

Bayesian forecasting of UEFA Champions League under alternative seeding regimes

Questa è la versione Post print del seguente articolo:

Original

Bayesian forecasting of UEFA Champions League under alternative seeding regimes / TENA HERRILLO, J; Corona, F; Forrest, D; Wiper, M.. - In: INTERNATIONAL JOURNAL OF FORECASTING. - ISSN 0169-2070. - 35:2(2019), pp. 722-732. [10.1016/j.ijforecast.2018.07.009]

Availability:

This version is available at: 11388/215448 since: 2021-03-22T20:56:27Z

Publisher:

Published

DOI:10.1016/j.ijforecast.2018.07.009

Terms of use:

Chiunque può accedere liberamente al full text dei lavori resi disponibili come "Open Access".

Publisher copyright

note finali coverpage

(Article begins on next page)



Bayesian forecasting of UEFA Champions League under alternative seeding regimes

Journal:	<i>International Journal of Forecasting</i>
Manuscript ID	INTFOR_1712582.R1
Manuscript Type:	Original Article
Keyword:	Bayesian methods, Forecasting competitions, Simulation, Sports forecasting, Uncertainty
Abstract:	Evaluation of seeding rules requires use of probabilistic forecasting models both for individual matches and the tournament. Prior papers have employed a match-level forecasting model and then applied Monte Carlo simulation of the tournament to estimate outcome probabilities, allowing, for example, an outcome uncertainty measure to be attached to each proposed seeding regime. But this approach takes no note of uncertainty surrounding parameter estimates in the underlying match-level forecasting model. We propose a Bayesian approach to address this problem. We illustrate its potential by simulating the UEFA Champions League under alternative seeding regimes. We find that changes in 2015 tended to increase uncertainty over progression from the Group Stage of the competition but made limited difference over which clubs would compete in the Final.

SCHOLARONE™
Manuscripts

1. Introduction

The most obvious purpose served by probabilistic sports forecasting models is to enhance the experience of fans whose enjoyment is derived not only from the event itself but also from anticipating it and talking about it in the run-up to the game. Again forecasts are of direct utility to those who trade on associated betting markets and perhaps also to sports managers who seek a basis for benchmarking the realised performance of their team.

Sports forecasting also has the potential to help researchers address wider issues in sport. For example, researchers modelling (stadium or television) audience size may require a variable to capture match significance, a concept operationalised as the probability that the result of the current match will affect the outcome of the tournament. However, this depends on the outcome of matches yet to be played. [1] are amongst those who have measured the importance of the current match by incorporating a match result probabilistic forecasting model, which is applied to those matches yet to be played, in their simulation methodology.

In general it can also be illuminating to employ a model and use it to generate both a set of forecasts given that the state of the world is as it is and also a set of forecasts given some counter-factual. The present paper adopts this approach. We produce forecasts for various outcomes in the UEFA Champions League (e.g. the probability that club i will reach the Round of 16 and the probability that club i will reach the Final) under both the actual and alternative seeding schemes. By comparing the forecasts for the actual and alternate states of the world, we are able to answer questions such as which types of club gained and which lost from the changes made to the seeding system. Clearly, the methodology could be extended to cover multiple alternative seeding systems, which would be of benefit of organisers of this and other sports tournaments

Embedded within a probabilistic forecasting model for a tournament, based on simulation, is a match-level forecasting model. An innovation in our paper is to use a Bayesian approach to match level forecasting and this enables us to calculate predictive probability distributions for the tournament outcome probabilities under both the current and alternative seeding schemes. In cases where the only purpose of the exercise is to generate forecasts of tournament outcomes, the approach could equally be employed to provide credible intervals as well as point estimates of the probability of a particular team winning, for example.

Our methodology is related to a strand of literature [for example 2, 3] in which simulation models are constructed to identify which teams have the greatest probability of winning a championship. [4] extended this approach (in the context of the Champions League) by using simulation to generate metrics for outcome uncertainty and other tournament characteristics under alternative, hypothetical seeding arrangements. [5] applied similar methodology in the context of the 2010 FIFA World Cup to derive probabilities of success for the competing teams not only for the tournament design which was actually in place but also for hypothetical contests where the same teams took part but with alternative seeding arrangements. Since the authors of these two papers used their results to compare outcome uncertainty under alternative seeding regimes, their papers were a natural starting point for developing our approach.

1
2
3
4 The possibility of employing simulation to examine and estimate the change in seeding arrangements in
5 the Champions League also attracted [6]. Their foci of interest were different from ours, concentrating on the
6 competitive balance of the Final of the competition and providing extended treatment of the probabilities
7 of different countries rather than individual clubs winning the championship. By contrast, our contribution
8 looks at probabilities of clubs progressing to each phase of the competition.
9

10
11 The methodology followed by [6] also differed from ours in that they, like the earlier papers cited above,
12 estimated a probabilistic forecasting model for individual football matches based on some measure or measures
13 of the strength of each team. Each of these papers then applied Monte Carlo simulation to the whole
14 tournament, where probabilities of outcomes of each game were represented by point estimates derived from
15 the individual match forecasting model.
16

17
18 Using this classical, plug-in approach, point estimates for the individual match probabilities can be used
19 to simulate the outcome of the tournament. However, these match-level probabilities themselves are subject
20 to uncertainty which is not accounted for by such approaches. Ignoring this uncertainty will not affect the
21 efficacy of the approach if the purpose is simply to generate a point estimate of each team's chance of winning
22 the competition. But our purpose is to compare probabilities of success under alternative seeding regimes.
23 For example, it might be that, following the standard approach adopted by previous authors, it was estimated
24 that, for a particular club, the probability of its winning the tournament would be 0.10 under one seeding
25 regime and 0.15 under the other. However, how can we tell if these probabilities are different from each other?
26 Previous papers have not addressed this issue. But, with our Bayesian approach, it is possible to obtain the
27 posterior distribution of the model parameters and then estimate, for example, the predictive distribution of
28 the probability that a given team wins the tournament. If the distribution of a team winning the tournament
29 under one seeding system is very different from the distribution under another, it is then almost impossible
30 that the team has the same chance of winning under both systems, whereas if the distributions overlap
31 significantly, it could be that the "true" underlying probability of winning the tournament is essentially the
32 same under both seeding systems.
33
34
35
36
37
38
39
40

41 Could we achieve the same objective by adapting the frequentist models of the earlier literature? It
42 would be possible in principle but might prove a very complex exercise. One would use some sort of bootstrap
43 procedure to try to re-sample the parameter distribution to give more reasonable predictions as an alternative
44 to the Bayesian approach used here. This would necessitate deciding on issues such as, for example, how
45 to define a multivariate distribution of t-statistics and replicate values from it and the number of bootstrap
46 replications to achieve sound critical values.
47
48
49

50 The paper proceeds as follows. In section 2, we describe the structure of the UEFA Champions League
51 and the changes made to the seeding rules in 2015-6. Section 3 describes the entropy measure we will employ
52 to as one of the ways of helping us to assess the degree of outcome uncertainty over which clubs will reach
53 various stages of the competition (including winning it). In section 4, we set out our Bayesian Poisson model
54 for probabilistic forecasting of individual match results and in section 5 we describe the execution and results
55
56
57
58
59
60

of the Monte Carlo simulation. Section 6 consists of conclusions and reflections.

2. Seeding rules in the UEFA Champions League

The structure of the competition was not changed when the new seeding rules we were to assess were introduced. It remained the case that 32 clubs participated in each annual edition of the tournament. One of these places was reserved for the current holders of the Champions League trophy. Twenty-one places were for clubs from relatively strong footballing countries which qualified automatically through their finishing positions in their domestic leagues. The final ten places were for clubs which were successful in a qualifying tournament held in the summer (where participation in the qualifiers was again conditional on finishing positions in domestic leagues).

Each year, by the end of August, the names of the 32 clubs taking part are known. At this stage, the competitive balance of the tournament (the variation in strength across the competing clubs) is fixed. But the outcome uncertainty of the tournament (the variation between clubs in the probabilities of winning) will be influenced by the seeding arrangements applied.

According to the taxonomy of tournament designs set out in [5], the Champions League is a “hybrid 1G-KO” competition. This means that, just as in the World Cup, there is a single “group stage” (played September to December) followed by a series of knock-out or elimination rounds (February to May) culminating in a grand final (June).

At the start of the competition, the 32 clubs are split into eight groups (A to H). Each group is a mini double round-robin league such that each club plays each other club both home and away. A win is rewarded with three points and a draw with one point. The top two clubs in each group proceed to the round of 16. From then on, the tournament is a straight knock-out competition with a fresh draw conducted for each round. However, a constraint imposed on the draw for the round of 16 is that top place finishers from the group stage play second-placed finishers from the group stage.

The point at which seeding is applied is in the draw which allocates clubs between the eight initial groups. For the draw, clubs are allocated to either pot 1, pot 2, pot 3 or pot 4 according to their seeding. The idea is that each group will then include one team from each pot. A slight complication to be noted is that each group is also constrained to consist of clubs from four different countries.

Seeding is based primarily on the “UEFA coefficients” of each club prior to the tournament. These are ranking points earned by wins and progression in the UEFA Champions League and in the UEFA Europa League in the preceding five seasons by the club itself and (with lesser weight) by clubs from the national association to which it belongs.

Prior to season 2015-6, allocation of clubs was according to UEFA coefficients. Thus pot 1 consisted of the eight strongest clubs according to UEFA rankings and none of these could play each other in the group stage. This is classic seeding which is intended to minimise the probability of an early exit by one of the “best” competitors. Similarly clubs were distributed across the other pots in order according to their ranking

1
2
3
4 by UEFA coefficient.

5 From season 2015-6, different seeding arrangements were put into place. Pot 1 would now be populated by
6 the reigning champion of the Champions League and the champion clubs from the seven strongest national
7 leagues according to UEFA national coefficients. For that particular edition of the Champions League,
8 Barcelona was entitled to a place in pot 1 both from being Champions League and Spanish League champion
9 in the previous season, so an extra place was available for the champion club, PSV Eindhoven, from the
10 eighth ranked national league (Netherlands). All the remaining clubs were then distributed across pots 2, 3
11 and 4 according to rank order of coefficient rankings.

12 This apparently minor change in the seeding rules for allocation to pot 1 had the potential to affect the
13 degree of outcome uncertainty surrounding the tournament since it does shift clubs between groups. Table 1
14 shows the changes in pot allocations in the first season of the new regime compared with how the allocation
15 would have looked under the old regime.

16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60

Table 1 about here

It is hard from intuition to be confident about how these changes will change outcome uncertainty. For example, according to the UEFA coefficients, Real Madrid was the strongest club of all but it was in pot 2 because it had failed to win the Spanish League. This might reduce the probability of progression for clubs like Chelsea (which would have been in pot 1 under either system) since Chelsea is not now protected from the possibility of playing the mighty Real Madrid in the group stage. On the other hand, Chelsea might be confident of finishing in the top two even in a group with Real Madrid and then it would be protected from playing Real Madrid in the round of 16. On the contrary, in the case of a weaker UEFA rated club like PSV Eindhoven, we would expect the change to benefit them as their chances of playing stronger clubs in the first stage of the competition are clearly reduced by their new status in pot 1. Clearly the effects of the seeding change are hard to work out and analysis therefore requires simulation of the tournament.

3. Entropy measure

We will measure outcome uncertainty from the perspective of the point in time just prior to the draw to allocate the 32 competing clubs between the eight groups which comprise the first stage of the tournament. It is at this stage that the seeding is implemented which influences the flow of events throughout the rest of the competition.

Entropy is a measure of unpredictability of information content. It has been employed to capture uncertainty in sports since [7] used it to investigate the apparent “improvement” in uncertainty over time in Major League Baseball. Here, in the context of the Champions League, we are interested in measuring uncertainty over several outcomes, for example which club will win the tournament and which clubs will emerge from the group stage to take places in the Round of 16.

To illustrate for the case of uncertainty over which club will win the tournament, let $p_{j0} = P(V_j|\xi_0)$ be the probability of victory for club j conditional on pre-tournament information concerning the 32 clubs in the competition.

Entropy is defined as follows:

$$e_0 = - \sum_{j=1}^{32} p_{j0} \log_2 p_{j0}. \quad (1)$$

Minimum entropy or maximum information occurs when some $p_{j0} = 1$ while the others are 0. In this case, $e_0 = 0$. At the other extreme, maximum entropy is when all probabilities p_{j0} are equal to $1/32$ so that there is maximum outcome uncertainty.

Note again that entropy is calculated before the start of the tournament. [8] and [9] estimate the entropy in different phases of sports tournaments. For each, their focus is on identifying the decisive matches in a tournament. But our interest in the entropy measure is not to identify decisive matches in a given tournament. Instead it is to compare tournament outcome uncertainty under alternative seeding policies.

4. Bayesian Poisson Model

The first requirement for simulating a tournament is a viable model which can be used for probabilistic forecasting of individual matches which might take place within the tournament. Two categories of model have been popular in the academic literature. Economists [such as 10] have tended to favour ordered probit models which directly generate probability estimates for match outcome (home win, draw, away win). Statisticians [since 11] have typically focused on variants of Poisson regression which yield for each team, the estimated probabilities of its scoring 0 goals, 1 goal, 2 goals etc. Combining these probabilities enables point estimation of winning probabilities for each side.

[12] found that the performance of each class of model was similar. Here, we adopt a Poisson regression approach because, sometimes, progression in the Champions League depends on goals, not just on match outcomes. Thus, if clubs are tied in points in a group at the end of the group stage, the tie break rules include comparison of goal difference. The model underlying the simulation must therefore generate probabilities of each possible scoreline in a match as well as win-draw-lose probabilities.

In the football forecasting literature, there have been a number of variations on the basic Poisson model proposed, for example, [13]. However, following [5], we apply the simpler Maher model given that our focus is on tournament simulation rather than on the match-level forecasting models themselves.

We apply a Poisson regression model (BP) for the goals scored by each team in a match as follows:

$$Y_{t,k} \sim \text{Poisson}(\lambda_{T,k}),$$

where $Y_{t,k}$ represents the number of goals scored by team T in a match at time k and

$$\log(\lambda_{T,k}) = \beta_{A_T}x_{A_T,k} + \beta_{A_O}x_{A_O,k} + \beta_{H_T}x_{H_T,k} + \beta_{Aw_T}x_{Aw_T,k} + \beta_{F_T}x_{F_T,k}, \quad (2)$$

where $x_{A_T,k}$ represents the strength of team T and $x_{A_O,k}$ is the strength of the opposing club. $x_{H_T,k}$ indicates if team T plays at home, $x_{Aw_T,k}$ if it plays away and $x_{F_T,k}$ if the match is the Final of the whole competition (the only match played at a neutral ground). The parameters β_{A_T} , β_{A_O} , β_{H_T} , β_{Aw_T} and β_{F_T} are coefficients that express the relationship between the explanatory variables and $\lambda_{T,k}$.

Here, we proxy the strength of a club by its tally of UEFA points (i.e its UEFA coefficient) prior to the start of each year's competition, as recorded on the UEFA website. Such points are used to rank clubs across Europe and are earned on the basis of previous performances in UEFA competitions by the subject club and by other clubs in the national league to which it belongs. Our use of official rankings mirrors the approach of [14] who used a Poisson model with FIFA ranking points in their simulation of the 1998 World Cup. In some applications of the Poisson model, in particular modelling of national leagues, goals scored and conceded in past matches are used as measures of strength. This produces satisfactory forecasts [11].

It may be argued that recent goals scored and conceded reflect form so that the team strength coefficients should vary over time. In the present case, the Champions League is played over a time span of eight months and simulations are from the perspective of before the start of the competition. A team's form in August, when the regular league season is just starting or about to start may not be relevant to predicting the results of possible Champions League matches in April. Thus a longer-run measure of strength, such as is provided by UEFA coefficients seems most appropriate. Further, because of heterogeneous quality across national leagues, it is difficult to gauge relative strengths of clubs on the basis of comparing their propensities to score and concede goals in domestic competition. This again points to use of information on past European performances, which is captured by UEFA coefficients.

It would be possible to fit the Poisson regression model using a classical maximum likelihood as in [14] but here we prefer to use Bayesian inference. In our context, we are interested in predicting the results of future matches and an advantage of Bayesian methods as opposed to classical, plug-in based prediction is that they directly take parameter uncertainty into account.

In order to complete the Bayesian formulation, we need to define a prior distribution for the regression coefficients $\beta = (\beta_0, \beta_{A_T}, \beta_{O_T}, \beta_{H_T}, \beta_{Aw_T}, \beta_{F_T})'$ in (2). One option in the presence of good prior information would be to consider a multivariate normal prior structure for β , but here we prefer to use an improper, uniform prior distribution, $f(\beta) \propto 1$ as recommended in [15].

Exact calculation of the posterior distribution under this model is impossible, but instead we can generate an (approximate) Monte Carlo sample of values from the posterior distribution using Markov chain Monte Carlo (MCMC) methods, see e.g. [16] for a good review. In particular, we use a random walk Metropolis algorithm to generate the successive elements of β from their conditional posterior distributions as in e.g.

[17].

After calculating the Monte Carlo samples, to obtain the win, draw and loss probabilities, $p_{W,k}$, $p_{D,k}$ and $p_{L,k}$ respectively, we follow the procedure used by [14], [3], [9], among many others. In this way, the probabilities for each game during the competition for teams T against the teams O become as follows:

$$p_{W,k} = \sum_{i_T=1}^{\infty} \sum_{i_O=1}^{i_T-1} P(y_{T,k} = i_T)P(y_{O,k} = i_O), \quad (3a)$$

$$p_{D,k} = \sum_{i_T=1}^{\infty} P(y_{T,k} = i_T)P(y_{O,k} = i_T), \quad (3b)$$

$$p_{L,k} = \sum_{i_O=1}^{\infty} \sum_{i_T=1}^{i_O-1} P(y_{T,k} = i_T)P(y_{O,k} = i_O), \quad (3c)$$

where i_T and i_O are the number of scored goals for teams T and O respectively, where $P(y_{T,k} = i_T)$ and $P(y_{O,k} = i_O)$ represent the corresponding Poisson probabilities of goals scored (with $\lambda_{T,k}$ and $\lambda_{O,k}$ as means) by the teams in the game. Similar to previous authors including [5] in their tournament simulation, we assume independence between the goals scored by the two teams.

To estimate the model, we employed a data set containing the results of every Champions League match played between seasons 2002-3 and 2014-5. This is the period since the current 1G-KO structure of the tournament was put into place.

Parameter values β , in (2) were sampled from the posterior distribution using 10,000 MCMC iterations with 5,000 iterations to burn in the chain and thinning to reduce autocorrelation. Figure 1 shows the traces of 1000 thinned parameter values and the associated estimated posterior parameter densities.

Figure 1 about here

It can be observed that all traces have a white noise behaviour which suggests that convergence has been achieved. Also, observing the fitted parameter densities, it can be seen that all densities, except for the away effect are concentrated away from 0 which suggests that these effects are important for determining match results.

Table 2 summarises the central values of these posterior distributions. To illustrate the interpretation of these values, consider a hypothetical match in which Real Madrid played at home versus Barcelona in the 2015-6 Champions League. At the start of the competition, the UEFA coefficients of the two clubs were 172 and 165 respectively but of course Real Madrid would have enjoyed home advantage in this imaginary fixture.

Table 2 about here

Using median values of the posterior distribution of the estimated parameters, the expected number of

goals for each team is then:

$$\lambda_{\text{Real Madrid}} = \exp(0.0057(172) - 0.0050(165) + 0.3520) = 1.66,$$

$$\lambda_{\text{Barcelona}} = \exp(0.0057(165) - 0.0050(172) + 0.0014) = 1.09.$$

To the nearest integer, the single most likely result was 2-1 to Real Madrid.

However, of course, modelling yields estimates for each team of the probability of each possible number of goals it might score. These two sets of probabilities are then combined to give probabilities of each possible match result (e.g. 0-0, 1-0, 0-2, etc). It is these probabilities that are employed in the simulation exercise.

Note that our approach is based on forecasting individual matches even though some rounds of the competition involve two-legged contests where the winner is determined by aggregate score. We model the two matches as independent of each other and this means that we abstract from the possibility that the result of the second match is influenced by the first. This implies we also abstract from any advantage there may be in playing at home in the second game. We chose not to incorporate the home-second advantage because there is evidence [18] that this has been diminishing (if not quite disappeared) and to use historic data would bring too great a correction to the forecasts.

The model yielded predictions (in the sense of 'most likely outcome') of a home win in 42.2% of the matches, of a draw in 25.0% of the matches and of an away win in 32.8% of matches. This reflected adequately well the actually observed proportions of match outcomes, 44.4%, 23.7% and 31.9% respectively.

When tested, our match-level model exhibited satisfactory performance in terms of predictive power. We took the (chronologically) first 70% of matches as the training set and estimated the parameters of the model from data on those 962 games. We then predicted match results in the test sample (the remaining 413 matches) using the median of the sample posterior distribution. In terms of home win/ draw/ away win outcome of each match, the most likely outcome from the model coincided with the actual outcome in 50.7% of matches.

In addition, we used the model estimated over the sample period 2003-15 to predict the home/ draw/ loss outcome of all matches which took place in the Champions League, 2015-16. In this out-of-sample testing, we again took the most likely outcome indicated by the model as our prediction and this time the success-rate was 51.2%. This is close to the success-rate (51.6%) reported by [8] for kernel regression modelling applied to one edition of the European Championships (a 1G-K0 football tournament for national teams). [8] notes that his was a superior success rate to that reported in some earlier literature (though it is hard to compare performance across competitions which differ in both format and heterogeneity of team strength).

Finally, we compared our model with bookmaker-odds (from Bet365) in terms of forecasting performance as captured by Brier Score. This is a challenging test because our model employs only information available before the start of the tournament whereas bookmakers-odds are as quoted immediately prior to the match. The mean values of the Brier Score were .301 for the model and .274 for Bet365 implicit probabilities. This

1
2
3
4 indicates greater forecasting ability from using bookmakers-odds but the empirical distribution of the two
5 Brier Scores are very similar and in fact cross three times. This similarity is encouraging given that betting
6 odds capture recent information which cannot be supplied to the model because the model forecasts are
7 needed prior to the draw for the first phase of the tournament
8
9

10 11 **5. Monte Carlo Simulation**

12
13 We carried out three separate simulations of the 2015-6 Champions League. Two of these related to
14 the old and new seeding regimes described in section 2 above. The third simulation assumed a completely
15 random draw to determine the allocation of the 32 clubs between the 8 initial groups. We term these three
16 possible seeding systems as the traditional (denoted by T), the new (denoted by N) and the random (denoted
17 by R).
18
19

20
21 [2] notes that it is important that procedures in a simulation exercise mimic the tournament in terms of
22 reflecting the competition rules. In Koning's case, this was a little more straightforward than in our case.
23 He was studying the FIFA World Cup and forecasting from a point subsequent to the allocation of national
24 teams between the groups in the initial group stage. By contrast, our analysis is from the perspective of the
25 time prior to the draw for the group stage. Further, there is no fresh draw in the FIFA World Cup after the
26 group stage. All pairings in the first knock-out round are pre-determined, e.g. winners of Group A versus
27 runners-up of Group B. By contrast there is a fresh draw for the Round of 16 in the Champions League
28 which randomly assigns clubs from the set of group winners with clubs from the set of runners-up (with the
29 additional constraint that clubs from the same country cannot play each other at this stage). Thus, in our
30 case, simulation includes simulation of both the group draw and the draw for the Round of 16 (and indeed
31 the draws for the quarter- and semi-finals) as well as simulation of the evolution of winners and losers as
32 the tournament progresses. Similarly, to mimic the tournament closely, we apply the same tie-break rules
33 as those set down by UEFA for determining which club proceeds to the next stage if two clubs are tied on
34 points in a group. In the subsequent knock-out rounds, ties up to and including the semi-final are two-legged.
35 In this case, where two teams are drawing after two legs, 30 minutes extra time is played and if the teams
36 are still tied after extra time, a penalty shoot-out is used to decide the result. In this case, we determine
37 the winner of the tie by a random process because of the small number of observations of extra time plus
38 penalties in the data set.
39
40
41
42
43
44
45
46
47

48 In the context of simulation of tournaments, an innovation in how we proceed is the use of MCMC
49 estimation (described in the previous section) instead of classical estimation. This means that we do not
50 just obtain an estimation of the model parameters and their standard deviations but an estimation of their
51 posterior distribution. This gives us a much fuller picture of the different values each parameter can take.
52 Knowledge of this uncertainty can then be exploited to evaluate differences in results from the three simulation
53 exercises.
54
55

56 In order to estimate the probability of a team winning the tournament, or surviving to the next round, we
57
58
59

1
2
3
4 use a Monte Carlo simulation procedure within the previous MCMC algorithm which simulates the outcome
5 of each match in the tournament. Thus, at each MCMC iteration, we simulate 1000 replications of the
6 tournament and use these to estimate, for example, the win probabilities for each team. Over the whole
7 MCMC run, the distribution of win probabilities for a given team is an approximation to the posterior
8 predictive distribution of the probability of a team winning the tournament.
9

10
11 Regarding the number of replications, we judged 1,000 adequate because the estimated tournament en-
12 tropy score stabilised after around 400 or fewer replications using the mean values of the estimated parameters.
13

14 Now we wish to examine tournament uncertainty under the different seeding systems. Given that under
15 the random seeding and new seeding systems, it is more likely for currently stronger UEFA rated teams to
16 meet each other in the group stage, then we might expect that under these systems, there is more possibility
17 of a big team being knocked out at this stage, and not qualifying for the round of 16, than under the
18 traditional system. Figure 2 shows kernel density estimates of the posterior predictive distributions of the
19 probability that each team qualifies under each the three seeding systems: the traditional (T), the new (N)
20 and a hypothetical random draw (R).
21
22
23

24 The results bear out this thesis. For example, in the case of the top eight UEFA rated teams (shown
25 in the first row of Figure 2), it can be seen that the posterior predictive distribution of the probability of
26 qualification under random seeding is concentrated to the left and that the distribution under the new system
27 is also centred at a slightly lower value than under the traditional system.
28
29
30

31 Figure 2 about here
32

33 For the lowest ranked teams, shown on the bottom row of Figure 2, both the traditional and new seeding
34 systems suggest a similar but much lower probability of qualification than under random seedings. Obviously,
35 these teams have little probability of qualification under usual seeding systems, but under random seeding,
36 it is possible that they could be drawn in a group containing only weaker teams and therefore have more
37 chance of qualification.
38
39
40

41 For more middle ranking teams, the results are more variable. In general, random seedings lead to very
42 different distributions to either of the other seeding systems. Two teams in particular: PSV Eindhoven and
43 Shakhtar Donetsk, show very noticeable differences in their predicted performance under the traditional and
44 new seedings as there is almost no overlap in their posterior predictive distributions under the traditional
45 and new seeding systems.
46
47
48

49 Comfortably the club most favoured by the change in seeding regime was PSV Eindhoven. Under the
50 new system, they qualify for pot 1 as champion of the UEFA country with the eighth highest ranking (in
51 reality ninth as Barcelona had won both the Champions League and the Spanish League and thus qualified
52 on two criteria, leaving PSV to fill the extra space). As we saw in Table 1, under the traditional system,
53 PSV would have been only in pot 3, leaving them with the possibility of facing much stronger teams in the
54 group stage and a much harder path to qualification.
55
56
57
58
59
60

1
2
3
4 It seems likely in fact that there will more often than not be a club, like PSV Eindhoven, propelled from
5 pot 3 to pot 1 by the rule changes. Over the ten years from 2007-8 to 2016-7, there were seven instances of
6 the club which won the Champions League also winning its domestic competition. On each such occasion,
7 a club from what was only the eighth-ranked league in Europe secured a place in pot 1 as a result. In the
8 2016-2017 season, an unlikely winner of the English League, Leicester City, benefited in a similar way: its lack
9 of European experience would have placed it in pot 3 under the old regime but it took its place in pot 1 (and
10 duly progressed to the quarter-finals). Thus any 'favourable' effect on outcome uncertainty within the 2015-6
11 tournament attributed to the boosting of the prospects of PSV Eindhoven might reasonably be expected to
12 be replicated in most years. After PSV Eindhoven, the next most favoured club was Juventus which, as
13 Italian champion, qualified for pot 1 whereas previously, on the basis of its UEFA coefficient, it would have
14 been in pot 2. To a lesser extent, both Paris St.-Germain and Zenit St. Petersburg were beneficiaries of the
15 changes and they too were able to claim places in pot 1 as national champions whereas, under the traditional
16 system, each would have been in pot 2.

17
18 In the case of Shakhtar Donetsk, under the traditional system they would have been placed in pot
19 2 whereas under the new system, they are lowered to pot 3. This change is reflected in the predictive
20 distribution of their probability of surviving the group stages which has been moved significantly to the left
21 under the new system, reflecting the fact that as a lower rated team they would typically have to play against
22 stronger opposing teams under the new system.

23
24 In order to get a Bayesian point estimate of the probability of qualifying for the next round, we can simply
25 calculate the mean of the posterior predictive distribution. The means under all three seeding systems are
26 given in Table 3.

27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60
Table 3 about here

38 The results confirm the previous conclusions. For the highest eight UEFA ranked teams, the mean
39 predictive probability of qualification is lowest under random seeding and then slightly higher under the
40 traditional scheme than under the new seeding scheme. For the lowest ranked teams, the mean probability
41 of qualification is higher under random seeding, with very little difference between the new and traditional
42 schemes. In the case of PSV Eindhoven, the mean probability of qualification increases substantially from
43 0.299 under the traditional scheme to 0.453 under the new scheme and for Shakhtar Donetsk the probability
44 is reduced from 0.633 to 0.483. Other results are also in line with what would be expected from Table 1.
45 Paris Saint German, Juventus and Zenit Saint Petersburg, the teams that are upgraded from pot 2 to pot 1
46 all show slight increases in the mean qualification probability under the new seeding scheme.

47
48 Similar results applied to the first knock out stage leading to qualification for the quarter finals. Again we
49 found that teams showing the most notable differences under the traditional and new systems were Shakhtar
50 Donetsk and PSV Eindhoven. High ranked clubs like Porto and Arsenal were likely to have lower chances of
51 reaching the quarter finals under the new system whereas in the case of Juventus, Zenit and Paris St.-Germain,
52

1
2
3
4 the advantage of the new seeding system was still reflected in the movement of the predictive density to the
5 right, although to a lesser extent than in the case of the group stage. The results were qualitatively similar
6 when we considered qualification for the semi finals but the effects were much reduced at this stage. Broadly,
7 the results of these exercises indicated a worthwhile impact on competitive balance in the early phases of the
8 Champion League but limited impact on which clubs would progress to the most lucrative rounds
9

10
11 With a more formal approach, we can capture the overall effect on uncertainty of outcome (in terms of
12 which club will win the competition). Figure 3 shows the densities of the entropies for each of our three
13 seeding regimes.
14

15
16 **Figure 3 about here**
17

18
19 The density associated with the random draw regime is more inclined to the right than for either the
20 traditional or old seeding systems, implying that failing to seed at all will add uncertainty as to which
21 competitor will win the tournament. The density of the entropy under the new seeding system is also slightly
22 to the right of that under the old system, again suggesting that it is more probable that there is slightly
23 higher uncertainty under the new system than the old.
24
25

26
27 But perhaps entropy does not capture the full story about competitive balance. The audience for the
28 competition may be interested mainly in the very strongest clubs which alone have a realistic chance of
29 winning the tournament. It has been argued that the first (group) stage lacks interest in that there is too
30 little risk of the most famous clubs failing to advance. We looked at the mean number of times in 1,000
31 repetitions that the eight clubs with the most UEFA points all advance to the knock-out stage. This was 43.3
32 times under the old system, 34.1 times under the new system and 14.6 times under the random draw system.
33
34 Thus, there are likely to be casualties under any system but the strong were expected to be less successful
35 as a group under the new seeding regime (to that extent the new exhibits superior competitive balance) and
36 would be much less successful under random draws.
37
38

39
40 Naturally the final at the end of the competition attracts considerable interest. We looked at the number
41 of times in 1,000 repetitions that both finalists were drawn from the eight strongest clubs according to pre-
42 tournament UEFA points. The results were 698 under the traditional seeding system, 670 under the new
43 seeding system and 643 under a random draw system. From the perspective of reaching the final, the change
44 in regime again moved the competition towards greater competitive balance though the effect seems rather
45 marginal and indeed even with a random draw it would happen more often than not that the eight strongest
46 clubs would produce both finalists.
47
48
49

50 51 **6. Concluding remarks** 52

53
54 The motives for seeding sports competitions may vary from case to case. Some organisers (including,
55 we suspect, UEFA) may seek to maximise commercial revenue. Others may desire to provide the highest
56
57
58
59

1
2
3
4 possible degree of outcome uncertainty or emphasise the somewhat intangible concept of “fairness” or even
5 offer enhanced protection to the politically most powerful competitors.

6
7 But, regardless of the goals of tournament organisers, the contribution of sports analytics is to provide
8 simulation methodologies to allow tournament organisers to evaluate changes in seeding arrangements (which,
9 realistically, could not even be evaluated ex-post because we would observe too few runs of the competition).
10 Previous approaches to simulation could be employed but would provide just mean probability values of final
11 outcomes under each seeding scheme. They could not show whether the distributions of these probabilities
12 differ substantially from each other. This is relevant for tournament organisers because pushing through
13 changes is likely to be politically costly and they will wish to know whether there is, for example, a clear
14 gain to be made in terms of improving the competitive balance of the competition.

15
16
17
18 In principle, the methods described could be applied across a variety of sports and tournaments. Our case
19 study examined changes in seeding arrangements introduced into football’s European Champions League in
20 2015. Those particular changes were found to have increased uncertainty over which club would win the
21 competition. They also slightly reduced the probability of each of the most powerful clubs surviving to the
22 knock-out stage while increasing the chances of some of the weakest clubs. Probably these impacts were
23 consistent with what UEFA hoped to achieve given adverse comments on how excessively predictable the
24 early stages of the competition had become.

25
26
27
28
29 The effects of the change in seeding have been illustrated here in the context of one particular edition of
30 the Champions League, that of 2015-6. From season to season, the 32 clubs taking part will of course vary.
31 But it is reasonable to suppose that our analysis of which particular clubs gained or lost from the reforms in
32 2015-6 illustrates also which types of clubs would be better or worse off in the general case. Although, for
33 reasons of space and to avoid repetition, we do not report the details, we in fact repeated the exercise set out
34 in this paper for the Champions League, 2014-5. That year the old seeding system was still in place but we
35 asked what differences there would have been had the new system been in effect instead. Broadly the effects
36 were similar to those for 2015-6 not only in terms of the difference made to entropy but also in the sorts of
37 clubs which tended to be winners or losers. Thus, in 2014-5, CSKA Moscow would have had a significantly
38 higher chance of progression had the new system applied: it was a national champion and would have been
39 in pot 1 instead of pot 3 (similar to PSV Eindhoven which was the biggest gainer from regime change in
40 2015-6). Again, Arsenal would have had a markedly lower chance of progression had the new system been in
41 place in 2014-5. Its previous record in UEFA competitions entitled it to a place in pot 1 whereas under the
42 later regime it would have been in pot 2 because it had not won its national championship. In fact, Arsenal
43 proved to be in a similar situation in 2015-6 and, as noted above, was one of the clubs most adversely affected
44 when the new system was actually introduced.

45
46
47
48
49
50
51
52
53 We have focused on the use of forecasting models to illuminate the choice of seeding rules in a particular
54 football tournament. More generally, our approach could be useful just in the context of a forecasting exercise
55 for a tournament. Traditional approaches yield point estimates of probabilities of each competitor winning
56
57
58
59
60

1
2
3
4 the competition but our methodology provides distributions around these point estimates.

5 Richer information will be valuable to organisers because of the commercial implications of setting different
6 seeding rules. For example, high television revenue for the Final of a tennis or football tournament might
7 require a high probability that it will be contested by the 'best' competitors. Other stakeholders may also
8 require such additional information from forecasts of individual matches. For example, traders on associated
9 betting markets would wish to consider not only the point estimates but also seek to evaluate risk attached
10 to possible trading decisions.
11
12
13

14 15 16 **References**

- 17
18 [1] P. A. Scarf, X. Shi, The importance of a match in a tournament, *Computers and Operations Research*
19 35 (2008) 2406–2418.
20
21 [2] R. Koning, M. Koolhaas, G. Renes, G. Ridder, A simulation model for football championships, *European*
22 *Journal of Operational Research* 142 (2) (2003) 268–276.
23
24 [3] A. K. Suzuki, L. E. B. Salazar, J. G. Leite, F. Lozada-Neto, A Bayesian approach for predicting match
25 outcomes: The 2006 (Association) Football World Cup, *Journal of the Operational Research Society* 61
26 (2010) 1530–1539.
27
28 [4] P. A. Scarf, M. M. Yusof, M. Bilbao, A numerical study of designs for sporting contests, *European*
29 *Journal of Operational Research* 198 (1) (2009) 190–198.
30
31 [5] P. A. Scarf, M. M. Yusof, A numerical study of tournament structure and seeding policy fo the soccer
32 World Cup finals., *Statistica Neerlandica* 65 (1) (2011) 43–57.
33
34 [6] D. Dagaev, V. Y. Rudyak, Seeding the UEFA Champions League participants: Evaluation of the reform,
35 NRU Higher School of Economics. Series EC “Economics”. No. WP BRP 129/EC/2016.
36
37 [7] I. Horowitz, The increasing competitive balance in Major League Baseball, *Review of Industrial Orga-*
38 *nization* 12 (1997) 373–387.
39
40 [8] G. Geenens, On the decisiveness of a game in a tournament, *European Journal of Operational Research*
41 232 (2014) 156–168.
42
43 [9] F. Corona, J. Tena, M. Wiper, On the importance of the probabilistic model in identifying the most
44 decisive games in a tournament, *Journal of Quantitative Analysis in Sports* 13 (1) (2017) 11–23.
45
46 [10] R. Koning, Balance in competition in Dutch soccer', *The Statistician* 49 (3) (2000) 419–431.
47
48 [11] M. J. Dixon, S. G. Coles, Modelling association football scores and inefficiencies in the football betting
49 market, *Journal of the Royal Statistical Society C* 46 (2) (1997) 265–280.
50
51
52
53
54
55
56
57
58
59
60

- 1
2
3
4 [12] J. Goddard, Regression models for forecasting goals and match results in association football, International Journal of Forecasting 21 (2) (2005) 331–340.
5
6
7 [13] G. Boshnakov, T. Kharrat, I. McHale, A bivariate weibull count model for forecasting association football
8 scores, International Journal of Forecasting 33 (2) (2017) 458–466.
9
10 [14] D. Dyte, S. R. Clarke, A ratings based poisson model for World Cup soccer simulation, Journal of the
11 Operational Research Society 51 (2000) 993–998.
12
13 [15] A. Martin, K. Quinn, J. Park, MCMCpack: Markov Chain Monte Carlo in R, Journal of Statistical
14 Software 43 (9) (2011) 1–29.
15
16 [16] C. Robert, G. Casella, Monte Carlo Statistical Methods, 2nd Edition, Springer, 2004.
17
18 [17] C. Sherlock, P. Fearnhead, G. Roberts, The random walk metropolis: Linking theory and practice
19 through a case study, Statistical Science 25 (2).
20
21 [18] R. Flores, D. Forrest, C. de Pablo, J. Tena, What is a good result in the first leg of a two-legged football
22 match?, European Journal of Operational Research 247 (2) (2015) 641–647.
23
24
25
26
27
28

29 **Tables and Figures**
30
31

32 Table 1: Summary of pot changes under the new rating system

Pot change	Teams
1 → 2	{Real Madrid, Atletico Madrid, Porto, Arsenal}
2 → 3	{Shakhtar Donetsk}
2 → 1	{Paris Saint German, Juventus, Zenit Saint Petersburg}
3 → 1	{PSV Eindhoven}

33
34
35
36
37
38
39
40
41
42
43
44 Table 2: Summary of BP model

Variable	Median	Mean	SD
AT	0.0057	0.0057	0.0004
AO	-0.0050	-0.0050	0.0004
Home	0.3520	0.3499	0.0568
Away	0.0014	0.0002	0.0566
Final	0.3804	0.3739	0.1861

Table 3: Predictive mean probabilities of qualifying from the group stage to the first knockout round.

Team	Country	UEFA coef	Seeding system		
			R	T	N
Real Madrid	ESP	171.999	0.946	0.983	0.976
Barcelona	ESP	164.999	0.944	0.976	0.970
Bayer Munich	GER	154.883	0.901	0.962	0.949
Chelsea	ENG	142.078	0.861	0.939	0.918
Atletico Madrid	ESP	120.999	0.772	0.873	0.852
Benfica	POR	118.276	0.747	0.862	0.837
Porto	POR	111.276	0.718	0.831	0.790
Arsenal	ENG	110.078	0.717	0.826	0.789
Manchester United	ENG	103.078	0.670	0.734	0.750
Paris Saint-Germain	FRA	100.483	0.657	0.729	0.743
Valencia	ESP	99.999	0.659	0.739	0.748
Juventus	ITA	95.102	0.588	0.679	0.710
Zenit Saint Petersburg	RUS	90.099	0.555	0.654	0.677
Bayer Leverkusen	GER	87.883	0.569	0.647	0.656
Manchester City	ENG	87.078	0.592	0.633	0.650
Shakhtar Donetsk	UKR	86.033	0.561	0.633	0.483
Sevilla	ESP	80.499	0.514	0.450	0.475
Lyon	FRA	72.983	0.475	0.400	0.395
Dynamo Kyiv	UKR	65.033	0.412	0.337	0.342
Olympiacos	GRE	62.380	0.415	0.325	0.326
PSV Eindhoven	NED	58.195	0.379	0.299	0.453
CSKA Moscow	RUS	55.599	0.352	0.281	0.286
Galatasaray	TUR	50.020	0.336	0.251	0.254
Roma	ITA	43.602	0.296	0.217	0.219
BATE Borisov	BLR	35.150	0.218	0.135	0.136
Borussia Monchengladbach	GER	33.883	0.219	0.130	0.131
Wolfsburg	GER	31.883	0.218	0.123	0.123
Dinamo Zagreb	CRO	24.700	0.173	0.098	0.099
Maccabi Tel Aviv	ISR	18.200	0.142	0.079	0.080
Gent	BEL	13.440	0.145	0.067	0.068
Malmo FF	SWE	12.545	0.135	0.065	0.066
Astana	KAZ	3.825	0.114	0.048	0.049

Figure 1: Trace (left panel) and density (right) of each posterior parameter distribution of a Poisson regression model in 1,000 MCMC.


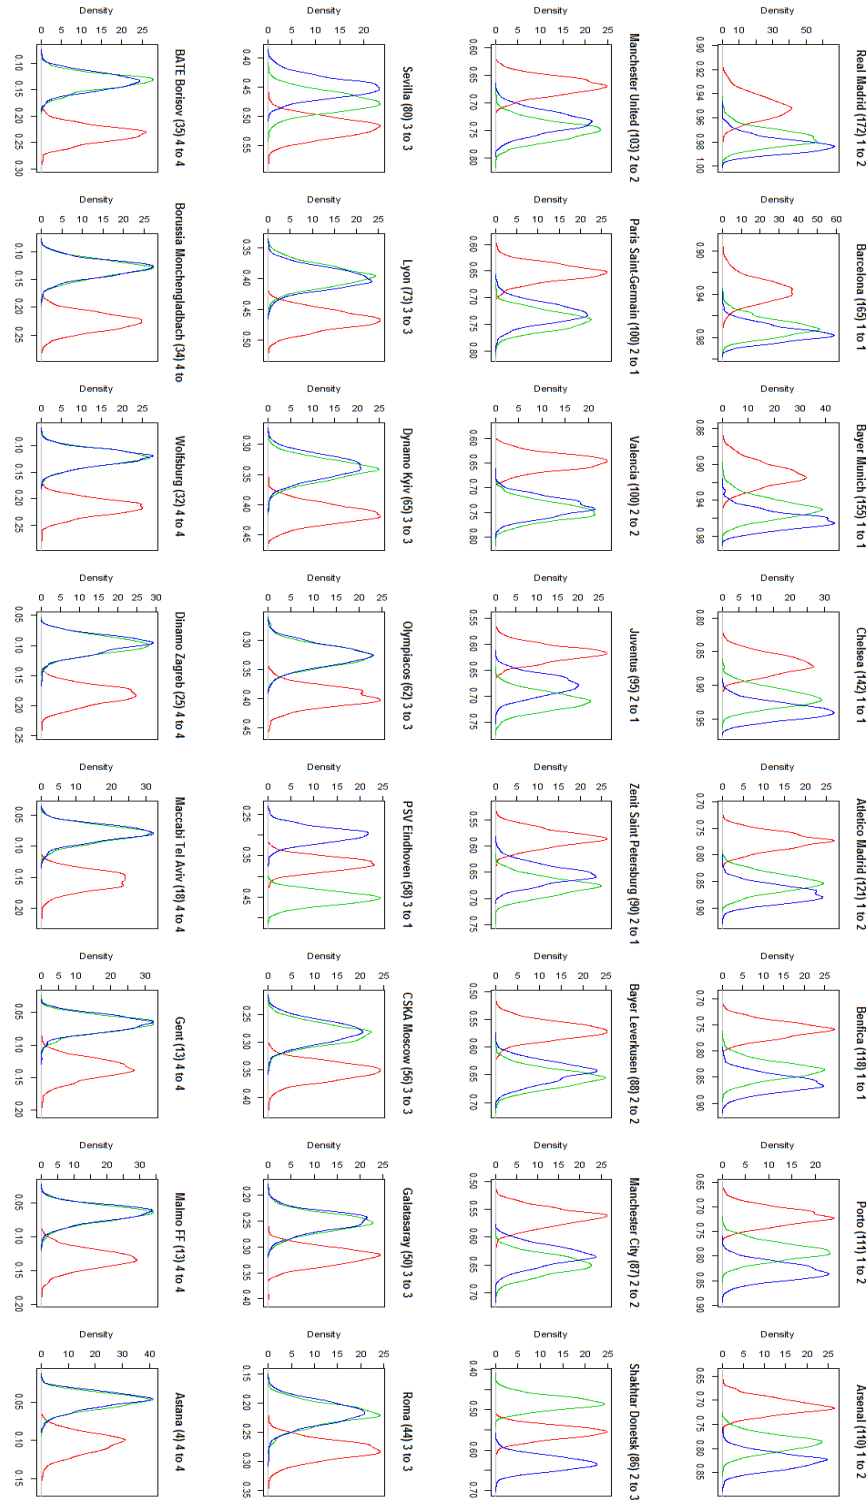



Figure 2: Kernel density estimates of the posterior predictive density of the probability of qualifying from the group stage for each team under the different seeding systems: T (blue), N (green) and R (red).



1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60

Figure 3: Predictive distributions of entropy of the distribution of the probability of the competition winner:
T (blue), N (green) and R (red).



RESPONSE TO COMMENTS BY THE GUEST EDITOR AND TWO REVIEWERS

Guest Editor's comments:

Both referees like the paper, as do I. Having said that, both authors are concerned that the contribution to the literature is "slight". I would like you to respond to that concern, and perhaps "sell" the paper, and your findings, a little better. In addition to responding to the referees' other points, I am particular keen for a few sentences to be added to the paper about the possibility of using a frequentist framework to do the same tasks you do here.

We have added a few sentences at the end of the manuscript to specify which stakeholders might be helped by the methods described in the paper; we note how commercially important it might be to choose the optimal seeding system. We have added extensive new text on a possible frequentist framework as an alternative to our Bayesian approach (see response to the first point of reviewer 1 below).

Reviewer: 1

Comments to the Author

I have some remarks to the authors:

- The main contribution of the paper is to allow for parameter uncertainty in the simulation model. One could do that in a frequentist framework as well, by drawing repeatedly from the multivariate distribution of the parameters, and do the simulation for each draw. This option is not discussed, nor compared to the Bayesian method proposed.

We have added some twenty lines of text towards the end of Section 1. Here we acknowledge that there is a frequentist alternative to our approach but point out that the complexity of the methodology required would be considerable: our Bayesian framework is in fact simpler to apply.

- I am not too convinced about the Maher model, can the authors provide a table with frequencies of results, and corresponding probabilities?

- I am worried about the lack of the time dimension in the knock-out stage. It is well know that it is advantageous to play the second leg at home.


We have added a new paragraph in Section 4 where we are transparent about the fact that we abstract from the issue of two-legged matches, treating each leg as an independent event. We note and give a reference to support that the home-second advantage has been diminishing over time and is now marginal.

- I find the entropy measure as a way to aggregate all match information not convincing. Some teams start with the expectation they will qualify for, say, at least the quarterfinals, and other teams are happy to be in some tournament after the winter break. I suggest to distinguish between likely contenders for the title and non contenders.

We entirely take the point and, in new text at the end of Section 5, we note that “entropy does not capture the full story”. We go on to quote figures focusing on the eight strongest clubs in the tournament (their chances of advancing to the knock-out stage and of supplying both clubs in the Final).

- How do the outcomes of the Maher model and the simulation model compare to betting quotes?

At the end of Section 4 we report Brier scores for match-level forecasts from our model and from implicit probabilities quoted by Bet365. We note that, while the point estimates of Brier score was lower (as was to be expected given that model forecasts are made pre-tournament), the empirical distributions of the two Brier scores are in fact remarkably similar. For information, this figure shows the two distributions (black is for model forecasts and red for Bet365 forecasts):



It is not possible to test the simulation model against betting odds. Typically a bookmaker will not publish odds for a club reaching a specific stage of a competition (such bets are possible but on special request). There are therefore no relevant odds archives on which to draw.

Reviewer: 2

Comments to the Author

This manuscript is mostly sound, but the contribution is not overly significant, and I think the abstract and the conclusion are rather inconclusive as a result. I think the authors should make a more positive and substantial claim about the value of the work, particularly in respect of its value to tournament administrators.

This comment was taken up by the guest editor: please see the response to the guest editor above.

1
2
3 My other points are minor:
4

- 5 1) It's not clear to me how the discussion of surprise and suspense (the JPE paper) is relevant.
6

7 *We have deleted this section of text.*

- 8 2) Page 1, line 51, states that the authors "use a Bayesian approach....to test for statistical
9 significance". This seems like a contradiction in terms.
10

11 *Thank-you for pointing out this careless usage- we have removed this error.*

- 12
13 3) Page 2 line 29 onwards makes the point that parameters are uncertain and the tournament
14 simulation should reflect this. The authors state "previous papers have not addressed this issue".
15 However, if I recall correctly, Scarf et al, in an EJOR paper on tournament design that is not cited, did
16 look at such an issue. The abstract makes the same point. So I think the authors need to check that
17 paper, probably cite it, and clarify exactly what is new and what is not here.
18

19 *We take it that the reference is to Scarf, Yusof & Ibao (2009). We were clearly amiss in not picking up
20 this paper before our first submission. It is now fully acknowledged in the text. However, it does not
21 address our crucial issue, namely uncertainty over parameter estimates in the match forecasting
22 model. It does test robustness of conclusions about rival seeding systems by imagining different sets of
23 covariate values (UEFA points) from those observed. Our equivalent is to report (albeit briefly) on
24 results obtained for a different year of the tournament (with therefore a different array of UEFA
25 points).
26
27*

- 28 4) The description of the Maher model is too long as this is all well known, and not really the
29 contribution of the manuscript.
30

31 *We agree that most readers will be familiar with the Maher model. On the other hand, the description
32 of the model takes only ten lines and even if we removed them we would have to put in a couple of
33 lines of linking text to introduce the model. On balance, we decided to leave the ten lines in the text.*

- 34 5) There are too many figures. I think this doesn't help the value to length ratio of this
35 manuscript.
36
37

38 *We accept that there were too many pictures and they were not strictly necessary to the narrative. We
39 have deleted old Figures 2,4,5 and 6, with therefore considerable saving of space.*

- 40 6) In the text, it is not always clear to the reader which seeding system the authors are referring
41 to at any point in the manuscript. It would be better if these were defined somewhere obvious and
42 given obvious labels.
43
44

45 *At the beginning of Section 6, we explicitly define the three seeding systems considered in the paper.*

- 46 7) Finally, the last line of the abstract raises questions... "The rule changes...increase OU". First
47 question is: by how much, a lot or a little? Second question is: OU of what exactly: the entire
48 tournament or the group stages? I suggest to rewrite to provide answers if possible.
49

50 *In the text itself, we have gone into more detail than before on different dimensions of outcome
51 uncertainty (see responses to Reviewer 1). We have tweaked the Abstract to take into account that we
52 found that the new seeding system had its largest effect in the case of progression to the knock-out
53 stage and lesser effects thereafter. This is also reflected in the (slightly) revised Abstract.*
54
55
56
57
58
59
60