

Are Sports Betting Markets Prediction Markets? Evidence from a New Test

Kyle J. Kain* and Trevon D. Logan†

January 2011

Abstract

Researchers commonly use sports betting lines as predictions of the outcome of sporting events. Betting houses set betting lines conditional on bettors ex ante beliefs about game outcomes, which implies that the predictive power of the sports betting market could be an unintended consequence of betting house profit maximization. Using this insight, we propose a new test of the predictive power of the sports betting market, which incorporates a seldom-used piece of complementary betting information: the over/under—the predicted sum of scores for a game. Since the over/under has the same market properties as the betting line, it should be similarly predictive about the actual outcome, while if bettors have different beliefs about this game feature it need not be predictive. Using the universe of betting lines and over/unders on NFL, NBA, NCAA college football, and NCAA college basketball games from 2004-2010, we test the predictive power of the sports betting market in a seemingly unrelated regression (SUR) structure that allows us to characterize both features of the betting market simultaneously. Our joint test reveals that while the betting line is an accurate predictor of the margin of victory the over/under is a poor predictor of the sum of scores. We consistently reject the hypothesis that the sports betting market overall functions well as a prediction market.

* Corresponding author. University of Michigan, e-mail: kylekain@umich.edu

† The Ohio State University and NBER, e-mail: logan.155@osu.edu

Jonathan Burns, Rodney Paul, and Andrew Weinbach graciously shared data sources. Michael Kovach and Chris Zukauckas provided superb research assistance. We thank Rodney J. Andrews, Paul W. Rhode, Michael J. Sinkey, the editor and two anonymous referees for helpful comments. The usual disclaimer applies.

1. Introduction

Researchers commonly use sports betting lines (the predicted margin of victory) as a measure of the expected outcome of sporting events [Pankoff 1968, Sauer 1998]. One use of the betting line is to measure the degree to which a particular outcome was unexpected—a proxy for shocks. This analysis extends beyond sports betting and includes a host of studies that use betting lines to determine unexpected outcomes and their effects on behavior. These studies include domestic violence [Card and Dahl 2011], riots on college campuses [Rees and Schnepel 2009], fundraising [Coate and Depken 2008], cognitive biases [Andrews, Sinkey and Logan 2011] and a host of other issues in economics, finance and public policy. The argument for this use of betting markets is due to the literature which shows that betting lines are extremely good predictors of the actual margin of victory—better than experts and econometric models [Song et al. 2007].

The use of sports betting lines as predictions is somewhat problematic on both theoretical and statistical grounds. Theoretically, the betting houses that set the betting lines have clear incentives to maximize profit by taking into account the *ex ante* beliefs of bettors.¹ If the goal of the betting house is to minimize the risk of having more winners than losers (and having to pay winners from its own sources), betting lines should reflect the wager-weighted median of bettor beliefs. Similarly, betting houses could adopt other, riskier, strategies depending on their risk tolerance and the potential profitability of alternative strategies. In either case, the betting line need not correspond to the expected outcome of the game. Indeed, any correspondence would be incidental depending on the profit function of the betting house. Statistically, existing studies usually regress the betting line on the actual game outcome and find that the relationship is quite strong, betting lines are well correlated with actual margins of victory [Zuber et. al 1985, Gandar et. al 1988, Sauer

¹ Some betting houses purchase their betting lines from professional odds-makers, but even in these situations the same profit motives would apply for the betting house.

et. al 1988, Golec and Tanarkin 1991, Dare and McDonald 1996, Dare and Holland 2004, Sinkey 2011]. While this may establish the general predictive accuracy of the betting line, it is not a powerful test against many reasonable alternatives. For example, these tests usually do not consider subsets of games where the difference between bettor beliefs and expected values may be acute.²

Recent research suggest that betting houses may systematically move betting lines to counteract bettor biases such as beliefs in the “hot hand.” In particular, Logan and Sinkey [2011] find that betting houses overprice favorites who “beat the spread” in their previous contest. Similarly, Levitt [2004] shows that betting houses attempt to manipulate bettors to have skewed distributions of wagers. While the betting house incurs substantial risk in this situation, the payoffs may be substantial. In general, the research to date has shown that the primary function of the betting house is *not* to accurately predict the outcome of a given contest.

While betting lines may be good predictors of margins of victory, it does not necessarily follow that sports betting markets serve as good markets of prediction. In this paper we propose a new test for the accuracy of the sports betting market. Our test incorporates a neglected, but related, piece of sports betting market information—the over/under, the predicted sum of scores in a game (also known as the “totals”). Bettors bet on over/unders just as they do betting lines, that the sum of scores will be greater (“over”) or less than (“under”) the given value. If the betting line and the over/under are accurate predictors then together they fully characterize the game—they predict the scores of each team, their sum and their difference [Evan and Noble 1992]. On the other hand, if both betting lines and over/unders are set to be medians of bettor beliefs or some other function of bettor beliefs about a given contest they need not coincide to the extent that bettors may have different beliefs about the margin of victory as opposed to the sum of scores, and

² As Sinkey and Logan [2010] note, these are tests of predictive power, but have usually been interpreted as tests of market efficiency. We make clear here that we are testing for predictive power, which differs fundamentally from efficiency. See Sinkey and Logan [2010] for more on the distinction between the two.

one piece may be more predictive than another [Evan and Noble 1992, Gandar, Zuber and Russo 1993, Paul and Weinbach 2002]. For example, bettors may believe that a football team will win by 7 (a touchdown) in either a high scoring game or a low scoring game. In other words, the combination of the betting line and the over/under allows us to analyze the predictive power of the sports betting *market* since the betting house faces the same profit function for both outcomes. The question is whether the betting house is actually predicting game outcomes or setting values as a function of bettor beliefs, or doing one for a given betting market and another for the other.

This distinction is important. While economists have concentrated on whether bets are fair we explicitly concentrate on whether betting market prices are *predictive*. We combine betting lines and over/unders to construct a stronger test for the predictive power of the sports betting market. Our approach is straightforward—the betting line and over/under allow us to determine a system of two equations for (1) the sum of scores and (2) their difference. As these two features are related if the betting market is predictive, we use a Seemingly Unrelated Regression (SUR) structure to define the system, allowing for correlations in the errors between the equations. While this need not hold for a betting market (bettor beliefs about one feature need not be related to the other), the correlation between the two should exist if the betting market is predictive as they would be correlated with each other.

Our test is a joint test of the coefficients in both equations—that they are equal to one another and equal to the actual game outcome ($\beta_{\text{BettingLine}} = \beta_{\text{Over/Under}} = 1$). A second test is that the intercept in the regressions are equal to one another and equal to zero ($\alpha_{\text{BettingLine}} = \alpha_{\text{Over/Under}} = 0$), which would be consistent with the betting line and over/under accurately predicting the outcome without systematic adjustment, an obvious indication of lack of prediction. Our more general test exploits the SUR framework—while a prediction market would require a mechanical relationship between the betting line and the over/under, beliefs about the two could be distinct. As such, testing

for the independence of the two underlying equations forms a test of whether or not the betting market is a prediction market. If the two are independent it holds that the sports betting market is most interested in setting lines and over/unders that ensure profitability as opposed to accurately predicting the outcomes of contests.

We use data on the universe of betting lines and over/unders from the NFL, NBA, NCAA football, and NCAA basketball from 2004 -2010 to test the predictive power of the sports betting market. In general, our joint test reveals that the sports betting market does not perform well as a prediction market. We consistently reject the hypothesis that the sports betting market is an accurate predictor of both game outcomes. Specifically, we find that over/unders are set with systematic errors. Moreover, we find that the results vary by sport- football-betting markets (both NCAA and NFL) are more predictive than basketball markets (either NCAA or NBA). More important, we fail to reject the hypothesis that the betting line and over/under equations are independent. While in a prediction market the two would be well correlated with one another by definition, our evidence is inconsistent with that explanation. We conclude that the predictive power of the sports betting market is relatively weak.

Our study makes several contributions to the literature. First, we extend the limited use of over/unders to the analysis of market prediction as opposed to market efficiency. Evan and Noble [1992] and Gander, Zuber, and Russo [1993] used NFL over/unders for one season, and more recently Paul and Weinbach [2002] used more than a decade of NFL data on over/unders. Those studies, however, were interested in the efficiency of the over/unders as opposed to their predictive accuracy. Similarly, all of those studies used NFL data, while we analyze four major sports for which totals markets exist. As such, we conclude that the lack of prediction of the sports betting market is a general trend and not confined to one sport. Second, our use of the SUR allows us to construct a simple test that incorporates a basic and intuitive feature of the totals market—if the sports betting

market is predictive the two metrics should be related to one another as they fully characterize a game.

2. Sports Betting Markets, Prediction, and Profit Maximization

Prediction markets have recently gained prominence due to their ability to forecast future events with less variance and more accuracy than other techniques, such as polling experts. Research on prediction markets has increased along with their use. Theoretical papers in this literature include Ottaviani and Sorensen [2007, 2008], and Wolfers and Zitzewitz [2004, 2007], while empirical papers cover a wide variety of issues, including macroeconomic events and political elections [Rhode and Strumpf 2005, Snowberg et. al 2007]. Many firms now use prediction markets to estimate the likelihood of future events³. Among prediction markets, gambling on sports is one of the oldest and largest—a key strength is that results of a particular contest are public and known with certainty.⁴ For this reason sports betting markets have been among the most popular of prediction markets studied by researchers.

The sports betting market is straightforward. Betting houses facilitate the process by setting the betting line that bettors wager on. The money that is bet on the losing side of the outcome is used to pay off the winning wagers and the betting house retains their fixed percentage, known as the vigorish. Intuitively, betting houses do not risk losing money if exactly half of the total amount bet is on one side of the betting line and the other half on the other side, regardless of the outcome. If this does not occur, the betting house incurs some risk.

³ For example, Google uses prediction markets to predict events relevant to company profitability, including the number of users for Gmail, the quality rating of Google Talk, whether or not Apple will release an Intel-based Mac, etc. [Cowgill, Wolfers, Zitzewitz 2009].

⁴ The public nature of the outcome is important. For example, previous Intrade contracts have been controversial. In 2006, the bet of whether or not North Korea would successfully fire a missile outside of its own airspace was improperly worded, leading to confusion and anger from bettors in the betting market.

Betting houses could profit handsomely if there were more losing bets than winning bets, but such a strategy would involve substantial risk to the betting house. A particularly strategic betting house could attempt to maximize profit over a range of contests, but they would also risk losing substantial sums of money. For this reason, models of betting houses typically begin by assuming that betting houses are primarily interested in setting betting lines that will guarantee equal betting on either side of the betting line (see Gray and Gray [1997] and others for use of this assumption)⁵.

The betting line is usually understood to be expected wager-weighted median of bettor beliefs about the outcome of the game in question as opposed to the expected value of the outcome. Kilby, Fox and Lucas [2002] and Roxborough and Roden [1998] point out in guides for sports book management that a bookmaker's primary objective is, in fact, to minimize risk. This need not be the case, however. For example, betting houses compete with each other for customers, and in doing so they could offer more attractive betting lines if they would result in a larger pool of wagers. Indeed, betting houses earn more in expectation, the more money that is wagered.

In general, the betting house is responsive to the beliefs of bettors. While the betting house's profit function must take into account the wager-weighted distribution of bettor beliefs (and pick a point in that distribution that they believe would maximize profit given their own beliefs about the outcome of the specific contest) the betting house does not have to consider the accuracy of the betting lines that they set. In fact, the predictive accuracy of the betting line is irrelevant to the betting house. For example if betting lines $L1$ and $L2$ would yield the same expected profit the betting house would be perfectly indifferent between offering either line to bettors as their expected profits would be the same.

⁵ Sinkey and Logan [2011] show that risk-minimization is equal to profit maximization if the betting house has the same belief about winning probabilities for a given line as the bettors. We show the proof in an appendix.

The same argument applies, directly, to the over/unders, the only difference being that the outcome is the sum of scores of the contest in question. In both cases the betting houses incentives are the same, to set a line that maximizes profit given the *ex ante* beliefs of bettors.⁶ We take the fact that betting houses set both betting lines and over/unders as evidence that both markets are profitable for the betting house in expectation.

Given this structure of the market, it is easy to see that the sports betting market need not be predictive. Indeed, the sports betting market is only predictive if the bettors happen to have beliefs that are predictive about the game outcome in question. As noted earlier, there may be particular cases where bettor beliefs deviate from expected outcomes, and in these situations betting lines and over/unders would cease to be predictors of game outcomes. Bettors may have accurate beliefs about one outcome but systematic bias about another. Betting houses profit when they take these biases into account, but it would leave the betting market a poor predictive market, at least for features where bettors have known biases. While previous research has established that bets placed on the over/unders of NFL games are indeed fair [Evan and Noble 1992, Paul and Weinbach 2002], the predictive power of over/unders has not been explored in any detail in the literature. A related open question is whether these markets are independently and jointly predictive.

3. Derivation of the Test

The most common method of assessing the predictive power of the sports betting market is to regress the actual outcome of a game on the associated betting line and a list of meaningful covariates, and then to test whether or not the betting line is statistically different from the outcome on average. Variations of this method have been used by Zuber et. al [1985], Gandar et. al [1988], Sauer et. al [1988], Golec and Tamarkin [1991], Dare and McDonald [1996], Dare and Holland

⁶ More precisely, this median would be weighted by the size of the bets placed.

[2004] and Fair and Oster [2007]. Most of these studies have found that the betting line (*SPREAD*) is not statistically different from the margin of victory (*MOV*), on average.

Our test adopts that general structure to the sum of scores (*SUM*) and the over/under (*Over/Under*). While this method may be problematic in testing for the *efficiency* of the over/under [Evan and Noble 1992, Gander, Zuber and Russo 1993, Paul and Weinbach 2002], this is the preferred test for the *predictive* power of the totals market.⁷ This yields a system of equations that characterize the predictive power of the betting market. In particular, the general structure is

$$MOV_i = \beta_0 + \beta_1 SPREAD_i + \varepsilon_{1i} \quad (1)$$

$$SUM_i = \beta_2 + \beta_3 Over/Under_i + \varepsilon_{2i} \quad (2)$$

The first test is that $\beta_0 = 0$, $\beta_2 = 0$ which is a necessary condition that the prediction be free of any systematic error. The second test is that $\beta_1 = 1$, $\beta_3 = 1$ which is the key to prediction—the betting line or over/under should be an accurate predictor.

We further propose a joint test of the sports betting market. When thinking of the joint test it is important to note that the errors of one test may be correlated with the errors of the other. Indeed, the sum of scores and the margin of victory have a mechanical relationship to one another. This naturally gives rise to a system of equations ((1) and (2) above) that are Seemingly Unrelated Regressions (SUR), as in Zellner [1962]. That is, shocks to one equation would have natural spillovers to the other. For example, games that are predicted to be blowouts are, by definition, high scoring games. This implies that $E[\varepsilon_1 \varepsilon_2 | SPREAD, Over/Under] \neq 0$, and as such the SUR is the appropriate specification.

⁷ Evan and Noble [1992] and Paul and Weinbach [2002] use a likelihood ratio tests to test for the efficiency of the totals market due to skewness in the forecast errors. Our position is that skewness of the forecast errors itself should be considered problematic for the totals market to be a prediction market, and throws into question the general predictive power of the betting market.

Our joint test combines (1) and (2) above into a standard linear SUR framework where we subsequently formally test the joint hypotheses that $\beta_0 = \beta_2 = 0$ and that $\beta_1 = \beta_3 = 1$ (each of these tests has a chi-square distribution with two degrees of freedom). Our conjecture is that if the market is predictive then we will fail to reject the two hypotheses given above—this would imply that (strong-form) the betting house’s accurately predict game outcomes and/or that (weak-form) bettor beliefs about all facets of game outcomes are equally accurate. If we reject either of these hypotheses, however, this implies that betting houses (and, by implication, bettors) are not as good about predicting some game outcomes as others. Naturally, this implies that the sports betting market does not function well as a prediction market overall, while it may function well as a predictor of certain outcomes.

In addition to this, we also use the Breusch-Pagan test of the independence of the two equations as a test of the relationship *between* these two features of the betting market. As we noted earlier, there is a mechanical relationship between the betting line and the over/under if the sports betting market also functions as a prediction market. Since in any prediction market the two features fully characterize the contest (they jointly solve for the scores of both teams, their sum, and their difference) they must be correlated.⁸ In a belief structure, which need not be predictive, this need not be the case. For example, a bettor may have beliefs about a contest being high or low scoring that is independent from their belief about the margin of victory. It is in this sense that the sports betting market ceases to be predictive—features of the contest can be subject to different beliefs that need not fully characterize the contest. Formally, if σ_{12} is the covariance between the two regression equations we test the hypothesis that $\sigma_{12} = 0$ (that is, that the regressions are

⁸ While one possibility is to estimate a system of four equations (adding equations for the predictions of the two teams’ scores in addition to equations 1 and 2), the system would be over-determined by definition since it would have four equations and two unknowns. As such, tests of the independence of the regressions would be rejected simply due to the functional form employed and would not constitute a robust test for the predictive power of the betting market.

independent, which would imply that the market is not predictive). The Breusch-Pagan test statistic in this instance is $\lambda = N \left(\frac{\sigma_{12}^2}{\sigma_{11}\sigma_{22}} \right)$, which has a chi-square distribution with one degree of freedom.

4. Data

We exploit a unique source of data to test for market efficiency in this betting market. We use data from four major betting markets, NCAA football, NCAA basketball, NBA, and the NFL from the 2004-2010 seasons in addition to game outcomes from the sports books at Pinnacle. Table 1 presents the summary statistics for the data of the four sports. Using all games for each sport gives us large sample sizes, with 2,245 NFL games, 4,902 college football games, 9,763 NBA games, and 23,462 college basketball games. The summary statistics reveal several interesting differences between the professional and college sports. In general, professional sports have smaller average margins of victory than their college counterparts (3.3 versus 5.2 for basketball, 2.6 versus 3.4 for football), NFL games are lower scoring than their college counterparts (42 versus 52 respectively), while the opposite is true in basketball (196 for the NBA versus 137 for college). We note that part of the difference in basketball could be due to the fact that NBA games are eight minutes longer than college contests.

There are also interesting features of the betting markets for these sports. Football, both college and professional, exhibits larger differences between opening and closing betting lines than basketball. Since both college and professional football usually have a week between games, a longer time than the few days between basketball games, we would expect the difference in opening and closing lines to be larger in football than in basketball as the longer fallow period would allow for more information to be gathered, which could affect a bettor's beliefs about the upcoming game. The NBA has no change from the opening line to the closing line and college basketball only sees a

.01 increase. The NFL closing line is .06 lower than the opening and college football has an even bigger differential of .12. Although this shows some movement, we note that overall movement of betting lines and over/unders is small. This is consistent with research that has shown that line movements respond to new information (Avery and Chevalier 1999). Since most contests will feature limited new information from the time the initial betting line is set, we would expect relatively small movements, on average.

5. Testing for the Predictive Power of the Sports Betting Market

5.1 Independent Tests of Sports Betting Market Predictive Power

We begin our analysis with the single equation test (equations 1 and 2) to see how accurately the opening and closing betting lines predict the actual margin of victory and to see how closely the opening and closing over/under predict the sum of scores for each game. Table 2 shows the opening/closing line regressed against the actual margin of victory as well as the opening/closing over under against the actual sum of scores. In general, the sports betting lines accurately predict the margin of victory in games. Indeed, the opening and closing lines predict with virtually the same accuracy. As we argued earlier, if line movements reflect bettor responses to new information, and if new information is rare, then opening and closing lines should not differ substantially and their relationship to the final outcomes would also be similar. For example, in Table 2 the predictive power of the opening betting line in the NBA is quite similar to the predictive power of the closing line (coefficients of .953 versus .945) and we fail to reject the hypothesis that the two are equally predictive ($p > .01$). College basketball follows a similar pattern; the closing line coefficient of 1.002 is closer to 1 than the opening line coefficient of .962.

We next test whether or not the coefficient on each betting characteristic, such as the opening/closing line or over under, is equal to 1 in the four different sports reported in Table 2. The results show that the opening line fails to accurately predict the margin of victory in every sport but

the NFL. We reject the null hypothesis that the opening betting line coefficient is equal to one in the NBA ($p < .05$), college football ($p < .01$), and college basketball ($p < .01$). Only for the NFL do we fail to reject the null hypothesis that the opening betting line coefficient is equal to one ($p > .1$). The closing line is slightly more accurate as it correctly predicts margin of victory in the NBA, NCAAB, and the NFL. College football, however, is an exception. This may be a result of a bettor bias towards certain teams. Bettors could behave irrationally and pick highly ranked teams or heavy favorites regardless of the betting line. Bettors might also lack sufficient information to make an accurate bet and instead rely on their loyalty to a team, past results, or an inherent favoritism towards certain teams. This could also reflect the fact that betting markets are uneven—teams with larger fan bases may place a disproportionate number of bets in the college football market.

The results for the predictive power of over/unders are reported in Table 2. In all of the sports the opening over/under does not accurately predict the sum of scores of games. There are significant differences by sport. The opening over/under is far and away the most accurate in college football (coefficient of .890) when compared to the other three sports. College basketball is the least accurate (.185) followed by the NFL (.226) and the NBA (.318). The closing over/under is more accurate in both the NFL and college football (coefficients of .886 and 1.016, respectively) but it has very low predictive power in college basketball and the NBA (.390 and .301, respectively). Both the opening and closing over/unders are very inaccurate in basketball generally, especially when compared with those of the NFL and college football.

In stark contrast to the betting line, we nearly always reject the hypothesis that the over/unders are predictive of the actual sum of scores. When we test that the coefficient for the opening and closing over/under is equal to 1 for the NFL we reject the null hypothesis ($p < .01$) for both the opening and closing. We also reject the null hypothesis that the opening over/under is predictive in college football ($p < .01$), but with the closing over/under we fail to reject the null. In

the NBA we reject the null hypothesis with both the opening and closing over/under ($p < .01$ in both). College basketball follows the same pattern—we reject both hypotheses at the 1% level.

Table 2 also shows the results for tests of systematic bias in the predictive power of the sports betting market. We test whether the intercepts are equal to 0. If we fail to reject the null hypothesis then this would be evidence that the betting market is accurately predicting the actual characteristics of the game. In the NFL both the opening and closing over/unders appear to have systematic differences with the actual sum of scores, as we reject the null hypothesis for both at the 1% level. College football follows a slightly different pattern. We reject our null hypothesis for both the opening line and over/under ($p < .01$), which also implies systematic differences with the actual outcome. We fail to reject the null hypothesis in regards to the opening and closing over/under. Basketball sums of scores, both NBA and college, are poorly predicted by both the opening and closing over/unders (rejected at the 1% level). Overall in basketball the over/under has a systematic difference with the actual sum of scores while the betting line seems to accurately predict the margin of victory. In college football both the opening line and over/under have systematic differences with the actual game outcomes. The NFL is slightly different, the opening and closing over/under appear to be accurate while neither line accurately predicts actual game outcomes.

5.2 Joint Test of the Predictive Power of the Sports Betting Market

We now address our new, stronger test using the betting line and over/under in a joint test of the predictive power of the sports betting market. By introducing the over/under as a second characteristic of the game this new test gives us a better idea of the predictive power of the betting market. If we reject our null hypothesis that $\beta_{spread} = \beta_{over/under} = 1$ and $\alpha_{spread} = \alpha_{over/under} = 0$, we have very clear evidence that the sports betting market does not seek to predict the actual outcomes of games, but a function of bettor beliefs. We also perform the Breusch-Pagan Independence test to determine whether or not the two equations are related. As we noted earlier, a prediction market

implies a tight correlation between the betting line and over/unders. Our null hypothesis is that the two equations are independent from one another, $\sigma_{12} = 0$, which we take as evidence that the market fails to be internally consistent with a prediction market.

As Table 2 shows, in the NFL we reject both our hypotheses for the joint test of the opening betting line and over/under ($p < .01$) and we fail to reject the hypothesis that the two equations are independent. We see the same results for the closing line and over/under ($p < .01$). College football is slightly different—we reject our joint test at the 1% level for the opening betting line and over/under and reject it for the closing as well at the 1% level. We also reject the test of independence between the two equations ($p < .01$) for both the opening and closing values. This shows that the two equations are seemingly related, such that the correlation between the betting line and over/unders is relatively strong in college football.

In the NBA for the opening line and over/under we reject our joint test hypothesis at the 1% level and we fail to reject our null hypothesis that the equations are independent from each other. Each of our test shows that the NBA betting market is not a prediction market. We find the same results in college basketball, where we reject our joint test's null hypothesis ($p < .01$) for both the opening and closing values. In addition we fail to reject the null hypothesis that the equations are independent from each other. Overall, the joint tests for all four sports are inconsistent with the sports betting market performing well as a prediction market once the over/unders are included in the analysis.

6. Conclusion

Rather than testing the predictive power of a single aspect of the market, we combined information on both the betting line and the over/under to construct a joint test of the predictive power of the market. Given that the betting house has the same incentives for all market features,

our test explicitly took advantage of the potential distinction between prediction and profitability in the betting market. We found that, overall, the sports betting market is not as predictive in some features as it is for others. Our joint tests revealed that while the betting line is an accurate predictor of the margin of victory, the over/under is a poor predictor of the sum of scores in a contest. As such, tests of the joint predictive power of the market revealed that it functions poorly as a prediction market. Furthermore, we found that the two sets of market information were independent—the mechanical relationship that would exist if the market were predictive was not found in the data.

There are several possible explanations for our main finding. First, it could be the case that bettors on over/unders are less informed than bettors on betting lines. This would leave the bettors on over/unders to have more skewed belief distributions that betting houses would need to account for. Second, there could be inherent differences in the ability of betting houses and bettors to predict the sum of scores as opposed to their difference. In essence, there could be more widespread agreement that a given contest will be close as opposed to agreement on it being high or low scoring. This would imply that the bettor beliefs about betting lines would have thinner tails than those for over/unders. Third, it could be the case that there is significantly less money bet on the over/unders as opposed to the betting line. In this instance, the relatively small potential losses from incorrectly setting the over/under would cause the betting house to under-invest in developing precise information that would lead to more accurate predictions. While each of these possibilities is plausible, we note that each hinges on the betting house functioning as a profit maximizer as opposed to an accurate predictor of outcomes. While the exact mechanism behind our main results remains elusive, every reasonable explanation requires that the betting house be a profit maximizer whose interest in predicting the actual outcome is important only to the extent that it coincides with

profit maximization. Predictive power of some outcomes of interest may be an unintended outcome of profit maximization by betting houses.

Future research on betting markets should seek to incorporate additional information to form stronger tests of features believed to apply to betting markets overall. As we noted earlier, the few studies that do employ over/unders [Evan and Noble 1992, Gander, Zuber and Russo 1993, Paul and Weinbach 2002] have used NFL data, and only one of those studies used data for more than one season. Future studies should add additional evidence to the literature as well as gauge the magnitude of the effects. For example, studies of the effects of information flows on prices should exploit movements in the over/unders in addition to movements in the betting line. If betting lines move and over/unders do not (or vice versa) this type of asymmetric price adjustment could be exploited to derive and test the nature of the effects of information on prices, cognitive biases, and other relevant topics.

References

- Andrews, Rodney, Michael Sinkey and Trevon D. Logan (2011) "Identifying Confirmatory Bias: A Regression-Discontinuity Approach." Working Paper, The Ohio State University.
- Avery, Christopher and Judith Chevalier (1999). "Identifying Investor Sentiment from Price Paths: The Case of Football Betting." *Journal of Business* 72: 493-521.
- Card, David and Gordon Dahl (2011). "Family Violence and Football: The Effect of Unexpected Emotional Cues on Violent Behavior." *Quarterly Journal of Economics* 126(1): 103-143.
- Cowgill, B., J. Wolfers and E. Zitzewitz (2009). "Using Prediction Markets to Track Information Flows: Evidence from Google." Working Paper, Dartmouth College.
- Dare, William and Steven Holland (2004). "Efficiency in the Betting Market: Modifying and Consolidating Research Methods". *Applied Economics*, 36(1): 9-15.
- Dare, William and Scott McDonald (1996). "A Generalized Model for Testing the Home and Favorite Team Advantage in Point Spread Markets". *Journal of Financial Economics*. 40: 295-318.
- Evan, W.E. and N.R. Noble (1992). "Testing Efficiency in Gambling Markets." *Applied Economics* 24: 85-88.
- Fair, Ray and John Oster (2007). "College Football Rankings and Market Efficiency." *Journal of Sports Economics*, 8(1): 3-18.
- Gandar, John, Richard Zuber, and Ben Russo (1993). "Testing Efficiency in Gambling Markets: A Comment." *Applied Economics*, 25: 937-943.
- Gandar, John, Richard Zuber, Thomas O'Brien, and Ben Russo (1998). "Testing Rationality in the Point Spread Betting Market." *Journal of Finance*, 43(4): 995-1008.
- Golec, Joseph and Maury Tamarkin (1991). "The Degree of Inefficiency in the Football Betting Market." *Journal of Financial Economics*, 30: 311-323.
- Gray, Phillip and Stephen Gray (1997). "Testing Market Efficiency: Evidence From the NFL Sports Betting Market." *Journal of Finance*. 52(4): 1725-1737.
- Kilby, Jim, Jim Fox, and Anthony Lucas (2002). *Casino Operations Management*, Second Edition. Wiley: Hoboken.
- Levitt, Steven D. (2004). "Why Are Gambling Markets Organised so Differently From Financial Markets?" *Economic Journal* 114(495): 223-246.
- Ottaviani, Marco and Peter Sorensen (2007). "Aggregation of Information and Beliefs in Prediction Markets." Working Paper, *Northwestern University*.

- Ottaviani, Marco and Peter Sorensen (2008). "Aggregation of Information and Beliefs in Prediction Markets: Lessons for Asset Pricing." Working Paper, *University of Copenhagen*.
- Pankoff, Lyn D. (1968). "Market Efficiency and Football Betting." *Journal of Business* 41 (April), pp. 103-114.
- Paul, Rodney, and Andrew Weinbach (2002). "Market Efficiency and a Profitable Betting Rule: Evidence from Totals on Professional Football". *Journal of Sports Economics*. 3(3): 256-263.
- Paul, Rodney, Andrew Weinbach, and Patrick Coate (2007). "Expectations and Voting in the NCAA Football Polls". *Journal of Sports Economics*. 8(4): 412-424.
- Paul, Rodney, Andrew Weinbach, and Chris Weinbach (2003). "Fair Bets and Profitability in College Football Gambling." *Journal of Economics and Finance*. 27(2): 236-242, 2003.
- Rees, Daniel I. and Kevin T. Schnepel (2009). "College Football Games and Crime." *Journal of Sports Economics*, Vol. 10, pp. 68-86.
- Rhode, Paul and Koleman Strumpf (2008). Historical Presidential Betting Markets. *Journal of Economic Perspectives*. 18(2): 127-142.
- Roxborough, R. and M. Rhoden (1998). *Sports Book Management: A Guide for the Legal Bookmaker*. Las Vegas: Las Vegas Sports Consultants.
- Sauer, R. (1998). "The Economics of Wagering Markets." *Journal of Economic Literature* 36: 2021-2064.
- Sauer, Raymond, Vic Brajer, Stephen Ferris, and M. Wayne Marr (1998). "Hold Your Bets: Another Look at the Efficiency of the Gambling Market for National Football League Games." *Journal of Political Economy*. 96(1): 206-213.
- Sinkey, Michael (2011). "How do Experts Use Bayes' Rule? Evidence from an Incentive-Free Environment." Unpublished Manuscript, The Ohio State University.
- Sinkey, Michael and Trevon D. Logan (2011). "Does the Hot Hand Drive the Market? Evidence from College Football." Working Paper, The Ohio State University.
- Snowberg, Erik, Justin Wolfers, and Eric Zitzewitz (2007). "Partisan Impacts on the Economy: Evidence from Prediction Markets and Close Elections." *Quarterly Journal of Economics*. 122(2): 807-829.
- Song, C., B.L. Boulier, and H.O. Stekler (2007). "The Comparative Accuracy of Judgmental and Model Forecasts of American Football Games." *International Journal of Forecasting* 23: 405-413.
- Wolfers, Justin and Eric Zitzewitz (2004). "Prediction Markets." *Journal of Economic Perspectives*. 18(2): 107-126.

- Wolfers, Justin and Eric Zitzewitz (2007). "Interpreting Prediction Market Prices as Probabilities."
Working Paper, *Stanford GSB*.
- Zellner, A. (1962). "An Efficient Method of Estimating Seemingly Unrelated Regressions and Tests
for Aggregation Bias." *Journal of the American Statistical Association* 57: 348-368.
- Zuber, Richard, John Gandar and Benny Bowers (1985). "Beating the Spread: Testing the Efficiency
of the Gambling Market for National Football League Games." *Journal of Political Economy*.
93(4): 800-806.

Appendix: Proof of the Profit Maximization of Risk Minimization

Given the risk minimization motive for betting houses, we start by establishing a minimum acceptable threshold for making profit in the betting market. With some probability, p , the bet will be successful, and since this market allows for only two outcomes, the remaining probability, $1-p$, captures the instance when the bet is not successful. The revenue of a bettor is multiplied by the bet size, B .⁹ Also every bet has a fixed amount of money that is given to the betting house, the *vigorish*. The threshold for any particular bet is thus a probability of success that exceeds the sum of the probability of failure and the transaction cost. After accounting for the vigorish, it must be the case that

$$pB - cB \leq (1 - p)B \quad (\text{A1})$$

and

$$(1 - p)B - cB \leq pB \quad (\text{A2})$$

If $c = 0$, there were no vigorish, then a risk-neutral bettor would be indifferent between betting that a team would beat (or, conversely, lose to) the line if the probability that a team beat the spread was equivalent to the probability that a team lost to the line. (This would represent a coin flip.)

Otherwise, the bettor would bet on the outcome that was more likely.¹⁰

The left hand side of (A1) represents the profit made by betting on a team to beat the betting line and the bet being successful. The left hand side of (A2) represents the profit made by betting against a team to beat the betting line and the bet being successful. The right hand sides of both equations represent the loss realized from an unsuccessful bet. Factoring out B and rearranging

terms yields the key relationship: $p \in \left\{ \frac{1-c}{2}, \frac{1+c}{2} \right\}$. A bettor will place a bet on a team to beat the

⁹ Betting houses may cap the size of the bet; for example, Kilby, Fox and Lucas [2002] suggest that gambling houses should cap the size of the bet at \$2,000 for college football, although betting houses may choose to pursue higher limits if they feel that bettors are particularly uninformed. However, this would not prevent bettors from having others bet for them. Thus, we model bet size to be arbitrarily large, although in principle the bet size is capped.

¹⁰ Note that since all participants in the betting market pay the vigorish, they are inherently risk-loving.

spread if she believes that the team has a greater than $\frac{1+c}{2}$ percent chance of beating the betting line. Conversely, a bettor will bet against a team beating the spread if she believes that the team has less than a $\frac{1-c}{2}$ percent chance of beating the betting line.¹¹

Betting houses aggregate this betting behavior to make profit. From the perspective of a betting house, a betting line l is efficient if it guarantees profit in expectation. More formally, the betting house chooses l to satisfy both

$$E(p_i^b | l_t, \Omega_t(l_t, I_t))\mathcal{B} - E(1 - p_i^b | l_t, \Omega_t(l_t, I_t))\mathcal{B} + cB \geq 0 \quad (\text{A3})$$

and

$$E(1 - p_i^b | l_t, \Omega_t(l_t, I_t))\mathcal{B} - E(p_i^b | l_t, \Omega_t(l_t, I_t))\mathcal{B} + cB \geq 0 \quad (\text{A4})$$

where l_t is the line at time t , Ω_t is the distribution of bettor beliefs about the likelihood of beating a betting line, and I_t is the information set of the bettors at time t , which includes all relevant information a bettor chooses to use when betting on a particular team to beat the betting line. This information set includes any information that bettors may use, such as injuries, whether a team has a strong tradition, opponent strength, and past results against the line. From a betting house's perspective, any $p_i^b \in \left\{ \frac{1-c}{2}, \frac{1+c}{2} \right\}$ guarantees that (A3) and (A4) hold. Note further that for common distributions of beliefs, such as uniform and normally distributed beliefs, setting a line close to the median is efficient—it guarantees that bettors will be indifferent between betting on either side of the betting line, which in expectation would give equal amounts of money on either side of the bet.

It is easy to show that risk-minimizing behavior is profit maximizing for the betting house. Suppose that, for a given line l_t the cumulative probability distribution associated with individuals

¹¹ For most betting markets, this number is 52.4 percent.

who believe the line is too low is $f(l_t)$ and the cdf associated with individuals who believe the line is too high is $g(l_t)$, such that $f(l_t) + g(l_t) = 1$. Furthermore, suppose that the betting house has the same beliefs, $q(l_t)$, as the bettors, such that $f(l_t) = q(l_t)$. For a given line, the betting house's profit function becomes:

$$f(l_t)B + g(l_t)B + c(f(l_t)B + g(l_t)B) - 2q(l_t)f(l_t)B - 2(1 - q(l_t))g(l_t)B \quad (\text{A5})$$

Here, setting l_t such that $f(l_t) = 0.5$ is both risk minimizing and profit maximizing.

Table 1
Summary Statistics for Betting Market Data

	NFL Sample				NCAA Football Sample		
	Mean	Std. Dev.	N		Mean	Std. Dev.	N
Score of Team	22.09	10.13	2,219	Score of Team	28.77	14.3	4,800
Score of Opponent	19.48	10.14	2,219	Score of Opponent	23.38	13.36	4,800
Sum of Scores	41.57	14.20	2,219	Sum of Scores	52.14	18.34	4,800
Margin of Victory	2.61	14.47	2,219	Margin of Victory	5.39	20.72	4,800
Over/Under - Opening	44.46	44.37	2,233	Over/Under - Opening	51.72	8.12	3,931
Point Spread - Opening	2.89	6.60	2,240	Point Spread - Opening	5.05	13.49	4,889
Over/Under - Closing	44.24	45.20	2,231	Over/Under - Closing	51.64	7.95	3,929
Point Spread - Closing	2.83	6.78	2,240	Point Spread - Closing	5.17	13.69	4,889

Table 1 (cont.)
Summary Statistics for Betting Market Data

	NBA Sample				NCAA Basketball Sample		
	Mean	Std. Dev.	N		Mean	Std. Dev.	N
Score of Team	99.51	12.85	9,646	Score of Team	71.23	12.41	23,421
Score of Opponent	96.18	12.66	9,646	Score of Opponent	66.06	12.14	23,421
Sum of Scores	195.69	21.75	4,800	Sum of Scores	137.29	20.49	23,421
Margin of Victory	3.33	13.32	9,646	Margin of Victory	5.17	13.53	23,421
Over/Under - Opening	196.54	20.71	8,687	Over/Under - Opening	136.99	23.62	19,227
Point Spread - Opening	3.43	5.68	9,672	Point Spread - Opening	5.32	8.26	23,520
Over/Under - Closing	196.53	21.69	8,681	Over/Under - Closing	136.87	16.66	19,221
Point Spread - Closing	3.43	5.88	9,672	Point Spread - Closing	5.33	8.19	23,520

Table 2
 OLS and SUR Tests of Betting Market Predictive Power
 Panel A -- NFL Sample

	<u>Regression Type and Dependent Variable</u>							
	OLS - Opening		OLS - Closing		SUR - Opening		SUR- Closing	
	I SUM	II MOV	III SUM	IV MOV	V SUM	VI MOV	VII SUM	VIII MOV
Over/Under	0.226*** [0.0301]		0.886** [0.0556]		0.226*** [0.0301]		0.886*** [0.0556]	
Point Spread		1.046 [0.0517]		1.043 [0.0493]		1.050*** [0.0516]		1.050*** [0.0493]
Constant	32.26*** [1.273]	-0.168 [0.314]	5.495** [2.281]	0.093 [0.308]	32.27*** (1.272)	-0.186 (0.314)	5.495 (2.280)	-0.126 (0.308)
Observations	2,216	2,219	2,214	2,219	2,216	2,216	2,214	2,214
<u>Hypothesis Tests (Test Statistics):</u>								
Breusch-Pagan Independence					0.195		0.367	
Over/Under=Point Spread=1					664.57***		5.2*	
Constants Both Equal 0					643.59***		5.98*	

Standard errors in brackets
 *** p<0.01, ** p<0.05, * p<0.1

Critical value for Breusch-Pagan Test (Chi square 1 degree of freedom) = 2.71 (p=0.1)
 Critical value for all other test (Chi square 2 degrees of freedom) = 4.61 (p = 0.1)

Table 2 (cont.)
 OLS and SUR Tests of Betting Market Predictive Power
 Panel B -- NCAA Football Sample

	<u>Regression Type and Dependent Variable</u>							
	OLS - Opening		OLS - Closing		SUR - Opening		SUR - Closing	
	I SUM	II MOV	III SUM	IV MOV	V SUM	VI MOV	VII SUM	VIII MOV
Over/Under	0.890*** [0.0334]		1.016 [0.0341]		0.884*** [-0.0333]		1.010*** [0.0340]	
Point Spread		-1.004*** [0.0167]		-1.012*** [0.0163]		1.002*** [-0.019]		1.018*** [-0.0186]
Constant	5.750*** [1.747]	0.255*** [0.241]	-0.703 [1.781]	0.047 [0.239]	6.096*** (1.744)	-0.153 (0.268)	-0.386 (1.779)	-.448* (0.265)
Observations	3,841	4,793	3,839	4,793	3,841	3,841	3,839	3,839
<u>Hypothesis Tests (Test Statistics):</u>								
Breusch-Pagan Independence					8.305***		8.152***	
Over/Under=Point Spread=1					12.23***		1.06	
Constants Both Equal 0					12.57***		2.9	

Standard errors in brackets
 *** p<0.01, ** p<0.05, * p<0.1

Critical value for Breusch-Pagan Test (Chi square 1 degree of freedom) = 2.71 (p=0.1)
 Critical value for all other test (Chi square 2 degrees of freedom) = 4.61 (p = 0.1)

Table 2 (cont.)
 OLS and SUR Tests of Betting Market Predictive Power
 Panel C -- NBA Sample

	<u>Regression Type and Dependent Variable</u>							
	OLS - Opening		OLS - Closing		SUR - Opening		SUR - Closing	
	I SUM	II MOV	III SUM	IV MOV	V SUM	VI MOV	VII SUM	VIII MOV
Over/Under	0.318*** [0.0107]		0.301*** [0.0103]		0.318*** (0.0107)		0.301*** [0.0103]	
Point Spread		0.953** [0.0218]		0.945*** [0.0210]		.962*** [.00964]		0.956*** (0.0217)
Constant	134.0*** [2.122]	0.0438 [0.145]	137.4*** [2.030]	0.0803 [0.143]	134.0*** (2.122)	-0.0352 (0.152)	137.4*** [2.030]	0.0152 (0.150)
Observations	8,648	9,621	8,639	9,621	8,648	8,648	8,639	8,642
<u>Hypothesis Tests (Test Statistics):</u>								
Breusch-Pagan Independence					0.166		0.434	
Over/Under=Point Spread=1					4033.07***		4639.50***	
Constants Both Equal 0					3989.42***		4584.28***	

Standard errors in brackets
 *** p<0.01, ** p<0.05, * p<0.1

Critical value for Breusch-Pagan Test (Chi square 1 degree of freedom) = 2.71 (p=0.1)
 Critical value for all other test (Chi square 2 degrees of freedom) = 4.61 (p = 0.1)

Table 2 (cont.)
 OLS and SUR Tests of Betting Market Predictive Power
 Panel D -- NCAA Basketball Sample

	<u>Regression Type and Dependent Variable</u>							
	OLS - Opening		OLS - Closing		SUR - Opening		SUR - Closing	
	I SUM	II MOV	III SUM	IV MOV	V SUM	VI MOV	VII SUM	VIII MOV
Over/Under - Opening	0.185*** [.0061]		.390*** [.0060]		0.185*** [.0061]		0.390*** (0.00847)	
Point Spread - Opening		0.973*** [0.0086]		1.01 [0.00847]		.962*** [.00964]		1.002*** (0.00962)
Constant	111.9*** [0.852]	-0.0032 [0.0845]	83.77*** [0.461]	-0.2 [1.168]	111.9*** (0.852)	0.058 (0.0930)	83.78*** (1.168)	-0.157* (0.0921)
Observations	19,105	23,346	19,100	19,099	19,101	19,101	19,095	19,095
<u>Hypothesis Tests (Test Statistics):</u>								
Breusch-Pagan Independence					1.352		0.625	
Over/Under=Point Spread=1					17721.98***		5184.75***	
Constants Both Equal 0					17245.27***		5145.61***	

Standard errors in brackets
 *** p<0.01, ** p<0.05, * p<0.1

Critical value for Breusch-Pagan Test (Chi square 1 degree of freedom) = 2.71 (p=0.1)
 Critical value for all other test (Chi square 2 degrees of freedom) = 4.61 (p = 0.1)