided all possible bets (252,690 on soccer and 111,976 on tennis) into deciles according to their decimal odds. The charts in Figure 1 show average payouts for a unit bet for each decile of odds. The data clearly suggest that these betting markets do not satisfy strong-form efficiency. Average payouts decline as the odds rise, dropping off particularly for the upper deciles of odds. This pattern of low average payouts for longshot bets is somewhat larger for tennis than for soccer. Standard *t* tests for differences of means across these deciles strongly reject the hypotheses of the mean payouts for the lower deciles of odds being the same as for the higher deciles.

The failure of the strong efficiency assumption implied by the favorite-longshot bias described Figure 1 does not, however, provide evidence against weak efficiency for these markets. Average payouts per unit bet for all deciles are below one, indicating that on average bets in all deciles lose money. This result still holds if we divide the data into smaller quantiles. While average payouts rise as odds fall, the average payout per unit bet for even the lowest one percent of odds is 0.997 for the soccer data and 0.985 for the tennis data.

³Bookmakers’ practices of customer profiling and stake restrictions are documented by Davies (2022a, 2022b).

Figure 1: Average payout rates for bets by deciles of the decimal odds (1 = lowest odds, 10 = highest odds)



4.2. Regressions

Table 1 reports estimation of equation 13 (for normalized probabilities) and equation 18 (for inverse odds) for both the soccer and tennis datasets. Dummy variables that equal one with a probability P have a variance of $P(1 - P)$, implying the error terms will feature heteroskedasticity. We follow Pope and Peel (1989) in using Weighted Least Squares (WLS) for estimation with the variances approximated by the $P_{ij}(1 - P_{ij})$ where P_{ij} is the relevant probability proxy used in the regression (either normalized probabilities or inverse odds). Because each match shows up multiple times in the regressions and there are correlations between the errors for outcomes of individual matches (i.e. the ϵ_{ij} terms are correlated for each specific value of i), we cluster standard errors at the match level.

The results confirm our theoretical prediction about how these two estimators should behave. The normalized probability test firmly rejects the hypothesis of market efficiency for both datasets with positive and significant estimates of $\hat{\gamma}$. In contrast, the inverse odds method does not reject the hypothesis of efficiency for the soccer market and produces a much weaker rejection for the tennis data, with the estimated $\hat{\gamma}$ still fairly close to zero. Based on the raw evidence in Figure 1 illustrating favorite-longshot bias in both our datasets, we think the normalized probability method delivers the correct answer and the inverse odds method is less able to detect the true pattern in the data.

As with the informal approach of charting average returns in Figure 1, while the favorite-longshot bias reported in these regressions implies strong-form efficiency does not hold, it does not provide any evidence against weak form efficiency. Recalling that the expected return on a one unit bet is $E(\pi_{ij}) = O_{ij}E(Y_{ij})$, one can calculate predicted values of Y_{ij} from the regression model in equation 12 and multiply by the odds to obtain an expected payout. The standard errors for the predicted values of Y_{ij} can also be used to construct confidence intervals for the expected payout.

For the soccer data, using the estimated regression with normalized probabilities, we find that only 343 bets out of 252,690 (0.14% of the sample) have predicted payouts on unit bets that are greater than one, implying an expected profit. And, of these, only 14 bets (0.01% of the sample) have 95% confidence intervals for predicted returns in which all values exceed one. Using the same method for the tennis data, the same exercise reveals 94 bets out of 111,976 (0.08% of the sample) have predicted payouts above one but none have 95% confidence intervals for predicted returns in which all values exceed one. If we instead use the inverse-odds regressions, equation 14, to construct predicted payouts, none of the predicted payouts exceed one.⁴ From this we can conclude that while bets on favorites lose less money than bets on longshots, betting on even extreme favorites is not a profit-making strategy.

⁴While the inverse-odds regressions produce biased estimates of β , the dependent variable being of the form $\frac{P_{ij}^N}{\mu_i}$ rather than P_{ij}^N means they can still produce noisy but unbiased forecasts of Y_i provided the μ_i terms are randomly distributed across the events indexed by i .

Table 1: Testing market efficiency with normalized probabilities and inverse odds
(Estimation of equations 13 and 18. Estimated via weighted least squares regression. Standard errors in brackets)

	Normalized Probabilities	Inverse Odds
Dependent Variable:	$Y_{ij} - P_{ij}^N$	$Y_{ij} - \frac{1}{O_{ij}}$
Explanatory Variable:	P_{ij}^N	$\frac{1}{O_{ij}}$
A. Soccer Data		
$\hat{\alpha}$	-0.025 (0.0024)	-0.026 (0.024)
$\hat{\gamma}$	0.075*** (0.007)	0.009 (0.007)
F test of $\alpha = \gamma = 0$	106.5***	1.85
N	252,690	252,690
B. Tennis Data		
$\hat{\alpha}$	-0.033 (0.003)	-0.036** (0.003)
$\hat{\gamma}$	0.066*** (0.006)	0.0166** (0.006)
F test of $\alpha = \gamma = 0$	186.61***	8.39**
N	111,976	111,976

5. Simulation Evidence

Here, we provide some additional evidence for the properties of the two methods by simulating cases in which the data generating process features an efficient betting market and cases in which it features a favorite-longshot bias pattern. Because we know “the truth” in these cases, we can see how well the two methods do in reflecting this truth in their estimates.

We create four simulated datasets, two calibrated to match features of our soccer data set and two calibrated against the tennis data. For each sport, our first simulated dataset is based on the assumption that betting markets are efficient. It takes a set of simulated probabilities \tilde{P}_{ij} and assumes that bookmakers set decimal odds as

$$\tilde{O}_{ij} = \frac{\mu_i}{\tilde{P}_{ij}} \quad (19)$$

where μ_i is a match-specific expected payout that applies to all bets in a match. The second simulated dataset assumes that bookmakers set odds as

$$\tilde{O}_{ij} = \frac{\mu_{ij}}{\tilde{P}_{ij}} \quad (20)$$

where μ_{ij} is a bet-specific expected payout that depends positively on P_{ij} . In other words, the bookmaker sets odds to have a favorite-longshot bias pattern with higher returns for favorites.

The simulations are calibrated to match features of our two datasets as follows.

1. We construct a set of outcome probabilities for a number of matches matching the number in our datasets. Specifically, we use the normalized probabilities from our two datasets to construct a sample of model probabilities, \tilde{P}_{ij} for each data set. We set the probabilities for simulated match i by randomly drawing from individual matches in the dataset using bootstrapping with replacement (e.g. if the bootstrapped random number is 254, then we use the normalized probabilities from match number 254 for the i th simulated match).
2. For each simulated match, we construct two sets of odds as described above. For the case where the market is efficient, we set μ_i using the formula in equation 10 for the corresponding match used to create the probabilities. For the case where there is favorite-longshot bias, we set μ_{ij} by sorting the normalized probabilities into deciles and then setting the bet-specific expected payout μ_{ij} equal to the average payoff for bets in the decile that the relevant \tilde{P}_{ij} is in.
3. Finally, for each simulated match, we generate random numbers between zero and one to get simulated match outcomes consistent with the specified probabilities. For example, for soccer, we divided $[0, 1]$ into three ordered subsets of size corresponding to the three simulated probabilities to determine the ex ante chance of a home win, away win and draw.

We use this method to construct 10,000 simulated datasets of each type. For each of the four

simulated datasets, we run regressions of the form of equations 13 and 18 and calculate t statistics testing $\gamma = 0$. Table 2 reports median t -statistics for these tests of $\gamma = 0$. Figures 2 to 5 display histograms of the t statistics for each of the four simulated cases: an efficient market calibrated to the soccer data (Figure 2), a favorite-longshot bias pattern calibrated to the soccer data (Figure 3), an efficient market calibrated to the tennis data (Figure 4) and a favorite-longshot bias pattern calibrated to the tennis data (Figure 5).

When the data generating process features our calibration of the favorite-longshot bias, the test based on normalized probabilities correctly rejects the null hypothesis of market efficiency in all simulations, with the median t -statistic closely matching the regression results in Table 1. When the market is efficient, the median t -statistic from the normalized probability method is approximately zero and these statistics reject the null at conventional Type 1 error rates (see the upper panels of Figures 2 and 4). In contrast, the inverse odds method nearly always rejects the hypothesis of market efficiency when it is true, suggesting instead there is a reverse favorite-longshot bias with large negative median t -statistics. When the data generating process features favorite-longshot bias, its median t -statistic is not statistically significant, so the tests fail to correctly reject the null hypothesis most of the time. For the soccer data calibration, this method rejects the null hypothesis of efficiency 13% of the time and does so for the tennis data 47% of the time (see the lower panels of Figures 3 and 5).

Table 2: Median t -statistics testing $H_0 : \gamma = 0$ in equations 13 and 18 for 10,000 simulations

	Normalized Probabilities	Inverse Odds
A. Soccer Data		
Data Generating Model: Efficiency	-0.013	-11.811
Data Generating Model: Favorite-Longshot Bias	11.382	0.52
B. Tennis Data		
Data Generating Model: Efficiency	0.009	-10.99
Data Generating Model: Favorite-Longshot Bias	12.79	1.78

Figure 2: Distribution of simulated t statistics testing $\gamma = 0$ in equation 13 from 10,000 simulations under the null hypothesis of market efficiency calibrated to the soccer data




Figure 3: Distribution of simulated t statistics testing $\gamma = 0$ in equation 13 from 10,000 simulations under the null hypothesis of favorite-longshot bias calibrated to the soccer data




Figure 4: Distribution of simulated t statistics testing $\gamma = 0$ in equation 13 from 10,000 simulations under the null hypothesis of market efficiency calibrated to the tennis data





Figure 5: Distribution of simulated t statistics testing $\gamma = 0$ in equation 13 from 10,000 simulations under the null hypothesis of favorite-longshot bias calibrated to the tennis data



6. Conclusions

Our findings point to three different conclusions that have relevance for the analysis of sports betting markets and their use as proxies in quantitative studies of sporting events.

First, from a market efficiency perspective, our results suggest that the fixed-odds betting markets for soccer and tennis that we have examined do not satisfy the strong-form definition of betting market efficiency which requires the expected return on each outcome in a contest to be the same. The markets feature a substantial favorite-longshot bias pattern such that expected returns on favorites are higher than expected returns on longshots. However, this bias does not allow bettors to earn profits so this pattern does not provide evidence for a failure of weak-form market efficiency.

Second, from the perspective of sports analytics, the failure of strong-form efficiency means betting odds in these markets cannot be used as unbiased proxies for the win probabilities of contestants or to construct unbiased measures of the uncertainty of outcomes.

Third, from a methodological perspective, we have shown that regressions testing strong-form market efficiency should follow the approach of Forrest and Simmons (2005, 2008) and others and use normalized probabilities rather than inverse-odds as the explanatory variable when explaining match outcomes.

References

- [1] Angelini, G. and L. De Angelis (2019). "Efficiency of online football betting markets," *International Journal of Forecasting*, 35, 712–721.
- [2] Davies, R. (2022a). *Jackpot*, Faber.
- [3] Davies, R. (2022b). Revealed: how bookies clamp down on successful gamblers and exploit the rest, *The Guardian*. www.theguardian.com/society/2022/feb/19/stake-factoring-how-bookies-clamp-down-on-successful-gamblers
- [4] Elaad, G. J. J. Reade and C. Singleton (2020). "Information, prices and efficiency in an online betting market," *Finance Research Letters*, Article 101291.
- [5] Fama, G. (1970). "Efficient capital markets: A review of theory and empirical work," *Journal of Finance*, Volume 25, pages 383-417.
- [6] Forrest, D. and R. Simmons (2002). "Outcome uncertainty and attendance demand in sport the case of English soccer," *Journal of the Royal Statistical Society. Series D (The Statistician)*, Vol. 51, pages 229-241.

- [7] Forrest, D. and R. Simmons (2005). Efficiency of the odds on English professional football matches. in Williams, L. V., editor, *Information Efficiency in Betting Markets*, Cambridge University Press.
- [8] Forrest, D. and R. Simmons (2008). "Sentiment in the betting market on Spanish football," *Applied Economics*, 40, 119-126.
- [9] Hegarty, T. and K. Whelan (2023). Forecasting soccer matches with betting odds: A tale of two markets. CEPR discussion paper No. 17949.
- [10] Lyócsaa, S. and I. Fedorko (2014). "What drives intermediation costs? A case of tennis betting market," *Applied Economics*, Volume 48, pages 2037–2053.
- [11] Mincer, J and A. Zarnowitz (1969). "The evaluation of economic forecasts." in *Economic Forecasts and Expectations: Analysis of Forecasting Behavior and Performance*. New York: National Bureau of Economic Research.
- [12] Paul, R. A. Weinbach, R. Borghesi and M. Wilson (2009). "Using betting market odds to measure the uncertainty of outcome in Major League Baseball," *International Journal of Sport Finance*. Volume 4, pages 255-263.
- [13] Pope, P. and D. Peel (1989). "Information, prices and efficiency in a fixed-odds betting market," *Economica*, 56, 223-241.
- [14] Štrumbelj, E. and M. Šikonja (2010). "Online bookmakers' odds as forecasts: The case of European soccer leagues". *International Journal of Forecasting*, 26, 482-488.
- [15] Thaler, R. and W. Ziemba (1988). "Anomalies: Parimutuel betting markets: racetracks and lotteries," *Journal of Economic Perspectives*, 2, 161-174.
- [16] Winekelmann, D., M. Ötting, C. Deutscher and T. Makarawicz (2023). "Are betting markets inefficient? Evidence from simulations and real data," *Journal of Sports Economics*, forthcoming.

A Leagues in the Soccer Dataset

Table 3: Description of the 22 football leagues included in the dataset

Nation	Number of Divisions	Division(s)
England	5	Premier League, Championship, League 1 & 2, Conference
Scotland	4	Premier League, Championship, League 1 & 2
Germany	2	Bundesliga 1 & 2
Spain	2	La Liga 1 & 2
Italy	2	Serie A & B
France	2	Ligue 1 & 2
Belgium	1	First Division A
Greece	1	Super League Greece 1
Netherlands	1	Eredivisie
Portugal	1	Primeira Liga
Turkey	1	Super Lig