Ranking soccer teams on basis of their current strength: a comparison of maximum likelihood approaches

Christophe Ley [*1], Tom Van de Wiele [†2], and Hans Van Eetvelde [‡3]

[1] *Ghent University, Ghent, Belgium*
[2] *DeepMind, London, UK*
[3] *Ghent University, Ghent, Belgium*

Abstract

We present ten different strength-based statistical models that we use to model soccer match outcomes. The models are of three main types: Bradley-Terry, Independent Poisson and Bivariate Poisson, and their common aspect is that the parameters are estimated via weighted maximum likelihood, the weights being a match importance factor and a time depreciation factor giving less weight to matches that are played a long time ago. We compare their predictive performance via the Rank Probability Score and the log loss at the level of both domestic leagues and national teams, and find that the best models are the Independent and Bivariate Poisson models with the fewest number of parameters per team. We compare the best model's predictive performance to the bookmakers' prediction over seventeen seasons of the English Premier League and for the EURO2016. We conclude the paper by giving our strength-based rankings for the current Premier League season as well as for European national teams at the beginning of 2018.

Key words: Bivariate Poisson model, Bradley-Terry model, Game forecasting, Independent Poisson model, Predictive performance, Weighted likelihood

1 Introduction

Football, or soccer, is undeniably the most popular sport worldwide. Predicting which team will win the next World Cup or the Champions League final are issues that lead to heated discussions and debates among football fans, and even attract the attention of casual watchers. Or put more simply, the question of which team will win the next match, independent of its circumstances, excites the fans. Bookmakers have made a business out of football predictions, and they use highly advanced models taking into account numerous factors (like a team's current form, injured players, the history between both teams, the importance of the game for each team, etc.) to obtain the odds of winning, losing and drawing for both teams.

One major appeal of football, and a reason for its success, is its simplicity as game. This stands somehow in contrast to the difficulty of predicting the winner of a football match. A help in this respect would be a ranking of the teams involved in a given competition based on

---

[*]Electronic address: `christophe.ley@ugent.be`
[†]Electronic address: `tomvandewiele@google.com`
[‡]Electronic address: `hans.vaneetvelde@ugent.be`

their current strength, as this would enable football fans and casual watchers to have a better feeling for who is the favourite and who is the underdog. However, the existing rankings, both at domestic leagues level and at national team level, fail to provide this, either because they are by nature not designed for that purpose or because they suffer from serious flaws.

Domestic league rankings obey the 3-1-0 principle, meaning that the winner gets 3 points, the loser 0 points and a draw earns each team 1 point. The ranking is very clear and fair, and tells at every moment of the season how strong a team has been since the beginning of the season. However, given that every match has the same impact on the ranking, it is not designed to reflect a team's current strength. A very recent illustration of this fact can be found in the German Bundesliga, where Borussia Dortmund had a tremendous start in the current season 2017-2018 with 19 out of 21 points after the first 7 rounds. In the next 7 rounds they only scored 3 points and were in a small crisis. In Round 15, Dortmund, ranked 6th with 22 points, received Werder Bremen, ranked 17th with only 11 points but who had the momentum on their side with 2 wins in the previous 3 matches. The ranking thus clearly favoured the home-team Dortmund, who however lost 1-2 to Bremen, a not so surprising result if one takes into account the recent performances of both teams.

Contrary to domestic league rankings, the FIFA/Coca-Cola World Ranking of national soccer teams is intended to rank teams according to their recent performances in international games. Bearing in mind that the FIFA ranking forms the basis of the seating and the draw in international competitions and its qualifiers, such a requirement on the ranking is indeed necessary. However, the current FIFA ranking fails to reach these goals in a satisfying way and is subject to many discussions (Cummings (2013); Tweedale (2015); The Associated Press (2015)). It is based on the 3-1-0 system, but each match outcome is multiplied by several factors like the opponent team's ranking and confederation, the importance of the game, and a time factor. We spare the reader those details here, which can be found on the webpage of the FIFA/Coca-Cola World Ranking (http://www.fifa.com/fifa-world-ranking/procedure/men.html). In brief, the ranking is based on the weighted average of ranking points a national team has won over each of the preceding four rolling years. The average ranking points over the last 12 month period make up half of the ranking points, while the average ranking points in the 13-24 months before the update count for 25% leaving 15% for the 25-36 month period and 10% for the 37-48 month period before the update. This arbitrary decay function is a major criticism of the FIFA ranking: a similar match of eleven months ago can have approximately twice the contribution as a match played twelve months ago. A striking example hereof was Scotland: ranked 50th in August 2013, it dropped to rank 63 in September 2013 before making a major jump to rank 35 in October 2013. This high volatility demonstrates a clear weakness in the FIFA ranking's ability of mirroring a team's current strength.

In this paper, we intend to fill the gap and develop a ranking that does reflect a soccer team's current strength. To this end, we consider and compare various existing and new statistical models that assign one or more strength parameters to each soccer team and where these parameters are estimated over an entire range of matches by means of maximum likelihood estimation. We shall propose a smooth time depreciation function to give more weight to more recent matches. The resulting ranking represents an interesting addition to the well-established rankings of domestic leagues and can be considered as promising alternative to the FIFA ranking of national teams.

The present paper is organized as follows. We shall present in Section 2 no less than 10 different strength-based models whose parameters are estimated via maximum likelihood. More precisely, via weighted maximum likelihood as we introduce two types of weight parameters: the above-mentioned time depreciation effect and a match importance effect for national team matches. In Section 3 we describe the exact computations behind our estimation procedures as

well as two criteria according to which we define a statistical model's predictive performance. Two case studies allow us to compare our 10 models at domestic league and national team levels: we investigate the English Premier League seasons from 2000-2017 (Section 4) as well as European national team matches between 2000 and 2016 (Section 5). Regarding national teams, we focus on European teams playing each other in order to have a balanced picture. Teams from different continents play different amounts of matches and the strength levels vary between continents. The FIFA ranking itself gives different weights to different continents, but these weights seem arbitrary, so we prefer not to take them up and just focus on European teams and inner-European matches. We illustrate the predictive strength of the best strength-based model for national teams by comparing our best model's prediction of the EURO2016 to the bookmaker's prediction (Section 5.2). On basis of the best-performing models, we then provide in Section 6 our new ranking for the current Premier League season as well as our ranking of European national teams at the beginning of 2018. We conclude the paper with final comments and an outlook on future research in Section 7.

2 The statistical strength-based models

2.1 Time depreciation and match importance factors

Our strength-based statistical models are of three main types: Bradley-Terry type models, independent Poisson models and bivariate Poisson models. Each model assigns strength parameters to all teams involved and models match outcomes via these parameters. Maximum likelihood estimation is employed to estimate the strength parameters, and the teams are ranked according to their resulting overall strengths. More precisely, we shall consider weighted maximum likelihood estimation, where the weights introduced are of two types: time depreciation (domestic leagues and national teams) and match importance (only national teams).

2.1.1 A smooth decay function based on the concept of Half period

A feature that is common to all considered models is our proposal of decay function in order to reflect the time depreciation. Instead of the step-wise decay function employed in the FIFA ranking, we rather suggest a continuous depreciation function that gives less weight to older matches with a maximum weight of 1 for a match played today. Specifically, the time weight for a match which is played x_i days back is calculated as

$$w_{time,i}(x_i) = \left(\frac{1}{2}\right)^{\frac{x_i}{\text{Half period}}},\tag{1}$$

meaning that a match played *Half period* days ago only contributes half as much as a match played today and a match played $3\times$*Half period* days ago contributes 12.5% of a match played today. Figure 1 shows a graphical comparison of our continuous time decay function versus the arbitrary FIFA decay function. In the sequel, $w_{time,i}$ will serve as weighting function in the likelihoods associated with our various models. This idea of weighted likelihood or pseudo-likelihood to better estimate a team's current strength is in line with the literature on modelling (mainly league) football scores, see Dixon and Coles (1997).

2.1.2 Match importance

While in domestic leagues all matches are equally important, the same cannot be said about national team matches where for instance friendly games are way less important than matches
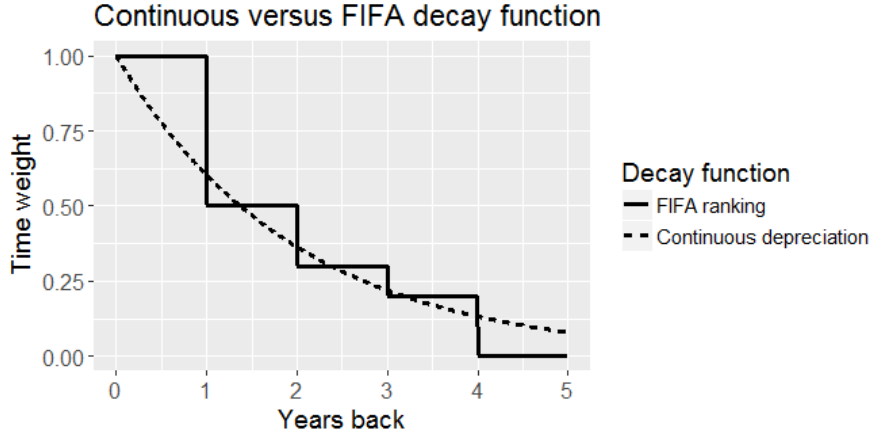
Figure 1: Comparison of the FIFA ranking decay function versus our exponential smoother (1). The continuous depreciation line uses a Half Period of 500 days.

played during the World Cup. Therefore we need to introduce importance factors. The FIFA weights seem reasonable for this purpose and will be employed whenever national team matches are analyzed. The relative importance of a national match is indicated by $w_{type,i}$ and can take the values 1 for a friendly game, 2.5 for a confederation or world cup qualifier, 3 for a confederation tournament (e.g. UEFA EURO2016 or the Africa Cup of Nations 2017) or the confederations cup, and 4 for World Cup matches.

2.2 The Bradley-Terry type models

Bradley-Terry (BT) models (Bradley and Terry, 1952) have been designed to predict the outcome of pairwise comparisons. Assuming that two individuals i and j possess strength parameters s_i and s_j, the probability of i winning over j is given by $\frac{s_i}{s_i+s_j}$ in BT models. Wang and Vandebroek (2013) modelled soccer match outcomes by means of an adapted BT model. We shall first describe their original idea and then propose two extended BT-type models.

2.2.1 Original Bradley-Terry model

The standard BT model only allows two possible outcomes (home win or home loss) whereas in football the probability of a draw must be modeled as well. The draw probability is modeled similarly to Davidson (1970) in such a way that the probability of a draw increases as the strengths of both teams are more similar. The adapted BT model also takes a home effect into account. Wang and Vandebroek (2013) have opted for a multiplicative home effect, following an analysis of 10 recent Italian Serie A seasons that revealed that home advantage is proportional to the strength of the home team. A different viewpoint, which we shall not follow here, was developed in Bassett (1997) who suggested an additive home constant.

From now on we denote the home team as team H and the away team as team A. A total of M team strengths need to be estimated when M teams are analyzed. If we call P_{iH} the probability of a home win in match i, P_{iD} the probability of a draw in match i and P_{iA} the

probability of an away win in match i, then the outcome probabilities are

$$P_{iH} = \frac{h\beta_{iH}}{h\beta_{iH} + d\sqrt{h\beta_{iH}\beta_{iA}} + \beta_{iA}};$$

$$P_{iD} = \frac{d\sqrt{h\beta_{iH}\beta_{iA}}}{h\beta_{iH} + d\sqrt{h\beta_{iH}\beta_{iA}} + \beta_{iA}};$$

$$P_{iA} = \frac{\beta_{iA}}{h\beta_{iH} + d\sqrt{h\beta_{iH}\beta_{iA}} + \beta_{iA}},$$

with home effect represented by h, draw effect by d, and β_{iH} and β_{iA} respectively stand for the strength parameters of home and away team in match $i \in \{1, \ldots, N\}$, N denoting the total number of matches in the considered time period. A home effect $h > 1$ inflates the strength of the home team and increases its modeled probability to win the match. This is typically the case since playing at home gives the benefit of familiar surroundings, the support of the home crowd and the lack of traveling. Matches on neutral ground are modeled by dropping the home effect h. The draw effect is best understood by assuming similar strengths in the absence of a home effect. In that case P_{iH} is similar to P_{iA} and the relative probability of P_{iD} compared to a home win or loss is approximately equal to d. All parameters are non-negative.

The strength parameters are estimated using maximum likelihood estimation on match outcomes. Let y_{ij} be 1 if the result of game i is j and $y_{ij} = 0$ otherwise, for $j = H, D, A$ as explained above. Under the common assumption that matches are independent, the likelihood for N matches corresponds to

$$L = \prod_{i=1}^{N} \prod_{j \in \{H,D,A\}} P_{ij}^{y_{ij} \cdot w_{type,i} \cdot w_{time,i}} \tag{2}$$

with $w_{type,i}$ and $w_{time,i}$ the weights described in Section 2.1. The model parameters p (here, the M team strengths and the parameters h and d) are replaced with $\exp(p)$ in the estimation procedure to make sure that the estimated parameters are positive.

2.2.2 Bradley-Terry model with Goal Difference weights

The basic BT model of the previous section does not use all of the available information. It only takes the match outcome into account, omitting likely valuable information present in the goal difference. A team that wins by 8-0 and loses the return match by 0-1 is probably stronger than the opponent team. Our first extension of the BT model modifies the basic BT model in the sense that matches are given an increasing weight when the goal difference grows. The likelihood function is calculated as

$$L = \prod_{i=1}^{N} \prod_{j \in \{H,D,A\}} P_{ij}^{y_{ij} \cdot w_{goalDiffscaled,i} \cdot w_{type,i} \cdot w_{time,i}}.$$

This formula slightly differs from (2) through the goal difference weight

$$w_{goalDiffscaled,i} = \begin{cases} 1 & \text{if draw} \\ \log_2(goalDiff_i + 1) & \text{else,} \end{cases}$$

with $goalDiff_i$ the absolute value of the goal difference in match i (both outcomes 2-0 and 0-2 thus give the same goal difference of 2). This way, a goal difference of 1 receives a goal difference weight of 1 and every additional increment in goal difference results in a smaller increase of the goal difference weight. A goal difference of 7 goals receives a goal difference weight of 3. Parameter estimation is achieved in the same way as in the basic BT model.

2.2.3 Bradley-Terry model with home and away strengths

Our second extension of the basic BT model contains separate team strength parameters for home and away teams instead of a single home effect parameter common to all teams. This way, the home team strength is not influenced by performances on the road and the away team strength is not influenced by results in home matches. A team is modeled as if it were a separate team playing at home or away from home. The draw effect is modeled as in the basic BT model, and parameter estimation follows again via maximum likelihood; we omit the expression here as it is highly similar to (2). The strength of a team on neutral ground is modeled in the same fashion as we define the overall strength of a team, namely by combining the home and away strengths. All three Pythagorean means were considered for this purpose (the arithmetic, geometric and harmonic means). An analysis of more than 100 seasons of Premier League data (1892-2014) revealed that the geometric mean is best suited for the combination of team strengths. The absolute difference in ranks between the ranking based on each of the Pythagorean means of the modeled team strengths (two for each team) and the actual final league ranking based on the 2-point-for-win-system was used to identify the preferred approach.

2.3 The Independent Poisson models

Poisson models were first suggested by Maher (1982) to model football match results. He assumed the number of scored goals by both teams to be independent Poisson distributed variables. Let G_{iH} and G_{iA} be the random variables representing the goals scored by the home team and the away team in match i, respectively. With those assumptions the probability function can be written as

$$\mathrm{P}(G_{iH} = x, G_{iA} = y) = \frac{\lambda_{iH}^x}{x!} \exp(-\lambda_{iH}) \cdot \frac{\lambda_{iA}^y}{y!} \exp(-\lambda_{iA}), \tag{3}$$

where λ_{iH} and λ_{iA} stand for the means of G_{iH} and G_{iA}, respectively.

Being a count-type distribution, the Poisson is a natural choice to model soccer matches. It bares yet another advantage when it comes to predicting matches. If $G_i = G_{iH} - G_{iA}$, then the probability of a win of team H over team A, the probability of a draw as well as the win of team A are respectively computed as $\mathrm{P}(G_i > 0)$, $\mathrm{P}(G_i = 0)$ and $\mathrm{P}(G_i < 0)$. The Skellam distribution, the discrete probability distribution of the difference of two independent Poisson random variables, is used to derive these probabilities given λ_{iH} and λ_{iA}. This renders the prediction of future matches via the Poisson model particularly simple.

2.3.1 Independent Poisson model with 1 strength parameter

Attributing again a single strength parameter to each team, denoted as before by β_{iH} and β_{iA} for match i, we define the Poisson means as $\lambda_{iH} = c \times h \times \frac{\beta_{iH}}{\beta_{iA}}$ and $\lambda_{iA} = c \times \frac{\beta_{iA}}{\beta_{iH}}$ with h the home effect, c a common intercept, and under the constraint that all parameters are positive like in the BT models. Matches on neutral ground are modeled by dropping the home effect h. With this in hand, the overall likelihood can be written as

$$L = \prod_{i=1}^{N} \left(\frac{\lambda_{iH}^{g_{iH}}}{g_{iH}!} \exp(-\lambda_{iH}) \cdot \frac{\lambda_{iA}^{g_{iA}}}{g_{iA}!} \exp(-\lambda_{iA}) \right)^{w_{type,i} \cdot w_{time,i}},$$

where g_{iH} and g_{iA} stand for the actual home and away goals scored in match i. Maximum likelihood estimation yields the values of the strength parameters. It is important to notice that the Poisson model uses two observations for each match (the goals scored by each team) while

using the same number of parameters (number of teams + 2). The BT approach, except for the model with Goal Difference Weight, only uses a single observation for each match.

2.3.2 Independent Poisson model with 2 strength parameters or Maher's original model

In the previous section we have defined a slightly simplified version of Maher's original model. Maher assumed the scoring rates to be of the form $\lambda_{iH} = c \times h \times \frac{o_{iH}}{d_{iA}}$ and $\lambda_{iA} = c \times \frac{o_{iA}}{d_{iH}}$, with o_{iH}, o_{iA}, d_{iH} and d_{iA} standing for offensive and defensive capabilities of both teams. The overall strength of a team is again a combination of both strength parameters, obtained in the same way as for the BT model with home and away strength parameters.

2.4 The Bivariate Poisson models

A potential drawback of the Independent Poisson models lies precisely in the independence assumption. Of course, some sort of dependence between the two playing teams is introduced by the fact that the strength parameters of each team are present in the Poisson means of both teams, however this may not be a sufficiently rich model to cover the interdependence between two teams.

Karlis and Ntzoufras (2003) suggested a bivariate Poisson model by adding a correlation between the scores. The home and away scores are modelled as $G_{iH} = X_{iH} + X_{iC}$ and $G_{iA} = X_{iA} + X_{iC}$, where X_{iH}, X_{iA} and X_{iC} are independent Poisson distributed variables with parameters λ_{iH}, λ_{iA} and λ_{iC}, respectively. The joint probability function of the home and away score is then given by

$$\mathrm{P}(G_{iH} = x, G_{iA} = y) = \frac{\lambda_{iH}^x \lambda_{iA}^y}{x!y!} \exp(-(\lambda_{iH} + \lambda_{iA} + \lambda_{iC})) \sum_{k=0}^{\min(x,y)} \binom{x}{k}\binom{y}{k} k! \left(\frac{\lambda_{iC}}{\lambda_{iH}\lambda_{iA}} \right)^k, \quad (4)$$

which is the formula for the bivariate Poisson distribution with parameters λ_{iH}, λ_{iA} and λ_{iC}. It reduces to (3) when $\lambda_{iC} = 0$. This parameter thus can be interpreted as the covariance between the home and away scores and might reflect the game conditions. The strength parameters λ_{iH} and λ_{iA} are similar as in the Independent model, but we attract the reader's attention to the fact that the means for the home and away scores are now given by $\lambda_{iH} + \lambda_{iC}$ and $\lambda_{iA} + \lambda_{iC}$, respectively.

Letting G_i again stand for the goal difference, we can easily see that the probability function of the goal difference for the bivariate case is the same as the probability function for the Independent model with parameters λ_{iH} and λ_{iA}, since

$$\begin{aligned} P(G_i = x) &= P(G_{iH} - G_{iA} = x) \\ &= P(X_{iH} + X_{iC} - (X_{iA} + X_{iC}) = x) = P(X_{iH} - X_{iA} = x), \end{aligned}$$

implying that we can again use the Skellam distribution for predicting the winner of future games.

2.4.1 Bivariate Poisson with constant covariance

The first bivariate Poisson models that we consider are parsimonious in terms of parameters since we assume constant covariance over all matches. The first model, termed Bivariate Poisson 00 with 1 strength parameter, defines $\lambda_{iH} = c \times h \times \frac{\beta_{iH}}{\beta_{iA}}$ and $\lambda_{iA} = c \times \frac{\beta_{iA}}{\beta_{iH}}$ and is a modified version of the simplest Karlis and Ntzoufras (2003) proposal. The second model, termed Bivariate Poisson

00 with 2 strength parameters, corresponds to that version by Karlis and Ntzoufras (2003) and is defined via $\lambda_{iH} = c \times h \times \frac{o_{iH}}{d_{iA}}$ and $\lambda_{iA} = c \times \frac{o_{iA}}{d_{iH}}$ where each team has an offensive and defensive strength parameter. In both cases the covariance is governed by $\lambda_{iC} = b_C$ for some constant covariance parameter b_C. Again all parameters should be positive and we apply the same weight factors as in the previous sections for the weighted maximum likelihood estimation.

2.4.2 Bivariate Poisson with non-constant covariance

Karlis and Ntzoufras (2003) added other team-specific parameters, apart from the offensive and defensive strength parameters. These new parameters are used to model the influence of the teams on the game conditions. In this regard we can investigate three new models, one where the home team influences the game conditions, one where the away team does and one model where both teams influence the conditions of the game. Let γ_H be 1 if the home team has influence on the covariance and 0 if not, and let γ_A be defined analogously. Then we consider the covariance as

$$\lambda_{iC} = b_C \times b_{iH}^{\gamma_H} \times b_{iA}^{\gamma_A} \tag{5}$$

where b_{iH} and b_{iA} are the covariance-influencing parameters for the home and away team, respectively. This means that, if for instance $\gamma_H = 1$ and $\gamma_A = 0$, we only allow the home team to affect the covariance. If $\gamma_H = \gamma_A = 0$ we are in the case of constant covariance between the home and away score, hence this corresponds to the Bivariate Poisson 00 model with two strength parameters. With an obvious notation, we term the three other possible models based upon (5) as Bivariate Poisson 10, Bivariate Poisson 11 and Bivariate Poisson 01, respectively. Each model has offensive and defensive strength parameters per team as well as the covariance-influencing parameter b_i, and the estimation procedure runs along the same lines as for the Bivariate Poisson 00 models.

One can think of many other ways to model dependent football scores. The dependence between the home and away scores can for instance be modelled by all kinds of copulas or adaptations of the Independent model. Incorporating them all in our analysis seems an impossible task, which is why we opted for the very prominent Karlis-Ntzoufras proposal. Notwithstanding, we mention some important contributions in this field: Dixon and Coles (1997) added an additional parameter to adjust for the probabilities on low scoring games (0-0, 1-0, 0-1 and 1-1), McHale and Scarf (2011) investigated copula dependence structures, and Boshnakov et al. (2017) recently proposed a copula-based bivariate Weibull count model.

3   Parameter estimation and model selection

In this section we shall briefly describe two crucial statistical aspects of our investigation, namely how we compute the maximum likelihood estimates and which criteria we apply to select the model with the highest predictive performance.

3.1   Computing the maximum likelihood estimates

Parameters in the BT, Independent and Bivariate Poisson models are estimated using maximum likelihood estimation. To this end, we have used the *optim* function in R (R Development Core Team, 2018) by specifying as preferred method the *BFGS* (Broyden-Fletcher-Goldfarb-Shanno optimization algorithm). We have opted for this quasi-Newton method because of its robust properties. Since all model parameters need to be positive by definition of our statistical model,

we have applied the *exp* transformation to all parameters during the likelihood calculation. All optimized parameters are initialized at 0 (1 in the estimation procedure after applying the *exp* transformation).

3.2   Measures of predictive performance

The studied models are built to perform three way outcome prediction (home win, draw or home loss). Each of the three possible match outcomes is predicted with a certain probability but only the actual outcome is observed. The predicted probability of the outcome that was actually observed is thus a natural measure of predictive performance. The ideal predictive performance metric is able to select the model which approximates the true outcome probabilities the best.

We consider two such metrics in our analysis. The first is the log loss criterion. The average log loss for N matches is defined as $log\ loss = -\frac{1}{N}\sum_{i=1}^{N}\sum_{j\in\{H,D,A\}} y_{ij} \cdot \log(P_{ij})$ where P_{ij} represents the modeled probabilities and y_{ij} is 1 if the result of game i is j and $y_{ij} = 0$ otherwise. The smaller the log loss, the better the predictive performance of the considered model. The log loss criterion is typically used as the performance metric in Data Science competitions when the goal is to predict probabilities for a set of possible outcomes.

The second metric is the Rank Probability Score (RPS) of Epstein (1969). It represents the difference between cumulative predicted and observed distributions via the formula

$$\frac{1}{2N}\sum_{i=1}^{N}\left((P_{iH} - y_{iH})^2 + (P_{iA} - y_{iA})^2\right)$$

where we keep the same notations as above. It has been shown in Constantinou and Fenton (2012) that the RPS is more appropriate as soccer performance metric than other popular metrics such as the RMS and Brier score. The reason is that, by construction, the RPS works at an ordinal instead of nominal scale, meaning that, for instance, it penalizes more severely a wrongly predicted home win in case of a home loss than in case of a draw.

4   Case study 1: Premier League

4.1   Data collection and comparison of the predictive models

The engsoccerdata package (Curley, 2015) contains results of all top 4 tier football leagues in England since 1888. The dataset contains the date of the match, the teams that played, the tier as well as the result. The number of teams equals 20 for each of the seasons considered (2000-2017). Matches are predicted for every season separately and on every match day of the second half of the season, using the first season half for training the models. So a total of 3230 matches are predicted. Accordingly, every season is analyzed separately from previous years. This avoids the issue of assigning strengths to promoted teams and takes team changes during the summer transfer season into account. Matches are predicted in blocks corresponding to each round, and after every round the parameters are updated. In all our models, the Half Period is varied between 20 and 600 days in steps of 20. We compare our models also to a simple method which we call the Majority method. Based only upon the percentage of home wins, draws and home losses in the season that is currently being evaluated, this method assigns to every game the same probabilities, regardless of the teams that are playing. It is therefore fair to say that models that perform worse than the majority method are valueless.

Table 1 summarizes the analysis by comparing the best performing models of each of the considered classes, i.e. the model with the optimal Half Period. As we can see, the Independent

Table 1: Comparison table for the best performing models of each of the considered classes with respect to the RPS criterion and the log loss criterion. All of the second season half English Premier League matches in the period between the seasons 2000-2001 and 2016-2017 are considered.

Model Class	Number of teamspecific parameters	Half Period best model RPS	Lowest RPS	Half Period best model log loss	Lowest log loss
Bradley-Terry	1	HP = 540	0.2019895	HP = 600	0.9948099
BT Goal Difference	1	HP = 360	0.2025017	HP = 440	1.0012263
BT Home & Away	2	HP = 600	0.2147502	HP = 600	1.0844735
Independent Poisson	1	HP = 200	0.1978573	HP = 240	0.9766120
Independent Poisson	2	HP = 240	0.1987503	HP = 300	0.9785288
Bivariate Poisson 00	1	HP = 180	0.1981067	HP = 240	0.9776285
Bivariate Poisson 00	2	HP = 240	0.1985977	HP = 380	0.9815752
Bivariate Poisson 10	3	HP = 260	0.1995125	HP = 400	0.9881019
Bivariate Poisson 01	3	HP = 240	0.1995247	HP = 300	0.9861578
Bivariate Poisson 11	3	HP = 280	0.1996435	HP = 200	0.9906929
Majority method	0	HP = ∞	0.2255210	HP = ∞	1.0387409

Poisson model with 1 strength parameter per team is the best according to both log loss and RPS, followed by the bivariate Poisson model with just one parameter per team. So, parsimony in terms of parameters to estimate is important. We also clearly see that all Poisson-based models outperform the BT models according to both criteria. This was to be expected since Poisson models use the goals as additional information. Considering the win margin in the BT model does not improve its performance. It is also noteworthy that the best two models have among the lowest Half Periods. Finally we remark that the BT model with home and away parameters does worse than the Majority method according to the log loss criterion.

4.2 Comparison with the bookmaker's prediction

The website http://www.football-data.co.uk/ hosts an excel file with match data and odds of each English Premier League season since the season 2000-2001. Decimal odds are available for thirteen major betting companies. Nine bookmakers (Bet365, Bet & Win, Gamebookers, Interwetten, Ladbrokes, Sportingbet, StanJames, VCBet and William Hill) contained odds for at least half of the matches, the other four were discarded. Decimal odds are translated to outcome probabilities assuming fixed overround margins on all three outcomes. If the observed odds are [2,3,3.5] for a home win, a draw and a home loss respectively, the transformed probabilities become $\left[\frac{1}{2}, \frac{1}{3}, \frac{1}{3.5}\right] / \left(\frac{1}{2} + \frac{1}{3} + \frac{1}{3.5}\right) \approx [0.447, 0.298, 0.255]$. The average bookmaker predictive performance is calculated by transforming the odds of nine major bookmaker firms to bookmaker probabilities.

The resulting RPS and log loss values are, respectively 0.1917733 and 0.9569953. This means that our best-performing model, the Independent Poisson model with 1 strength parameter, lies only 3.17% respectively 2.05% above the bookmaker scores, which is quite remarkable given the additional information that these gold standard predictions take into account. It underlines that our parsimonious model, based on sound statistical estimation, is able to achieve very good predictions.

5 Case study 2: National teams

As for the Premier League, we shall first compare our various models in terms of log loss and RPS, and analyze which model has the best predictive performance. For the sake of illustration, we then use this best-performing model to make a post-hoc analysis of the UEFA EURO2016 tournament by comparing our predictions to the bookmaker ratings. We will achieve this goal by simulating repeatedly all matches of the tournament, hence our predictions are based on 100.000 simulations of the EURO2016.

5.1 National teams data collection and comparison of the predicting models

National team match results were scraped from the website http://eu-football.info/. The platform contains a complete archive of all European national football team results since 1872 where at least one of the playing teams was European. The website also contains historical results of all European and domestic club competitions. Results are organized by national team and are extracted by parsing through all the national team pages using the R package rvest (Wickham, 2015). We predicted the outcome of 3868 games between two European teams in the period from 2000 to 2016. The last game in our analysis is played on 2016-11-14. The parameters are estimated by maximum likelihood on a period of four years. The half period is varied from 100 to 1500 days in steps of 100.

The results of our model comparison are provided in Table 2. This time we notice differences between the log loss and RPS rankings. Both rankings agree on preferring Poisson-like models over BT-type models. However, while the RPS criterion favors the Bivariate Poisson model with 2 strength parameters followed by the Bivariate Poisson model with 1 strength parameter and the Independent Poisson model with 1 strength parameter, the log loss criterion ranks the Bivariate Poisson model with 2 strength parameters only third, its 1-parameter competitor being classified first. Quite remarkably, the difference between the Bivariate and Independent Poisson models with 1 strength parameter is very small in both performance metrics. This is due to the small covariance we find in the bivariate model: it lies nearly all the time under 0.1. If we add up the log loss and RPS scores, the Bivariate Poisson model with 1 strength parameter turns out to be the best, which is why we consider it in the next section when simulating the EURO2016.

5.2 Predicting the EURO2016: simulation details and comparison to the bookmaker ranking

The UEFA EURO2016 tournament was played in France by 24 teams. Six groups of four teams were drawn on December 12, 2015 in Paris. The first two ranked teams progress to the knockout rounds as well as the four best third ranked teams of the six groups. Four knockout stages were played to select a winner. Portugal, who could not win any of its group matches, won the final against host team and tournament favorite France. On their way to the final, Portugal had to defeat Poland and Wales in the quarter- and semi-finals, while France won against Iceland and world champion Germany.

In all simulated matches with our Bivariate Poisson model, team strengths are converted to the expected number of goals scored by both teams according to (4). The home effect is only applied to France. During the simulations, we did not refit the Bivariate Poisson model after each of the three group stage rounds nor after the first three of the four knockout stages. This allowed us to simulate the entire tournament with the initial strengths of teams, and hence avoided introducing a potential simulation-based estimation bias. This choice is further corroborated by

Table 2: Comparison table for the best performing models of each of the considered classes with respect to the average log loss criterion and the RPS criterion. All of the matches between the European national teams in the period 2000-2016 are considered.

Model Class	Number of teamspecific parameters	Half Period best model RPS	Lowest RPS	Half Period best model log loss	Lowest log loss
Bradley Terry	1	HP = 1500	0.1729464	HP = 1400	0.8915426
BT Goal Difference	1	HP = 1500	0.1728042	HP = 1400	0.8973620
BT Home & Away	2	HP = 1500	0.1850369	HP = 1500	0.9998043
Independent Poisson	1	HP = 1400	0.1690249	HP = 1400	0.8612584
Independent Poisson	2	HP = 1500	0.1709804	HP = 1500	0.8691197
Bivariate Poisson 00	1	HP = 1100	0.1690129	HP = 1100	0.8609054
Bivariate Poisson 00	2	HP = 1100	0.1688754	HP = 1100	0.8646181
Bivariate Poisson 10	3	HP = 1200	0.1697885	HP = 1500	0.8778760
Bivariate Poisson 01	3	HP = 1300	0.1693430	HP = 1500	0.8779135
Bivariate Poisson 11	3	HP = 1400	0.1693875	HP = 1500	0.8838847
Majority method	0	HP=∞	0.2340603	HP = ∞	1.0725820

the high Half Period of 1100 days, meaning that matches played four years ago are considered about half as influential as matches played today. Consequently, the estimates of the outcomes are not expected to be biased heavily by not refitting the model.

Articles 16-19 of the rules (UEFA.com, 2016) for the UEFA EURO2016 tournament were used to determine the ranking of the teams based on the group match results. The rules were also used to calculate the best four thirds of the six groups and pair the correct teams in the knockout rounds. Knockout matches that end in a draw after ninety minutes are extended by two 15-minute periods. The simulation study handles this by simulating a new match with the expected goals scored set to one third of the main part of the match. Penalty shootouts are simulated by assigning a random winner.

The UEFA EURO2016 tournament was simulated 100,000 times. A bar chart of the tournament win frequency of the participating teams is shown in Figure 2. France is -also due to the home effect- the main favorite to win the tournament according to the simulation study. It wins in about 22% of the simulations, while Germany wins in 13.7%. Spain and Belgium are expected to win the tournament with a probability close to 10% and are followed by England that is ascribed a win probability of about 9%. Ukraine and Portugal, the actual winner, are the main outsiders according to the simulation study.

Bookmaker odds are assumed to be the gold standard and the quality of our predicted ranking is judged by the similarity with the bookmaker ranking. Bookmaker odds were collected from the website talksport.com on June 7, 2016. The ranks of the simulation study are strongly correlated with the bookmaker ranks ($\rho = 0.892$). Looking at the actual outcome, we see that our model performs as well as the bookmakers in predicting Portugal as winner of the EURO. France, the favorite of our model and of the bookmakers, nearly fulfilled the expectations as they only lost the final against the Portuguese. According to the last FIFA ranking before the start of the EURO2016, France was only ranked 10[th] among European countries, underlining the fallacies of the FIFA ranking to reflect the current strengths of national teams.
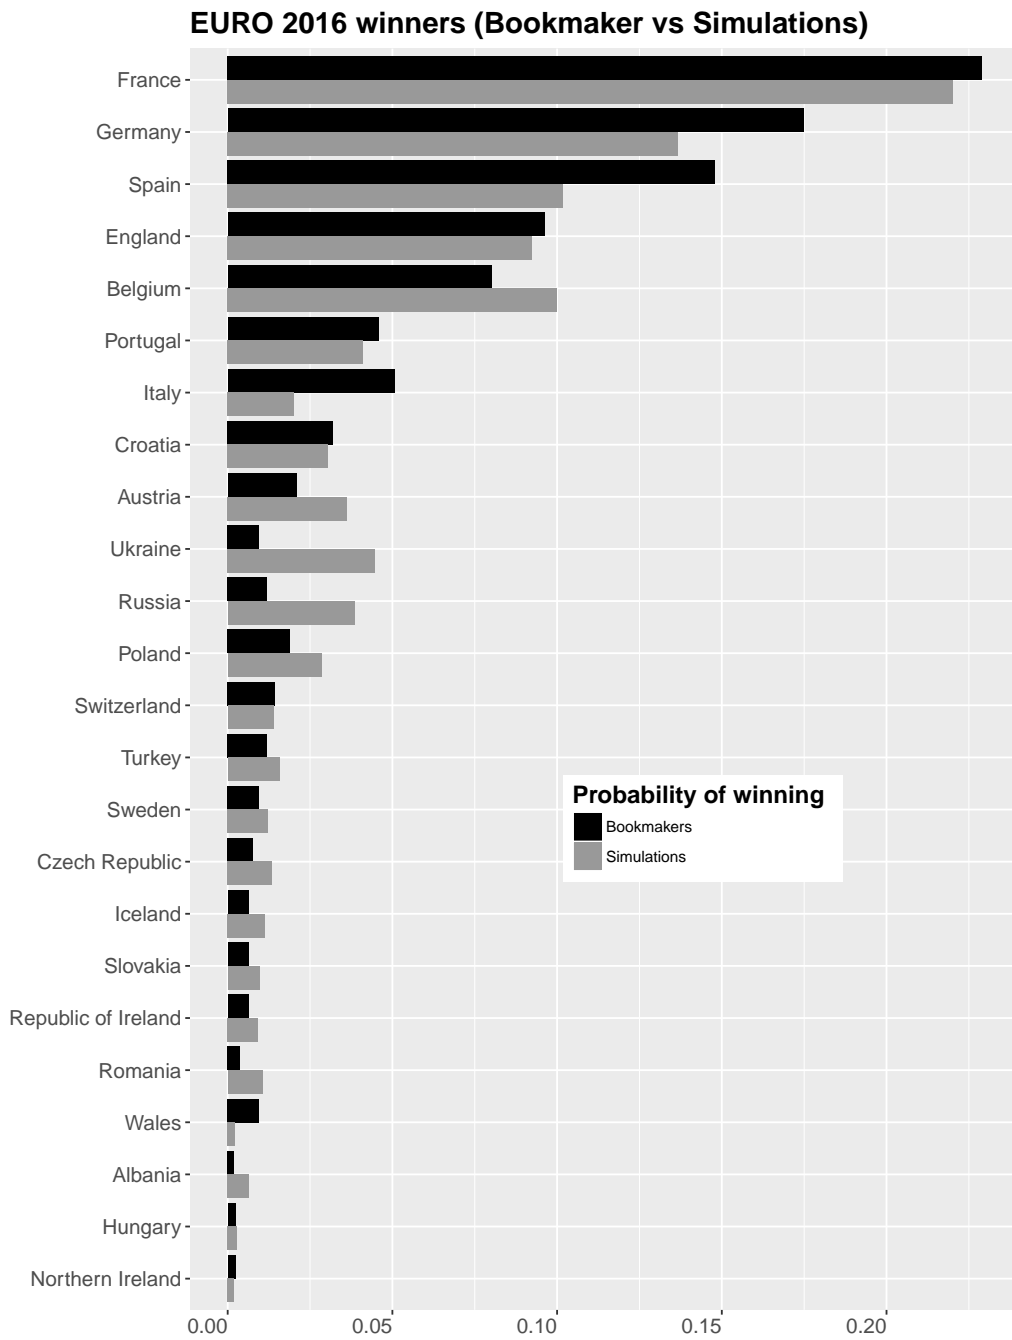
Figure 2: Comparison of winning probabilities of the bookmakers on June 7, 2016 and the probabilities of our simulation study based on all matches played up till the start of the EURO2016 tournament.

6 Rankings

We have determined in the previous sections which strength-based models have the best predictive power at both domestic leagues and national teams levels. On basis of the respective best-performing models, we now provide the ranking of the current Premier League season and the ranking of European national teams at the beginning of 2018.

Table 3 contains two rankings for the English Premier League on February 1 2018, after 25 rounds. The right-hand side shows the official ranking where every match has the same importance, while the left-hand side showcases our current-strength ranking based on the Independent Poisson model with Half Period 200, the best-performing model according to the RPS criterion. Both rankings have Manchester City in a dominating position, which is not surprising. A clear difference is the ranking of Liverpool, second according to our ranking but only fourth in the official ranking. This reflects well the current strong performance of Liverpool, who even beat Manchester City in round 23, but who had a difficult start into the season with only a nineth place after the first nine rounds. Liverpool is thus a striking example of the usefulness of our new ranking. Another notable difference is Huddersfield Town who had a strong start in the season with 7/9 points but who struggle recently.

Table 4 displays a comparison between our strength-based model (left-hand side) and the official FIFA/Coca-Cola ranking (right-hand side) on January 1 2018 (to be precise, the last calculation of the FIFA ranking was done on December 21 2017). Our ranking is based on the Bivariate Poisson 00 model with Half Period 1100, chosen according to the RPS. We attract the reader's attention to the fact that "Position" in the FIFA ranking means the world-wide position, which explains the missing numbers such as 2, 4, 10, etc. A first notable difference lies in the readability of the two rankings: while one can understand the values of the strength parameters as ratios leading to the average number of goals that one team will score against the other, the same cannot be said about the FIFA points which do not allow making predictions. In our ranking, the current strong performance of Spain is underlined. A major difference occurs with Russia: ranked 16^{th} in our ranking, it is only the 33^{rd} European team in the FIFA ranking (worldwide 64). This difference is due to the fact that Russia is hosting the World Cup 2018 and hence plays friendly games instead of the qualifiers, a fact that is penalized much more severely in the FIFA ranking than in our ranking.

7 Conclusion and outlook

We have compared 10 different statistical strength-based models according to their predictive performance. Our analysis clearly demonstrates that Poisson-type models outperform Bradley-Terry type models, and that the best models are those that assign the fewest parameters to teams. At domestic team level, the Independent Poisson model with one strength parameter per team was found to be the best in terms of both RPS and log loss criteria, while the Bivariate Poisson model with one strength parameter per team was overall the best at national team level. However, the difference between that model and the Independent Poisson with one strength parameter is very small, which is explained by the fact that the covariance in the bivariate Poisson model is close to zero. This is well in line with recent findings of Groll et al. (2017) who used the same Bivariate Poisson model in a regression context. Applying it to the European Championships 2004-2012, they got a covariance parameter close to zero. Besides in the context of regression models, bivariate Poisson models have also been used in a dynamic way. Such dynamic bivariate Poisson models were applied to analyze English Premier League matches by Crowder et al. (2002) and Koopman and Lit (2015).

Our final conclusion is that the Independent Poisson model with one strength parameter

Table 3: Ranking of the Premier League teams on February 1 2018 according to the Independent Poisson model with Half Period 200 compared to the official Premier League ranking at that date.

Position	Team	Strength	Team	Points
1	Manchester City	2.06	Manchester City	68
2	Liverpool	1.58	Manchester United	53
3	Manchester United	1.52	Chelsea	50
4	Tottenham Hotspur	1.49	Liverpool	50
5	Chelsea	1.42	Tottenham Hotspur	48
6	Arsenal	1.21	Arsenal	42
7	Leicester City	1.09	Burnley FC	35
8	Burnley FC	0.98	Leicester City	34
9	Bournemouth	0.95	Everton	31
10	Everton	0.88	Bournemouth	28
11	Crystal Palace	0.86	Watford	27
12	West Ham	0.86	West Ham	27
13	Southampton	0.84	Crystal Palace	26
14	Watford	0.82	Brighton and Hove Albion	24
15	Newcastle United	0.82	Huddersfield Town	24
16	West Bromwich Albion	0.80	Newcastle United	24
17	Swansea City	0.79	Stoke City	24
18	Brighton and Hove Albion	0.76	Southampton	23
19	Stoke City	0.67	Swansea City	23
20	Huddersfield Town	0.66	West Bromwich Albion	20

Table 4: Ranking of the European national teams on January 1 2018 according to the Bivariate Poisson 00 model with Half Period 1100 compared to the Official FIFA/Coca-Cola World Ranking on December 21 2017.

Position	Team	Strength	Position	Team	Points
1	Spain	2.23	1	Germany	1602(1601.91)
2	Germany	2.18	3	Portugal	1358(1357.89)
3	Belgium	2.08	5	Belgium	1325(1325.29)
4	France	2.01	6	Spain	1231(1230.83)
5	Portugal	1.84	7	Poland	1209(1208.96)
6	Poland	1.81	8	Switzerland	1190(1190.44)
7	England	1.68	9	France	1183(1183.3)
8	Denmark	1.63	12	Denmark	1099(1098.82)
9	Sweden	1.57	14	Italy	1052(1052.49)
10	Croatia	1.56	15	England	1047(1046.98)
11	Italy	1.55	17	Croatia	1018(1017.91)
12	Netherlands	1.55	18	Sweden	998(997.82)
13	Iceland	1.43	19	Wales	985(984.56)
14	Wales	1.40	20	Netherlands	952(952.49)
15	Slovakia	1.39	22	Iceland	910(909.55)
16	Russia	1.36	24	Northern Ireland	867(866.92)
17	Switzerland	1.36	28	Slovakia	817(816.92)
18	Serbia	1.32	29	Austria	815(815.41)
19	Scotland	1.27	32	Republic of Ireland	798(798.06)
20	Repupblic of Ireland	1.26	32	Scotland	798(797.76)
21	Ukraine	1.26	35	Ukraine	781(781.25)
22	Romania	1.23	36	Serbia	756(755.6)
23	Northern Ireland	1.23	37	Bosnia and Herzegovina	753(752.52)
24	Austria	1.22	41	Romania	737(736.67)
25	Czech Republic	1.22	42	Turkey	735(735.28)
26	Turkey	1.19	43	Bulgaria	719(718.88)
27	Bosnia and Herzegovina	1.17	46	Montenegro	681(681.41)
28	Montenegro	1.09	47	Greece	680(679.66)
29	Slovenia	1.08	48	Czech Republic	677(676.87)
30	Albania	1.07	53	Hungary	630(630.01)
31	Greece	0.97	59	Norway	573(573.29)
32	Hungary	0.95	62	Albania	553(552.62)
33	Finland	0.93	64	Russia	534(533.51)
34	FYR Macedonia	0.93	66	Finland	531(531.42)
35	Norway	0.93	69	Slovenia	522(521.72)
36	Bulgaria	0.90	76	FYR Macedonia	446(445.65)
37	Israel	0.86	83	Luxembourg	407(407.47)
38	Georgia	0.85	89	Estonia	397(396.53)
39	Belarus	0.74	90	Armenia	383(382.62)
40	Estonia	0.72	91	Cyprus	373(373.47)
41	Cyprus	0.71	92	Belarus	372(372.25)
42	Armenia	0.67	95	Faroe Islands	364(363.64)
43	Latvia	0.63	98	Israel	355(355.18)
44	Azerbaijan	0.62	104	Georgia	322(321.71)
45	Lithuania	0.60	117	Azerbaijan	281(281.31)
46	Moldova	0.60	132	Latvia	233(233.01)
47	Kazakhstan	0.59	137	Kazakhstan	220(220.1)
48	Faroe Islands	0.59	138	Andorra	215(214.8)
49	Luxembourg	0.58	149	Lithuania	179(178.96)
50	Malta	0.51	167	Moldova	111(110.7)
51	Andorra	0.41	177	Kosovo	97(97.49)
52	Liechtenstein	0.36	181	Liechtenstein	86(86.1)
53	Gibraltar	0.26	184	Malta	66(65.87)
54	San Marino	0.25	204	San Marino	11(11.09)

per team is a simple, easy-to-compute, interpretable and strongly predictive model for soccer games, both at domestic and national team level. Its best competitor is the Bivariate Poisson model with one strength parameter. Both models lead to rankings that reflect well a soccer team's current strength. Further exploration of these models at national team level including all teams worldwide is currently under investigation, with the aim of proposing an alternative FIFA ranking.

References

Bassett, G. (1997). Robust sports ratings based on least absolute errors. *The American Statistician*, 51:99–105.

Boshnakov, G., Kharrat, T., and McHale, I. G. (2017). A bivariate Weibull count model for forecasting association football scores. *International Journal of Forecasting*, 33(2):458–466.

Bradley, R. A. and Terry, M. E. (1952). Rank analysis of incomplete block designs: I. The method of paired comparisons. *Biometrika*, 39(3/4):324–345.

Constantinou, A. C. and Fenton, N. E. (2012). Solving the problem of inadequate scoring rules for assessing probabilistic football forecast models. *Journal of Quantitative Analysis in Sports*, 8(1):1559–0410.1418.

Crowder, M., Dixon, M., Ledford, A., and Robinson, M. (2002). Dynamic modelling and prediction of English Football League matches for betting. *Journal of the Royal Statistical Society Series D*, 51:157–168.

Cummings, M. (2013). FIFA World Rankings place Brazil 18th, reinforce flawed nature of system. *Bleacherreport.com*.

Curley, J. (2015). *engsoccerdata: Soccer Data 1871-2015*. R package version 0.1.4.

Davidson, R. (1970). On extending the Bradley-Terry model to accommodate ties in paired comparison experiments. *Journal of the Royal Statistical Society Series D*, 65:317–328.

Dixon, M. J. and Coles, S. G. (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.

Epstein, E. S. (1969). A scoring system for probability forecasts of ranked categories. *Journal of Applied Meteorology*, 8(6):985–987.

Groll, A., Kneib, T., Mayr, A., and Schauberger, G. (2017). On the dependency of soccer scores — a sparse bivariate Poisson model for the UEFA European Football Championship 2016. *Proceedings of the MathSport International 2017 Conference*, pages 161–175.

Karlis, D. and Ntzoufras, I. (2003). Analysis of sports data by using bivariate poisson models. *Journal of the Royal Statistical Society: Series D (The Statistician)*, 52(3):381–393.

Koopman, S. J. and Lit, R. (2015). A dynamic bivariate Poisson model for analysing and forecasting match results in the English Premier League. *Journal of the Royal Statistical Society: Series A (Statistics in Society)*, 178(1):167–186.

Maher, M. (1982). Modelling association football scores. *Statistica Neerlandica*, 36:109–118.

McHale, I. and Scarf, P. (2011). Modelling the dependence of goals scored by opposing teams in international soccer matches. *Statistical Modelling*, 11(3):219–236.

R Development Core Team (2018). *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria.

The Associated Press (2015). Romania, Wales set to be surprise World Cup top seeds. *USA Today*.

Tweedale, A. (2015). Belgium rise to No1 in Fifa world rankings after they beat Israel - despite playing one tournament in 13 years. *The Telegraph*.

UEFA.com (2016). Regulations of the UEFA European Football Championship 2014-2016. Available at http://www.uefa.com/MultimediaFiles/Download/Regulations/uefaorg/Regulations/02/03/92/81/2039281_DOWNLOAD.pdf.

Wang, C. and Vandebroek, M. (2013). A model based ranking system for soccer teams. *Available at SSRN 2273471*.

Wickham, H. (2015). *rvest: Easily Harvest (Scrape) Web Pages*. R package version 0.3.1.