The Wisdom of No Crowds:
The Reaction of Betting Markets to Lockdown Soccer Games

Tadgh Hegarty* Karl Whelan†

University College Dublin

April 2022

Abstract

The support of home spectators is one of the contributing factors to the home advantage effect in sports matches. The Covid-19 pandemic led to European soccer matches being played without spectators. We show that betting markets adjusted swiftly to account for a reduced home advantage in both goal difference and the probability of a win. These adjustments proved accurate over a large sample of soccer matches subsequently played without spectators even though the earliest games appeared to suggest a much bigger change in home advantage.

Keywords: Market Efficiency, Home Advantage, Soccer, COVID-19

JEL Classification: G14, L83, Z20, Z21

*tadgh.hegarty@ucdconnect.ie

†karl.whelan@ucd.ie. We are grateful to Chris Jepsen and Martina Lawless for comments on an earlier draft

1. Introduction

Since the pioneering work of Eugene Fama (1970), economists have debated whether financial market prices are efficient, in the sense of incorporating all relevant information. Many empirical studies have used sports betting markets to examine information efficiency (Williams, 1999). Like traditional financial markets, there are large number of participants and the total amounts at stake in betting markets are generally large. Unlike many financial markets, there is a definite outcome that determines the settlement of each transaction and the data on betting market prices and the outcomes of sports events are readily available. These characteristics make betting markets an ideal natural setting to test market efficiency (Thaler and Ziemba, 1988).

The Covid-19 pandemic provided a natural experiment for assessing how betting markets adapt to process new information. In this case, betting markets had to adjust to the fact that the pandemic led to a period in which most sporting events took place with no spectators present. The absence of spectators mattered for betting markets because the existence of a "home advantage" effect in sports has been well documented. While the precise contribution to this home advantage effect of the majority of spectators supporting the home team was not known, market participants are likely to have assessed it as an important contributing factor (Goumas, 2014; Ponzo and Scoppa, 2018). For this reason, the absence of spectators was something that betting markets needed to consider when setting odds during this period.

There have been some initial empirical studies investigating efficiency in soccer betting markets during the lockdown period and their findings have pointed towards the possibility of market inefficiencies in the pricing of home advantage. Winkelmann et al. (2021) argue that bookmakers mispriced the reduction in home advantage for initial Bundesliga matches played behind closed doors. Meier et al. (2021) argue temporary inefficiencies existed for early lockdown matches, concluding that in the short run, betting markets were not efficient. Fischer and Haucap (2021) suggest that the strength of the home advantage effect changed over the course of the lockdown period but that betting markets did not adjust their pricing of this element, also suggesting a possible inefficiency.

In this paper, we using data from the top four European soccer leagues to estimate full-sample regressions and also recursive regressions for outcomes and for betting odds during the period in which there were no spectators present. The recursive regression approach allows us to see how the evidence evolved on the impact of the absence of crowds on home advantage and also to see how market makers reacted to this evidence.

We find that betting markets made a swift adjustment to account for a reduced home advantage and that this adjustment—a reduction of about half of the goal difference previously associated with home advantage—proved to be very close to the size of the reduction in home advantage implied by the data once a reasonably large number of games had been played. We find that the data from the earliest games (most notably the early lockdown games in the German Bundesliga) suggested a far larger reduction in home advantage than would later prevail but betting markets did not react to the extreme outcomes in these early games. Overall, the evidence presented here is consistent with the patterns reported in earlier papers but is more suggestive of market efficiency. In particular, market makers appear to have had a good idea from the start of the underlying impact of the removal of crowds, with the evidence eventually supporting this initial assessment.

The paper is structured as follows. Section 2 describes our methodology and presents results for betting on goal difference in soccer matches as well as goal difference outcomes. Section 3 presents our analysis for wins and losses instead of goal difference. Section 4 presents more direct tests of market efficiency and reconciles our findings with other studies. Section 5 presents some concluding thoughts on our results.

2. Home Goal Advantage and Asian Handicap

The first outcome we examine is the goal difference in matches. These outcomes can be compared with the Asian Handicaps which are driven by the market's estimate of the expected goal difference between the two sides.

The Asian Handicap betting market is the most popular form of betting for European soccer matches, accounting for approximately 70% of total betting turnover (Kerr, 2018). This form of betting attempts to equalise the chances of a winning bet across the two teams by adding a hypothetical goal advantage in favour of the weaker team before kick off. To give an example, an Asian handicap of -1.5 gives 1.5 goals to the away team pre-match, requiring the home team to win by 2 goals or more for a bet on the home side to be successful. The handicaps are offered in increments of 0.25 goals, such that a handicap of -1.25 represents half the wager at -1.0 and half at -1.5. With the chance of each bet roughly equalised, bookmakers offer odds close to even money (meaning the winner profits via an amount equal to their bet) for both teams. This means the Asian handicap for each game is driven by the market estimate of the expected goal difference between home and away sides, up to the closest quarter of a goal. This market is popular with professional betting syndicates and offers the lowest margin to bettors (Hassanniakalager and Newall, 2019). Despite the popularity and relative size of the Asian handicap market, only a handful of papers have previously analyzed it (Constantinou, 2020).

Our dataset contains all 4,338 matches from the 2018/19 season to the 2020/21 season in the top divisions of the four major European soccer leagues (Germany, Italy, England and Spain). Data on the outcomes of the matches and betting market pricing are obtained from www.football-data.co.uk. We combine the market information for each game with the corresponding pre-match values for FiveThirtyEight's Soccer Power Index (SPI) ratings for both teams. These ratings are generated from a complex model and are recursively updated throughout the season, providing a time-varying measure of the fundamental strength of each team (Boyce, 2018).[1]

In early 2020, due to the COVID-19 pandemic, decisions were made to restart the competitions without spectators. Later, small numbers of attendees were permitted for certain matches, as restrictions eased over the period of our study. Table 1 displays the dates each league moved to behind closed doors with column *No Crowd* displaying the number of matches played with no spectators. The *Partial* column shows the number of matches with greatly reduced partial attendances for each league. Across these four leagues, there were a total of 1862 matches with either no spectators or else a small number.

[1] SPI ratings are available for download from https://projects.fivethirtyeight.com.

Table 1: Number of Lockdown Matches Per League

League	Behind closed doors from	No Crowd	Partial
Germany Bundesliga 1	March 11th 2020	355	34
Italy Serie A	March 8th 2020	501	9
England Premier League	June 17th 2020	441	31
Spain La Liga Primera	March 10th 2020	486	5
Total		1783	79

We estimate the effect of the absence of spectators on the actual goal difference between home and away sides using the following specification.

$$GD_j = \beta_0 + \beta_1 LD_j + \beta_2 SPIdiff_j + \sum_{k=1}^{3} \beta_{k+2} L_k + \epsilon_h \tag{1}$$

where GD_j is the number of goals in game j scored by the home team minus the number of goals scored by the away team. The variable LD_j indicates whether game j took place during lockdowns with either no spectators or a greatly reduced number. To measure the effect of the fundamental strength of each team, we incorporate $SPIdiff_j$, which is the home team's *SPI* rating minus the away team's rating. Finally, to control for potential differences in patterns of home advantage across leagues, we use dummy variables L_k for three of the leagues with the Premier League omitted as the reference league.

We report results for estimating (1) over a full sample and we also use a recursive approach, increasing window size and re-estimating the model. The coefficient of interest is β_1, which represents the change in home goal advantage for lockdown games. We start estimating β_1 after the first 40 games without spectators and then recursively add windows of 40 games sequentially, re-estimating the coefficient β_1 each time. We choose 40 matches because this corresponds to approximately equivalent to one full round of weekly matches in the four leagues. The results were robust to using either a smaller or larger window.

Separately, we model the Asian Handicap market prediction. MP_j is the Asian Handicap related to the expected number of goals the home team will score in match j minus the expected number of goals the away team will score. We explain this variable using the same model used for goal difference outcomes.

$$MP_j = \beta_0 + \beta_1 LD_j + \beta_2 SPIdiff_j + \sum_{k=1}^{3} \beta_{k+2} L_k + \epsilon_h \tag{2}$$

We estimate this model over the full sample and also use the same recursive method, estimating β_1

after the first 40 games without spectators and then recursively add windows of 40 games sequentially, re-estimating the coefficient β_1 each time. Both models are estimated via OLS using robust standard errors.

Table 2 shows a comparison of the estimation results for models (1) and (2) on all 4338 matches. The value of the constant coefficient (β_0) in model (1) represents the estimated home advantage prior to the Covid lockdown. We estimate an approximate advantage to the home side of one quarter of a goal on average. The coefficient β_1 for model (1) of -0.128 is statistically significant and substantively significant since it represents a reduction in home goal advantage for games played behind closed doors of almost 50%. The estimates of (2) show that over the lockdown period, betting markets priced a reduction in home advantage of -0.153 per match, also a reduction of about a half. The sampling variation in betting odds is considerably smaller than in match outcomes, so this estimate has a much lower standard error than the estimate of β_1 from the goal difference outcome regressions. However, the difference between the point estimate of the reduction in home advantage estimated from outcomes in model (1) and the estimate of this effect in betting odds is below half of the standard error of β_1 estimated in model (1). These results imply we cannot reject that betting odds priced the reduction in home advantage in a manner consistent with the actual outcomes.

Table 2: Estimation results for goal difference outcome and Asian handicap. Premier League dummy excluded.

	(1)		(2)	
	Goal difference		Asian Handicap	
Constant (β_0)	0.259***	(0.0534)	0.313***	(0.00970)
Lockdown (β_1)	-0.128*	(0.0503)	-0.153***	(0.00915)
SPIdiff (β_2)	0.0593***	(0.00180)	0.0600***	(0.000381)
Bundesliga	0.0689	(0.0795)	0.0135	(0.0133)
Serie A	0.0206	(0.0679)	0.00338	(0.0121)
La Liga	0.101	(0.0663)	0.0595***	(0.0131)
N	4338		4338	
R^2	0.212		0.892	

Standard errors in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Figure 1 shows the β_1 coefficients obtained from our recursive regression approach to illustrate how the evidence on home advantage evolved and how betting markets reacted. It shows there was an immediate reaction from the betting market (the blue dots) with the estimated effect of home advantage on goal difference odds for the first 40 games falling by about 0.10 per game. This reduction was consistent with the subsequent estimate of the actual effect during these initial games (the first red dot). However, this was followed by a period in which the estimated reduction in home advantage based on the games played without spectators jumped to over 0.5 for samples based on the first 80 and 120 games.

Adding more data sees the estimated impact of the absence of spectators on goal difference outcomes converge over the next few weeks towards the full-sample estimate of about -0.13. Notably, however, the betting odds did not adjust much at all in response to the early results, with the estimated β_1 settling down pretty quickly at close to its final estimated figure. Overall, throughout the lockdown period, betting markets did not adjust their estimate of the home change in home advantage by a substantial amount: The mean estimate in our recursive samples is -0.14, the minimum is -0.16 and maximum is -0.10.

Figure 2 shows the two sets of estimated β_1 coefficients while also incorporating 95% confidence intervals.[2] The difference in scale of the two charts should be noted. The confidence intervals tighten as the sample sizes increase but the intervals for the estimates based on goal difference outcomes are far wider than those based on the betting odds. Indeed, in only one of the goal difference outcome samples would the hypothesis of β_1 equalling the final betting odds estimate be rejected at a 95% confidence level.

How should these results be interpreted? The initial estimate of the reduced home advantage effect by betting markets in the opening round of games proved relatively close to the final estimates. While results from the early rounds of games without crowds pointed to a much larger reduction in home advantage, there were a number of reasons why betting markets may have largely ignored these outcomes.

First, these early estimates were based on small samples and thus subject to significant sampling error. Second, the large point estimates from the initial games played perhaps strained credulity. An estimate of $\beta_1 < -0.5$ implies a complete over-turning of home advantage so that the away teams would have an advantage larger than the estimated pre-Covid home advantage. Given the other elements beyond crowd support that likely influence home advantage (such as the home team having less far to travel and their greater familiarity with the home ground) this complete reversal was perhaps unlikely to reflect the underlying reality. Finally, the large initial values for β_1 from the early goal difference outcome regressions were heavily influenced by results in one league, the

[2]The standard error bands were based on adding and subtracting 1.96 times the standard error. Confidence intervals based on bootstrap procedures produced similar results.

German Bundesliga, which had a very ususual set of results over this period: In the first six rounds of games without spectators, there were only 11 home wins from 54 games. Whatever the relative merits of these different factors, the subsequent evidence suggests betting markets were probably correct to not over-react to these early events.

Figure 3 shows, separately for each of the four leagues, the point estimates of β_1 in both the goal difference outcome and the Asian handicap recursive regressions. In general, data for each of the leagues support the hypothesis that betting markets moved relatively quickly to price the impact of the absence of spectators with the data for outcomes in each league converging to close to the estimate settled on in betting markets. For three of the four leagues, the estimated β_1 values associated with goal difference outcomes converge closely to the estimate implied from early betting market prices (though betting markets were a little slower to adjust for the English Premier League, where the initial outcomes suggested a possible strengthening of home advantage.) For the Italian Serie A, the final estimates suggested no impact of the absence of spectators on home advantage, with an estimated β_1 close to zero. Interestingly, however, the betting markets continued pricing these games using a similar adjustment for reduced home advantage to those made for the other leagues. Overall, the data suggest that betting markets were using a full range of evidence across leagues to determine pricing, rather than reacting to events on a league-by-league basis.

Figure 1: Asian handicap β_1 coefficients compared with goal difference outcome β_1 coefficients

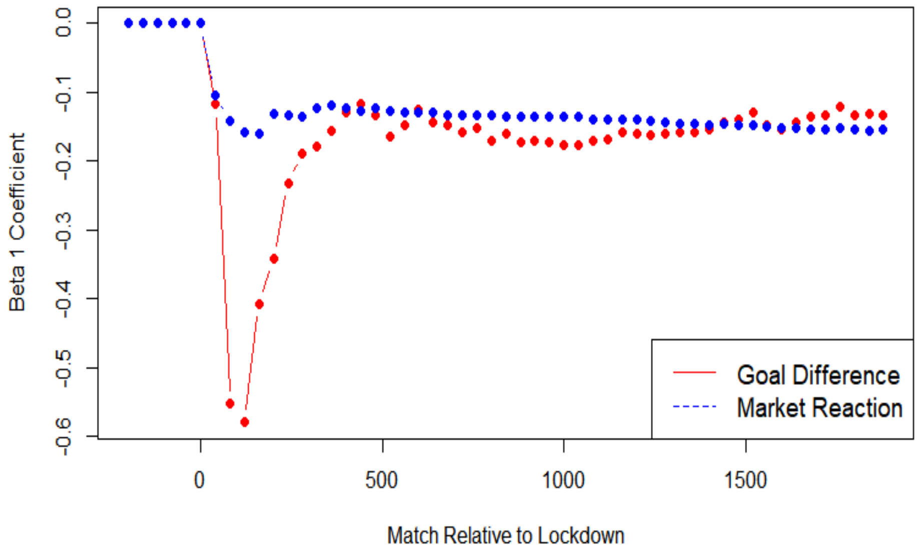

Figure 2: Confidence bounds for estimates of β_1

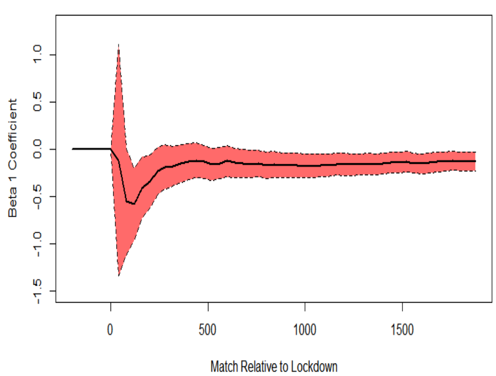

(a) Goal Difference

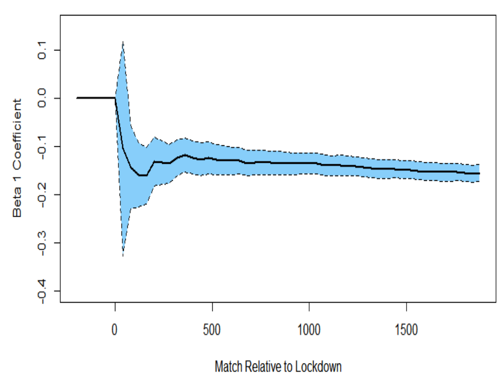

(b) Market Reaction

Figure 3: Asian handicap β_1 coefficients compared with goal difference outcome β_1 coefficients for each league

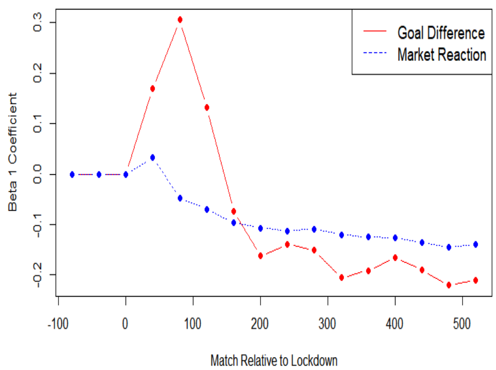

(a) English Premier League

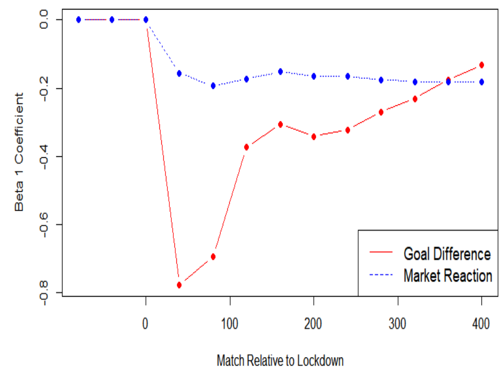

(b) Germany Bundesliga

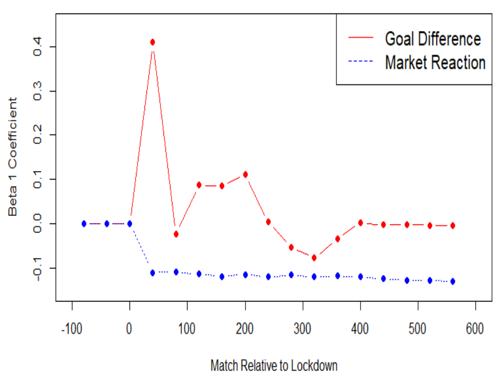

(c) Italy Serie A

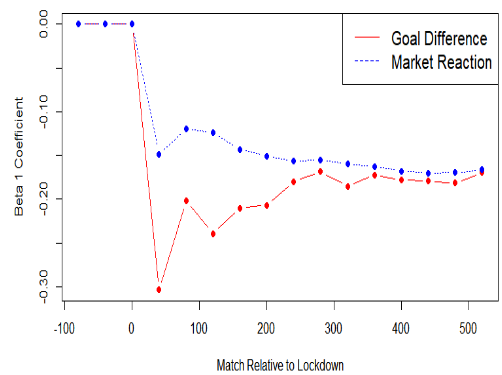

(d) Spain La Liga

3. Home Wins and Betting Win Probabilities

We also model the effect of the absence of spectators on how home advantage influences wins and losses. Following the established literature (Ponzo and Scoppa, 2018; Meier et al., 2021; Winkelmann et al., 2021) we use two data points for each match, W_{ij}, the first data point equalling one if the first team won and zero otherwise and the second data point equalling one if the second team won and zero otherwise. To give a concrete example, if Team A play Team C, we enter a data point W_{AC} which is one if Team A won and zero otherwise and a data point W_{CA} which is one if Team C won and zero otherwise. The following specification is estimated.

$$W_{ij} = \beta_0 + \beta_1 Home_i + \beta_2 Home_i \times LD_{ij} + \beta_3 SPIdiff_{ij} + \sum_{k=1}^{3} \beta_{k+3} L_k + \epsilon_i \tag{3}$$

$Home_i$ indicates if Team i played at home and an interaction term $Home_i \times LD_{ij}$ indicates if Team i played on home field during lockdown conditions. The associated coefficient, β_2, is a measure of the reduction in the likelihood of home wins over the lockdown period. As before, we use $SPIdiff_{ij}$ to control for both the fundamental strength of each participating team i. The specification is estimated via a Probit model with robust standard errors and clustering at the game level.

For betting markets, we estimate the implied probability the market assigns to wins for the home and away teams. Bookmaker odds can be represented as a positive number. For example, decimal odds of 2.50 for a team to win will return 250 for a successful 100 wager if that team wins (a profit of 150). These odds can be used to derive initial probability estimates by calculating their inverse. So, for example, the decimal odds of 2.50 imply an initial probability estimate of 0.4 that the team will win. Because bookmaker odds are affected by profit margins, these initial estimated probabilities do not sum to one. To address this issue, we calculate new probability estimates P_{ij} for both home and away wins for each match by dividing the initial probability estimates by the sum of these initial estimates for the three possible match outcomes.

We model the portion of probability the market assigns to team i beating team j with a specification identical to that for home and away wins.

$$P_{ij} = \beta_0 + \beta_1 Home_i + \beta_2 Home_i \times LD_j + \beta_3 SPIdiff_{ij} + \sum_{k=1}^{3} \beta_{k+3} L_k + \epsilon_i \tag{4}$$

Here, the coefficient of interest is β_2, which measures the change in proportion of market odds given to the home team during lockdown games. Because these are "proportions data" taking values between zero and one, the model is estimated via GLM with a logit link and binomial family as recommended for proportions data by Papke and Wooldridge (1996).[3] As with the goal difference models,

[3] See also the Stata FAQ on this topic by Allen McDowell and Nicholas Cox, available at https://www.stata.com/support/faqs/statistics/logit-transformation/

we estimate (3) and (4) for the full sample and also using a recursive method, estimating β_2 after the first 40 games without crowds and then adding windows of 40 games sequentially, re-estimating the coefficient β_2 each time. Table 3 shows that the full-sample estimates. Figure 4 reports the point estimates of β_2 from the recursive regressions and Figure 5 displays the confidence intervals.

Overall, the results are similar to those for goal difference. Table 3 shows that the full-sample estimate for the reduction in home advantage as estimated by betting markets is close to that estimated from the regression for win outcomes. Figure 4 shows a similar pattern in terms of the evolution of evidence on the home advantage effect and the betting market's reaction. The early data on home and away wins suggest a complete reversal of the home advantage effect so that the advantage switches to away teams. Again, the betting market does not react to these extreme outcomes and the estimates based on home and away wins outcomes converge relatively quickly to be close to the betting market estimates.

One small difference from the goal difference analysis is that the initial reduction in the odds of a home win in the first 40 games is not as proportionally large as the initial reduction in Asian Handicap estimates shown in Figure 1. By the second set of 40 games, however, the betting market estimate of this effect is close to its final level.

Table 3: Estimation results for win outcomes and win probabilities. Premier League dummy excluded.

	(3)		(4)	
	Match Outcome		Win Probability	
Constant (β_0)	-0.479***	(0.0260)	-0.862***	(0.00453)
Home (β_1)	0.373***	(0.0400)	0.612***	(0.00816)
Home x Lockdown (β_2)	-0.116**	(0.0408)	-0.138***	(0.00808)
SPIdiff (β_3)	0.0384***	(0.00139)	0.0599***	(0.000308)
Bundesliga	-0.0411	(0.0280)	0.0200***	(0.00304)
Serie A	-0.0665*	(0.0269)	-0.0192***	(0.00312)
La Liga	-0.0885***	(0.0267)	-0.0214***	(0.00311)
N	8676		8676	
AIC	10034.1		6846.6	

Standard errors in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Figure 4: Win probability β_2 coefficients compared with win outcome β_2 coefficients

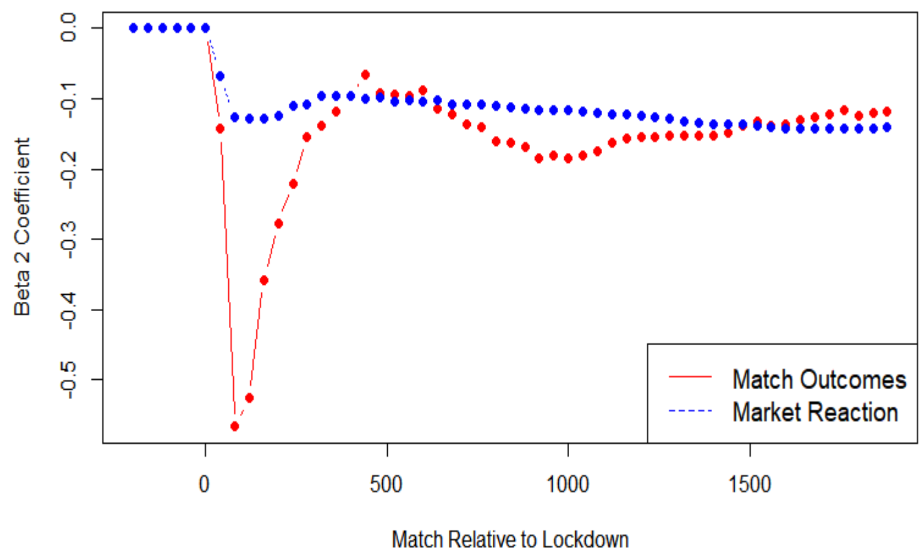

Figure 5: Confidence bounds for estimates of β_2

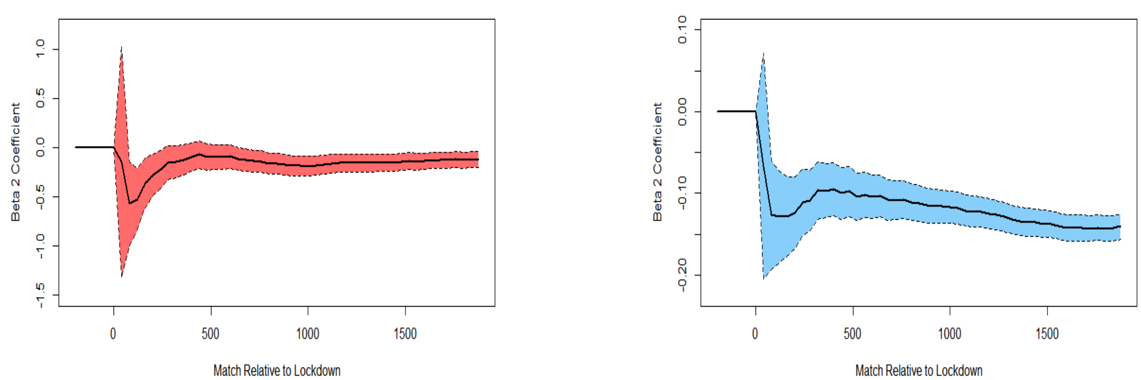

(a) Home Wins

(b) Market Reaction

4. Efficiency Tests

Much of the traditional literature on sports betting has focused on testing the efficiency of the odds using regressions of match outcomes on the relevant betting odds. If the betting markets are efficient, then it should not be possible to use other publicly available variables to improve on the forecast implicit in the odds. Here we present some efficiency tests of this sort, again using our full sample and recursive samples. Specifically, we add the betting market probabilities to the win-loss specification, equation (4):

$$W_{ij} = \beta_0 + \beta_1 P_{ij} + \beta_2 Home_i + \beta_3 Home_i \times LD_j + \beta_4 SPIdiff_{ij} + \sum_{k=1}^{4} \beta_{k+4} L_k + \epsilon_i \tag{5}$$

As before, the specification is estimated via a Probit model with robust standard errors and clustering at the game level.

Table 4 reports the results from estimating this model with the full sample (the last column), with the first 40 observations (the first column) and then a set of recursive samples that each add 300 additional matches. If betting markets systematically under-estimated the impact of the absence of spectators, then this would show up as a negative value for β_3 in this specification. However, Table 4 shows that only two of the recursive samples produce estimates of β_3 that are statistically significant at the 5 percent level. The first sample, with by far the largest estimated coefficient, relates to the first 40 games. The other sample is the one using the first 940 games. Adding more data after that point, the estimated coefficient gets smaller and the t-statistics decline. Table 5 repeats this analysis excluding the German Bundesliga. It shows no statistically significant values for β_3 in any of the recursive samples.

These results generally argue against there being an important informational inefficiency relating to home advantage during this period. How can these results be reconciled with previous findings? One explanation is the longer data set used here. Winkelmann et al. (2021) argue for mis-pricing based on games played before August 2020. Similarly, Meier et al. (2021) argue only for temporary inefficiencies existed for early lockdown matches, concluding that in the short run, betting markets were not efficient. In a specification explaining wins and losses, Fischer and Haucap (2021) use a time trend variable interacted with home advantage to demonstrate that there were changes over time, with home advantage strengthening after the initial lockdown period. They note that betting markets did not feature a similar strengthening in their pricing of home advantage. However, as we demonstrate here, this can be interpreted as betting markets having a good estimate of the underlying effect and the data converging towards this estimate.

Table 4: Testing Efficiency of Betting Odds

	Dependent Variable: Win/Loss Match Outcome						
	(I)	(II)	(III)	(IV)	(V)	(VI)	(VII)
	+40	+340	+640	+940	+1240	+1540	ALL
Bookmaker Prob (β_1)	2.384***	2.290***	2.246***	2.260***	2.346***	2.471***	2.495***
	(0.395)	(0.367)	(0.344)	(0.330)	(0.318)	(0.308)	(0.289)
Home (β_2)	0.0823	0.0882	0.0852	0.0810	0.0704	0.0597	0.0555
	(0.0723)	(0.0676)	(0.0634)	(0.0605)	(0.0582)	(0.0561)	(0.0534)
Home x Lockdown (β_3)	-0.540*	-0.0686	-0.0789	-0.122*	-0.0882	-0.0604	-0.0429
	(0.253)	(0.0781)	(0.0595)	(0.0516)	(0.0469)	(0.0438)	(0.0413)
SPIdiff (β_4)	0.00794	0.00973*	0.0104*	0.0102*	0.00915*	0.00769	0.00731
	(0.00517)	(0.00483)	(0.00454)	(0.00433)	(0.00416)	(0.00402)	(0.00377)
Bundesliga	-0.0235	-0.0317	-0.0411	-0.0554	-0.0409	-0.0444	-0.0492
	(0.0366)	(0.0346)	(0.0333)	(0.0319)	(0.0304)	(0.0294)	(0.0282)
Serie A	-0.0573	-0.0609	-0.0519	-0.0656*	-0.0554	-0.0555*	-0.0606*
	(0.0359)	(0.0340)	(0.0319)	(0.0307)	(0.0293)	(0.0282)	(0.0270)
La Liga	-0.0697	-0.0652	-0.0600	-0.0687*	-0.0712*	-0.0690*	-0.0711**
	(0.0357)	(0.0336)	(0.0319)	(0.0306)	(0.0296)	(0.0285)	(0.0272)
Constant (β_0)	-1.259***	-1.222***	-1.201***	-1.194***	-1.227***	-1.272***	-1.275***
	(0.127)	(0.119)	(0.112)	(0.108)	(0.105)	(0.102)	(0.0964)
N	5016	5616	6216	6816	7416	8016	8676
AIC	5762.0	6437.4	7135.4	7817.4	8501.1	9169.4	9933.8

Standard errors in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table 5: Testing efficiency of betting odds: Excluding Bundesliga

Dependent Variable: Win/Loss Match Outcome						
	(I)	(II)	(III)	(IV)	(V)	(VI)
	+40	+340	+640	+940	1240	ALL
Bookmaker Prob (β_1)	2.654***	2.367***	2.324***	2.443***	2.622***	2.597***
	(0.428)	(0.392)	(0.368)	(0.353)	(0.339)	(0.319)
Home (β_2)	0.0697	0.106	0.0863	0.0761	0.0511	0.0514
	(0.0802)	(0.0737)	(0.0690)	(0.0657)	(0.0628)	(0.0600)
Home x Lockdown (β_3)	-0.154	0.0343	-0.0548	-0.0723	-0.0468	-0.0400
	(0.230)	(0.0787)	(0.0609)	(0.0534)	(0.0490)	(0.0465)
SPIdiff (β_4)	0.00462	0.00924	0.00934	0.00820	0.00622	0.00642
	(0.00565)	(0.00522)	(0.00490)	(0.00465)	(0.00445)	(0.00420)
Serie A	-0.0567	-0.0554	-0.0605	-0.0544	-0.0526	-0.0607*
	(0.0360)	(0.0333)	(0.0311)	(0.0297)	(0.0281)	(0.0270)
La Liga	-0.0723*	-0.0678*	-0.0704*	-0.0658*	-0.0687*	-0.0705**
	(0.0358)	(0.0334)	(0.0310)	(0.0299)	(0.0285)	(0.0273)
Constant (β_0)	-1.356***	-1.268***	-1.226***	-1.273***	-1.327***	-1.313***
	(0.137)	(0.127)	(0.120)	(0.116)	(0.112)	(0.106)
N	3958	4558	5158	5758	6358	6840
AIC	4514.6	5186.3	5912.9	6569.4	7238.3	7797.1

Standard errors in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

5. Conclusion

We have examined the reaction of betting markets to the removal of spectators from high-profile European soccer matches during the COVID-19 pandemic. We find evidence of a swift and accurate adjustment in market pricing to reduce the allowance made for home advantage. This finding contrasts with some recent studies that emphasised the possibility of inefficiencies in the pricing of odds for the early games played without spectators.

The data for the early games played under lockdown suggested a much larger reduction in home advantage than was factored into the betting market's calculations. It is possible that the results during early lockdown games represented a genuine failure by betting markets and that away teams were temporarily given an outright advantage by the new conditions. However, our preferred interpretation is that these early games likely represented a statistical anomaly driven by small samples. As noted above, for almost all of our recursive samples for goal difference and win outcomes, the data could not reject the hypothesis that the lockdown period's home advantage effect equalled the final estimate of this effect derived from betting odds.

Our findings raise the question of how betting markets managed to have such an accurate estimate of the absence of spectators on outcomes. One possibility is that bookmakers had consulted previous academic research on this topic. For example, Ponzo and Scoppa (2018) examined derby matches where both teams usually played in the same stadium but one team had more support due to being designated the home team. They find a significant impact of being supported by home fans but their sample size was small (124 games), the estimated effect had a large standard error and the situation being described was different from the one occurring in lockdown, since fans were present at these matches. More relevant is the possible use of data from previous matches played without spectators. Reade et al. (2020) examine 160 European professional soccer matches which took place without spectators after 2002/2003 but prior to April 2020, usually due to punishment bans related to mis-behaviour by fans. They found that, on average, home advantage eroded for games played behind closed doors but the effects were highly uncertain. For example, their estimated reduction in the effect of home advantage on goal difference was -0.11 with a standard error of -0.15.

So while there was some evidence to draw on to isolate the potential effects of the absence of spectators, it was of limited usefulness for use in formulating betting odds. The limited availability of information makes the performance of betting markets during this period all the more remarkable. Rather than evidence of the failure of markets to adapt, we think the evidence points towards the ability of betting markets to incorporate new information relatively efficiently even in a situation where there is considerable uncertainty.

References

Boyce, J. (2018). *How our club soccer predictions work*. https://fivethirtyeight.com/methodology/how-our-club-soccer-predictions-work/.

Constantinou, A. (2020). Asian handicap football betting with rating-based hybrid bayesian networks. *arXiv preprint arXiv:2003.09384*.

Fischer, K. and Haucap, J. (2021). COVID-19, home advantage in professional soccer, and betting market efficiency. *COVID Economics*, 80:211–231.

Goumas, C. (2014). How does crowd support contribute to home advantage in soccer? *Journal of Sport Behavior*, 37(3):236.

Hassanniakalager, A. and Newall, P. W. (2019). A machine learning perspective on responsible gambling. *Behavioural Public Policy*, pages 1–24.

Kerr, J. (2018). How can legislators protect sport from the integrity threat posed by cryptocurrencies? *The International Sports Law Journal*, 18(1):79–97.

Meier, P. F., Flepp, R., and Franck, E. (2021). Are sports betting markets semistrong efficient? Evidence from the COVID-19 pandemic. *International Journal of Sport Finance*, 16(3):111–126.

Papke, L. E. and Wooldridge, J. (1996). Econometric methods for fractional response variables with an application to 401(k) plan participation rates. *Journal of Applied Econometrics*, 11:619–632.

Ponzo, M. and Scoppa, V. (2018). Does the home advantage depend on crowd support? Evidence from same-stadium derbies. *Journal of Sports Economics*, 19(4):562–582.

Reade, J., Schreyer, D., and Singleton, C. (2020). Echoes: what happens when football is played behind closed doors? *Available at SSRN 3630130*.

Thaler, R. H. and Ziemba, W. T. (1988). Anomalies: Parimutuel betting markets: Racetracks and lotteries. *Journal of Economic Perspectives*, 2(2):161–174.

Williams, L. V. (1999). Information efficiency in betting markets: A survey. *Bulletin of Economic Research*, 51(1):1–39.

Winkelmann, D., Deutscher, C., and Ötting, M. (2021). Bookmakers' mispricing of the disappeared home advantage in the German Bundesliga after the COVID-19 break. *Applied Economics*, pages 1–11.