

Informational efficiency and behaviour within in-play prediction markets

Article

Accepted Version

Creative Commons: Attribution-Noncommercial-No Derivative Works 4.0

Angelini, G., De Angelis, L. and Singleton, C. ORCID: <https://orcid.org/0000-0001-8247-8830> (2022) Informational efficiency and behaviour within in-play prediction markets. *International Journal of Forecasting*, 38 (1). pp. 282-299. ISSN 0169-2070 doi: 10.1016/j.ijforecast.2021.05.012 Available at <https://centaur.reading.ac.uk/98329/>

It is advisable to refer to the publisher's version if you intend to cite from the work. See [Guidance on citing](#).

To link to this article DOI: <http://dx.doi.org/10.1016/j.ijforecast.2021.05.012>

Publisher: Elsevier

All outputs in CentAUR are protected by Intellectual Property Rights law, including copyright law. Copyright and IPR is retained by the creators or other copyright holders. Terms and conditions for use of this material are defined in the [End User Agreement](#).

www.reading.ac.uk/centaur

CentAUR

Central Archive at the University of Reading

Reading's research outputs online

Informational efficiency and behaviour within in-play prediction markets*

Giovanni Angelini[†], Luca De Angelis[†], Carl Singleton[‡]

[†]Department of Economics, University of Bologna

[‡]Department of Economics, University of Reading

May 2021

Accepted at International Journal of Forecasting

Abstract

Studies of financial market informational efficiency have proven burdensome in practice, because it is difficult to pinpoint when news breaks and is known by some or all the participants. We overcome this by designing a framework to detect mispricing, test informational efficiency and evaluate the behavioural biases within high-frequency prediction markets. We demonstrate this using betting exchange data for association football, exploiting the moment when the first goal is scored in a match as major news that breaks cleanly. There are pre-match and in-play mispricing and inefficiency in these markets, explained by reverse favourite-longshot bias (favourite bias). The mispricing tends to increase when the major news is a surprise, such as a goal scored by a longshot team late in a match, with the market underestimating their chances of going on to win. These results suggest that, even in prediction markets with large crowds of participants trading state-contingent claims, significant informational inefficiency and behavioural biases can be reflected in prices.

Keywords: Market efficiency; Favourite-longshot bias; Mispricing; Behavioural bias; Betting strategy

JEL codes: G14, G41, L83

*l.deangelis@unibo.it, +39 051 20 9 8149, g.angelini@unibo.it, +39 0541 434 168, Dipartimento di Scienze Economiche, Piazza Scaravilli 2, Bologna, Italy; c.a.singleton@reading.ac.uk, +44 118 378 7864 (corresponding author), Department of Economics, University of Reading, Whiteknights Campus, RG6 6EL, UK.

With thanks to three reviewers, the associate editor, Iliyan Georgiev, Romain Gauriot and J. James Reade for comments and advice, as well as participants at the 12th International Conference on Computational and Financial Econometrics (Pisa, December 2018), the 39th International Symposium on Forecasting (Thessaloniki, June 2019), the 1st Reading Football Economics Workshop (Reading, September 2019), and the 13th International Conference on Computational and Financial Econometrics (London, December 2019). Special thanks to Wasim Rehman for sharing the Betfair Exchange data.

I Introduction

In the past fifty years, many researchers have attempted to test Fama's (1965; 1970) Efficient Market Hypothesis (EMH), using a variety of methods and contexts. Studies have not only looked at whether asset prices reflect all relevant historical information (weak-form efficiency) but also whether the arrival of new information is immediately and fully incorporated (semi-strong efficiency).¹ The answers to these questions are important both practically and theoretically. Not least, if markets are inefficient, then it implies that better informed economic agents can gain at the expense of the less well informed. Moreover, if the pricing mechanism in a market is inefficient, then an asset's price will not reflect its fundamental value, complicating any economic analysis. Investigations of financial market efficiency and related behavioural biases have proved somewhat burdensome in practice because, among other things, it is problematic to identify precisely the point at which news breaks and is known by some or all of the agents involved. It is typically hard to believe that efficiency tests on such markets are not affected by information leakages and asymmetries, which are unknown to the econometrician.

To overcome these challenges and gain insight on the validity of the EMH, a large literature has studied prediction or betting markets, typically those relating to sports events.² Unlike conventional financial markets, sports betting provides 'real world laboratories' in which to test the EMH and study departures from it, as participants are generally regarded as being well-informed, motivated, experienced and, most importantly, breaking news is usually reported cleanly, in a way that is easy for the participants to share and process.³ The assets (bets) in these markets have defined end points upon which their values become certain, which is typically not the case when evaluating financial securities pricing (Thaler and Ziemba, 1988; Ziemba, 2020). The main findings from the

¹See Malkiel (2003); Vaughan Williams (2005); Lim and Brooks (2011) for comprehensive reviews of this literature.

²See among many others: Thaler and Ziemba (1988); Pope and Peel (1989); Golec and Tamarkin (1991); Kuypers (2000); Levitt (2004); Smith et al. (2006); Snowberg and Wolfers (2010); Page and Clemen (2012); Franck et al. (2013); Brown (2013); Croxson and Reade (2014); Brown (2015); Deutscher et al. (2018); Angelini and De Angelis (2019).

³Some studies have attempted to replicate these conditions in the laboratory, such as Plott and Sunder (1988); Plott et al. (2003); List (2004) and Koessler et al. (2012), though these naturally lack realism and are open to the standard critiques of the settings being artificial.

literature studying betting markets, however, show mixed evidence both on the degree to which these markets are efficient and, in some cases, the potential behavioural biases which might account for why they are inefficient.

The main contribution of this paper is a new approach to test the semi-strong efficiency of prediction markets when ‘in-play’ trading is allowed, i.e. after the event has begun (e.g. between kick-off and the final whistle in a game of football). This approach is based on the [Mincer and Zarnowitz \(1969\)](#) forecast evaluation framework, and, in two ways, substantially extends a previous application by [Angelini and De Angelis \(2019\)](#), which only applied to testing the weak-form efficiency of bookmaker betting markets. First, we show how the approach applied to bookmakers in [Angelini and De Angelis \(2019\)](#) can be extended to the general case of prediction markets, to both detect mispricing and test whether prices are efficient just before an event begins. Second, we extend this to provide ways of detecting bias and testing market efficiency in the aftermath of in-play news which is major, relevant and plausibly arrives cleanly to all participants. We also discuss how this framework of regression models and hypothesis testing provides a practical approach to describe and evaluate some of the possible behavioural biases present in these markets, such as the well-know favourite-longshot bias, sometimes just referred to as the longshot bias (e.g. [Ottaviani and Sørensen, 2008](#); [Snowberg and Wolfers, 2010](#); [Newall and Cortis, 2021](#)), the home bias (e.g. [Levitt, 2004](#)), under or overreaction to major news (e.g. [De Bondt and Thaler, 1985, 1990](#)) and confirmation bias (e.g. [Wason, 1960](#); [Rabin, 1998](#)).

The secondary contributions of this paper come from an illustrative application of our methodological framework. Using a high-frequency data set of pre-match and in-play odds (prices), we study the final result markets of 1,004 English Premier League (EPL) association football matches. These data come from the Betfair Exchange, which is the world’s largest online betting exchange and prediction marketplace. As our primary focus is on the impact of news arriving on the market, betting exchange markets, where customers can bet against each other directly, are natural candidates to test for informational efficiency and to evaluate behavioural biases among the participants. Not least, betting exchange odds have more predictive power than corresponding bookmaker odds ([Smith et al., 2009](#); [Franck et al., 2010](#); [Reade, 2014](#)). The efficiency of these exchange markets has been studied before,

notably by [Croxson and Reade \(2014\)](#), who investigated the market reaction to goals scored just before half-time in a football match. This approach allowed the authors to separate the major news of a goal being scored from the continual flow of minor news. [Croxson and Reade \(2014\)](#) found that these markets were semi-strong efficient, as prices generally updated swiftly and fully following a goal. Our methodological framework allows us to add to this previous analysis and other literature in a number of ways.

First, we study the efficiency of the pre-match result market, finding significant evidence of mispricing, which can be explained by a *reverse* favourite-longshot bias (or favourite bias). In other words, the team which the market did not expect to win was significantly underpriced. Studies of the football match result odds offered by bookmakers have tended to find the opposite, and are generally consistent with the majority of other examples from sports wagering analyses since the seminal study on horse-racing by [Ali \(1977\)](#).⁴ This contrasting result between betting exchange markets and bookmakers is novel, and could suggest that too much in the past has been inferred about the preferences and biases of risk-takers from studying bookmaker prices. The behaviour of bookmakers has largely not been modelled in that literature, with [Levitt \(2004\)](#) being a notable exception, whereas betting exchanges involve the more general trade of state-contingent claims between market participants. Other explanations of the favourite bias that we find could include bettors on exchange markets being more risk averse compared with those using bookmakers (less risk loving, e.g. [Ottaviani and Sørensen, 2015](#)) and being better informed (fewer casual bettors, e.g. [Smith et al., 2006, 2009](#); [Bruce et al., 2009](#)).⁵ We use our findings to carry out a simple strategy of systematically betting only on significantly underpriced matches, both in sample and out of sample. In this way, substantial gross returns on investment of around 50% could have been earned by market participants who exploited the presence of the reverse favourite-longshot bias, both in the sample period and over 5 years later, out-of-sample, in the 2019/20 English Premier League season.

⁴See [Cain et al. \(2000\)](#); [Deschamps and Gergaud \(2007\)](#) for English football, as well as [Angelini and De Angelis \(2019\)](#) more generally for European professional football.

⁵Ours is by no means the only study to find a reverse favourite-longshot bias. For example, [Woodland and Woodland \(1994\)](#) found something similar in prediction markets for US baseball, though insufficient in that case to imply significant market inefficiency. [Newall and Cortis \(2021\)](#) suggest in a review of the literature that betting markets with fewer outcomes tend to produce a favourite bias (e.g., team sports), whereas a longshot bias appears in markets with many outcomes (e.g., horse racing or golf). Our results roughly fit within this dichotomy but studies of association football match odds from bookmakers generally do not.

Second, we evaluate whether the in-play prices of a football match result are efficient following the first goal, which is major news that implies a large change in the probability of one team winning. Specifically, we identify the combinations of when a first goal was scored and the pre-match odds which implied that the price after was significantly mispriced. Consistent with [Croxson and Reade's \(2014\)](#) study of goals around the half-time break, we find that the win odds for the team playing at home (away) were semi-strong efficient when the first goal in the match was scored by the away (home) team. In other words, at whatever point in the game the first goal arrived and independent of the pre-match odds, the prices afterward fully incorporated the new information, responded immediately, and did not drift. However, we find significant evidence that the home (away) win odds were mispriced in the period at least 5 minutes after the first goal was scored by the home (away) team, such that the markets exaggerated the decisiveness of the first goal in determining the final match outcome when it was scored by a favourite team, especially early in the match, but underestimated its decisiveness when scored by a longshot team, especially late in the match. This mispricing was strongest twenty seconds after the goal, but still remained significant as much as five minutes later. Again, applying a simple betting strategy based on the pre-match odds and the time when the goal was scored, we show that market participants could have systematically exploited these facts to make substantial returns on investment. Thus, there is evidence that these markets were inefficient in the aftermath of common instances of major news.

Third, we test for and evaluate the behavioural biases suggested by how the betting exchange markets reacted to major news. Depending both on whether the pre-match odds suggested favourite bias, longshot bias or no bias at all for the team playing at home or away, and depending on which team scored the first goal, we test whether this major news constituted a significant change in the degree and nature of mispricing. News which arrived early in the event and which reflected expectations did not cause significant revisions in the bias implied by market prices. However, when the first goal was scored later, in cases when the pre-match odds reflected favourite bias, the market significantly adjusted toward unbiasedness. Conversely, when the odds were generally priced correctly before the first goal, mispricing occurred after. If a longshot scored the first goal, particularly if this happened

toward the end of a match, then the initial mispricing tended to be amplified, i.e. the reverse favourite-longshot bias was increased.⁶ Overall, we find that how prices and expectations changed following major news on the Betfair Exchange was consistent with a pattern of markets responding towards efficiency following expected major news but inefficiently when this news came as a surprise. In Premier League football terms, this is equivalent to bettors generally underestimating the probability of Burnley F.C., a less-fancied team, going on to win after scoring the first goal late in the game against Manchester City F.C., a more fancied team.

The remainder of the paper is organised as follows: Section II outlines a general approach to testing prediction market efficiency, both for the outcome of events before they have begun and in the aftermath of major in-play news; Section III describes a data set of football match prediction markets; Section IV applies the testing approach to these data and analyses the degree of market efficiency on the Betfair exchange; and Section V concludes.

II Testing the Informational Efficiency of In-play Prediction Markets

In this section, we outline a forecast-based approach to test the efficiency of prediction markets. This approach is directed toward addressing the following two main questions in Sections II.i and II.ii:

1. Are market prices efficient just before an event begins?
2. Are market prices efficient in the aftermath of relevant news, which occurs between the beginning of an event and its end?

The former question aims at discovering whether pre-event prediction market odds are mispriced, where the specific events we will apply this approach to later are the outcomes of association football matches. The latter question focuses on the reaction of market participants to the arrival of new and important news, which should almost certainly affect

⁶There is mixed evidence from conventional financial markets on whether markets under or overreact to ‘surprise’ news, depending on the type of news. For example, [Brooks et al. \(2003\)](#) show that markets overreact to industrial disasters or the death of a CEO, whereas [Chan \(2003\)](#) finds evidence that investors underreact to headline-making news about a company.

expectations about an event’s final outcome. In our application, the major news we will study are the instances of the first goal being scored by one of the teams playing in a match; football is a relatively low-scoring game, with the most common outcome of a match being 1-1, i.e. one goal scored by each of the two teams, and with typical goal-scoring rates for each team lying between one and two goals per match.⁷ These prediction markets for final event outcomes are active throughout an event’s duration, with in-play trading. Therefore, we study the new price equilibrium reached by the market after news arrives and the evolution of prices afterward. Moreover, in Section II.iii we discuss the potential behavioural biases of prediction market participants, such as the well-known favourite-longshot bias (e.g. Ottaviani and Sørensen, 2008), and we interpret how deviations from no bias following major news could be related to any underreaction or overreaction by the participants. We also address the duration of mispricing in prediction markets, studying whether any deviation from market efficiency persists or is absorbed quickly.

II.i Are market prices efficient just before an event begins?

To evaluate whether prices in prediction markets (or betting exchanges) are set efficiently, or whether there is evidence of bias, we extend the analysis of betting markets by Angelini and De Angelis (2019). Let $p_{i,0}$ be the implied probability of an outcome of event i observed pre-event (i.e. at time $t = 0$; or at kick-off in football terms) and let $(p_{i,0})^{-1}$ be the corresponding pre-event prediction market price (decimal odds).⁸ For a given outcome (e.g. a win by the home side in a football match), we consider the market’s pre-event forecast error, computed as $e_{i,0} = y_i - p_{i,0}$, where $y_i = 1$ if i ended with that specific outcome (e.g. a home win) and 0 otherwise (e.g. a draw or an away win). Then, following an approach akin to the Mincer and Zarnowitz (1969) forecast evaluation regression, consider the following model:

$$e_{i,0} = \gamma_0 + \beta p_{i,0} + u_{i,0} , \quad (1)$$

⁷Author calculations with thanks to J. James Reade, using the entire history of football matches listed on Soccerbase.com, i.e. from 511,759 recorded matches up to 8 January, 2019.

⁸See Wolfers and Zitzewitz (2006); Manski (2006) for discussions on the interpretation of prediction market prices as probabilities.

where $u_{i,0}$ is an i.i.d. error term. As [Ioannidis and Peel \(2005\)](#) show that forecast errors can exhibit heteroskedasticity under the null hypothesis of market efficiency, we estimate Equation (1) by Weighted Least Squares (WLS), where the $n \times n$ weighting matrix is diagonal with elements $\sigma_{1,0}^2, \dots, \sigma_{n,0}^2$ and n denotes the total number of events studied. Since y_i is a Bernoulli random variable, its variance, $\sigma_{i,0}^2$, can be approximated by $p_{i,0}(1 - p_{i,0})$.

The estimation results of (1) are then used to assess whether prediction markets are unbiased just before events began. In particular, a rejection of the null hypothesis,

$$H_0 : \gamma_0 = \beta = 0 , \quad (2)$$

implies that, conditional on all the information available to market participants regarding event i , the expected value of the forecast error is not zero. Specifically, $E(e_{i,0}|\mathcal{I}_{i,0}) \neq 0$, where $\mathcal{I}_{i,0}$ is the general information set that participants are using, including the implied probabilities ($p_{i,0}$). If the null hypothesis (2) is not rejected, then the odds are set efficiently by the market participants and no bias is detected. If the null is rejected because of $\gamma_0 \neq 0$, then the odds imply significant forecast errors on average, perhaps because they are biased toward one outcome type over another. Similarly, we would anticipate rejecting the null if the markets were not competitive, such that on average one side of the market is earning significant profits in expectation, as is implied in the case of traditional bookmakers, for example, whereby γ_0 would then capture their expected profit margin (overround or vigourish in betting terms). A rejection of the null hypothesis (2) can also imply a significant relationship between the forecast error and the odds, $\beta \neq 0$, which in turn implies that the forecasts made by the market participants are biased for certain values of the pre-event implied probability. In other words, the odds before an event starts are mispriced, suggesting the presence of informational inefficiency.

In the spirit of [Angelini and De Angelis \(2019\)](#), we investigate whether any biases implied by the rejection of the null hypothesis (2) are large enough to generate market inefficiency, i.e. we test whether these deviations from no bias are significantly different from zero. In particular, consider the estimated parameter values of Equation (1), $\hat{\theta}_0 = (\hat{\gamma}_0, \hat{\beta})'$, interpolate over all possible probability values, $p_G \in (0, 1)$, and derive the ‘efficiency curve’

as:

$$\hat{G}(p_G) = \hat{\gamma}_0 + \hat{\beta}p_G . \quad (3)$$

The related confidence bands are then computed as:

$$\begin{aligned} CI_0 &= [\underline{CI}_0, \overline{CI}_0] , \\ &= \left[\hat{G}(p_G) - z_{\alpha/2} \text{ s.e. } \left(\hat{G}(p_G) \right), \hat{G}(p_G) + z_{\alpha/2} \text{ s.e. } \left(\hat{G}(p_G) \right) \right] , \end{aligned} \quad (4)$$

where $z_{\alpha/2}$ is the $100(1 - \alpha/2)$ -th percentile of the standard normal distribution, $\text{s.e.} \left(\hat{G}(p_G) \right) = \left[\nabla \hat{G}(p_G)' V_{WLS} \nabla \hat{G}(p_G) \right]^{1/2}$, $\nabla \hat{G}(p_G) = (1, p_G)'$ is the gradient, V_{WLS} is the variance of the WLS estimator, and the error term in (1) is assumed to be i.i.d.

The confidence intervals in (4) are useful as a procedure to test market efficiency and to evaluate prediction market bias. We define the probability ranges where either the lower bound of the confidence interval is larger than zero or the upper bound of the confidence interval is smaller than zero:

$$\underline{Q}_0 = \{p_{i,0} \in P : \underline{CI}_0 > 0\} , \quad (5)$$

$$\overline{Q}_0 = \{p_{i,0} \in P : \overline{CI}_0 < 0\} , \quad (6)$$

where $P = \{p : 0 < p < 1\}$, whereas \underline{CI}_0 and \overline{CI}_0 denote the lower and upper confidence bounds reported in (4), respectively.

These ranges, if any, define the values of $p_{i,0}$ to which a bias in the pre-event prices correspond. Specifically, \underline{Q}_0 (\overline{Q}_0) defines the range of implied probabilities which corresponds to an underpricing (overpricing) in the prediction market. To test whether these biases were large enough that profitable opportunities existed for participants who could have exploited them, which in turn would imply pre-event market inefficiency, we use a simple betting strategy. We systematically imagine wagers only on all the events within the estimation sample in which odds are identified as having been generally underpriced, i.e. in all cases where the implied probabilities belong to \underline{Q}_0 in (5). According to Fama's EMH, a positive return on investment (ROI) from such a strategy implies that the prediction market

is (weak-form) inefficient.

II.ii Are market prices efficient in the aftermath of in-play news?

In the language of online in-play prediction and financial markets, we consider the time since the event began as a number of discrete ‘ticks’, which corresponds to discrete multiples of some amount of time, for example ten seconds, which is the interval that we observe prices on the Betfair Exchange in our later application. In terms of a football match, these ticks translate to the amount of time played, or ‘on the clock’, where the clock in football only stops for the half-time interval between the beginning (‘kick-off’) and the end of the match (‘final whistle’).

Consider some type of major and relevant news about an event i ’s outcome that arrives after it has begun (in-play) at tick t , and let $p_{i,t+h}$ be the implied probability of an outcome observed after the news arrives, i.e. at tick $t + h$, for $h = 1, 2, \dots, H$, where H in each case is constrained by the end of the event. Therefore, $(p_{i,t+h})^{-1}$ represents the new equilibrium price that the market sets h ticks after the arrival of the specific piece of new information.

To evaluate whether the new equilibrium prices are set efficiently, or whether there is evidence of bias in how the market processes information, we extend the method described above to deal with high frequency data and in the style of an event study, where the event being studied in this case is the arrival of in-play news. Mimicking the approach described in Section II.i for the case of pre-event odds, the market forecast error h ticks after the in-play news is given by $e_{i,t+h} = y_i - p_{i,t+h}$, and we consider the following model:

$$e_{i,t+h} = \gamma_0 + \gamma_1 t + \gamma_2 t^2 + \beta p_{i,\tau} + u_{i,t+h} , \quad (7)$$

where $p_{i,\tau}$ is the probability of the outcome at tick τ , for $\tau = 0, 1, \dots, t - 1$. For example, $p_{i,0}$ denotes the pre-event probability of the outcome whereas $p_{i,t-1}$ is the probability of the outcome one tick before the news arrives and $u_{i,t+h}$ is an i.i.d. error term. By including the tick count in the model, we allow for the possibility that the forecast errors evolve in-play and after the major news event according to how much time has elapsed, which is equivalent to how much time is remaining if the market has a defined end point. We specify quadratic

terms for the tick in Equation (7) because we find this is appropriate for our application later in the paper; as an example, see the evolution of the odds-implied probabilities in Figure 2 below, which is clearly non-linear and can be adequately approximated by a quadratic form.

As for the case of pre-event odds in Section II.i, to account for the heteroskedasticity of the forecast errors, we estimate Equation (7) by Weighted Least Squares (WLS), where the $n_j \times n_j$ weighting matrix is diagonal with elements $\sigma_{1,\tau}^2, \dots, \sigma_{n_j,\tau}^2$, where n_j denotes the number of events where the in-play news of type j occurs (e.g. a goal scored by the home team in a football match). Again, since y_i is a Bernoulli random variable, we can approximate $\sigma_{i,\tau}^2$ with $p_{i,\tau}(1 - p_{i,\tau})$, thus yielding a weighting matrix $W_\tau = \text{diag}[\sigma_{i,\tau}^2] = \text{diag}[p_{i,\tau}(1 - p_{i,\tau})]$ in the WLS estimation.

The results from estimating Equation (7) are used to assess whether the prediction markets are generally unbiased after in-play news. In particular, a non-rejection of the null hypothesis,

$$H_0 : \gamma_0 = \gamma_1 = \gamma_2 = \beta = 0 , \quad (8)$$

implies that the expected value of the forecast error is zero, conditional on all the information available to market participants on event i until tick t , i.e. including all other in-play news before and related to the particular news studied, which arrives at tick t . This would in turn imply that the prediction market is efficient h ticks afterward. More specifically, we would have $E(e_{i,t+h}|\mathcal{I}_{i,t}) = 0$, where $\mathcal{I}_{i,t}$ is the general information set that participants use to make their forecasts and decisions, which also incorporates the regressors t and $p_{i,\tau}$, for $\tau = 0, \dots, t - 1$. Conversely, a rejection of the null hypothesis (8) implies a significant relationship between the forecast error and (at least) one of the regressors. This would in turn imply that the forecast of the market participants is biased for certain values of t and p . In other words, the arrival of news creates informational inefficiency as mispricing is observed in the market. Similarly to the case of pre-event prices, we investigate whether these deviations from no bias are large enough to generate market inefficiency. Using the estimated parameter values of Equation (7) ($\hat{\theta} = (\hat{\gamma}_0, \hat{\gamma}_1, \hat{\gamma}_2, \hat{\beta})'$), we interpolate over all possible values of $p_G \in (0, 1)$ and $t_G \in (0, \bar{t})$, where \bar{t} gives the end of the event, and derive

the efficiency curve as:

$$\hat{G}(t_G, p_G) = \hat{\gamma}_0 + \hat{\gamma}_1 t_G + \hat{\gamma}_1 t_G^2 + \hat{\beta} p_G, \quad (9)$$

as well as the related confidence bands as:

$$\begin{aligned} CI &= [\underline{CI}, \overline{CI}] , \\ &= \left[\hat{G}(t_G, p_G) - z_{\alpha/2} \text{ s.e. } \left(\hat{G}(t_G, p_G) \right), \hat{G}(t_G, p_G) + z_{\alpha/2} \text{ s.e. } \left(\hat{G}(t_G, p_G) \right) \right] , \end{aligned} \quad (10)$$

where $\text{s.e.} \left(\hat{G}(t_G, p_G) \right) = \left[\nabla \hat{G}(t_G, p_G)' V_{WLS} \nabla \hat{G}(t_G, p_G) \right]^{1/2}$, $\nabla \hat{G}(t_G, p_G) = (1, t, t^2, p_G)'$ is the gradient and V_{WLS} is the variance of the WLS estimator, and the error term in (7) is assumed to be i.i.d.

The purpose of this is to define regions where either the lower bound of the confidence interval is larger than zero or the upper bound of the confidence interval is smaller than zero. More specifically, for the case of post-news efficiency:

$$\underline{Q} = \{(t, p_{i,\tau}), t \in T, p_{i,\tau} \in P : \underline{CI} > 0\} , \quad (11)$$

$$\overline{Q} = \{(t, p_{i,\tau}), t \in T, p_{i,\tau} \in P : \overline{CI} < 0\} , \quad (12)$$

where $T = \{t : 0 < t < \bar{t}\}$, $P = \{p : 0 < p < 1\}$, and \underline{CI} and \overline{CI} denote the lower and upper confidence bounds reported in (10), respectively. These regions, if any, define the combinations of when the major in-play news arrives, t , and the event outcome probability prior to this, $(p_{i,\tau})$ which correspond to bias in the prediction markets. \underline{Q} (\overline{Q}) defines the combinations of the news occurring at tick t and the prior event outcome probability $(p_{i,\tau})$ that in general correspond to an underpricing (overpricing) in the in-play odds of the studied markets. To test whether these biases are large enough to imply in-play market inefficiency, we evaluate whether positive returns can be achieved. We do this by systematically betting only on all the events within sample, when the combinations of implied probabilities of some final outcome and the arrival ticks of particular types of in-play news are in the region \underline{Q} , i.e. in all cases where the in-play odds are underpriced. A positive ROI would then imply that the set of prediction markets studied are generally not efficient in semi-strong form.

Moreover, the analysis of the prediction market forecast errors in Equation (7) can be repeated for different values of h . This allows us to evaluate not only whether the market is inefficient but also how long any mispricing lasts, and how much time is required by the participants to accurately process the news, i.e. how much time is needed to absorb any biases and potentially adjust or re-adjust toward efficiency.

II.iii Detection of bias in prediction markets

The methods described above will potentially provide evidence of deviations from no bias within prediction markets. In the following, we provide an interpretation of possible biases and a test of whether the arrival of major news on the market provokes a significant change in participants' beliefs about an event's outcome.

The well-known *favourite-longshot bias* postulates that the odds on expected winners are underpriced while the odds on unlikely winners are overpriced, which typically implies that wagering on favourites is more profitable than wagering on longshots (e.g. [Ali, 1977](#); [Thaler and Ziemba, 1988](#); [Ottaviani and Sørensen, 2008](#)). In a prediction market containing bettors with heterogeneous beliefs, who are risk-neutral price takers, [Manski \(2006\)](#) showed formally that the overpricing (upward bias) of the longshot would arise in equilibrium due to the combination of budget constraints and skewed payoffs. [Ottaviani and Sørensen \(2015\)](#) and [He and Treich \(2017\)](#) generalised this result to broader sets of risk preferences, demonstrating sufficient conditions such that the favourite-longshot bias would emerge in prices. For example, the latter authors showed that this occurs when twice the degree of absolute risk aversion of participants is less than the degree of absolute prudence. In the case of constant relative risk aversion, this occurs when bettors are less risk averse than implied by logarithmic utility. If bettors are more risk averse, then the direction of the bias in prices could be reversed. [Ottaviani and Sørensen \(2015\)](#) also showed that the favourite-longshot bias would emerge among risk averse bettors with bounded wealth, or among bettors with unbounded wealth but decreasing risk aversion with wealth, as an underreaction to public information. However, a dynamic version of that model also predicts that this bias ought to be reversed over time. Besides these predictions from neoclassical theory, there is a competing set of behavioural explanations for the favourite-longshot bias, which emphasises

the misperception of probabilities. [Snowberg and Wolfers \(2010\)](#) looked to distinguish the behavioural and neoclassical explanations using exotic bets on US horse racing. They found evidence suggesting that bettors' inability to distinguish between different low probabilities, rather than risk-love, appears to explain why longshots are overbet. [Vaughan Williams et al. \(2018\)](#) also found that misperception rather than risk-love provided the best explanation for the favourite-longshot bias observed in behaviour during online poker games.

The presence of the favourite-longshot bias in prediction markets can be evaluated by testing whether the slope of the $p_{i,0}$ or $p_{i,\tau}$ regressors in Equations (1) and (7), respectively, are zero against the following two alternatives:

$$\left\{ \begin{array}{ll} H_0 : \beta = 0 & \text{no bias} \\ H_{1A} : \beta > 0 & \text{favourite-longshot bias} \\ H_{1B} : \beta < 0 & \text{reverse favourite-longshot bias .} \end{array} \right. \quad (13)$$

Further, we can compare the degree of bias in the market pre-event with the aftermath of major types of in-play news. Consider the initial or pre-event forecast errors ($e_{i,0}$), and the post-news forecast errors ($e_{i,t+h}$) to investigate whether the same biases apply before an event begins and after in-play news changes participants' prior expectations. In particular, define the variable:

$$\xi_{i,t} = \hat{\boldsymbol{\theta}}' \mathbf{x}_{i,t} , \quad (14)$$

which measures the mispricing at tick t for $\mathbf{x}_{i,t} = (1, p_{i,0}, t, t^2)'$, and where $\hat{\boldsymbol{\theta}} = (\hat{\gamma}_0, \hat{\beta}, \hat{\gamma}_1, \hat{\gamma}_2)'$ are the estimated parameters from (3) or (9). In the former case for $t = 0$: $(\hat{\gamma}_1, \hat{\gamma}_2) = \mathbf{0}$. The difference between the post-news and pre-event bias is defined as:

$$\Xi_{i,t} = \xi_{i,t} - \xi_{i,0} , \quad (15)$$

and we test the following null hypothesis against the alternatives:

$$\begin{cases} H_0 : E(\Xi_{i,t}) = 0 \\ H_{1A} : E(\Xi_{i,t}) > 0 \\ H_{1B} : E(\Xi_{i,t}) < 0 . \end{cases} \quad (16)$$

Not rejecting the null in (16) implies that market participants' degree of bias does not significantly react to the arrival of new information. A rejection of the null in favour of either alternative in (16) suggests that the news significantly changes the participants' bias about the expectations of an event's outcome.

In Table 1, we summarise possible combinations of the pre-event degree of bias ($\xi_{i,0}$) and the post-news market reaction ($\Xi_{i,t}$). If we observe no pre-event deviation from no bias, $\xi_{0,t} \approx 0$, (top panel of Table 1) and we reject the null in (16), then the arrival of new information on the market creates mispricing, as participants adjust their expectations and deviate from no bias. The middle panel of Table 1 shows that, starting from a situation of pre-event positive mispricing ($\xi_{i,0} > 0$) the in-play news at tick t may lead to: (i) the same degree of bias on the market as before, i.e. we do not reject the null in (16); (ii) an amplification of the positive mispricing when H_0 is rejected in favour of H_{1A} ; (iii) a significant reduction of the mispricing when H_0 is rejected in favour of H_{1B} . The latter may result in the market completely absorbing or even reversing the previous bias in the aftermath of the new information. The interpretation of the market reaction to new information in the case of pre-event negative mispricing ($\xi_{i,0} < 0$), is reported in the lower panel of Table 1 and is opposite to the case of $\xi_{i,0} > 0$.

The idea that market participants overreact to salient new information, suggested by [Kahneman and Tversky \(1973\)](#), has been extensively studied and, for example, has been demonstrated in practice by participants in the stock market (e.g. [De Bondt and Thaler, 1985, 1990](#)). For betting markets on professional football matches, [Choi and Hui \(2014\)](#) found that market reaction to particularly surprising in-play outcomes overcompensates the typical underreaction in these markets to news, and, therefore, creates an opposite mispricing on the market after the surprise. We also investigate whether *unexpected news* during an event, i.e. in-play outcomes that are characterised by a low probability, lead to

Table 1: Combinations of pre-event market mispricing, results of the hypotheses testing in (16) and interpretations of the market reaction to in-play news at tick t .

Pre-event	Result of the test	Interpretation of the market reaction
No bias ($\xi_{i,0} \approx 0$)	Accept H_0 Reject H_0 for H_{1A} Reject H_0 for H_{1B}	No change in beliefs: still no bias Creating positive mispricing Creating negative mispricing
Positive mispricing ($\xi_{i,0} > 0$)	Accept H_0 Reject H_0 for H_{1A} Reject H_0 for H_{1B}	No change in beliefs: still positive mispricing Amplifying mispricing Absorbing mispricing
Negative mispricing ($\xi_{i,0} < 0$)	Accept H_0 Reject H_0 for H_{1A} Reject H_0 for H_{1B}	No change in beliefs: still negative mispricing Absorbing mispricing Amplifying mispricing

overreaction. In our particular football betting application, we look for evidence of this by comparing the reaction and evolution of in-play odds after goals are scored by teams which are either less or more fancied to win the match.

We also empirically investigate under which conditions a *confirmation bias* (Wason, 1960; Rabin, 1998) can be suggestively found in prediction markets. We interpret any cases in which mispricing was detected before the event began ($\xi_{i,0} \neq 0$) and where we do not reject the null in (16), $E(\Xi_{i,t}) = 0$, as evidence of a confirmation bias. Such a case would suggest that participants stick to their prior beliefs after new information breaks on the market, even though those beliefs were biased. In that sense, the new information is perceived by the market as not informative enough to provoke a reaction to compensate for the previous mispricing.

Finally, we analyse a further well-documented bias in sports betting markets. Among others, Levitt (2004) and Vlastakis et al. (2009) show that bettors tend to overestimate the probability of the home team winning, i.e. *home bias*. Notwithstanding the fact that playing at home significantly increases the probability of winning the match (e.g. Nevill and Holder, 1999), bettors tend to overvalue the actual chance of the home team winning, and this bias is amplified when the home team is the favourite to win.

III Data & Estimation

We use a sample of $n = 1,004$ matches played in the English Premier League (EPL) from 15 August, 2009, to 11 May, 2014, from the total number of 1,900 that took place in this

period.⁹ For each match, we observe the in-play odds (prices) collected every 10 seconds on the Betfair Exchange market for the final result outcome, i.e. whether the game finishes in a draw (tie) or a win for either the team playing at home or away. Thus, in this application our definition of a tick is a period of ten seconds.¹⁰

A betting exchange operates as a limit order-driven market, which matches the ‘back’ and ‘lay’ orders, that is the bets on and against an outcome, respectively. Essentially, this allows individuals to bet against each other directly. The back and lay odds are equivalent to the bid and ask prices in financial markets. In betting exchange markets, the prices (odds) are not dictated by market makers (bookmakers), but the bettors can buy (back) or sell (lay) bets both pre-match and during the game. Moreover, as Betfair charged a commission of up to 5% on net winnings ex post in the UK, falling to 2% for heavy bettors and in the 2019/20 season for our out-of-sample test, this is not reflected in the price data. In a nutshell, the matched bets on each outcome are zero-sum games between the back and lay bettors and, given that there is no bookmaker, any biases we observe from the betting exchange odds should derive from the behaviour of market participants.

Betfair operates the world’s biggest exchange by volume traded and claims to have millions of customers. The arrival of the online betting exchanges in 2000 in the UK is credited with revolutionising the betting industry, driving down bookmaker profit margins (overrounds) and increasing competition (Forrest et al., 2005). As outlined by Croxson and Reade (2014), the number of daily trades on the Betfair Exchange has historically been greater than all the European Stock Exchanges combined. There is no liquidity issue in the prediction markets we study, as also described by Croxson and Reade (2014). This is the case both pre-match and during the match itself. It is generally a feature of the most popular betting exchange markets that the volume of trading multiplies after the event has begun. To illustrate this, we present in Figure 1 scraped data from a recent EPL match, showing the cumulative amount of money (pounds sterling) matched on the market for the

⁹We obtained the data set from a third party who had purchased it from Betfair. We looked into obtaining data for the other matches in these seasons but discovered this was not feasible. The selection of matches into the sample is random within seasons, with all teams in the EPL in this period being represented home and away. 47% of matches are represented in the 2009/10 season, 37% in 2010/11, 54% in 2011/12, 61% in 2012/13 and 65% in 2013/14. We can provide a list of the sample matches on request.


¹⁰Standard participants on the Betfair Exchange can observe live price and volume updates as frequently as every 1 second, but we were only able to obtain regular frequency historical data for markets at 10 second intervals.

final result (i.e. home, away win or draw), from 90 minutes before kick-off up to the point the market closed after the final whistle. The vertical dashed line shows when the match began. At this point, £4.7 million of bets had been matched, but by market close this figure was £12.3 million. Also shown in Figure 1 is what the exchange terms the ‘book percentage’ (right axis) throughout the duration of the market. From the perspective of the backer, this is the sum over all possible event outcomes of the odds-implied probabilities, or in other words one plus the exchange market equivalent of a bookmaker’s overround, and thus it gives a measure of competitiveness. For the vast majority of the event this measure was less than 101%, implying that the prices being offered were competitive. The horizontal dashed line in Figure 1 traces out the goal difference as the example match progressed and ended 3-1. This illustrates that the competitiveness of the market and the implied liquidity did not fluctuate wildly when a goal was scored.¹¹ This is why we chose Premier League football match result markets to illustrate and apply the methodological framework in Section II. An application with less liquid markets is likely to significantly limit the chances of detecting significant in-play mispricing and inefficiency or interpreting behavioural biases. In such markets, the back and the lay odds can be persistently some distance apart and will not necessarily imply the market expectations of outcomes (see [Flepp et al., 2017](#), for a discussion of how less liquid betting exchange markets allow bookmakers (sportsbooks) to remain successful).

The data we analyse concern the actual prices at which trades were made, rather than the back or lay prices being offered at any point in time. Figure 2 shows an example of these in-play betting exchange data, using the Southampton vs. Manchester United match played on 11th May, 2014. The time series of the odds-implied probabilities are depicted for each of the possible final result outcomes and for all ticks from 1 to 550, where each additional tick corresponds to 10 seconds of the match, allowing for a short amount of injury time in all matches in either half. We can observe two distinct jumps in the patterns of the implied probability series in this match, caused by two separate instances of major news. The first of these is a goal scored at tick 163 by Southampton, which provoked an abrupt change in the three outcome probabilities and, subsequently, a new market equilibrium was reached. Specifically, the implied probability of a Southampton win increased from 0.35 to

¹¹Unfortunately, we do not have data on the market volumes traded or competitiveness of the matches in our analysis sample. However, in addition to the example shown, the data employed by [Croxxson and Reade \(2014\)](#) for an earlier period demonstrate that these markets are heavily traded, liquid and competitive.

Figure 1: An example of the in-play liquidity and competitiveness of Betfair Exchange English Premier League match result markets



Notes: author calculations from Betfair Exchange: time series from 90 minutes before kick-off to the market close for the final result outcome of Liverpool vs. Manchester City, 10th November, 2019. The vertical dashed line shows the time of kick-off. The horizontal dashed line traces out the goal difference between the home and away team during the match - the final scoreline was 3-1.

0.62 after the goal, while the draw and Manchester United win probabilities dropped to 0.23 and 0.14 from 0.31 and 0.33, respectively. The second major news is a goal by Manchester United at tick 322, which promptly increased the implied probabilities of the draw and the away win, and dramatically decreased the home win probability. Thereafter, since no other major news arrived on the market, such as more goals or a player being dismissed, the draw probability tended to 1 toward the end of the match at an increasing rate, while both the home and away win probabilities shrank toward 0.


Importantly for the application of our methodology in what follows, Betfair briefly suspends the final result markets at kick-off and whenever a ‘Material Event’ is deemed to have occurred. The latter includes the awarding of red cards or penalty kicks, but the most significant material events in terms of price reactions are goals being scored, especially the first one. In these cases, the market is suspended only so long as it takes for a goal to be awarded with certainty to one team or the other, which is typically just a few seconds, at least before the introduction of video assistant referees, which postdated our sample. This delay is put in place in case the referee or another official rules the goal out, for example because of a player being deemed offside or a foul being spotted in the build-up to the goal. When the market is suspended, all unfilled orders are cancelled, clearing out the market.¹² As soon as the market re-opens we observe an immediate jump in prices, as shown in Figure 2 and previously demonstrated by Croxson and Reade (2014). We observe the tick and prices just before and just after a goal is scored.

III.i Estimation

We apply the methodology described in Section II to the data set of $n = 1,004$ matches (events), to evaluate the efficiency of the exchange betting market. We use the prices backing either a home or an away win throughout the analysis. The odds on ties in football are generally tightly bounded above and below over matches, both pre-match and after the first goal, in which latter case the draw outcome becomes even more unlikely. As a consequence, the efficiency curves in Equations (3) and (9) cannot be computed over all

¹²In fact, Betfair actively voids any bets that were ‘unfairly’ matched after a material event if the market was not suspended on time. See Betfair Rules and Regulations, Part B, 1.3 (November 2019): <https://www.betfair.com/aboutUs/Rules.and.Regulations/>.

Figure 2: In-play match result probabilities from Betfair Exchange: Manchester United vs. Southampton, 11th May 2014



Notes: author calculations from Betfair Exchange, time series from tick 0 to tick 550 of the probabilities of a home win (red), a draw (blue) and an away win (green) implied by in-play odds. The match ended 1-1.

values of $p_G \in (0, 1)$. Therefore, we do not consider the draw outcome explicitly in our analysis.

First, we concentrate on the pre-event market, as per Section II.i. Second, we study the markets after the arrival of major in-play news, as per Section II.ii. In particular, we focus on the ‘first goal’ of a match as major news, in the sense that it significantly affects the in-play odds (see Figure 2). Therefore, for this part of the analysis, we exclude all matches which ended with no goals scored, such that $\tilde{n} = 882$, of which $\tilde{n}_H = 513$ are ‘home team goal’ matches and $\tilde{n}_A = 369$ are ‘away team goal’ matches. In estimating Equation (7), we consider $p_{i,\tau}$ with $\tau = 0$, i.e. we consider the pre-match probability as the regressor. We can identify $\tau = 0$, the first tick of the match and starting odds, because Betfair always briefly suspends the market at kick-off. Focusing on the first goal of a match and $p_{i,0}$ as the regressor in (7) allows us to consider an exhaustive range of possible combinations of tick and implied probabilities, which enables us to study the evolution of prices after events and investigate the market participants’ reaction in four different scenarios: (i) home odds after the first goal is scored by the home team (HH); (ii) home odds after the first goal is scored by the away team (HA), (iii) away odds after the first goal is scored by the home team (AH), (iv) away odds after the first goal is scored by the away team (AA).¹³

IV Empirical Analysis of Betfair Exchange Prediction Markets



In this section, we show the results from applying the bias and efficiency testing approach described above to our sample of Betfair Exchange prediction markets. By doing so, we look to answer the questions posed in Section II within this particular context. First, in Section IV.i, we investigate the market efficiency for the prices of final result outcomes set just before the beginning of a football match (event). Second, in Section IV.ii, we address the semi-strong form of market efficiency, focusing on how these prediction markets reacted after the first goal was scored (major news). Finally, in Section IV.iii, we analyse and interpret the presence of the behavioural biases described in Section II.iii.

¹³We also estimate the model using the probability of the match outcome prior to the goal, i.e. $p_{i,t-1}$, in (7), and describe the results later.

IV.i Are market prices efficient just before a match kicks off?

We evaluate whether the exchange market participants set prices (odds) efficiently before the beginning of a match. The results are reported in Table 2. The top panel of the table reports the estimates of Equation (1), the middle panel depicts the derived efficiency curves ($\hat{G}(p_G)$), as per Equation (3) and over all possible values of $p_G \in (0, 1)$, and the bottom panel shows the returns on investment from applying an in-sample betting strategy based on Equation (5), i.e. betting the same amount in all cases where the implied probabilities at the start of a match were significantly underpriced.

Table 2: Pre-match analysis of prediction market mispricing and efficiency

Estimates	
<i>Home Odds</i>	<i>Away Odds</i>
$\hat{\gamma}_0 = 0.0942^{***}, \hat{\beta} = -0.2599^{***}$ (0.0031) (0.0001)	$\hat{\gamma}_0 = 0.0606^{***}, \hat{\beta} = -0.2530^{***}$ (0.0048) (0.0006)
$F\text{-test} = 7.94^{***} (0.0003)$	$F\text{-test} = 5.85^{***} (0.0030)$
Efficiency curves	
<i>Home Odds</i>	<i>Away Odds</i>
	
$\underline{Q}_0 = (0, 0.24], \bar{Q}_0 = [0.46, 1)$	$\underline{Q}_0 = (0, 0.14], \bar{Q}_0 = [0.33, 1)$
Betting strategy	
<i>Home Odds</i>	<i>Away Odds</i>
Odds > 4.16	Odds > 7.14
Matches = 183	Matches = 195
Correct bets = 20.77%	Correct bets = 14.36%
ROI = 40.05%	ROI = 56.41%

Notes: author calculations using pre-match Betfair result odds from $n = 1,004$ matches. Top panel: WLS estimates of Equation (1), p -values in parentheses. *** indicates significance at the 1% level, two-sided tests. F -test displays the test statistic for the test of the null hypothesis in (2). Middle panel: efficiency curves as per Equation (3). Bottom panel: application of a simple betting strategy described in the text. ROI is gross, before commission or charges are applied to profits by Betfair Exchange.

The F -tests reported in the top panel of Table 2 show a rejection of the null hypothesis in (2) for both the home and away win odds. Therefore, the pre-match odds for the winning outcomes on the Betfair Exchange markets were generally not set efficiently, and thus there is

evidence of mispricing (bias). The estimated slopes, $\hat{\beta}$, are significantly negative. This result suggests a *reverse* favourite-longshot bias, as we reject the null hypothesis ($H_0 : \beta = 0$), in favour of the alternative (H_{1B}) in (13). Therefore, bettors operating in the exchange betting markets backed too strongly the teams which were expected to win, such that wagering on longshots tended to offer greater expected returns. This evidence is opposite to the result typically found for fixed-odds bookmaker markets on football matches (e.g. Angelini and De Angelis, 2019). In line with Smith et al. (2009) and Snowberg and Wolfers (2010), in addition to the theoretical predictions in Ottaviani and Sørensen (2015) and He and Treich (2017), this evidence of a reverse favourite-longshot bias in exchange betting markets could be explained by a higher degree of risk aversion among the bettors operating in these markets compared with those using fixed-odds bookmakers. The estimated intercept of Equation (1) is significantly positive, though not enough on average to offset the favourite-longshot bias, for both the home and away win outcomes, implying on average negative forecast errors. This would suggest that the market participants were biased toward a win outcome and away from the draw, which would be consistent with a ‘splitting’ bias, or ‘black and white thinking’. The significant tendency of individuals to under-predict draws in football matches has been documented before by Na et al. (2019) in an experimental setting.

The estimated efficiency curves $\left(\hat{G}(p_G)\right)$, and the related confidence bands, computed as per (4), define the probability ranges (\underline{Q}_0) in (5) for home and away odds (middle panel of Table 2). Due to the estimated negative slope (and positive $\hat{\gamma}_0$), these probability ranges are defined by the lowest probabilities (i.e. the highest odds), as \underline{Q}_0 is given by $0 < p_G \leq 0.24$ and $0 < p_G \leq 0.14$ for home and away odds, respectively. Therefore, consistent with the reverse favourite-longshot bias, higher probabilities (smaller odds) were more likely overpriced, while smaller probabilities (higher odds) were more likely underpriced.

To evaluate the forecasting performance of these markets and their efficiency, the simple betting strategy described in Section II.i is applied in-sample (*ex post*) to all n matches in our data set, i.e., betting one unit on all pre-match odds which were identified as generally significantly underpriced by the regression model. The results reported in the bottom panel of Table 2 show that by systematically wagering the same amount on all the home (away) odds of matches whose pre-match probability was in the range $\underline{Q}_0 = (0, 0.24]$

($\underline{Q}_0 = (0, 0.14]$), i.e. on each of the 183 home and 195 away matches where the odds were larger than 4.16 and 7.14, respectively, a bettor could have earned a substantial gross ROI of 40% in the case of home wins and 56% for away wins, before paying any commission.¹⁴ From this evidence, we can conclude that the reverse favourite-longshot bias detected within pre-match odds, for final results on the Betfair Exchange, was large enough to create profitable opportunities for bettors in expectation, using a relatively simple betting strategy. Therefore, these markets were (weak-form) inefficient. This result is in spite of the apparent overpricing of the win outcomes relative to the draw outcome in these football matches.

We also apply the same betting strategy implied by the results in Table 2 out-of-sample, to all 380 matches in the 2019/20 season of the English Premier League, using the pre-match prices from Betfair Exchange. We find a gross ROI of 51.9% over home and away wins combined, from a total of 175 bets with average decimal odds of 10.3, implying an average implied probability of 9.8% against an actual winning percentage of 17.7%.¹⁵ Despite applying the results of the mispricing testing framework at least some four years after the last match in our estimation sample, it appears as though the reverse favourite-longshot bias and relatively straightforward source of inefficiency has remained in these markets. Given the high liquidity of the elite-level football result markets on the Betfair Exchange (Figure 1), substantial profits could have been made.

At first glance, the gross ROIs we find from this application, in and out-of-sample, are staggeringly high. But such high ROIs only come from selectively betting where there is a statistically significant ‘edge’ over the market, from the systematic mispricing we have identified. This contrasts with evidence from online bookmakers on these markets, including on the Premier League for the same matches, where there is a small favourite-longshot bias that is not enough to overcome the bookmakers’ profit margin or overround (see Angelini and De Angelis, 2019). Taking the Betfair Exchange and bookmaker results together, there are likely to be arbitrage opportunities, both pre-match and in-play. We can speculate that these occasional arbitrage opportunities are able to persist for two reasons. First, profitable

¹⁴The standard commission rate on profits applied by Betfair is 5% in the UK, though with the basic My Betfair Rewards plan most participants would pay at most 2%. There are other factors affecting the potential net ROI, meaning this would differ not only based on the country of the participant but also their exchange markets history. See for details: <https://www.betfair.com/aboutUs/Betfair.Charges/>.

¹⁵Details of the odds and which matches and outcomes bets were made are available on request. We do not have in-play price data since the 2013/14 season and our model estimation sample.

bettors and arbitragers are quickly identified and have their activity blocked and limited by online sportsbooks, leaving behind only the more casual or less sophisticated punters, who favour long-odds accumulators, in spite of their huge overround, or markets with low liquidity on the exchanges (see [Flepp et al., 2017](#)). Second, bookmakers participate in the exchange football match result markets themselves, not only to hedge but also to profit from the mispricing and any arbitrage.

IV.ii Are market prices efficient after the first goal is scored?

We now consider whether the odds set on the Betfair Exchange in the aftermath of the first goal scored were unbiased. As a preliminary analysis of the in-play data, Appendix Figure [A1](#) reports the mean of the jumps in the implied probabilities $h = 2$ ticks after the first goal in each of the four cases considered, namely $\{HH, HA, AH, AA\}$. The behaviour of the implied win probabilities ($p_{i,t+2}$) in the HH and AA cases was similar. The magnitude of the jump increases along with the tick when the first goal was scored. The higher jumps in the probability of a win were concentrated around cases where the pre-match probability was 0.50. However, for the cases of HA and AH , it is less easy to observe regularity in the pattern of probability jumps following the arrival of the first goal, and overall the mean changes are lower in absolute value.

We estimate the model in Equation [\(7\)](#) with the pre-match probability ($p_{i,0}$) as a regressor (i.e. $\tau = 0$) and then consider different horizons after the major news arrived, namely $h = \{2, 5, 30\}$ ticks after the first goal.¹⁶ Table [3](#) reports the estimated coefficients and the efficiency tests for the null hypothesis in [\(8\)](#). The F -tests reject the null of market efficiency for the cases of HH and AA in the top-left and bottom-right panels, respectively, for all horizons (h) considered, except for $h = 5$ in the AA case. Conversely, for the two cases of HA and AH (top-right and bottom-left of Table [3](#), respectively) the exchange market prices evolved efficiently after the arrival of major news for all horizons considered, as the F -tests for the null in [\(8\)](#) do not reject market efficiency.¹⁷ Because the book percentage

¹⁶We exclude from the analysis all matches in which further major news arrives after the first goal and before tick $t + h$, e.g., a second goal.

¹⁷We control for misspecification in the functional form of [\(7\)](#) with a set of Ramsey RESET tests. The results indicate that the model in [\(7\)](#) is correctly specified. We also considered adding a dummy variable to the model, interacted with the odds-implied pre-match probability, for when this latter value was greater than 0.5 (i.e. $d_{i,0}p_{i,0}$, where $d_{i,0}$ takes the value 1 if $p_{i,0} > 0.5$ and zero otherwise). This was done to test the assertion in [Newall and Cortis \(2021\)](#) that a favourite bias is more likely when the favourite has a > 0.5

(competitiveness) of these markets is mostly constant (Figure 1), the results in Table 3 suggest that the draw outcome of a football match is on average mispriced after the first goal is scored, in the opposite direction to the win outcome odds of the team that scored. Nevertheless, we still omit the draw outcome from the analysis here as the mispricing results and efficiency curves do not provide any additional insights – when forecasting football match results or analysing their betting markets, it is practical to focus on the home and away win outcomes and to regard the draw as the residual outcome.

Table 3: In-play analysis of market mispricing when the first goal is scored

		<i>Home Goal</i> ($\tilde{n}_H = 513$)			<i>Away Goal</i> ($\tilde{n}_A = 369$)		
		$h = 2$	$h = 5$	$h = 30$	$h = 2$	$h = 5$	$h = 30$
<i>Home Odds</i>	$\hat{\gamma}_0$	-0.0170 (0.7746)	-0.0638 (0.2895)	-0.0491 (0.4764)	0.0386 (0.3473)	0.0697* (0.0959)	0.0274 (0.5121)
	$\hat{\gamma}_1$	0.0011*** (0.0055)	0.0010*** (0.0098)	0.0010** (0.0298)	-0.0001 (0.5955)	-0.0002 (0.4591)	-0.0002 (0.6119)
	$1000\hat{\gamma}_2$	-0.0015** (0.0238)	-0.0013* (0.0509)	-0.0013 (0.1203)	0.0003 (0.5191)	0.0003 (0.5313)	0.0003 (0.5679)
	$\hat{\beta}$	-0.1784** (0.0244)	-0.1209 (0.1391)	-0.1248 (0.1833)	-0.1121* (0.0705)	-0.1119* (0.0965)	-0.0584 (0.2407)
	F -test	4.1158*** (0.0027)	3.4686*** (0.0083)	2.7496** (0.0279)	0.8916 (0.4690)	0.9737 (0.4219)	0.4455 (0.7756)
<i>Away Odds</i>	$\hat{\gamma}_0$	0.0289 (0.2442)	0.0442* (0.0851)	0.0289 (0.2279)	0.1146* (0.0932)	0.0220 (0.7427)	0.0574 (0.4435)
	$\hat{\gamma}_1$	-0.0002 (0.2658)	-0.0003 (0.1324)	-0.0003* (0.0835)	0.0007 (0.1967)	0.0006 (0.2017)	0.0006 (0.2931)
	$1000\hat{\gamma}_2$	0.0003 (0.2864)	0.0004 (0.1870)	0.0005* (0.0879)	-0.0008 (0.3887)	-0.0008 (0.3635)	-0.0007 (0.5339)
	$\hat{\beta}$	-0.0423 (0.3425)	-0.0185 (0.6497)	-0.0025 (0.9587)	-0.3752*** (0.0020)	-0.2112* (0.0741)	-0.2661** (0.0471)
	F -test	0.5164 (0.7238)	0.7769 (0.5405)	0.9078 (0.4593)	4.9844*** (0.0006)	1.7154 (0.1460)	2.4393** (0.0475)

Notes: author calculations using within match Betfair result odds from $\tilde{n}_H = 513$ ‘home team goal’ matches and $\tilde{n}_A = 369$ ‘away team goal’ matches. Presents WLS estimates of Equation (7). p -values in parenthesis. ***, **, * indicate significance at the 1%, 5% and 10% levels, respectively, two-sided tests. F -test displays the test statistic for the test of the null hypothesis in (8).

When the home team scored first, the estimated coefficients reported in Table 3 show that the average market bias is mainly explained by the time when the first goal was scored, i.e. $\hat{\gamma}_1$ is significantly greater than zero at the 5% level for all horizons considered, while $\hat{\gamma}_2$ is only so for $h = 2$. Taken together, the home win is on average increasingly underpriced after the first goal is scored as the match progresses, though $\hat{\gamma}_2$ being negative makes this a decreasing rate of increase. Conversely, the pre-match implied probability is (negatively) significant at the 5% level only twenty seconds after the first goal was scored, i.e. $h = 2$, but not for $h = 5$ (one minute) or $h = 30$ (five minutes). Therefore, we find evidence of *reverse* probability of winning. This variable was insignificant in all specifications.

favourite-longshot bias, i.e. we reject H_0 in (13) in favour of H_{1B} , though only for the case of $h = 2$. As such, this bias tended to be absorbed by the market as h increased. However, the markets generally still did not achieve efficiency as long as five minutes after the first goal was scored, mainly because of the significance of $\hat{\gamma}_1$. This is also the case when the away team scored the first goal and the reverse favourite-longshot bias is more evident in this case. The results in the bottom-right panel of Table 3 show that the reaction of away odds is only significantly explained by the pre-match probability of an away team win when they scored the first goal, as $\hat{\beta}$ is significantly negative for $h = 2$ and $h = 30$, at least at the 5% level, and for $h = 5$ at the 10% level.

Appendix Table A1 shows results of a robustness check, considering the odds-implied probability observed one tick (at most 10 seconds) before a goal was scored, $p_{i,t-1}$, in Equation (7), instead of the initial pre-match probability, $p_{i,0}$. Comparing Appendix Table A1, with Table 3 indicates that our main results are robust to whether we use the pre-match or the pre-goal odds-implied probability in the model. There is mispricing and a reverse favourite-longshot bias for football result prediction markets after the first goal is scored, which can be predicted by either the pre-match or pre-goal odds. Despite the informative content provided by the pre-goal regressor, we prefer and recommend using the pre-match regressor in this particular application to football result markets. The issue is that considering the probability before the goal, in practice, rules out several combinations of (t, p) , especially cases of large t and large p . As a matter of fact, it is unlikely to observe cases in which no goals have yet to be scored as a match nears its end but where one of the teams has a high win probability. In these cases the draw is the most likely outcome and, in general, odds are set accordingly by the market.

Appendix Figure A2 plots in blue the efficiency curves $\left(\hat{G}(t_G, p_G)\right)$ according to (9), and in red the related 90% confidence bands as per (10), for $h = 2$ ticks after the first goal was scored. From these plots we can identify the combinations of tick (t) and pre-match probability, $(p_{i,0})$ where the conditions $\underline{CI} > 0$ and $\overline{CI} < 0$ in (11) and (12), respectively, were satisfied. For instance, we observe that the combination of large (small) t and small (large) $p_{i,0}$ satisfied condition $\underline{CI} > 0$ ($\overline{CI} < 0$) for the cases of HH and AA . In line with the results in Table 3, we cannot distinguish such combinations in the cases of HA and AH ,


especially for the condition $\underline{CI} > 0$.

To better identify the combinations $(t, p_{i,0})$ which implied inefficiency, Figure 3 extrapolates from Appendix Figure A2 the *underpriced* areas, \underline{Q} , (in green) and the *overpriced* areas, \overline{Q} , (in yellow) for $h = \{2, 5, 30\}$ ticks after the first goal. It shows that the evolution of the in-play odds afterward was similar in the cases of HH and AA . The underpriced odds were concentrated in the area where the pre-match probability was roughly below 0.7 and 0.5 for HH and AA , respectively. One implication suggested by these results is that a goal scored by a pre-match longshot team led to mispricing of the team's win probability. In particular, this "surprise" news was unexpected by the market participants, and thus the underpriced odds set after the goal was scored could be explained by a lack of confidence in the now greater possibility that the longshot would eventually win the match. Conversely, the overpriced odds were concentrated in the area where the pre-match probability was larger and the first goal was scored in the early stages of a match, namely t approximately smaller than 150 ticks (the first twenty-five minutes). This suggests that a match's first goal scored early by the favourite team was generally expected by the market participants, and thus the price after the goal arrived was set with excess confidence, since it reaffirmed beliefs that the favourite team would win.

To evaluate the efficiency of the market, we apply in-sample the betting strategy outlined in Section II.ii (*ex post*) to the sub-sample of $\tilde{n} = 882$ matches. The results for the underpriced odds after the first goal was scored, i.e. those satisfying the condition in \underline{Q} for the HH and AA cases, are reported in the left panel of Table 4. These results show that, by systematically wagering the same amount on all matches in the region \underline{Q} for the HH (AA) case, gross ROIs of 35.27%, 10.74% and 11.97% (70.42%, 67.16% and 24.58%) could have been earned for $h = 2, 5$, and 30 ticks after the first goal, respectively. As a robustness check, we present in the right panel of Table 4 the results from similarly and systematically betting in-sample on the region of overpriced odds, \overline{Q} . From these results, it is evident that betting in this way would generally have led to negative returns.¹⁸

¹⁸We also performed an out-of-sample betting strategy exercise for the last 100 matches in the sample, for both pre-match and in play odds; i.e. we did not use these matches to identify the regions of mispricing. The results from this are qualitatively similar to the in-sample ones reported in Tables 2-4, and are available on request.

Figure 3: Estimated regions of under and overpricing following the first goal scored



Notes: shows the estimated regions Q and \bar{Q} in (11) and (12), respectively, for $h = \{2, 5, 30\}$ ticks after the first goal event, extrapolating from Appendix Figure A2. The green areas depict combinations of t and $p_{i,0}$ when the Betfair Exchange in-play odds are in general *underpriced*. The yellow areas depict the combinations of t and $p_{i,0}$ when the Betfair Exchange in-play odds are in general *overpriced*.

In summary, we find significant and substantial evidence that the in-play odds in football match result Betfair Exchange markets were (semi-strong form) inefficient in the aftermath of the first goal, if we consider the reaction of home (away) win odds to a goal scored by the home (away) team. However, this is not the case for other combinations of the outcome and identity of the team scoring the first goal, i.e. the HA and AH cases.

IV.iii Detection of changes to bias on the Betfair Exchange after major news

In this section, we evaluate whether and how new information impacts on market participants' previous expectations. Table 5 shows the results from testing whether these expectations (biases) changed for all three of the different pre-match cases reported in Table 1, for $h = 2$ ticks after the first goal was scored. In particular, we consider the odds of the overpriced favourite teams, with an implied pre-match win probability larger than or equal to 0.75 ($p_{i,0} \geq 0.75$), and the underpriced longshot teams, $p_{i,0} \leq 0.20$. We also consider the odds when the pre-match market was unbiased, as previously described in Table 2: $0.25 \leq p_{i,0} \leq 0.45$ and $0.15 \leq p_{i,0} \leq 0.32$ for the home and away odds, respectively.

There are two cases where it appears as though the previous bias was generally not updated after the first goal arrived in a match. Both relate to the favourite team scoring in the first fifteen minutes ($t \leq 90$), i.e. cases of pre-match upward biased prices for either the home or away win. In these two cases, we do not reject the null hypothesis in (16). This emphasises that an early goal scored by a favourite team was somehow expected by market participants, and did not alter the pre-match mispricing. In other words, the probability that the bettors ascribed to the favourite team remained higher than it should have been. This is evidence that prices were affected by confirmation bias, i.e. individuals processed the implications of the major news in a way which confirmed their prior expectations (Wason, 1960; Rabin, 1998). However, from the results reported in the first row of each panel of Table 5, we observe an absorption of the pre-match mispricing, rejecting H_0 in favour of H_{1A} in (16), as a reaction when the first goal was either scored later in a match by a favourite or scored at any time by a longshot. These latter results suggest that in some circumstances the market participants did correctly process the arrival of new information,

and the exchange markets then tended to adjust toward unbiasedness.¹⁹ In fact, one can infer whether the adjustment toward unbiasedness is large enough to reach in-play market efficiency by looking at the white areas in Figure 3. For example, consider the reaction of home odds when the first goal in a match is scored by the home team in the second half ($t \geq 275$), with a pre-match implied home win probability of $p_{i,0} \geq 0.75$, which would have suggested favourite bias. Figure 3 shows that the reaction in such a case is on average large enough to fully absorb the pre-match favourite bias.

Another interesting case concerns the effect of the first goal arriving when the market was generally unbiased at kick-off. The second row of each panel in Table 5 shows that we reject H_0 in (16) in all these cases, thus losing market unbiasedness after the major news arrived. The first goal being scored by the away team generally produced a downward (upward) bias on the home (away) odds, suggesting that the market participants did not believe sufficiently that the away team would eventually win. This result is in support of a *home bias* (e.g. Levitt, 2004). In the case of the first goal being scored by the home team, the results are mixed. We observe a downward (upward) bias for home (away) odds if the first goal occurred in the first fifteen minutes of a match ($t \leq 90$), while the opposite market reaction is observed when the first goal occurred late in a match ($t \geq 275$). These results are difficult to interpret and this could be due to the large range of $p_{i,0}$ considered.

The last row of each panel in Table 5 reports results relating to the pre-match upward bias of longshot odds. In this case, the mispricing tended to be absorbed after the favourite scored the first goal, both for the home and away win odds. Conversely, this mispricing was amplified when the first goal was scored by the longshot, except for the case of an early goal scored when the longshot was playing at home. Therefore, we find evidence that the market participants reacted toward no bias following expected news (i.e. the first goal scored by the favourite team), while the reverse favourite-longshot bias was amplified when there was surprise news (i.e. the first goal scored by the longshot).

¹⁹Page and Clemen (2012) found that the favourite long-shot bias in prediction market prices of political events tended to be absorbed by markets over long periods of time and as the expiration date approached. However, this study did not address whether the absorption was driven by major news arriving on the market, but rather instead modelled it as a mechanical effect of time discounting by the participants, given that the length of time studied was considerably longer.

Table 4: In-play betting strategy results on final result outcomes after the first goal is scored

	Betting Strategy: underpriced odds				Robustness check: overpriced odds			
	Home Odds/Home Goal	Away Odds/Away Goal	h = 2	h = 5	h = 30	h = 2	h = 5	h = 30
ROI (%)	35.27	10.74	11.97	70.42	67.16	24.58	-24.12	-10.66
Correct bets (%)	67.95	74.14	73.02	48.44	80.00	57.69	64.71	70.42
Mean winning odds	1.99	1.49	1.53	3.52	2.09	2.16	1.17	1.22
Matches	78	58	63	64	15	26	34	71
							46	5

Notes: see Figure 3. Shows results from a simple betting strategy, systematically wagering the same amount on all cases identified as underpriced or overpriced (for robustness), following the first goal scored during an in-sample match. ROI is gross, before commission or charges are applied to profits by Betfair Exchange.

Table 5: Change in market participants expectations and bias, $\Xi_{i,t}$, following the first goal scored in a match, depending on pre-match conditions

	Home Goal ($\tilde{n}_H = 513$)				Away Goal ($\tilde{n}_A = 369$)			
	$t \leq 90$	$t \geq 275$	$0 \leq t \leq 550$		$t \leq 90$	$t \geq 275$	$0 \leq t \leq 550$	
Home Odds	Favourite bias, $p_{i,0} \geq 0.75$	-0.0069 (0.2223)	0.1396*** (0.0000)	0.0860*** (0.0000)	0.0517*** (0.0000)	0.0412*** (0.0000)	0.0437*** (0.0000)	
	No bias $0.25 \leq p_{i,0} \leq 0.45$	-0.0342*** (0.0000)	0.1087*** (0.0000)	0.0390*** (0.0003)	-0.0061*** (0.0001)	-0.0156*** (0.0000)	-0.0136*** (0.0000)	
	Longshot bias $p_{i,0} \leq 0.20$	-0.0511*** (0.0000)	0.0893*** (0.0000)	0.0263*** (0.0003)	-0.0360*** (0.0000)	-0.0457*** (0.0000)	-0.0421*** (0.0000)	
Away Odds	Favourite bias, $p_{i,0} \geq 0.75$	0.1067*** (0.0000)	0.0812*** (0.0000)	0.0963*** (0.0000)	-0.0045 (0.4716)	0.1004*** (0.0000)	0.0578*** (0.0000)	
	No bias $0.15 \leq p_{i,0} \leq 0.32$	0.0101*** (0.0000)	-0.0118*** (0.0000)	-0.0028*** (0.0034)	0.0482*** (0.0000)	0.1561*** (0.0000)	0.1043*** (0.0000)	
	Longshot bias $p_{i,0} \leq 0.20$	-0.0085*** (0.0000)	-0.0340*** (0.0000)	-0.0237*** (0.0000)	0.0627*** (0.0000)	0.1676*** (0.0000)	0.1188*** (0.0000)	

Notes: see Tables 2 & 3. Shows test results of the null hypothesis (16), i.e. whether the nature of the pre-match *reverse* favourite-longshot bias changes after the first goal is scored. *p*-values in parentheses. *** indicates significance at the 1% level, two-sided tests. See Table 1 for interpretation. Results shown in **green** are the cases in which the pre-match bias tends to be absorbed following the goal. Results shown in **red** are the cases in which the pre-match bias is amplified; favourites were negatively mispriced and longshots were positively mispriced.

V Concluding Remarks

In this paper, we proposed a practical framework which can be used to investigate the behaviour of participants in prediction markets. We demonstrated this using a high-frequency data set of sports betting exchange prices (odds) on the final result markets of football matches. The methodology could be easily applied to other prediction markets with high-frequency data and the clean arrival of major news, beyond sports, such as those run within major companies among employees (e.g. [Cowgill and Zitzewitz, 2015](#)) and public markets on political or financial events (e.g. [The Iowa Electronic Markets \(IEM\)](#), [PredictIt](#) and the now defunct [Intrade.com](#)).

In our application, we tested for weak-form market efficiency, by analysing pre-match exchange odds, and semi-strong form efficiency, by focusing on the in-play odds after the arrival of the major news that the first goal of a match had been scored. The results suggested a reverse favourite-longshot bias for both pre-match and in-play odds. This is opposite to the findings from fixed-odds bookmaker markets, where the evidence in favour of the favourite-longshot bias has been widely documented (e.g. [Kuypers, 2000](#); [Direr, 2013](#); [Angelini and De Angelis, 2019](#)). The reverse bias on the betting exchange created profitable opportunities that could have been exploited by simple betting strategies. By wagering on longshots, we showed that substantial positive returns were possible both from betting before and during matches.

From our analysis of in-play pricing, we also tested for the presence of behavioural biases, focusing on how market participants reacted to major news. We found evidence that prior bias was not significantly updated only in the case when a favourite team scored at the beginning of a match. Conversely, when the first goal was scored by either a longshot or a favourite team later in a match, the response tended to either amplify or absorb the initial pre-match mispricing, depending on conditions. In particular, if the news was somehow expected by the market, then prices reacted toward no bias, i.e. when a favourite team scored first. On the other hand, when the news was unexpected or a surprise, then the market biases tended to increase, i.e. when the first goal was scored by a longshot. Moreover, in the case that no pre-match bias was detected in prices, the arrival of the first goal created mispricing. Empirical evidence of home bias and confirmation bias in these markets was

also found.

We hope that the methodological framework proposed here is simple to apply, and that other researchers will do so, either on larger datasets of football match final result markets, thus corroborating - or possibly overturning - the results here, or on entirely different in-play markets. The results from our application of the methodology suggest that more caution may be needed when inferring behavioural biases of bettors from bookmaker pricing. At the very least, further research is needed to uncover how far and why the suggested behavioural biases of the participants in betting exchanges differ from what bookmaker pricing suggests.

References

- Ali, M. M.** 1977. “Probability and Utility Estimates for Racetrack Bettors.” *Journal of Political Economy*, 85(4): 803–815.
- Angelini, G., and L. De Angelis.** 2019. “Efficiency of online football betting markets.” *International Journal of Forecasting*, 35(2): 712–721.
- Brooks, R. M., A. Patel, and T. Su.** 2003. “How the Equity Market Responds to Unanticipated Events.” *The Journal of Business*, 76(1): 109–134.
- Brown, A.** 2013. “Information Processing Constraints and Asset Mispricing.” *The Economic Journal*, 124(575): 245–268.
- Brown, A.** 2015. “Information Acquisition in Ostensibly Efficient Markets.” *Economica*, 82(327): 420–447.
- Bruce, A. C., J. E. V. Johnson, J. D. Peirson, and J. Yu.** 2009. “An Examination of the Determinants of Biased Behaviour in a Market for State Contingent Claims.” *Economica*, 76(302): 282–303.
- Cain, M., D. Law, and D. Peel.** 2000. “The Favourite-Longshot Bias and Market Efficiency in UK Football Betting.” *Scottish Journal of Political Economy*, 47(1): 25–36.
- Chan, W. S.** 2003. “Stock price reaction to news and no-news: drift and reversal after headlines.” *Journal of Financial Economics*, 70(2): 223–260.
- Choi, D., and S. K. Hui.** 2014. “The role of surprise: Understanding overreaction and underreaction to unanticipated events using in-play soccer betting market.” *Journal of Economic Behavior & Organization*, 107(PB): 614–629.
- Cowgill, B., and E. Zitzewitz.** 2015. “Corporate Prediction Markets: Evidence from Google, Ford, and Firm X.” *The Review of Economic Studies*, 82(4): 1309–1341.
- Croxson, K., and J. J. Reade.** 2014. “Information and Efficiency: Goal Arrival in Soccer Betting.” *The Economic Journal*, 124(575): 62–91.
- De Bondt, W. F. M., and R. Thaler.** 1985. “Does the Stock Market Overreact?” *Journal of Finance*, 40(3): 793–805.
- De Bondt, W. F. M., and R. Thaler.** 1990. “Do Security Analysts Overreact?” *American Economic Review*, 80(2): 52–57.
- Deschamps, B., and O. Gergaud.** 2007. “Efficiency in Betting Markets: Evidence from English Football.” *Journal of Prediction Markets*, 1(1): 61–73.
- Deutscher, C., B. Frick, and M. Ötting.** 2018. “Betting market inefficiencies are short-lived in German professional football.” *Applied Economics*, 50(30): 3240–3246.
- Direr, A.** 2013. “Are betting markets efficient? Evidence from European Football Championships.” *Applied Economics*, 45(3): 343–356.
- Fama, E. F.** 1965. “The Behavior of Stock-Market Prices.” *The Journal of Business*, 38(1): 34–105.
- Fama, E. F.** 1970. “Efficient Capital Markets: A Review of Theory and Empirical Work.” *Journal of Finance*, 25(2): 383–417.

- Flepp, R., S. Nüesch, and E. Franck.** 2017. “The liquidity advantage of the quote-driven market: Evidence from the betting industry.” *The Quarterly Review of Economics and Finance*, 64 306–317.
- Forrest, D., J. Goddard, and R. Simmons.** 2005. “Odds-setters as forecasters: The case of English football.” *International Journal of Forecasting*, 21(3): 551–564.
- Franck, E., E. Verbeek, and S. Nüesch.** 2010. “Prediction accuracy of different market structures – bookmakers versus a betting exchange.” *International Journal of Forecasting*, 26(3): 448–459.
- Franck, E., E. Verbeek, and S. Nüesch.** 2013. “Inter-market Arbitrage in Betting.” *Economica*, 80(318): 300–325.
- Golec, J., and M. Tamarkin.** 1991. “The degree of inefficiency in the football betting market: Statistical tests.” *Journal of Financial Economics*, 30(2): 311–323.
- He, X.-Z., and N. Treich.** 2017. “Prediction market prices under risk aversion and heterogeneous beliefs.” *Journal of Mathematical Economics*, 70(C): 105–114.
- Ioannidis, C., and D. Peel.** 2005. “Testing for market efficiency in gambling markets when the errors are non-normal and heteroskedastic an application of the wild bootstrap.” *Economics Letters*, 87(2): 221–226.
- Kahneman, D., and A. Tversky.** 1973. “On the psychology of prediction.” *Psychological Review*, 80(4): 237–251.
- Koessler, F., C. Noussair, and A. Ziegelmeyer.** 2012. “Information aggregation and belief elicitation in experimental parimutuel betting markets.” *Journal of Economic Behavior & Organization*, 83(2): 195 – 208.
- Kuypers, T.** 2000. “Information and efficiency: an empirical study of a fixed odds betting market.” *Applied Economics*, 32(11): 1353–1363.
- Levitt, S. D.** 2004. “Why are gambling markets organised so differently from financial markets?” *The Economic Journal*, 114(495): 223–246.
- Lim, K.-P., and R. Brooks.** 2011. “The Evolution Of Stock Market Efficiency Over Time: A Survey of the Empirical Literature.” *Journal of Economic Surveys*, 25(1): 69–108.
- List, J. A.** 2004. “Testing Neoclassical Competitive Theory in Multilateral Decentralized Markets.” *Journal of Political Economy*, 112(5): 1131–1156.
- Malkiel, B. G.** 2003. “The Efficient Market Hypothesis and Its Critics.” *Journal of Economic Perspectives*, 17(1): 59–82.
- Manski, C.** 2006. “Interpreting the predictions of prediction markets.” *Economics Letters*, 91(3): 425–429.
- Mincer, J., and V. Zarnowitz.** 1969. “The evaluation of economic forecasts.” In *Economic Forecasts and Expectations: Analysis of Forecasting Behavior and Performance*. NBER, 1–46.
- Na, S., Y. Su, and T. Kunkel.** 2019. “Do not bet on your favourite football team: the influence of fan identity-based biases and sport context knowledge on game prediction accuracy.” *European Sport Management Quarterly*, 19(3): 396–418.
- Nevill, A. M., and R. L. Holder.** 1999. “Home advantage in sport.” *Sports Medicine*, 28(4): 221–236.

- Newall, P. W. S., and D. Cortis.** 2021. “Are Sports Bettors Biased toward Longshots, Favorites, or Both? A Literature Review.” *Risks*, 9(1): .
- Ottaviani, M., and P. N. Sørensen.** 2008. “The Favorite-Longshot Bias: An Overview of the Main Explanations.” In *Handbook of Sports and Lottery Markets*. Eds. by D. B. Hausch, and W. T. Ziemba, San Diego Elsevier, 83 – 101.
- Ottaviani, M., and P. N. Sørensen.** 2015. “Price Reaction to Information with Heterogeneous Beliefs and Wealth Effects: Underreaction, Momentum, and Reversal.” *American Economic Review*, 105(1): 1–34.
- Page, L., and R. T. Clemen.** 2012. “Do Prediction Markets Produce Well-Calibrated Probability Forecasts?” *The Economic Journal*, 123(568): 491–513.
- Plott, C. R., J. Wit, and W. C. Yang.** 2003. “Parimutuel betting markets as information aggregation devices: experimental results.” *Economic Theory*, 22(2): 311–351.
- Plott, C., and S. Sunder.** 1988. “Rational Expectations and the Aggregation of Diverse Information in Laboratory Security Markets.” *Econometrica*, 56(5): 1085–1118.
- Pope, P. F., and D. A. Peel.** 1989. “Information, Prices and Efficiency in a Fixed-Odds Betting Market.” *Economica*, 56(223): 323–341.
- Rabin, M.** 1998. “Psychology and Economics.” *Journal of Economic Literature*, 36(1): 11–46.
- Reade, J.** 2014. “Information And Predictability: Bookmakers, Prediction Markets And Tipsters As Forecasters.” *Journal of Prediction Markets*, 8(1): 43–76.
- Smith, M. A., D. Paton, and L. Vaughan Williams.** 2006. “Market Efficiency in Person-to-Person Betting.” *Economica*, 73(292): 673–689.
- Smith, M. A., D. Paton, and L. Vaughan Williams.** 2009. “Do bookmakers possess superior skills to bettors in predicting outcomes?” *Journal of Economic Behavior & Organization*, 71(2): 539–549.
- Snowberg, E., and J. Wolfers.** 2010. “Explaining the Favorite-Long Shot Bias: Is it Risk-Love or Misperceptions?” *Journal of Political Economy*, 118(4): 723–746.
- Thaler, R., and W. T. Ziemba.** 1988. “Parimutuel Betting Markets: Racetracks and Lotteries.” *Journal of Economic Perspectives*, 2(2): 161–74.
- Vaughan Williams, L.** 2005. *Information Efficiency in Financial and Betting Markets*. Cambridge University Press.
- Vaughan Williams, L., M. Sung, P. A. F. Fraser-Mackenzie, J. Peirson, and J. E. V. Johnson.** 2018. “Towards an Understanding of the Origins of the Favourite–Longshot Bias: Evidence from Online Poker Markets, a Real-money Natural Laboratory.” *Economica*, 85(338): 360–382.
- Vlastakis, N., G. Dotsis, and R. N. Markellos.** 2009. “How efficient is the European football betting market? Evidence from arbitrage and trading strategies.” *Journal of Forecasting*, 28(5): 426–444.
- Wason, P. C.** 1960. “On the failure to eliminate hypotheses in a conceptual task.” *Quarterly Journal of Experimental Psychology*, 12(3): 129–140.
- Wolfers, J., and E. Zitzewitz.** 2006. “Interpreting Prediction Market Prices as Probabilities.” NBER Working Papers 12200, National Bureau of Economic Research, Inc.

- Woodland, L. M., and B. M. Woodland.** 1994. “Market Efficiency and the Favorite-Longshot Bias: The Baseball Betting Market.” *Journal of Finance*, 49(1): 269–279.
- Ziemba, W. T.** 2020. “Parimutuel betting markets: racetracks and lotteries revisited.” SRC Discussion Paper 103, Systemic Risk Centre, London School of Economics.


Appendix A Additional Tables and Figures

Table A1: In-play analysis of market mispricing when the first goal is scored, robustness check: using $p_{i,t-1}$ instead of $p_{i,0}$ in Equation (7)

		<i>Home Goal</i>			<i>Away Goal</i>		
		$h = 2$	$h = 5$	$h = 30$	$h = 2$	$h = 5$	$h = 30$
<i>Home Odds</i>	$\hat{\gamma}_0$	0.0076 (0.9099)	-0.00379 (0.5789)	-0.0427 (0.5557)	0.0767 (0.1032)	0.0988** (0.0409)	0.0586 (0.2120)
	$\hat{\gamma}_1$	0.0011** (0.0133)	0.0010** (0.0306)	0.0010** (0.0231)	-0.0001 (0.7968)	-0.0002 (0.5174)	-0.0002 (0.5736)
	$1000\hat{\gamma}_2$	-0.0018** (0.0323)	-0.0014 (0.1093)	-0.0016** (0.0649)	0.0001 (0.8725)	-0.0001 (0.8736)	0.0002 (0.7047)
	$\hat{\beta}$	-0.2251** (0.0203)	-0.1633* (0.0972)	-0.1422 (0.1822)	-0.2610** (0.0112)	-0.2174** (0.0398)	-0.1649* (0.0643)
	F -test	4.0970*** (0.0028)	3.5161*** (0.0076)	2.7521** (0.0278)	1.6810 (0.1538)	1.3385 (0.2553)	0.9581 (0.4311)
<i>Away Odds</i>	$\hat{\gamma}_0$	0.0384 (0.1608)	0.0510* (0.0771)	0.0260 (0.3108)	0.1023 (0.1612)	-0.0072 (0.9196)	0.0372 (0.6347)
	$\hat{\gamma}_1$	-0.0002 (0.3592)	-0.0003 (0.1784)	-0.0003* (0.0827)	0.0006 (0.3117)	0.0006 (0.2748)	0.0006 (0.3026)
	$1000\hat{\gamma}_2$	0.0002 (0.4945)	0.0004 (0.3010)	0.0005* (0.0859)	-0.0008 (0.4762)	-0.0008 (0.4444)	-0.0008 (0.4734)
	$\hat{\beta}$	-0.1069 (0.2045)	-0.0560 (0.5190)	0.0143 (0.8489)	-0.3239** (0.0217)	-0.1253 (0.3572)	-0.2083 (0.1667)
	F -test	0.7042 (0.5893)	0.8334 (0.5043)	0.9162 (0.4543)	3.9701*** (0.0036)	1.1262 (0.3439)	1.9207 (0.1074)


Notes: author calculations using within match Betfair result odds from $\tilde{n}_H = 513$ ‘home team goal’ matches and $\tilde{n}_A = 369$ ‘away team goal’ matches. Presents WLS estimates of Equation (7). p -values in parenthesis. ***, **, * indicate significance at the 1%, 5% and 10% levels, respectively, two-sided tests. F -test displays the test statistic for the test of the null hypothesis in (8).

Figure A1: Mean of the odds-implied outcome probability jump after the first goal is scored at tick t and for pre-match probability $p_{i,0}$



Notes: author calculations using the $\hat{n}_H = 513$ 'home team goal' matches and $\hat{n}_A = 369$ 'away team goal' matches. Black cells indicate that there are no events in the data set for the corresponding combination of tick and pre-match outcome probability.

Figure A2: Estimated efficiency curves for the first goal scored at tick t and for pre-match probabilities $p_{i,0}$



Notes: author calculations using the $\tilde{n}_H = 513$ 'home team goal' matches and $\tilde{n}_A = 369$ 'away team goal' matches. The blue plane shows the estimated efficiency curves as per Equation (9) and in red the related 90% confidence intervals as per (10). In black the zero plane is depicted, and in yellow the pre-match efficiency curve is also shown (see also Table 2).