Estimation of the attack and defence teams' abilities

Figure 1a: Estimated posterior 50% confidence bars for the attack (red) and the defense (blue) effects along the nine considered seasons used as training set (from 2007-2008 to 2015-2016) and predicted bars for the teams belonging to the German Bundesliga 2016/2017, the test set season. According to the model specification, these teams' specific parameters are included in a log-linear predictor built for modeling the scoring strength of each team and given an exchangeable prior distribution. Their interpretation is then the following: the greater/lower is the attack/defense team effect and the better is evaluated that specific team. As it results evident, in the summer 2016 Bayern Munich is predicted to lead the Bundesliga of the upcoming season - as it happened - and Borussia Dortumund is predicted to perform quite well. Wider posterior bars are in general associated with teams with fewer observations, i.e. teams that have been relegated along the nine seasons.

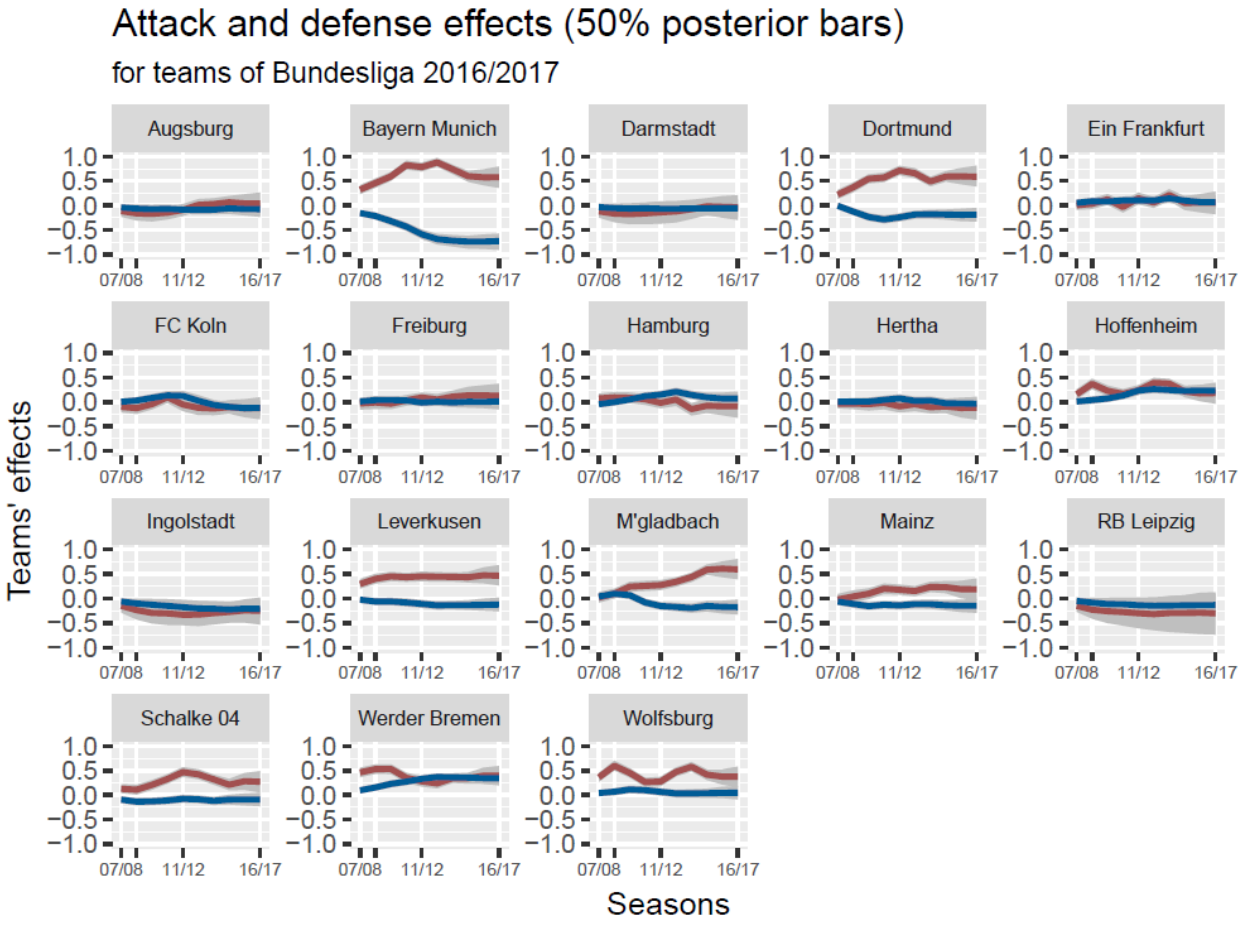

Figure 1b: Estimated posterior 50% confidence bars for the attack (red) and the defense (blue) effects along the nine considered seasons (from 2007-2008 to 2015-2016) and predicted bars for the teams belonging to the English Premier League 2016/2017, the test set season. According to the model specification, these teams' specific parameters are included in a log-linear predictor built for modeling the scoring strength of each team and given an exchangeable prior distribution. Their interpretation is then the following: the greater/lower is the attack/defense team effect and the better is evaluated that specific team. At the end of the 2015-2016 season Manchester City, Arsenal and Chelsea are predicted to perform quite well in the upcoming season: Chelsea actually won the English Premier League 2016-2017. Whereas Middlesbrough and Hull City are predicted to perform badly, as it actually happened with their relegation in the Football League Championship. Wider posterior bars are in general associated with teams with fewer observations, i.e. teams that have been relegated along the nine seasons.

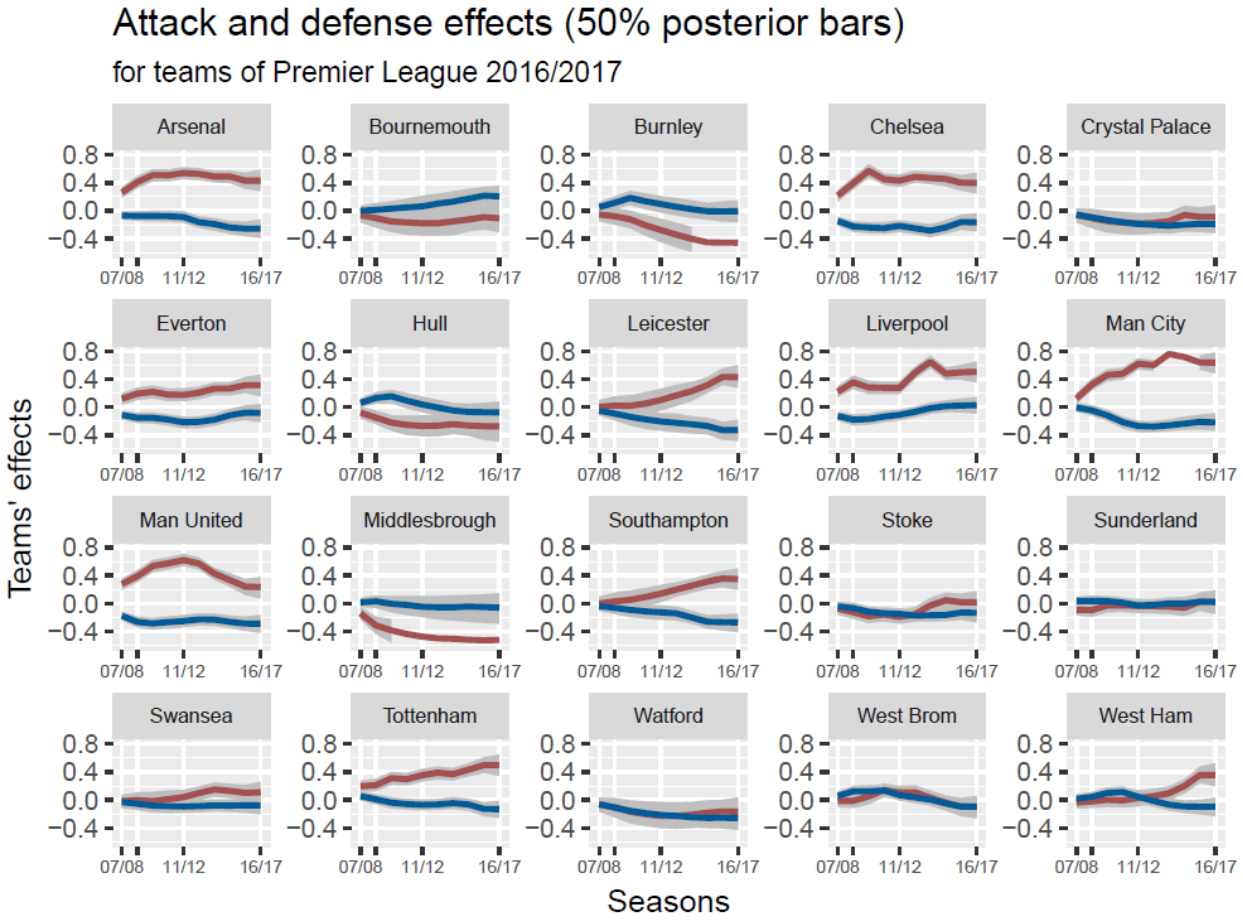

Figure 1c: Estimated posterior 50% confidence bars for the attack (red) and the defense (blue) effects along the nine considered seasons (from 2007-2008 to 2015-2016) and predicted bars for the teams belonging to the Spanish Liga 2016/2017, the test set season. According to the model specification, these teams' specific parameters are included in a log-linear predictor built for modeling the scoring strength of each team and given an exchangeable prior distribution. Their interpretation is then the following: the greater/lower is the attack/defense team effect and the better is evaluated that specific team. Liga is usually lead by Real Madrid, Barcelona and less often by Atletico Madrid. This pattern emerges clearly from this graph. Whereas the three relegated teams - Sporting Gijon, Granada and Osasuna - were correctly predicted to perform poorly at the beginning of the 2016-2017 season. Wider posterior bars are in general associated with teams with fewer observations, i.e. teams that have been relegated along the nine seasons.

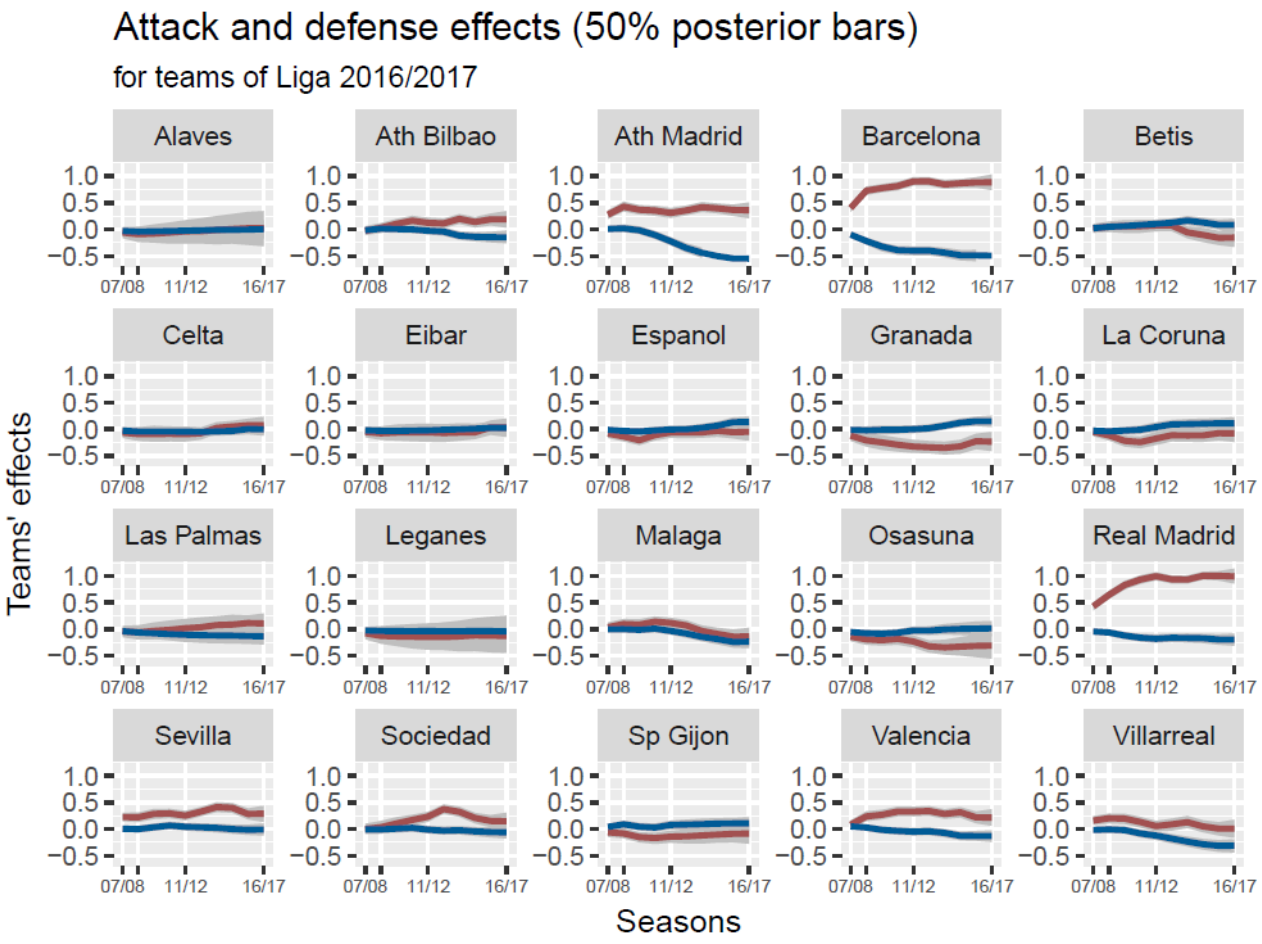

Figure 1d: Estimated posterior 50% confidence bars for the attack (red) and the defense (blue) effects along the nine considered seasons (from 2007-2008 to 2015-2016) and predicted bars for the teams belonging to the Italian Serie A 2016/2017, the test set season. According to the model specification, these teams' specific parameters are included in a log-linear predictor built for modeling the scoring strength of each team and given an exchangeable prior distribution. Their interpretation is then the following: the greater/lower is the attack/defense team effect and the better is evaluated that specific team. Juventus, the leading team in the previous five seasons, is predicted to perform quite well; also Roma and Napoli, respectively the second and the third team at the end of the test set season, are predicted to perform well, but with an associated defense effect lower than Juventus. Pescara, the least in the final rank, is correctly associated with the worst performances. Wider posterior bars are in general associated with teams with fewer observations, i.e. teams that have been relegated along the nine seasons.

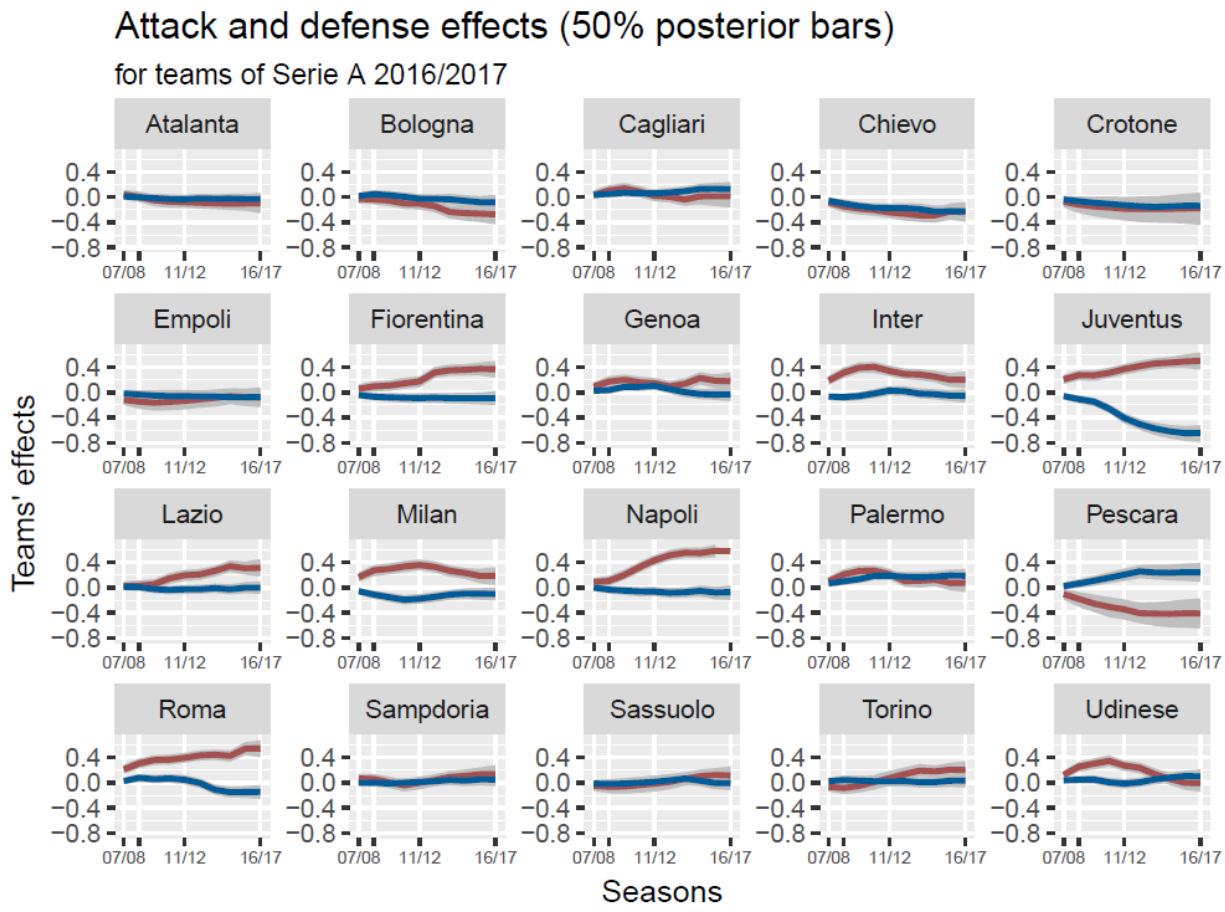

Model fit

Figure 2a: Graphical posterior predictive check: we generate the home and the away goals for every match of our training set (nine seasons, from 2007-2008 to 2015-2016) and we check whether the replicated goals' difference under the model fits the observed goals' difference. The top-row plots are probability plots for the German Bundesliga and the Spanish Liga, where on the x-axis we have the observed goals' difference y1 - y2, whereas on the y-axis we have the replicated goals' difference y1(rep) - y2(rep). Darker regions are associated with higher posterior probabilities. In correspondence of an observed goals' difference on the x-axis, this graphical tool provides the most likely replicated goals' difference. A perfect - maybe overfitting - model should display the black regions along the red bisector, but such a model would probably be inappropriate in terms of future predictions. In fact, let consider the bottom-row graphs that display the observed distributions of the goals' difference for the leagues plotted above. These distributions are slightly asymmetric, but the maximum is always concentrated at zero. In our dataset a goals' difference amounting at zero is then sensitively more likely than a goals' difference of, say, three. Perhaps, the top-row checks are comforting, since they suggest a good fit in correspondence of likely goals' difference - e.g. -1,0,1 - and a poorer fit when the goals difference turns out to be rare - say, 8.

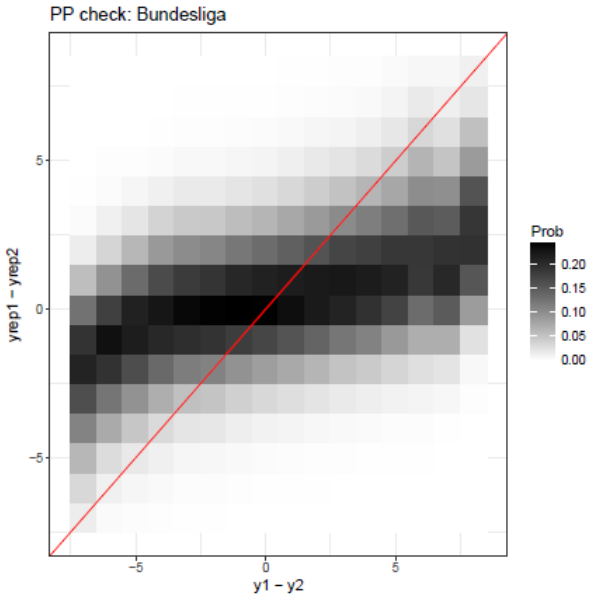
PP check: Bundesliga
yrep1 − yrep2
y1 − y2
Prob
0.20
0.15
0.10
0.05
0.00

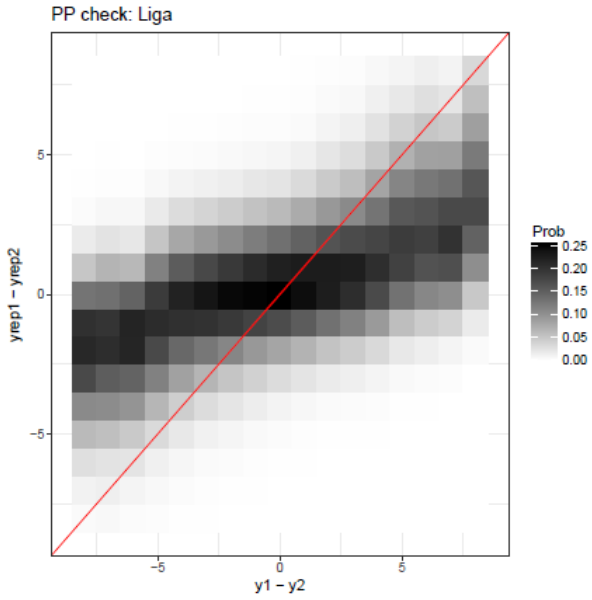
PP check: Liga
yrep1 − yrep2
y1 − y2
Prob
0.25
0.20
0.15
0.10
0.05
0.00

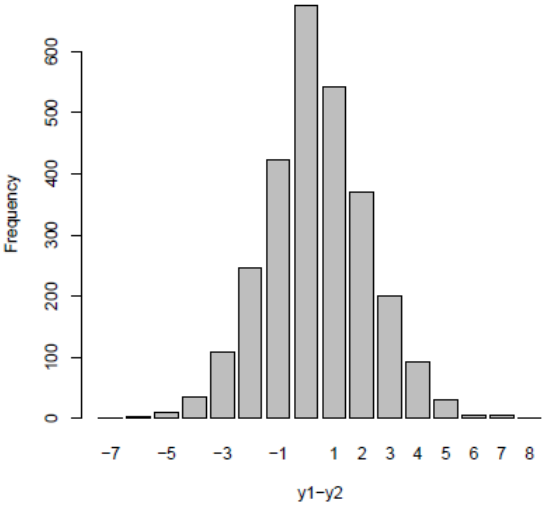
Frequency
y1−y2

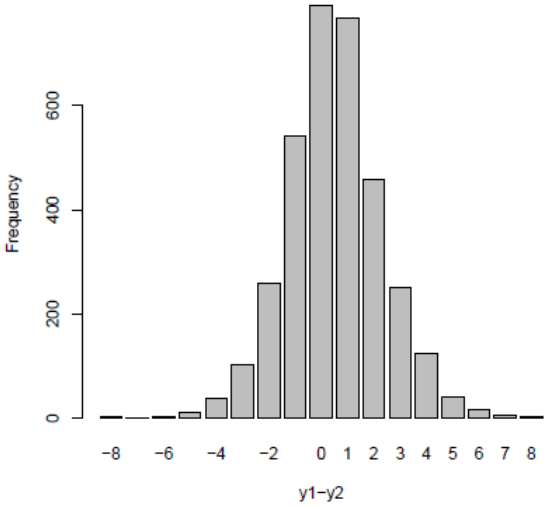
Frequency
y1−y2

Figure 2b: Graphical posterior predictive check: we generate the home and the away goals for every match of our training set (nine seasons, from 2007-2008 to 2015-2016) and we check whether the replicated goals' difference under the model fits the observed goals' difference. The top-row plots are probability plots for the English Premier League and the Italian Serie A, where on the x-axis we have the observed goals' difference y1 - y2, whereas on the y-axis we have the replicated goals' difference y1(rep) - y2(rep). Darker regions are associated with higher posterior probabilities. In correspondence of an observed goals' difference on the x-axis, this graphical tool provides the most likely replicated goals difference. A perfect - maybe overfitting - model should display the black regions along the red bisector, but such a model would probably be inappropriate in terms of future predictions. In fact, let consider the bottomrow graphs that display the observed distributions of the goals' difference for the leagues plotted above. These distributions are slightly asymmetric, but the maximum is always concentrated at zero. In our dataset a goals' difference amounting at zero is then sensitively more likely than a goals' difference of, say, three. Perhaps, the top-row checks are comforting, since they suggest a good fit in correspondence of likely goals' difference - e.g. -1,0,1 - and a poorer fit when the goals difference turns out to be rare - say, 8.

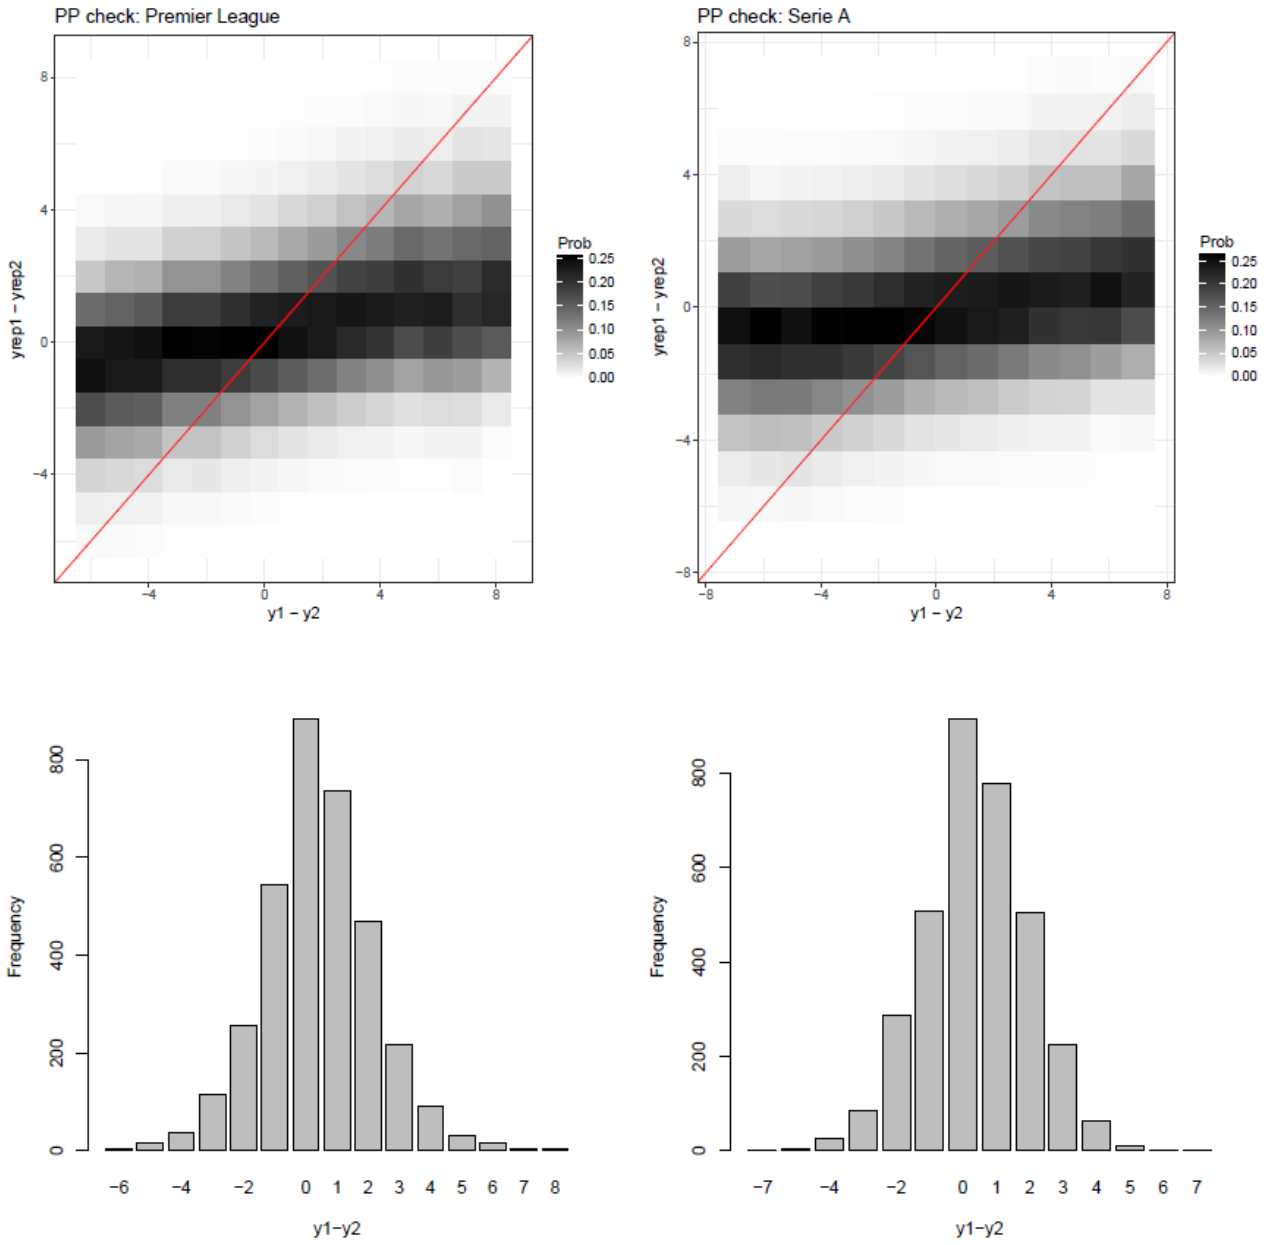

Posterior predictive distribution of the matches outcomes

Figure 3: Posterior predictive distribution of the possible results for the match Real Madrid-Barcelona, Spanish Liga 2016/2017, and Sampdoria-Juventus, Italian Serie A 2016-2017. Both the top-row and bottom-row plots display the posterior uncertainty related to the exact predicted outcome. In the top-row plots the results are ordered from the most likely to the least likely; the dashed blue line is in correspondence of the actual final result. In the bottom-row plots the home goals are on the x-axis and the away goals on the y-axis. Darker regions are associated with higher posterior probabilities, whereas a red square is in correspondence of the actual final result.

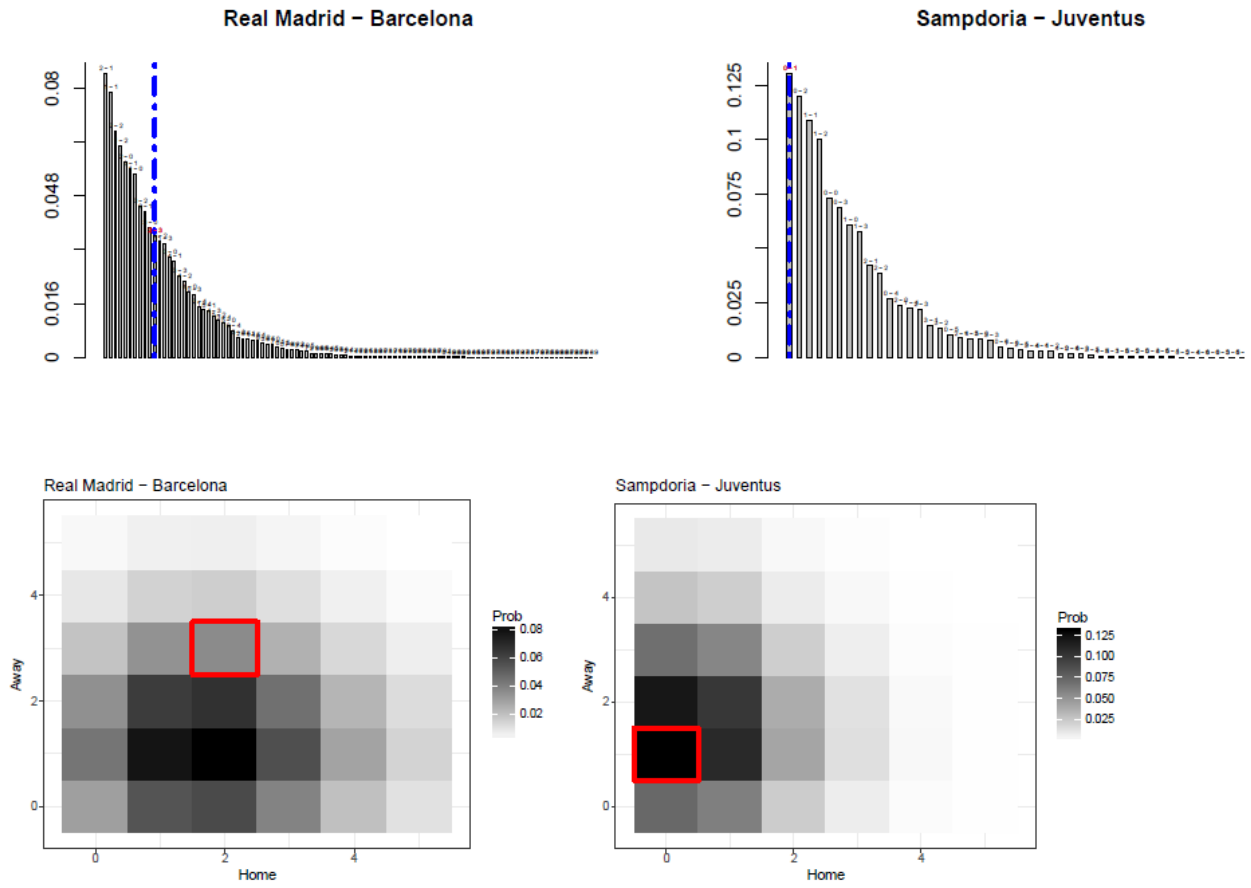

Posterior predictive distribution for the achieved points

Figure 4: Posterior 50% confidence bars (gray ribbons) for the achieved final points of the top four European leagues 2016-2017. Black points are the observed points. Black lines are the posterior medians. At a first glance, the pattern of the predicted ranks appears to match the pattern of the observed ranks, and the model appears to be well calibrated.

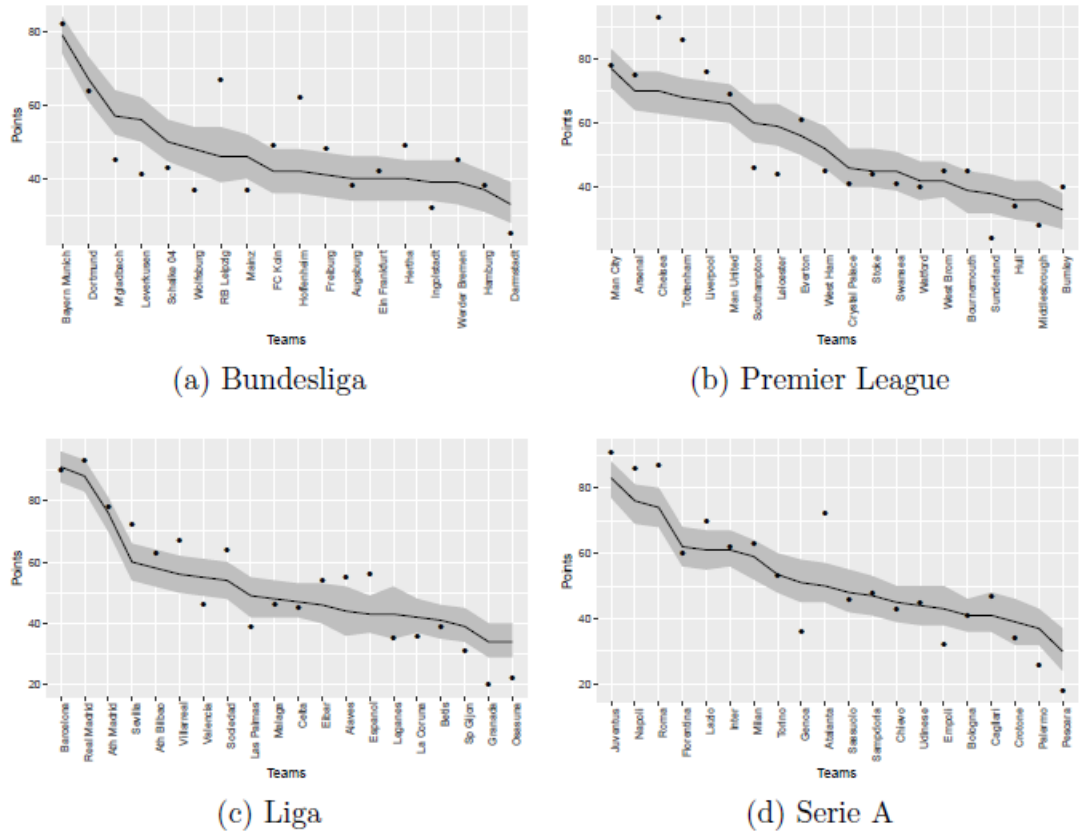

(a) Bundesliga (b) Premier League

(c) Liga (d) Serie A

Betting strategy

Table 1: Betting strategy: bet one unit on the three-way match outcome - home win, draw or away win - with the highest expected return. Here we provide the expected profits (%/100) for seven considered bookmakers, for each of the top-four European leagues. All the expected profits turn out to be positive.

	Bet365	Bwin	Interwetten	Ladbrokes	Sportingbet	VC Bet	W. Hill
Bundesliga	0.180	0.157	0.135	0.148	0.226	0.169	0.177
Premier League	0.242	0.188	0.157	0.198	0.247	0.216	0.248
Liga	0.100	0.085	0.07	0.073	0.138	0.095	0.104
Serie A	0.180	0.154	0.088	0.14	0.228	0.156	0.199