



Research Repository UCD

Title	Football is becoming more predictable; network analysis of 88 thousand matches in 11 major leagues
Authors(s)	Maimone, Victor Martins, Yasseri, Taha
Publication date	2021-12-15
Publication information	Maimone, Victor Martins, and Taha Yasseri. "Football Is Becoming More Predictable; Network Analysis of 88 Thousand Matches in 11 Major Leagues" 8, no. 12 (December 15, 2021).
Publisher	The Royal Society
Item record/more information	http://hdl.handle.net/10197/12723
Publisher's statement	Published by the Royal Society under the terms of the Creative Commons Attribution License http://creativecommons.org/licenses/by/4.0/ , which permits unrestricted use, provided the original author and source are credited.
Publisher's version (DOI)	10.1098/rsos.210617

Downloaded 2024-04-18 19:07:18

The UCD community has made this article openly available. Please share how this access benefits you. Your story matters! (@ucd_oa)



© Some rights reserved. For more information

Football is becoming boring; Network analysis of 88 thousands matches in 11 major leagues

Victor Martins Maimone^a and Taha Yasseri^{*,a,b}

^aOxford Internet Institute, University of Oxford, Oxford OX1 3JS, UK; ^bAlan Turing Institute, London NW1 2DB, UK

August 27, 2019

Abstract

Football is a major sport with worldwide popularity. In recent years excessive monetization of the game has been argued to have affected the quality of the match in different ways. In one hand playing football has become a high income profession and the players are highly motivated to perform well; on the other hand stronger teams have higher income and therefore they can afford better and more expensive players leading to even stronger appearance in tournaments that can make the game more imbalanced and hence predictable. In this work we take a data-heavy network science approach to measure predictability of football over 26 years in major European leagues. We benchmark our model against betting house predictions and after establishing its robustness we show that over time, the games in major leagues have become more predictable. We provide further support for this observation by showing that inequality between teams has increased in accord with the trends in predictability and the home field advantage has been vanishing ubiquitously. This is a first attempt to study football at a large scale and within a historical framework.

Football | Network | Prediction | Centrality

Playing football is, arguably, fun. So is watching it in stadiums or via public media, and to follow the news and events around it. Football is worthy of extensive studies, as it is played by roughly 250 million players in over 200 countries and dependencies, making it the world's most popular sport (1). European leagues alone are estimated to be worth more than £20 billion (2). The sport itself is estimated to be worth almost thirty times as much (3). It has been argued that the surprise element and unpredictability of football is the key to its popularity (4). A major question in relation to such a massive entertainment enterprise is if it can retain its attractiveness through surprise element or, due to significant recent monetizations, it is becoming more predictable and hence at the risk of losing popularity?

There has previously been a fair amount of research in statistical modelling and forecasting in relation to football. The Prediction models are generally either based on detailed statistics of actions in the pitch (5–8) or on a prior ranking system which estimates the relative strengths of the teams (9–11). Some models have considered team pairs attributes such as the geographical locations (6, 12, 13), and some models have mixed these approaches (14, 15).

A rather new approach in predicting performance is based on machine learning and network science (16, 17). Such methods have been used in relation to sport (18–20) and particularly football (21, 22). Most of the past research in this area however either focuses on inter-team interactions and modelling player behaviour rather than league tournament's results prediction, or are limited in scope—particularly they rarely take a historical approach in order to study the game as an evolving phenomenon (23–25). In the present work we use a network science approach to quantify predictability of football in a simple and robust way and by calculating the measures in 26 years of 11 major European leagues we examine if predictability of football has changed over time.

Modelling predictability

To measure the predictability of football we need to build a simple prediction model. For the sake of simplicity, we limit our analysis to the matches that have a winner and we eliminate the ties from the entirety of this study. Our prediction task therefore is simplified to predicting a home win versus an away win that refer to the events of the host, or the guest team wins the match respectively. To predict the results of a given match, we consider the performance of each of the two competing teams in their past N matches preceding that given match. To be able to compare leagues with different number of teams and matches, we normalize this by the total number of matches played in each season T , and define the model training window as $n = N/T$.

For each team we calculate its accumulated *Dyadic Score* as the fraction of points the team has earned during the window n to the maximum number of the points that the team could win during that window. Using this score as a proxy of the team's strength, we can calculate the difference between the strengths of the two teams prior to their match and train a predictive model which considers this difference as the input.

However, this might be too much of a naive predictive model, given that the two teams have most likely played against different sets of teams and the points that they collected can mean different levels of strength depending on the strength of the teams that they collected the points against. We propose a network based model in order to improve the naive dyadic model and account for this scenario.

*To whom correspondence should be addressed. E-mail: taha.yasseri@oii.ox.ac.uk

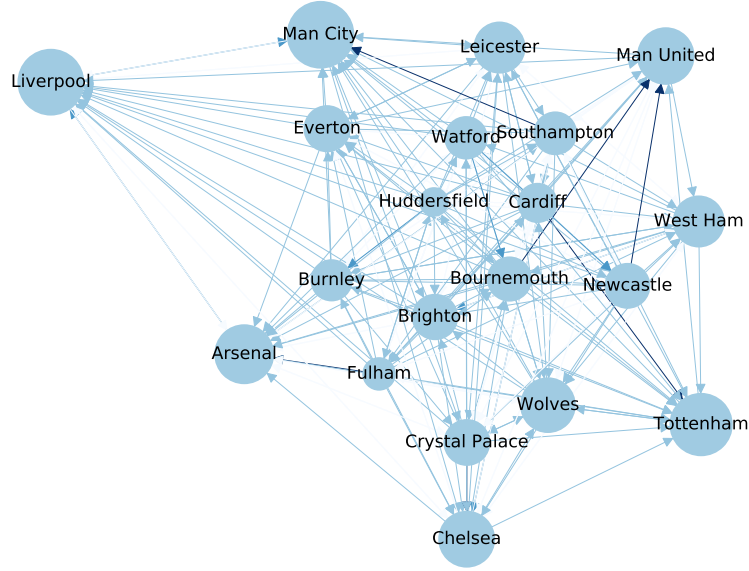


Fig. 1. The network diagram of the 2018-2019 English Premier League after 240 matches have been played for $n = 0.5$ (calculating centrality scores based on the last 190 matches).

Network Model. To overcome the above mentioned limitation and come up with a scoring system that is less sensitive to the set of teams that each team has played against, we build a directed network of all the matches within the training window, in which the edges point from the loser to the winner, weighted by the amount of points the winner earned. In the next step, we calculate the network *eigenvector centrality* score for all the teams. The recursive definition of eigenvector centrality, that is that the score of each node depends on the score of its neighbours that send a link to it, perfectly solves the problem of the dyadic scoring system mentioned above (16, 26, 27). An example of such network and calculated scores are presented in Figure 1.

We can calculate the score difference between the two competing teams for any match after the N th match. We will have $(T - N)$ matches with their respective outcomes and *score differences*, for both models. We then fit a *logistic regression model* of the outcome (as a categorical y variable) over the score differences according to Eq. 1. Without loss of generality, we always calculate the point score difference as $x = \text{home team score} - \text{away team score}$ and assume $y = 1/0$ if the home/away team wins.

$$y = F(x|\mu, s) = \frac{1}{1 + \exp(-\frac{x-\mu}{s})}, \quad [1]$$

where μ and s are model parameters than can be obtained by ordinary least squares methods.

Finally, the logistic regression model will provide the probabilistic assessment of each score system (Dyadic and Network) for each match, allowing us to understand how correctly the outcomes are being split as a function of the pre-match score difference. See Figure 2 for an example.

Benchmark Model. We benchmark our models against the implicit prediction model provided by the betting market payoffs (*Bet 365*). When betting houses allow players to bet on an event outcome, they are essentially providing an implicit probabilistic assessment of that event through the payoff values. We calculate the probabilities of different outcomes as the following:

$$\begin{cases} \text{Prob}(\text{HomeWin}) = \frac{\frac{1}{h}}{\frac{1}{h} + \frac{1}{a}} \\ \text{Prob}(\text{AwayWin}) = \frac{\frac{1}{a}}{\frac{1}{h} + \frac{1}{a}} \end{cases}, \quad [2]$$

where h is monetary units for a home team win, a units for an away team win. The experienced gambler knows that in the real world, on any given betting house's game, the sum of the outcomes' probabilities never does equal to 1, as the game proposed by the betting houses is not strictly *fair*. However, we believe this probabilities give us a good estimate of the state of the art predictive models in the market. In Figure 2 the data and corresponding fits to all the 3 models are provided.




Fig. 2. Logistic regression model example for the England Premier League 2018-2019 for the two models and the benchmark model.




Fig. 3. Accuracy of the models measured through a) average Brier Score and b) area Under the Curve (AUC), per league.

Model Performance. To compare our two models with the benchmark model, we calculate a loss function in the form of the Brier score for all the $(1 - n)T$ matches in different leagues and years from 2005 to 2018 and $n = 0.5$ (See Data and Methods for details). The distributions of scores are reported in Figure S1 and the average scores are shown in Figure 3(a) (Note that the smaller the Brier score, the better fit of the model). Progressing on the loss function to actual prediction accuracy, for each match, we considered a home team win prediction whether $Prob(HomeWin) \geq \pi$ and an away team win prediction otherwise. We applied this prediction rule for all three models (Dyadic, Network and the Betting Houses Market benchmark) and compared the predictions to the outcomes, calculating how many predictions matched the outcome. We change the threshold of π and calculate the average Area Under Curve (AUC) score of Receiver Operating Characteristic (ROC Curve) for each league, shown in Figure 3(b). The complete set of distributions for the accuracy indicators is depicted in Figures S1 and S2. By increasing n , naturally the accuracy of predictions increases. For the rest of the paper however, we focus on $n = 0.5$. The comparisons for different values of n is presented in Figures S3-S6.

Even though our model underperforms the market on average, t-tests for the difference between the Brier score and AUC distributions of our models and the betting houses market are statistically indistinguishable at the 2% significance level for the majority of year-leagues (See Tables S1 and S2 for details). One should note that, the goal here is not to develop a better prediction model than the one provided by the betting houses market, but to obtain a comparable, consistent and simple-to-calculate model. Market models evolve throughout time, which would bias the assessment of a single prediction tool throughout the years. In that capacity, we are looking for models that provide the time-robust tool we need to assess whether football is becoming more predictable over the years.

Results and Discussion

Predictability over Time. After having assessed the models' validity and robustness, we can address if predictability of football has been changing over time. From now on we only use the Network model as the marginally better one for assessing predictions and the AUC values as a measure of predictability (The results of Brier scores are shown in Figure S7). In Figure 4 we see the AUC scores and a smoothed fits to them per league over time for the last 26 years.

A positive trend in predictability is observed in most of the cases (England, France, Germany, Netherlands, Portugal, Scotland and Spain), whereas for Greece and Turkey we see a deterioration of prediction power. For both leagues, however, we see a recent increase in predictability. However, these two leagues are the one for which the predictability was already relatively high in the past. The remaining leagues (Belgium and Italy) display a somewhat stable predictability throughout time. All leagues, notwithstanding, tend to converge towards 0.75 AUC.

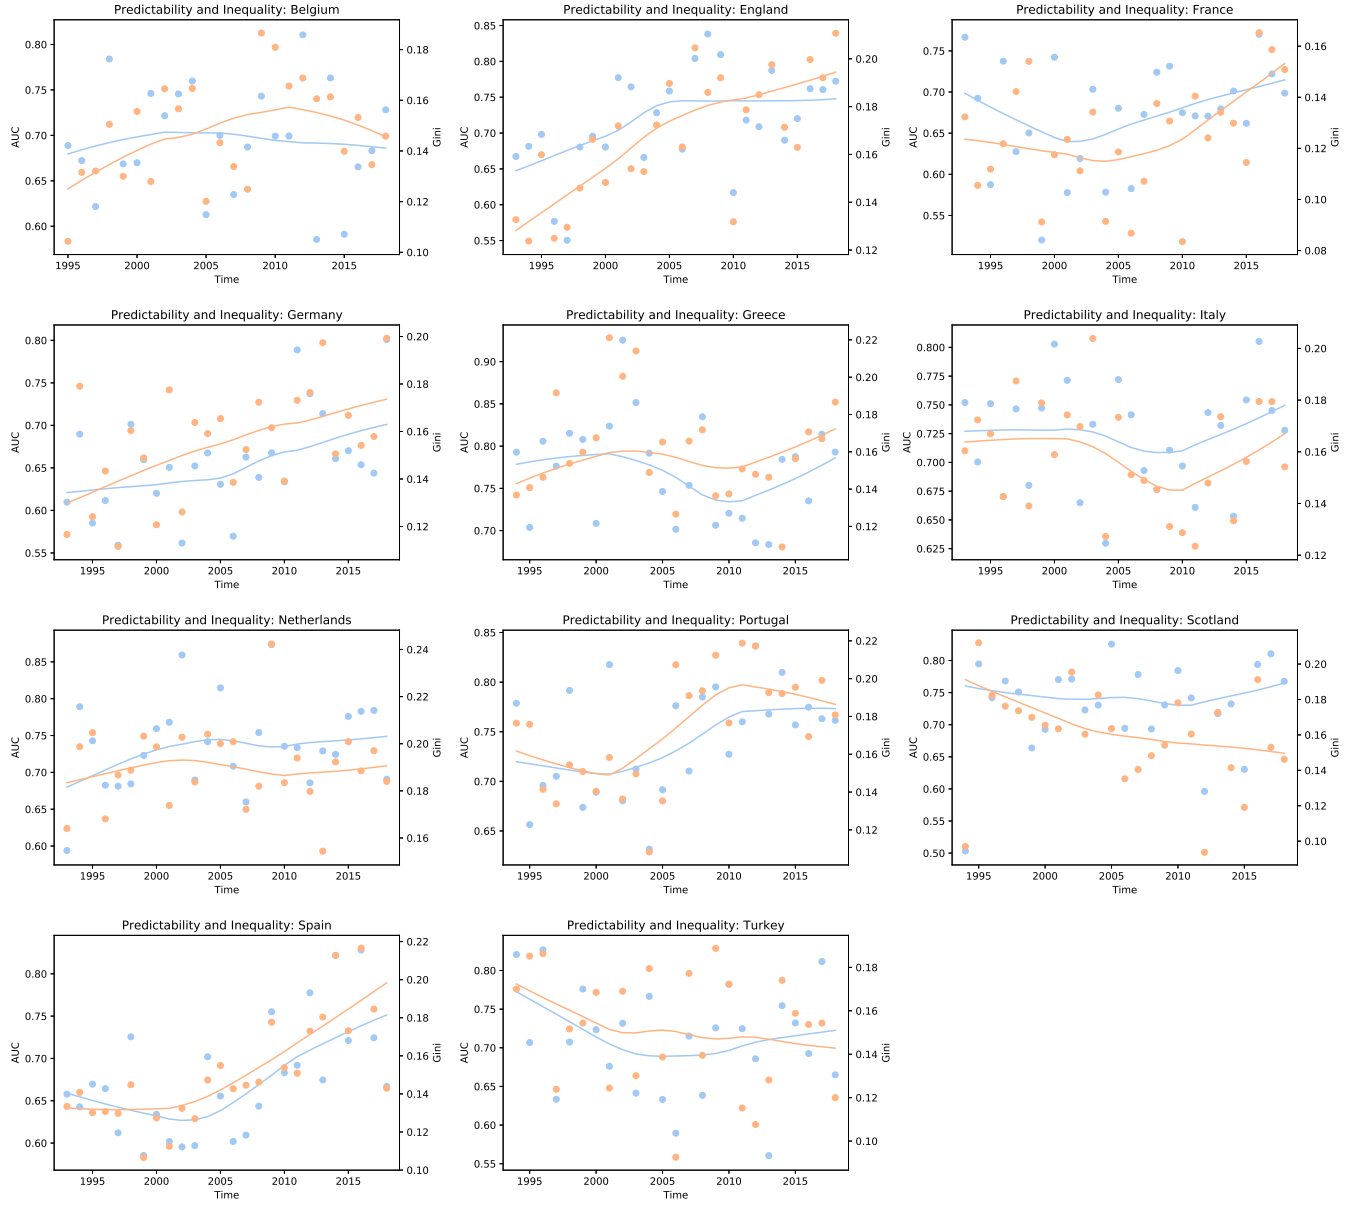


Fig. 4. Time Trends In AUC and *Gini* Coefficient. Blue dots/lines depict AUC and are marked in the primary (left) *y*-axis; Orange dots/lines depict the *Gini* coefficient and are marked in the secondary (right) *y*-axis; Both lines are fitted through a *lowess* model.

Table 1. AUC and Gini Correlation by league

League	Correlation	League	Correlation
Spain	0.874	Netherlands	0.676
England	0.823	Italy	0.622
Germany	0.805	France	0.563
Scotland	0.723	Greece	0.561
Portugal	0.694	Belgium	0.413
Turkey	0.686		

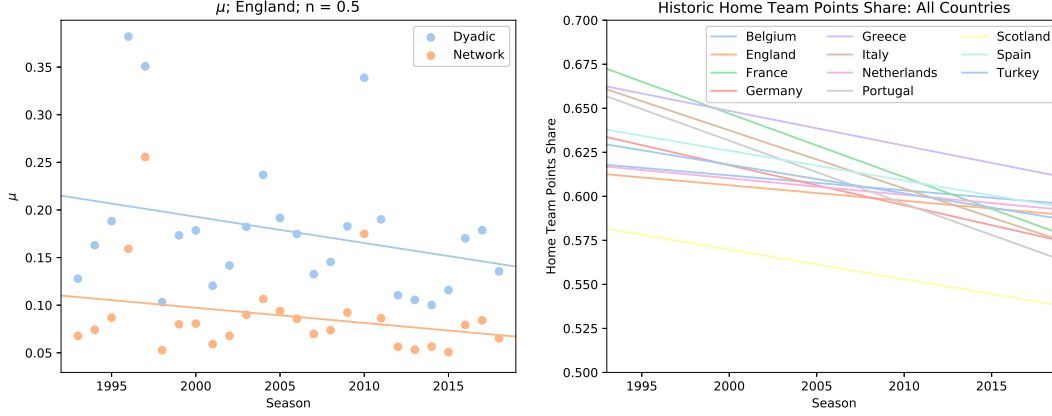


Fig. 5. The Decrease in Home Field Advantage. Left: Home field advantage calculated by the two models. Right: The share of home points measured on historical data. The straight lines are linear fits. See the detailed graphs for each league in Figures S8–S11.

Increasing Inequality. In analysing the predictability of different leagues, we observe that predictability has been increasing for the richer leagues in Europe, whereas the set for which the indicator is deteriorating is composed mainly by peripheral leagues.

It seems football as a sport is emulating society in its somewhat “gentrification” process, i.e., the richer leagues are becoming more deterministic because: better teams win more often; consequentially, becoming richer; allowing themselves to hire better players (from a talent pool that gets internationally broader each year); becoming even stronger; and, closing the feedback cycle, winning even more matches and tournaments; hence more predictability in more professional and expensive leagues.

To illustrate this trend, we use the *Gini* coefficient, proposed as a measure of inequality of income or wealth (28–30). We calculate the Gini coefficient of a given league-season’s distribution of points each team had at the end of the tournament. Figure 4 depicts the values for all the leagues in the database, comparing the evolution of predictability and the evolution of inequality between teams for each case. There is a high correlation between predictability in a football league and inequality amongst teams playing in that league, uncovering evidence in favour of the “gentrification of football” argument. The correlation values between these two parameters are reported in Table 1.

Home Advantage and Predictability. In Figure 2 we see that the fitted Sigmoid is not symmetric and a shift in the probable outcome (i.e., when $\text{Prob}(y = 1) \geq 0.50$) happens at an x value smaller than zero, meaning even for some negative values of x , which represents cases where the visiting team’s score is greater than the home team’s score, we still observe a higher chance for a home team victory. This is not surprising considering the existence of the notion of *Home Field Advantage*: a significant and structural difference in outcome contingent on the locality in which a sports match is played —favouring the team established in such locality, has been exhaustively uncovered between different team sports and between different countries both for men’s and women’s competitions (31). Home field advantage was first formalised for North-American sports (32), and then documented specifically for association football (33).

We calculated the average amount of home field advantage as measured in the prediction models (μ in Eq. 1) for each season of English Premier League and plotted it in Figure 5. We see a decrease in home field advantage over time (for the results for other leagues see Figures S8–S11).

However, one can calculate the home advantage directly from historical data by counting the total number of points that the home and away teams gained in each season. The trends in the share of home teams for different leagues are shown in Figure 5 (For detailed diagrams of each league see Figures S8–S11).

It is clear that the home field advantage is still present, however it has been decreasing throughout time for all the leagues under study. In explaining this trend, we should consider factors involved in home field advantage. Pollard counts the following as driving factors for the existence of home field advantage (34): crowd effects; travel effects; familiarity with the pitch; referee bias; territoriality; special tactics; rule factors; psychological factors; and the interaction between two or more of such factors.

Table 2. Database – Matches per league

Country	Total Matches	Country	Total Matches
Belgium	6620	Netherlands	7956
England	10044	Portugal	7122
France	9510	Scotland	5412
Germany	7956	Spain	10044
Greece	6470	Turkey	7616
Italy	9066		

While some of these factors are not changing over time, increase in the number of foreign players, diminishing the effects of territoriality and its psychological factors, as well as observing that fewer people are going to stadiums, traveling is becoming easier, teams are camping in different pitches and players are accruing more international experience, can explain the reported trend; stronger (richer) teams are much more likely to win, it matters less where they play.

Conclusion

Relying on large-scale historical records of 11 major football leagues, we have shown that, throughout time, football is dramatically changing; the sport is becoming more predictable; teams are becoming increasingly unequal; and home field advantage is steadily and consistently decreasing for all the leagues in the sample.

The prediction models in this work were designed to be self-contained, i.e., to solely consider the past results as supporting data to understand future events. Despite the strength in terms of time consistency, they do ignore the myriad of data that is available nowadays, which is being extensively used by cutting edge data science projects (35) and, obviously, the betting houses. In this sense, this work is limited by the amount of data and the sophistication of the models it considers, however, this is by design and in order to focus on the nature of the system under study rather than the predictive models.

Furthermore, this work is also limited by the sample size; Football has been played in over 200 countries, out of which only 5.5% were included in our sample. Although the main reasons for that are of practical matters, given the hardship in obtaining reliable and time-consistent data, the events become very sparse for most of the leagues going back in time.

Future work should, as speculated in this work, try and assess: the role money is playing in removing the surprise element of the sport; expanding the sample barriers beyond the European continent; and ultimately, but not exhaustively, should test the money impact over predictability on different sports and leagues that – theoretically – should not be affected by it, namely leagues and sports that impose salary caps over their teams, such as the United States of America’s Professional Basketball League (the NBA).

Materials and Methods

Data. In this study, each datapoint depicts a football match and, as such, contains: the date in which the match took place, which determines the season for which the match was valid; the league; the home team; the away team; the final score (as in, the final amount of goals scored by each team); and the pay-offs at *Bet 365* betting house.

The database encompass eleven different European countries and their top division leagues (Belgium, England, France, Germany, Greece, Italy, Netherlands, Portugal, Scotland, Spain and Turkey), ranging from the 1993-1994 to the 2018-2019 seasons (some leagues only have data starting in the 1994-1995 and 1995-1996 seasons). The data was extracted from <https://www.football-data.co.uk/>. Considering all leagues and seasons, the final database encompasses 87,816 matches (See Table 2). These 87,816 matches had a total of 236,323 goals scored, an average of 2.7 goals per match.

Performance Assessment. Assessing performance of a probabilistic model is both challenging and controversial in the literature, as the different measures present strengths and weaknesses. Notwithstanding, the majority of research done in the area applies (at least one of) two methods, namely: a loss function; and/or a binary accuracy measure. We analyse our models within both approaches.

Loss Function: The Brier Score. Proposed in (36), the Brier score is a loss function that measures the accuracy of probabilistic predictions. It applies to tasks in which predictions must assign probabilities to a set of mutually exclusive discrete outcomes. The set of possible outcomes can be either binary or categorical in nature, and the probabilities assigned to this set of outcomes must sum to one (where each individual probability is in the range of 0 to 1).

Suppose that on each of t occasions an event can occur in only one of r possible classes or categories and on one such occasion, i , the forecast probabilities are $f_{i1}, f_{i2}, \dots, f_{ir}$, that the event will occur in classes $1, 2, \dots, r$, respectively. The r classes are chosen to be mutually exclusive and exhaustive so that $\sum_{j=1}^r f_{ij} = 1, i = 1, 2, 3, \dots, t$. Hence, the Brier score is defined as:

$$P = \frac{1}{t} \sum_{j=1}^r \sum_{i=1}^t (f_{ij} - E_{ij})^2 \quad [3]$$

where E_{ij} takes the value 1 or 0 according to whether the event occurred in class j or not. In our case, we have: $r = 2$ classes, i.e., either the home team won or not; and $(1 - n)T$ matches to be predicted, transforming Equation 3 into:

$$P = \frac{1}{(1 - n)T} \sum_{j=1}^2 \sum_{i=1}^{(1-n)T} (f_{ij} - E_{ij})^2, \quad [4]$$

We see, from Equation 4, how the Brier Score is an averaged measure over all the predicted matches, i.e., each league-season combination will have its own single Brier Score. The lower the Brier score, the better (as in the more assertive and/or correct) our prediction model.

Accuracy: Receiver Operating Characteristic (ROC Curve). A receiver operating characteristic curve, or ROC curve, is a graphical plot that illustrates the diagnostic ability of a binary classifier system as its discrimination threshold is varied, which allows us to generally assess a prediction method's performance for different thresholds. When using normalised units, the area under the ROC curve (often referred to as simply the AUC) is equal to the probability that a classifier will rank a randomly chosen positive instance higher than a randomly chosen negative one (37).

The machine learning community has historically used this measure for model comparison (38), though the practice has been questioned given AUC estimates are quite noisy and suffer from other problems (39–41). Nonetheless, the coherence of AUC as a measure of aggregated classification performance has been vindicated, in terms of a uniform rate distribution (42), and AUC has been linked to several other performance metrics, including the above mentioned Brier score (43).

ACKNOWLEDGMENTS. The authors thank Luca Pappalardo for valuable discussions. TY was partially supported by the Alan Turing Institute under the EPSRC grant no. EP/N510129/1.

References

- Giulianotti R, Rollin J, Weil E, Joy B, Alegi P (2017) Football. encyclopedia britannica.
- Wilson B (2018) European football is worth a record £22bn, says deloitte.
- Ingle S, Glendenning B (2003) Baseball or football: which sport gets the higher attendance?
- Choi D, Hui SK (2014) The role of surprise: Understanding overreaction and underreaction to unanticipated events using in-play soccer betting market. *Journal of Economic Behavior & Organization* 107:614–629.
- Maher MJ (1982) Modelling association football scores. *Statistica Neerlandica* 36(3):109–118.
- Rue H, Salvesen O (2000) Prediction and retrospective analysis of soccer matches in a league. *Journal of the Royal Statistical Society: Series D (The Statistician)* 49(3):399–418.
- Dixon MJ, Coles SG (1997) Modelling association football scores and inefficiencies in the football betting market. *Journal of the Royal Statistical Society: Series C (Applied Statistics)* 46(2):265–280.
- Crowder M, Dixon M, Ledford A, Robinson M (2002) Dynamic modelling and prediction of english football league matches for betting. *Journal of the Royal Statistical Society: Series D (The Statistician)* 51(2):157–168.
- Hvattum LM, Arntzen H (2010) Using elo ratings for match result prediction in association football. *International Journal of forecasting* 26(3):460–470.
- Buchdahl J (2003) *Fixed odds sports betting: Statistical forecasting and risk management*. (Summersdale Publishers LTD-ROW).
- Constantinou AC, Fenton NE, Neil M (2012) pi-football: A bayesian network model for forecasting association football match outcomes. *Knowledge-Based Systems* 36:322–339.
- Goddard J, Asimakopoulos I (2004) Forecasting football results and the efficiency of fixed-odds betting. *Journal of Forecasting* 23(1):51–66.
- Kuypers T (2000) Information and efficiency: an empirical study of a fixed odds betting market. *Applied Economics* 32(11):1353–1363.
- Goddard J (2005) Regression models for forecasting goals and match results in association football. *International Journal of forecasting* 21(2):331–340.
- Baboota R, Kaur H (2019) Predictive analysis and modelling football results using machine learning approach for english premier league. *International Journal of Forecasting* 35(2):741–755.
- Fraiberger SP, Sinatra R, Resch M, Riedl C, Barabási AL (2018) Quantifying reputation and success in art. *Science* 362(6416):825–829.
- Williams OE, Lacasa L, Latora V (2019) Quantifying and predicting success in show business. *Nature communications* 10(1):2256.
- Cintia P, Coscia M, Pappalardo L (2016) The haka network: evaluating rugby team performance with dynamic graph analysis in *Proceedings of the 2016 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining*. (IEEE Press), pp. 1095–1102.
- Radicchi F (2011) Who is the best player ever? a complex network analysis of the history of professional tennis. *PloS one* 6(2):e17249.
- Yucesoy B, Barabási AL (2016) Untangling performance from success. *EPJ Data Science* 5(1):17.
- Rossi A, et al. (2018) Effective injury forecasting in soccer with gps training data and machine learning. *PloS one* 13(7):e0201264.
- Pappalardo L, Cintia P, Ferragina P, Pedreschi D, Giannotti F (2018) Playerank: Multi-dimensional and role-aware rating of soccer player performance. *arXiv preprint arXiv:1802.04987*.
- Cintia P, Rinzivillo S, Pappalardo L (2015) A network-based approach to evaluate the performance of football teams in *Machine learning and data mining for sports analytics workshop, Porto, Portugal*.
- Lazova V, Basnarkov L (2015) Pagerank approach to ranking national football teams. *arXiv preprint arXiv:1503.01331*.
- Milanlouei S (2017) Fall in love with network science at first sight!
- Newman M (2018) *Networks*. (Oxford university press).
- Bonacich P (1987) Power and centrality: A family of measures. *American journal of sociology* 92(5):1170–1182.
- Gini C (1912) Variabilità e mutabilità. *Reprinted in Memorie di metodologica statistica* (Ed. Pizetti E, Salvemini, T). Rome: Libreria Eredi Virgilio Veschi.
- Gini C (1921) Measurement of inequality of incomes. *The Economic Journal* 31(121):124–126.
- Gini C (1936) On the measure of concentration with special reference to income and statistics. *Colorado College Publication, General Series* 208:73–79.
- Pollard R, Prieto J, Gómez MÁ (2017) Global differences in home advantage by country, sport and sex. *International Journal of Performance Analysis in Sport* 17(4):586–599.
- Schwartz B, Barsky SF (1977) The home advantage. *Social forces* 55(3):641–661.
- Pollard R (1986) Home advantage in soccer: A retrospective analysis. *Journal of sports sciences* 4(3):237–248.
- Pollard R (2008) Home advantage in football: A current review of an unsolved puzzle. *The open sports sciences journal* 1(1).
- Pappalardo L, Cintia P (2017) Quantifying the relation between performance and success in soccer. *Advances in Complex Systems* p. 1750014.
- Brier GW (1950) Verification of forecasts expressed in terms of probability. *Monthly Weather Review* 78(1):1–3.
- Fawcett T (2006) An introduction to roc analysis. *Pattern recognition letters* 27(8):861–874.
- Hanley JA, McNeil BJ (1983) A method of comparing the areas under receiver operating characteristic curves derived from the same cases. *Radiology* 148(3):839–843.
- Hanczar B, et al. (2010) Small-sample precision of roc-related estimates. *Bioinformatics* 26(6):822–830.
- Lobo JM, Jiménez-Valverde A, Real R (2008) AUC: a misleading measure of the performance of predictive distribution models. *Global ecology and Biogeography* 17(2):145–151.
- Hand DJ (2009) Measuring classifier performance: a coherent alternative to the area under the roc curve. *Machine learning* 77(1):103–123.
- Ferri C, Hernández-Orallo J, Flach PA (2011) A coherent interpretation of auc as a measure of aggregated classification performance in *Proceedings of the 28th International Conference on Machine Learning (ICML-11)*. pp. 657–664.
- Hernández-Orallo J, Flach P, Ferri C (2012) A unified view of performance metrics: translating threshold choice into expected classification loss. *Journal of Machine Learning Research* 13(Oct):2813–2869.

Supporting Information

Table S1. The t-Test values comparing the Brier scores for the prediction models and the betting house benchmark for $n = 0.5$.

Country	Model	t-stat	p-value	Country	Model	t-stat	p-value
Belgium	Network	2.679	0.011	Belgium	Dyadic	2.386	0.023
England	Network	2.801	0.008	England	Dyadic	2.903	0.007
France	Network	0.549	0.586	France	Dyadic	0.562	0.578
Germany	Network	2.055	0.048	Germany	Dyadic	1.894	0.067
Greece	Network	1.054	0.299	Greece	Dyadic	0.918	0.365
Italy	Network	3.422	0.002	Italy	Dyadic	3.021	0.005
Netherlands	Network	2.420	0.021	Netherlands	Dyadic	2.455	0.019
Portugal	Network	2.847	0.007	Portugal	Dyadic	2.686	0.011
Scotland	Network	1.097	0.282	Scotland	Dyadic	1.335	0.192
Spain	Network	2.382	0.024	Spain	Dyadic	2.404	0.023
Turkey	Network	0.459	0.649	Turkey	Dyadic	0.454	0.653

Table S2. The t-Test values comparing the AUC scores for the prediction models and the betting house benchmark for $n = 0.5$.

Country	Model	t-stat	p-value	Country	Model	t-stat	p-value
Belgium	Network	-3.450	0.001	Belgium	Dyadic	-3.379	0.002
England	Network	-3.868	0.000	England	Dyadic	-3.911	0.000
France	Network	-2.582	0.014	France	Dyadic	-2.551	0.015
Germany	Network	-4.576	0.000	Germany	Dyadic	-4.284	0.000
Greece	Network	-1.994	0.054	Greece	Dyadic	-1.852	0.073
Italy	Network	-4.418	0.000	Italy	Dyadic	-4.021	0.000
Netherlands	Network	-3.544	0.001	Netherlands	Dyadic	-3.679	0.001
Portugal	Network	-4.413	0.000	Portugal	Dyadic	-4.389	0.000
Scotland	Network	-1.929	0.062	Scotland	Dyadic	-2.049	0.048
Spain	Network	-4.418	0.000	Spain	Dyadic	-4.430	0.000
Turkey	Network	-1.324	0.194	Turkey	Dyadic	-1.281	0.208

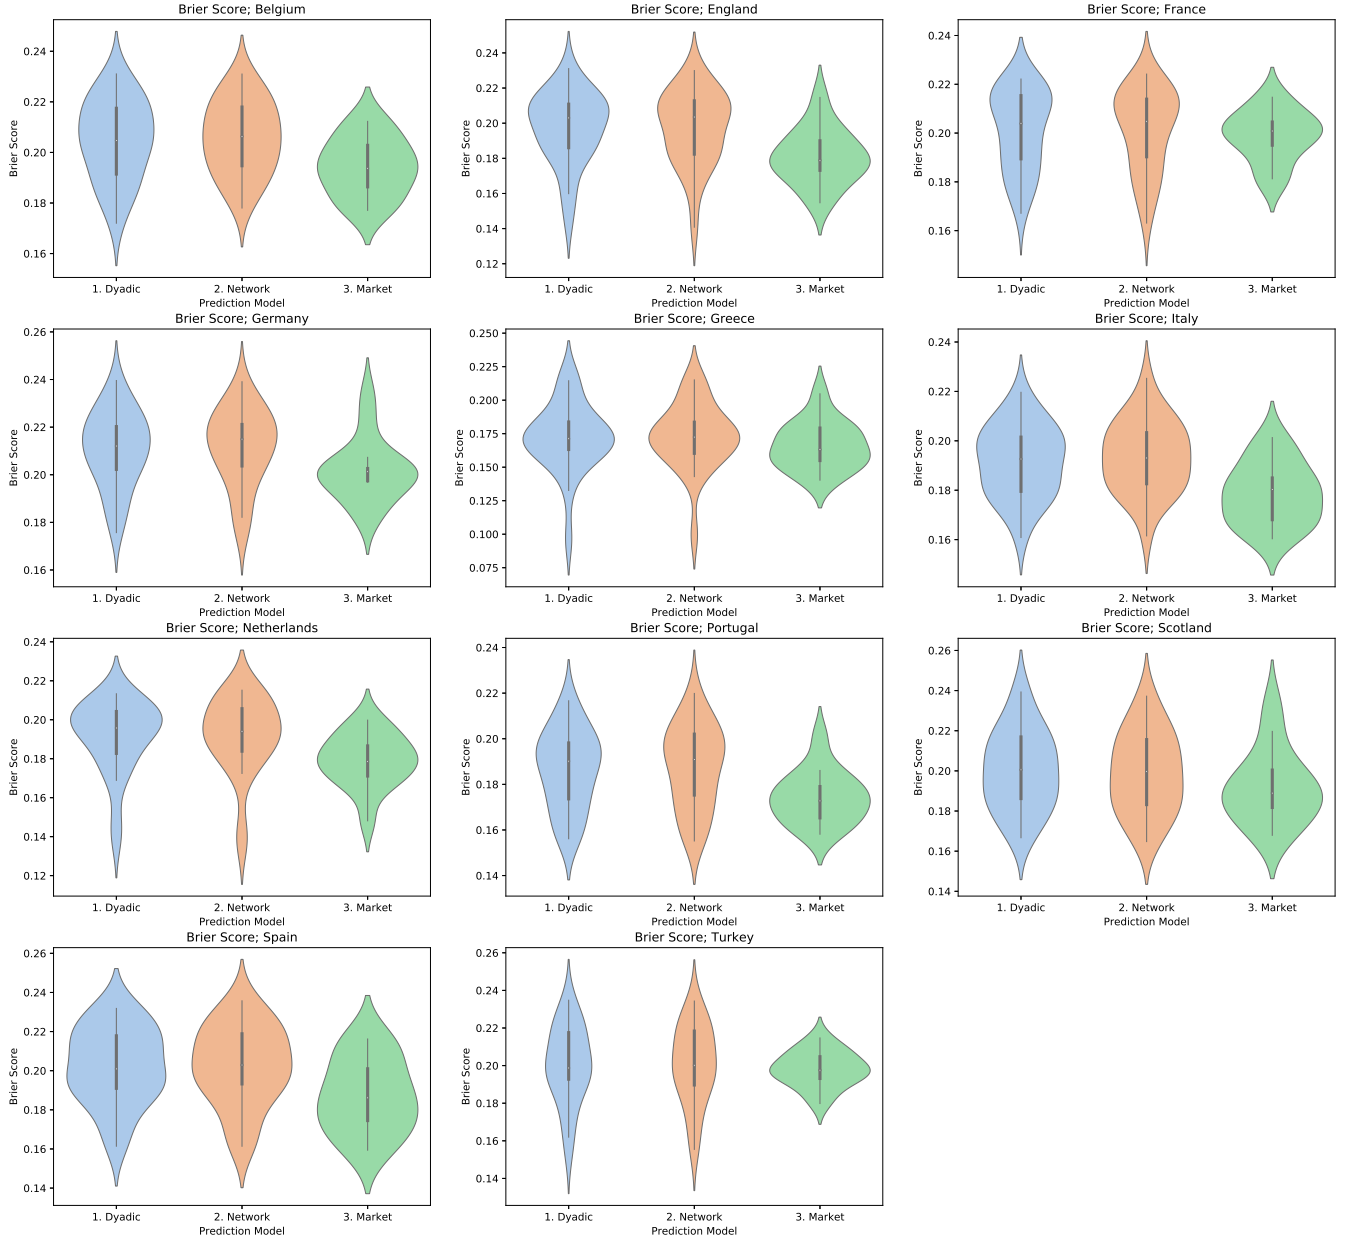


Fig. S1. Brier score distributions per country for the benchmark and prediction models (for $n = 0.5$).

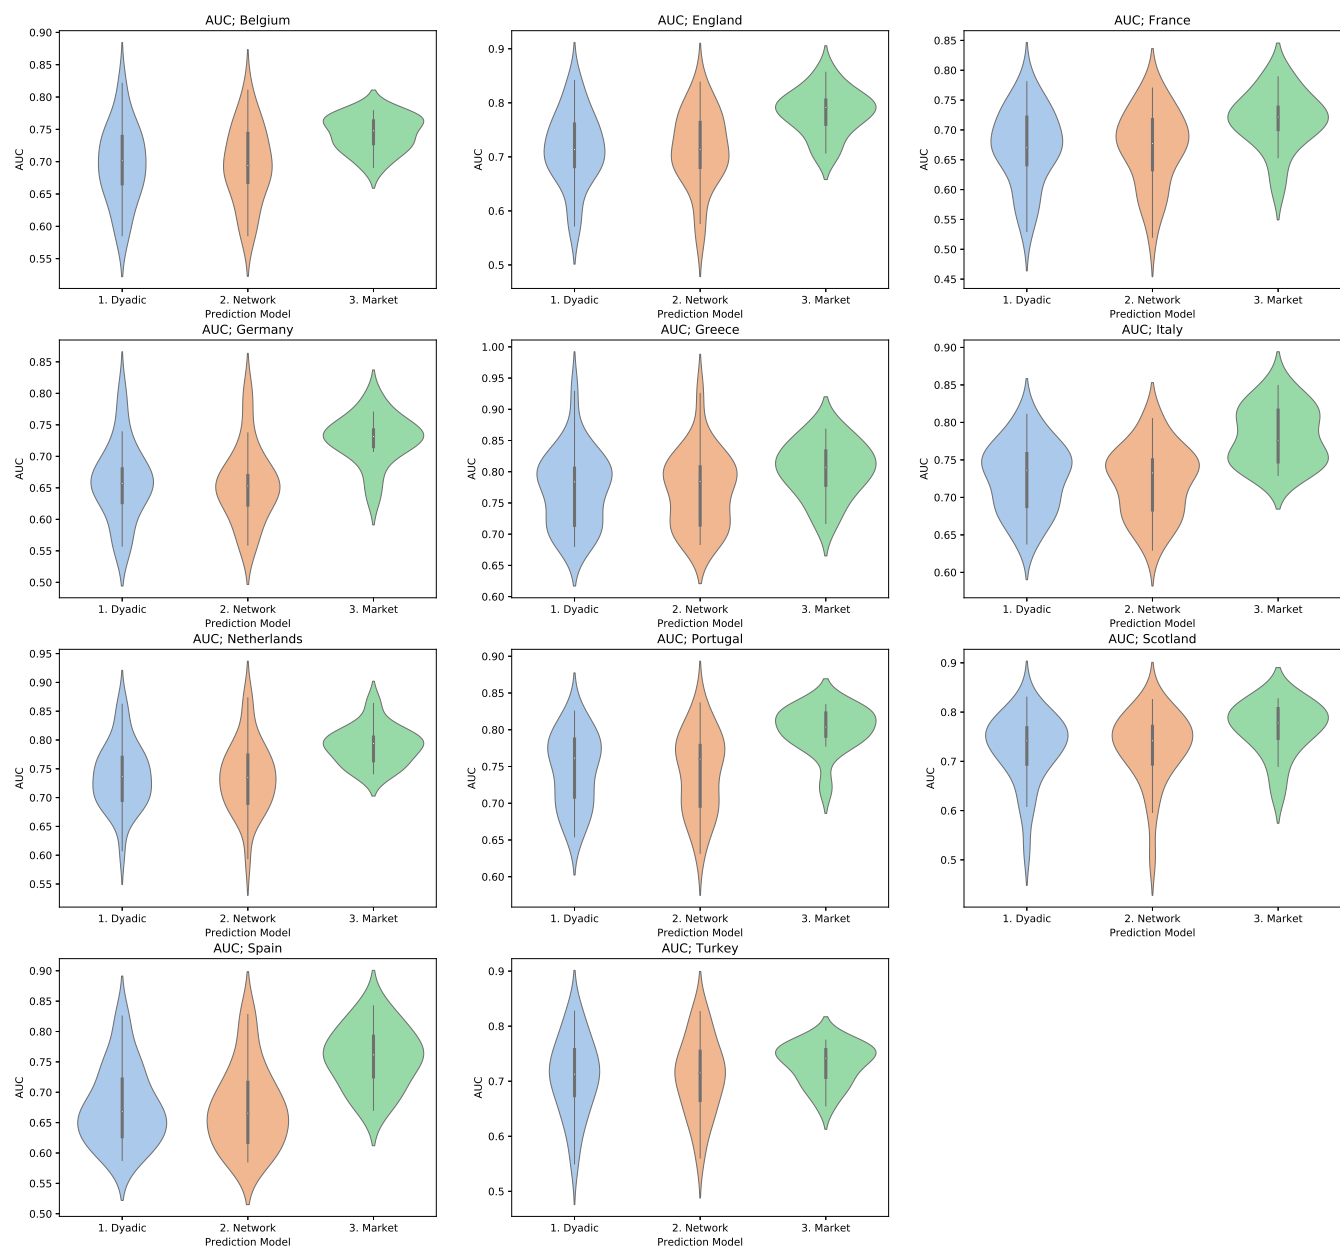


Fig. S2. Area Under the Curve (AUC) distributions per country for the benchmark and prediction models (for $n = 0.5$).





Fig. S3. Average Accuracy per Country: $n = 0.10$





Fig. S4. Average Accuracy per Country: $n = 0.30$





Fig. S5. Average Accuracy per Country: $n = 0.70$





Fig. S6. Average Accuracy per Country: $n = 0.90$

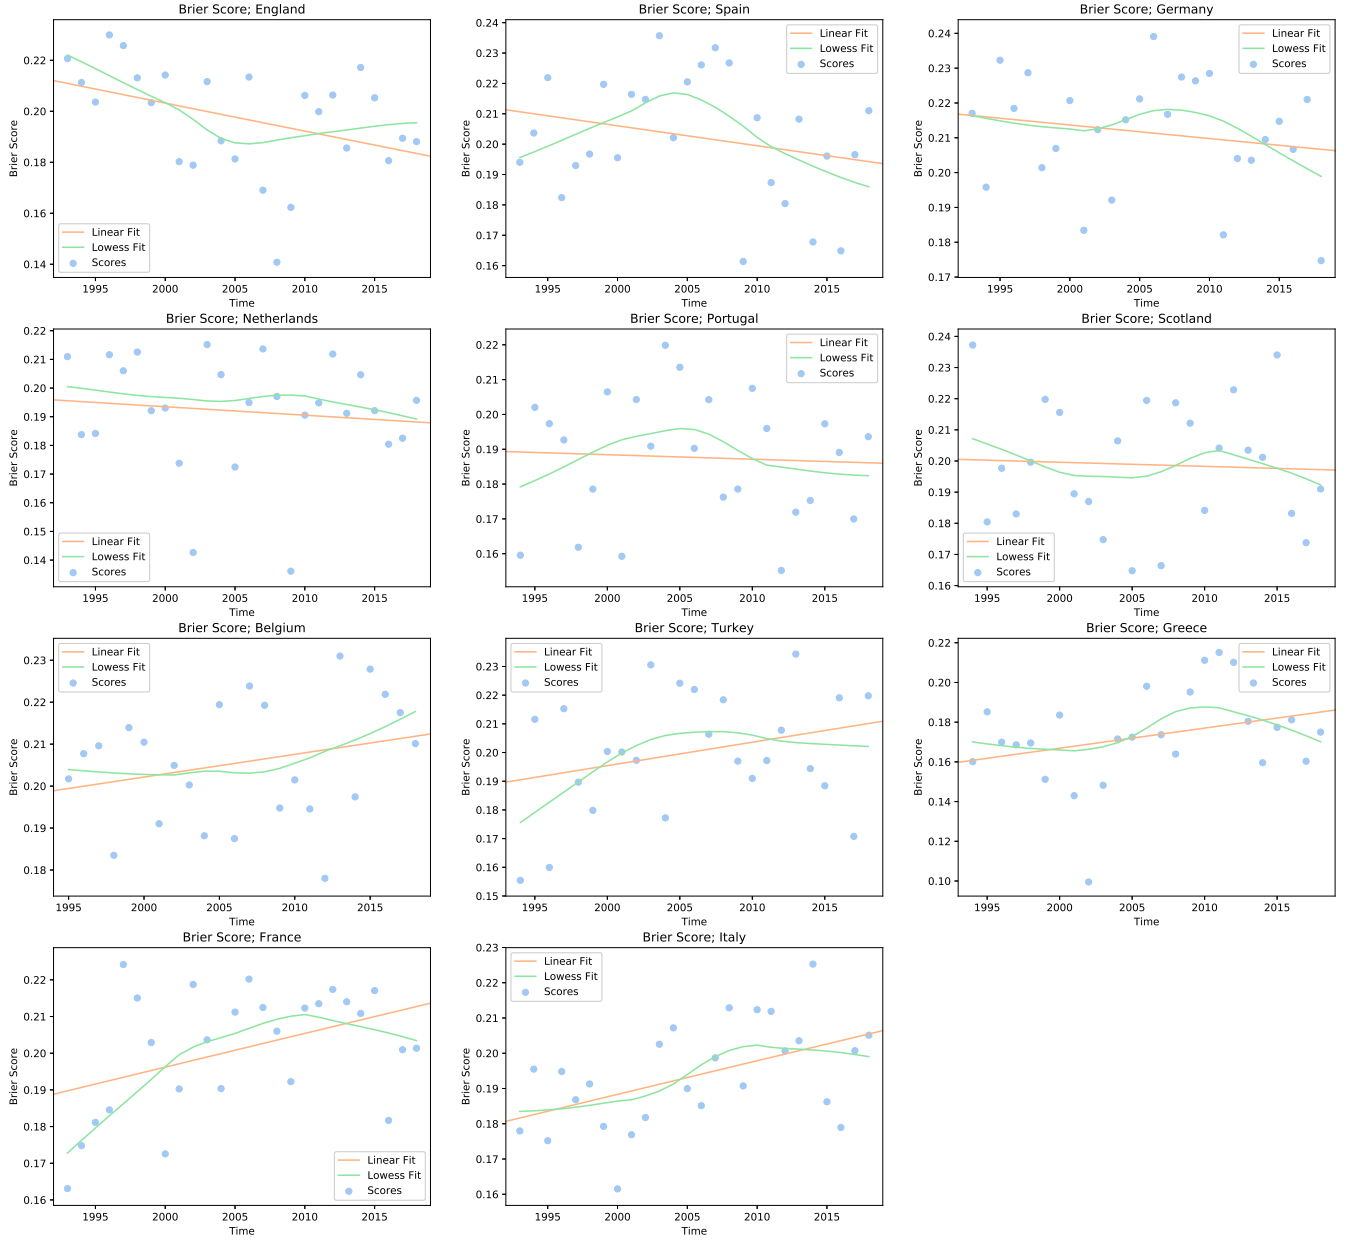


Fig. S7. Time Trends In The Brier Score per Country. The two lines show a linear regression fit and a non-parametric *lowess* fitting model.

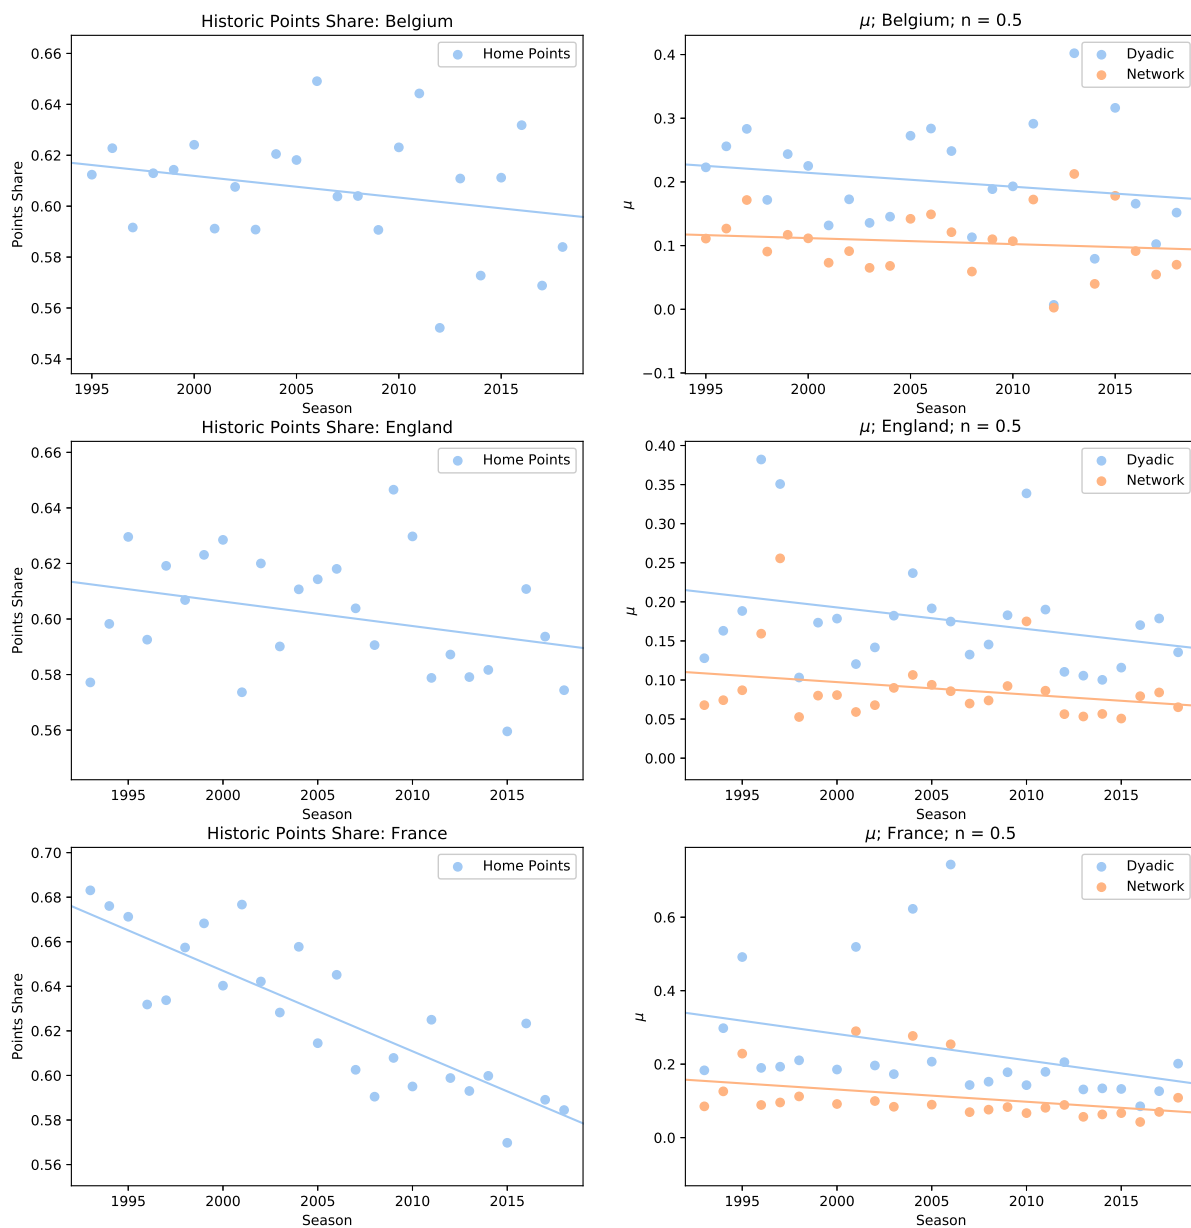


Fig. S8. The Decrease in Home Field Advantage: Selected Countries (1/4)

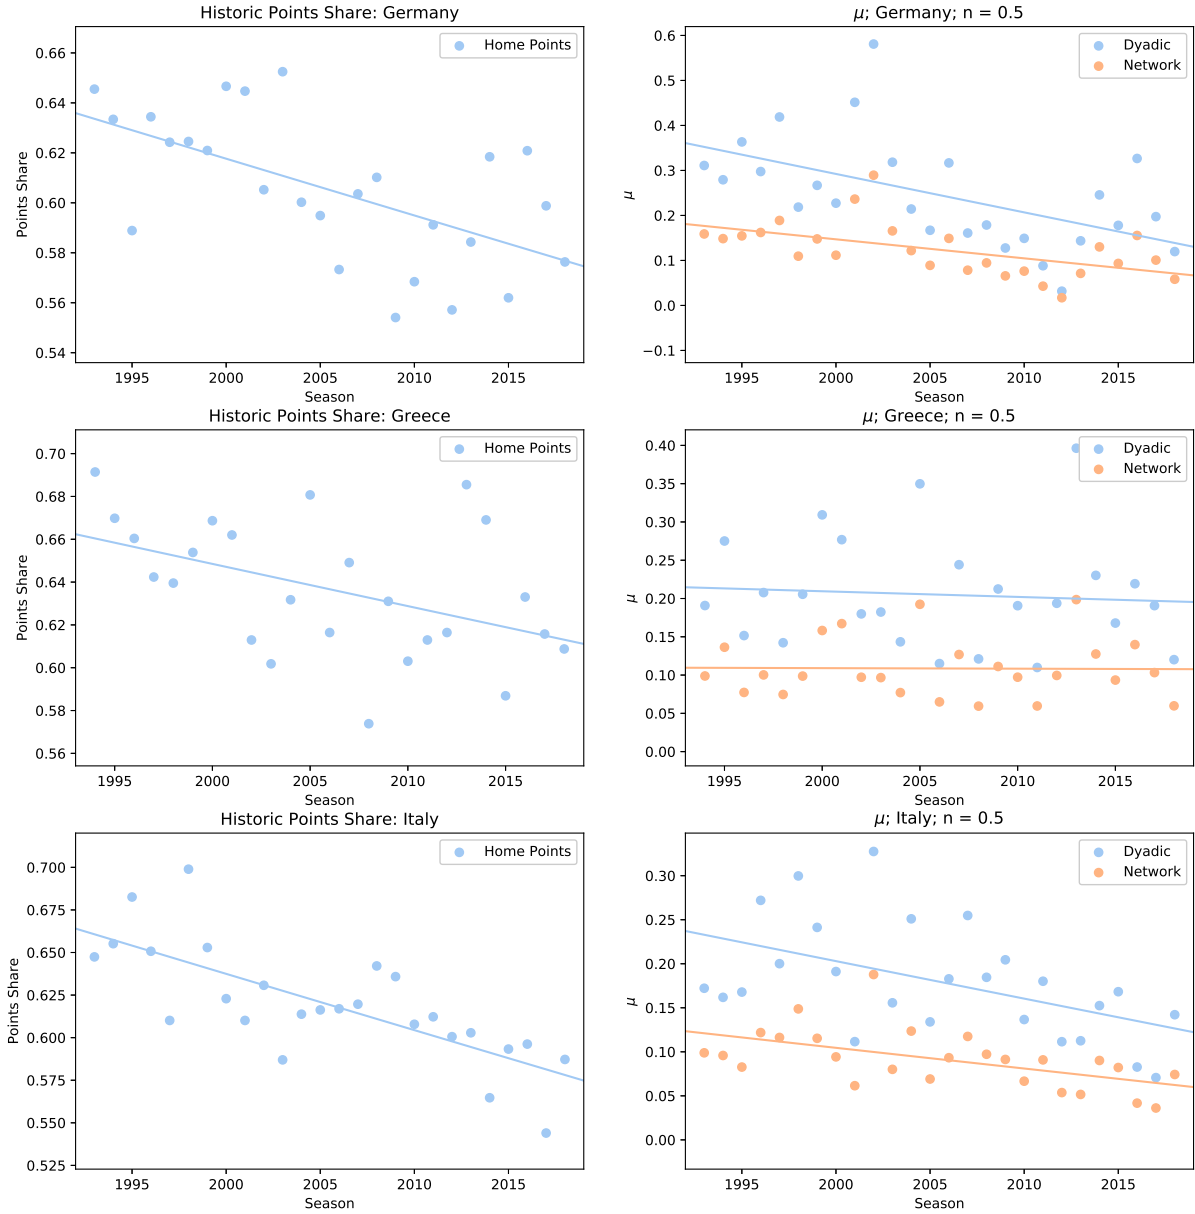


Fig. S9. The Decrease in Home Field Advantage: Selected Countries (2/4)

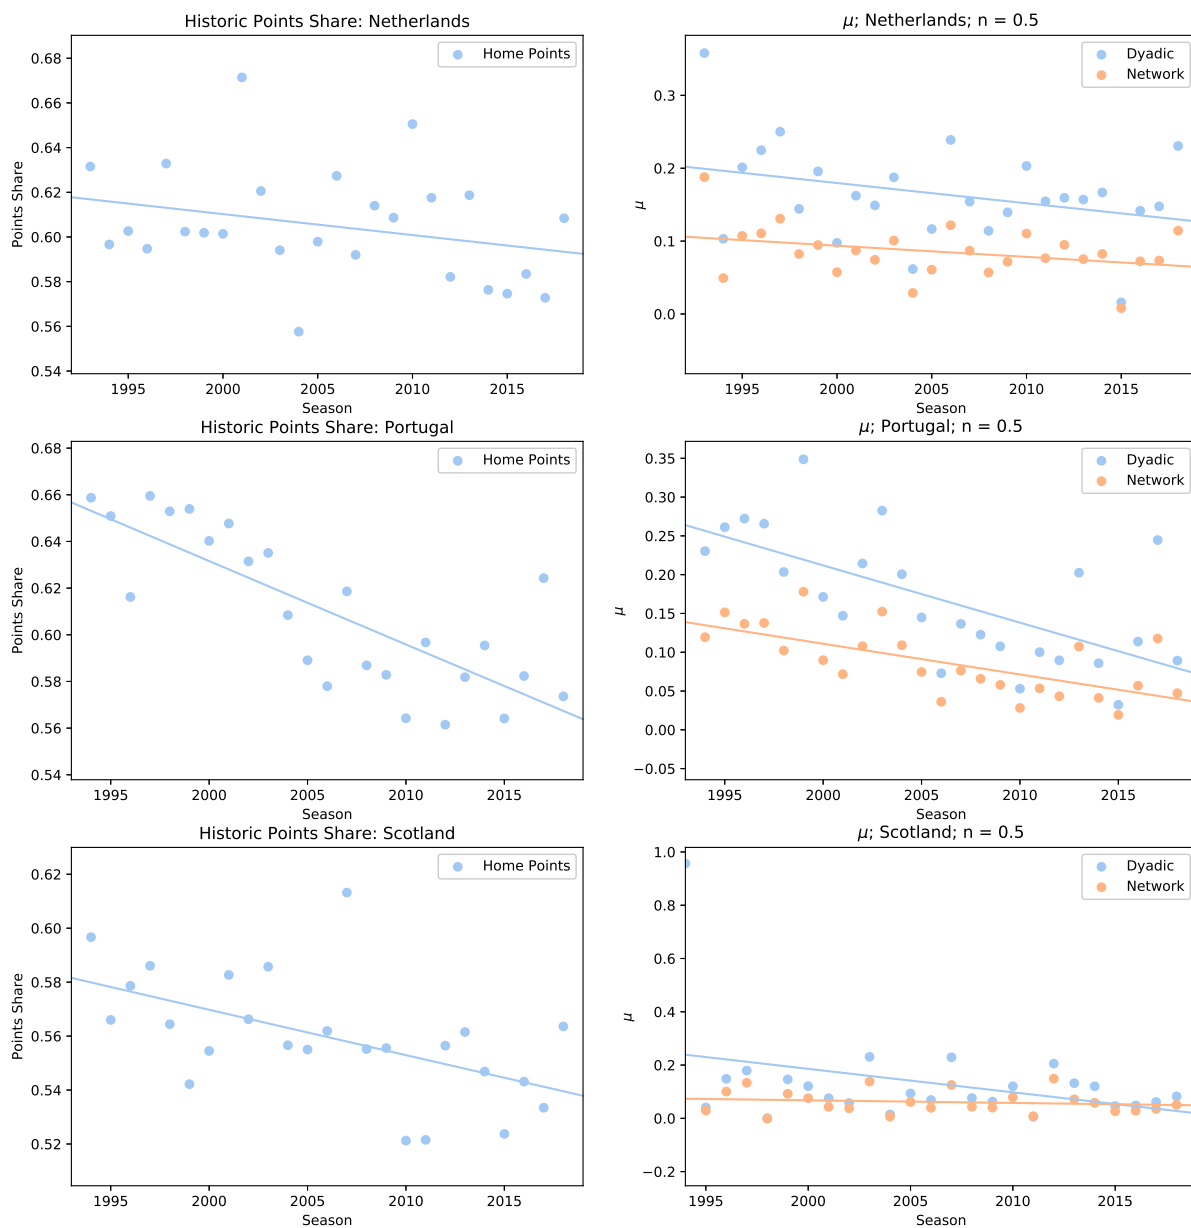


Fig. S10. The Decrease in Home Field Advantage: Selected Countries (3/4)




Fig. S11. The Decrease in Home Field Advantage: Selected Countries (4/4)