Newcastle University

MMathStat project

Statistical Models for the Prediction of Football Matches with Applications to Betting

Author:
Scott Lillico

Supervisor:
Dr. Phil Ansell

Contents

Abstract

We look at the history of research done in football statistics, considering the methods used to draw conclusions and how the methods have evolved over the years. We then examine a couple of models in detail before applying one of them to a set of results consisting a season. We also suggest a potential change to this model to see if we can improve it in any way and compare the results we get. When calculating the results, we contemplate both the parameters considering all games and how parameters change over time considering a certain number of games. We apply these results to betting on a weekends' fixtures and seeing how much of a profit is made before drawing our conclusions.

Chapter 1

Introduction

Football is one of the most popular sports in the world, if not the most popular. The 2013/14 English Premier League had an average viewership of over 1 million people in the United Kingdom for every game shown live on TV. [1]

With the popularity of football, it has a huge betting scene. From something as simple as betting on the outcome of a match to betting on the number of corners both teams will have over a particular match, there are a high number of possibilities to profit from betting on football matches. When betting, punters will want to try and maximise their profit while at the same time trying to minimise the risk of losing and getting the best possible odds they can on an outcome. It is here where we can use statistics to model outcomes of football matches based on the strengths and weaknesses of the teams playing.

There are a high number of factors to consider when trying to predict the results of a football match, such as the relative strengths and weaknesses of the two teams playing, the form of the respective teams, whether key players are missing through either injury or suspension or indeed key players coming back into the team. Additionally, we can also look at key personnel changes and see the impact they have on a team's performance, such as bringing in new players after a transfer window or changing manager partway through the season.

When considering football statistics, defining the ability of a team is usually determined by two factors; the number of goals they score and the number of goals they concede. Stronger teams will, naturally, score more goals and concede fewer goals than weaker teams. However, it is not always the case that stronger teams will beat weaker teams, since football can not be simplified in this way. The statistics calculated can only determine the probability that there will be a certain outcome in a match, whether it is simply the outcome of win, lose or draw or the specific score of a match.

Chapter 2

Literature Review

A substantial amount of research has been done into modelling various aspects of football. One such paper, by Reep and Benjamin (1968) [2], gives a reflection on how much modelling football over the years has changed. They modelled the number of passes a team would make until possession was switched over, via either a shot at goal, being tackled or an infringement, with a Negative Binomial distribution. They also made further observations into each passing move, such as the proportion of shots in a so called "shooting area" that led to goals compared to all goals scored and the proportion of goals conceded when losing the ball in the teams' own half compared to all goals conceded. One particularly interesting statistic they noted was the average number of shots it took to score a goal for a particular team for different seasons, as well as for various World Cups. In general, teams would take an average of 8-10 shots to score a goal, although they do point out that just because one team has more shots than the opposing team does not mean that they will win the game. They come to the conclusion that chance dominates the game in deciding who will win.

Karlis and Ntzoufras (2003) [3] considered the bivariate Poisson distribution, not widely used due to the difficulty of implementing it, as a way of modelling scores in a couple of sports, water polo and football. The bivariate Poisson distribution allows a parameter to act as a dependence parameter between the two random variables, X and Y, which model home and away goals respectively for a football match. A purpose of this approach was to obtain a better estimate of the number of draws that occur since, typically, the number of draws estimated by two independent Poisson models is underestimated in comparison to how many draws occur in reality. To implement this, Karlis and Ntzoufras inflated the probability of draws occurring by adding on some discrete distribution to their proposal of the bivariate Poisson distribution if the scores were equal, i.e.

$$P_D(x,y) = \begin{cases} (1-p)BP(x,y \mid \lambda_1,\lambda_2,\lambda_3), & x \neq y, \\ (1-p)BP(x,y \mid \lambda_1,\lambda_2,\lambda_3) + pD(x,\boldsymbol{\theta}), & x = y. \end{cases}$$

In this case, $D(x,\boldsymbol{\theta})$ is a discrete distribution with $\boldsymbol{\theta}$ as a parameter vector. Karlis and Ntzoufras used the 1991-92 Italian Serie A season to test the bivariate Poisson approach and compared it to using 2 Poisson distributions. It was found that using 2 Poisson distributions severely underestimated the number of 1-1 draws, and using the bivariate Poisson overestimated the number of 0-0 draws.

Baio and Blangiardo (2010) [4] considered a Bayesian approach when modelling football results, using the 1991/92 and 2007/08 Italian Serie A seasons to implement this. Like what will eventually form the basis of our model, they defined each team having various attacking and defensive parameters, as well as a home effect parameter. The attacking and defensive parameters for each team were calculated by considering the number of home and away goals scored in a match between two teams and modelling them based on a Poisson distribution, similarly to how our model will be defined. Since the Bayesian approach was being used, suitable priors needed to be defined, with the home parameter having a Normal prior with zero mean and a very small precision of 0.0001, reflecting the uncertainty of how the home parameter will shape up. Additionally, the attacking and defensive parameters also received Normal priors each with differing means μ and differing precisions τ. Finally, the two means had a prior distribution with the same parameters as the prior for the home effect, whilst the precisions each had a prior distribution of $Gamma(0.1, 0.1)$. The outcome of this analysis shows how well the Bayesian approach fits to the actual data and it is compared to the bivariate Poisson approach done in the past by Karlis and Ntzoufras.

Another approach considered was done by Hvattum and Arntzen (2010) [5], who used a measure known as the ELO rating to predict match results. Originally developed for chess in the 1970's by Arpad Elo, it has seen a widespread use in other sports, including football. The ELO rating uses the results of previous contests to measure a team's current strength, in an attempt to give an overall rating across all teams from differing leagues. Given the result of a match, we define a new rating of a team as

$$R_n = R_o + K(W - W_e).$$

R_n and R_o are the new and old ratings respectively, K is a weight constant depending on the tournament the match is played at, with further adjustments depending on the goal difference in the game, W is the result of the game, with a win corresponding to 1, a draw equaling 0.5 and a loss giving a value of 0, and W_e is the expected result of the game, calculated by

$$W_e = 1/\left(10^{(-dr/400)} + 1\right)$$

where dr is the difference in ratings between the teams playing plus 100 points for the home team. Hvattum and Arntzen compared two ELO based methods (with and without accommodating for goal difference) with a variety of other pre-established methods, such as probabilities calculated based on bookmakers odds. It was found that the ELO ratings were an efficient method of encoding past results over a short time span.

Chapter 3

Models

For our predictions, we will be using an adaptation of two different models proposed by two different authors; one by Maher (1982) and one by Dixon and Coles (1997).

3.1 Assumptions

We first define assumptions that both papers share in common. One such assumption is that the number of goals scored in a football match is defined by a Poisson distribution. In any given match, when a team has possession, they have a chance to score a goal. While this chance of scoring a goal is low, the number of times a team has possession in a match is very high. This leads to the number of goals in a game being Binomial which, since possession is very high, leads to a good Poisson approximation.

3.2 Maher

M.J. Maher [6] proposed that, to model the score in a football match between teams i and j, with team i being the home team, two independent Poisson distributions should be used, X_{ij} and Y_{ij}, for the number of home and away goals respectively. The observed score of such a match is (x_{ij}, y_{ij}), so we assume

$$X_{ij} \sim Poisson(\alpha_i \beta_j)$$

and

$$Y_{ij} \sim Poisson(\gamma_i \delta_j).$$

Here, α_i is the strength of team i's attack when they play at home, β_j the weakness of team j's defense when playing away, γ_i the weakness of team i's defense when playing at home and, finally, δ_j is the strength of team j's attack when playing away.

To ensure that we have a unique set of parameters, we impose the constraints that

$$\sum_i \alpha_i = \sum_i \beta_i$$

and

$$\sum_i \gamma_i = \sum_i \delta_i.$$

Since the number of home goals $\mathbf{X}$ and the number of away goals $\mathbf{Y}$ are independent, we can obtain maximum likelihood estimates for the α's and β's entirely from the $\mathbf{x}$ and, likewise, for the γ's and δ's entirely from the $\mathbf{y}$.

We define the likelihood function for the number of home goals in a match as

$$L(\alpha, \beta) = \prod_i \prod_{i \neq j} \frac{(\alpha_i \beta_j)^{x_{ij}} e^{-\alpha_i \beta_j}}{x_{ij}!}.$$

Taking logs of this function, we obtain the log-likelihood function as

$$\log(L(\alpha, \beta)) = \sum_i \sum_{i \neq j} (-\alpha_i \beta_j + x_{ij} \log(\alpha_i \beta_j) - \log(x_{ij}!))$$

Taking the derivative of the log-likelihood with respect to each of α_i and β_j, we get

$$\frac{\partial \log L}{\partial \alpha_i} = \sum_{i \neq j} \left(-\beta_j + \frac{x_{ij}}{\alpha_i} \right)$$

and

$$\frac{\partial \log L}{\partial \beta_j} = \sum_{i \neq j} \left(-\alpha_i + \frac{x_{ij}}{\beta_j} \right).$$

This culminates in our maximum likelihood estimates for each parameter being

$$\hat{\alpha}_i = \frac{\sum_{i \neq j} x_{ij}}{\sum_{i \neq j} \hat{\beta}_j}$$

and

$$\hat{\beta}_j = \frac{\sum_{i \neq j} x_{ij}}{\sum_{i \neq j} \hat{\alpha}_i}.$$

Since the maximum likelihood estimate for each parameter requires the same statistic for the other parameter, we will need to use an iterative scheme to obtain numerical values for our maximum likelihood estimates. One such way to do this is to define a starting value for each estimate and then use the $\hat{\alpha}$'s to estimate the $\hat{\beta}$'s, the $\hat{\beta}$'s to estimate the $\hat{\alpha}$'s and then continue on in this way until the estimates converge.

We can run similar calculations for the number of away goals in a match, leaving our maximum likelihood estimates for γ and δ as

$$\hat{\gamma}_i = \frac{\sum_{i \neq j} y_{ij}}{\sum_{i \neq j} \hat{\delta}_j}$$

and

$$\hat{\delta}_j = \frac{\sum_{i \neq j} y_{ij}}{\sum_{i \neq j} \hat{\gamma}_i}.$$

This model is a good basis to start from, although there are ways we can improve it. For example, having four parameters, for home and away and then attack and defense, may not be necessary, so we could find some means to explain the same amount of information whilst defining fewer parameters. More importantly, this model assumes that every possible score is independent, which while true for a lot of scores, investigating this point further shows that not all scores are independent.

3.3 Dixon-Coles

Dixon and Coles [7] proposed an adaptation to Maher's model in 1997. Like Maher, they proposed two independent Poisson distributions for the number of home and away goals scored in a particular football match, i.e.

$$X_{ij} \sim Poisson(\alpha_i\beta_j\gamma)$$
$$Y_{ij} \sim Poisson(\alpha_j\beta_i)$$

where the α_i's are the attacking parameters, with higher values signifying a better attack, the β_i's are the defensive parameters, where a lower value means a better defense, and γ is a home effect parameter. While this follows a similar structure to what Maher proposed, Dixon and Coles suggested a further modification to the model in light of an investigation into independence of scores they conducted. They found that, while most scores in a football match are independent between scores, the scores 0-0, 1-0, 0-1 and 1-1 were not considered as being independent between their respective scores. Taking into account this change, we define a new probability mass function as

$$Pr(X_{ij} = x, Y_{ij} = y) = \tau_{\lambda,\mu}(x, y)\frac{\lambda^x e^{-\lambda}}{x!}\frac{\mu^y e^{-\mu}}{y!} \tag{3.1}$$

where $\lambda = \alpha_i\beta_j\gamma$, $\mu = \alpha_j\beta_i$ and

$$\tau_{\lambda,\mu}(x, y) = \begin{cases} 1 - \lambda\mu\rho & \text{if } x = y = 0, \\ 1 + \lambda\rho & \text{if } x = 0, y = 1, \\ 1 + \mu\rho & \text{if } x = 1, y = 0, \\ 1 - \rho & \text{if } x = y = 1, \\ 1 & \text{otherwise.} \end{cases}$$

Here, ρ is a dependence parameter with constraint

$$\max\left(-\frac{1}{\lambda}, -\frac{1}{\mu}\right) \le \rho \le \min\left(\frac{1}{\lambda\mu}, 1\right).$$

If $\rho = 0$, then we have independence across all scores; otherwise, we lose independence for the above mentioned scores.

To prevent over-parameterisation, we impose the constraint of

$$\frac{1}{n}\sum_{i=1}^{n}\alpha_i = 1.$$

Similar to our calculations for Maher's model, we define the likelihood of this probability function. We index each respective match as $k = 1, \ldots, N$ and define the scores of each match as (x_k, y_k), giving our likelihood up to proportionality as

$$L(\alpha_i, \beta_i, \rho, \gamma; i = 1, \ldots, n) \propto \prod_{k=1}^{N} \tau_{\lambda_k,\mu_k}(x_k, y_k)e^{-\lambda_k}\lambda_k^{x_k}e^{-\mu_k}\mu_k^{y_k}, \tag{3.2}$$

where

$$\lambda_k = \alpha_{i(k)}\beta_{j(k)}\gamma,$$
$$\mu_k = \alpha_{j(k)}\beta_{i(k)}.$$

Here, $i(k)$ and $j(k)$ are the indices for the home and away teams, respectively, playing in match k. Taking the log-likelihood of Equation 3.2, we obtain

$$\log L = \sum_{k=1}^{N} \log\left(\tau_{\lambda_k,\mu_k}(x_k, y_k)\right) + x_k \log(\lambda_k) + y_k \log(\mu_k) - \lambda_k - \mu_k. \tag{3.3}$$

Because of our lack of independence, we cannot obtain a system of linear equations for our maximum likelihood estimates like we did for Maher's model. Instead, we must use direct numerical maximisation of the log-likelihood, although this is easily implemented using R.

3.4 Proposed Changes

Whilst the Dixon and Coles model is good for modelling the abilities of football teams, there could still be room for a potential improvement in it. A noteworthy property of the Dixon and Coles model is that it is assumed that the home effect parameter γ is constant across all teams, meaning that this leaves no room for some teams perhaps being stronger at home in comparison to some other teams. Therefore, we propose a slight modification to this model to accommodate for the possibility of some teams being stronger at home than others, so, assuming independence, the distributions are

$$X_{ij} \sim Poisson(\alpha_i\beta_j\gamma_i)$$
$$Y_{ij} \sim Poisson(\alpha_j\beta_i)$$

where the parameters are defined as before, with the exception of γ_i now being the home effect parameter for team i. However, since the fundamental structure of how we model results has not changed, we still have a loss of independence for the aforementioned scores. The only change this would make to Equation 3.1 is that now $\lambda = \alpha_i\beta_j\gamma_i$. Similarly, our definition of λ_k used in Equations 3.2 and 3.3 now becomes $\lambda_k = \alpha_{i(k)}\beta_{j(k)}\gamma_{i(k)}$.

Chapter 4

Modelling a Football Season

We consider the 2013/14 English Premier League season for using our respective models. For the purpose of calculating scores, since there are technically an infinite number of possible scores, we set a maximum of 8 goals per game between both teams, with the added restriction that one team can score no more than 7 goals. Since, out of 380 games in this season, only 2 games had a result that would fall outside of these constraints, we can safely assume that results outside of the set constraints are a very rare occurrence and not representative of a typical result. All results are taken from the Rec. Sports Soccer Statistics Foundation (RSSSF). [8]

4.1 Common γ

4.1.1 Static Model

We first consider the static model that Dixon and Coles proposed to estimate each of our parameters. Taking into account all of the results over the course of the season with no time dependence, we obtain the attacking and defensive parameter estimates for each team seen in Table 4.1.

The "home effect" parameter γ for the parameters in Table 4.1 is 1.31, whilst the corresponding dependence parameter ρ is 0.143. We can use these parameters to simulate seasons and compare to how well the simulation matches up with what happened in reality.

Using R, a simulation of a number of seasons with the same teams can be implemented using the respective parameters. Running a simulation of 1000 seasons and taking the average points for each team, we can view a simulated position in the league for each team and compare them to where the team actually finished, seen in Table 4.2. Additionally, a box plot of the amount of points teams earned over the 1000 simulations can be seen in Figure 4.1.

The simulation appears to match up reasonably well with how the season turned out in reality, meaning that this model is sensible at the very least. There are a few anomalies in this simulation compared to what happened in reality though. Some teams that were simulated to finish much lower than what they actually did include Crystal Palace and Newcastle. One possible explanation is the number of goals both teams scored and

Team	α	β
Arsenal	1.28	0.94
Aston Villa	0.74	1.35
Cardiff	0.61	1.64
Chelsea	1.32	0.62
Crystal Palace	0.63	1.09
Everton	1.14	0.88
Fulham	0.79	1.93
Hull	0.73	1.19
Liverpool	1.93	1.21
Manchester City	1.91	0.87
Manchester United	1.21	0.98
Newcastle	0.83	1.34
Norwich	0.53	1.37
Southampton	1.02	1.06
Stoke	0.86	1.16
Sunderland	0.78	1.35
Swansea	1.04	1.22
Tottenham	1.06	1.18
West Brom	0.83	1.33
West Ham	0.75	1.12

Table 4.1: Parameter Estimates with Common γ

conceded over the course of the season could be seen as being worse than other teams around them. In reality, Crystal Palace scored the third fewest number of goals over the season, with only Norwich and Cardiff, two teams who were relegated, scoring fewer. This culminates in a lacklustre attacking parameter for where they actually finished. For Newcastle, despite the fact that they finished 10th, only 5 teams conceded more goals than them, since they tended to lose heavily when they lost games. This ends up impacting their defensive parameter while not boosting their attacking parameter to compensate since, when Newcastle won games, they tended to be by small margins.

On the other end of scale, namely teams doing far better in the simulation compared to how they did in reality, we have Swansea and West Brom. West Brom in particular actually scored and conceded the same number of goals as Newcastle despite finishing 7 places lower and scoring 13 points fewer. Comparing the parameters between the two teams, West Brom actually have marginally better parameters by way of having slightly better defensive ability. This ends up impacting the simulation, making West Brom seem better than what they actually turned out to be. One potential reason that Swansea's position could have been overestimated is that their goal difference was much better than the teams that finished around them, being at an even zero compared to the teams around them, who were generally in the negatives.

One curious trend that should be noted is how the simulated positions seem to be decided for teams that appear to be close in ability to each other. Towards the top half of the

Team	Simulated Position	Actual Position
Arsenal	4th	4th
Aston Villa	17th	15th
Cardiff	20th	20th
Chelsea	2nd	3rd
Crystal Palace	15th	11th
Everton	5th	5th
Fulham	19th	19th
Hull	14th	16th
Liverpool	3rd	2nd
Manchester City	1st	1st
Manchester United	6th	7th
Newcastle	13th	10th
Norwich	18th	18th
Southampton	7th	8th
Stoke	10th	9th
Sunderland	16th	14th
Swansea	9th	12th
Tottenham	8th	6th
West Brom	12th	17th
West Ham	11th	13th

Table 4.2: Simulated Positions for Common γ

table, the simulation seems to suggest that having a better defense is superior to having a better attack. A potential reason for this is that teams who are near the top of the table are generally going to score a lot of goals in the first place, since they are winning more games, and the deciding factor comes from being able to stop goals being scored against you, since that can make the difference between a win, draw or loss. This can be seen with Chelsea, who conceded the fewest amount of goals in reality and have by far the best defensive parameter, and Liverpool, who have the best attacking parameter, but only eight teams have worse defensive parameters. Liverpool, who actually finished ahead of Chelsea with a superior goal difference, are still simulated to finish lower than Chelsea, potentially due to amount of goals they conceded and the teams they conceded them against influencing their defensive parameter negatively.

For the bottom half of the table, more importance seems to be placed in scoring goals rather than the ability to defend against them. Similarly in concept to the teams in the top half of the table needing a better defense, the potential logic here is that teams at the bottom half of the table are conceding a lot of goals to begin with, since they're more likely to be losing a fair amount due to their position, meaning they are needing to score goals to perhaps get marginal results rather than consolidate an outcome. Coming back to Crystal Palace, this may be the reason why they are simulated to finish as low as they are despite their similar goal difference to Newcastle, West Brom and Hull, due to their paltry attack. Another two teams to compare are the two bottom teams, Fulham and Cardiff. The simulation does match where they finished in reality, although Fulham did

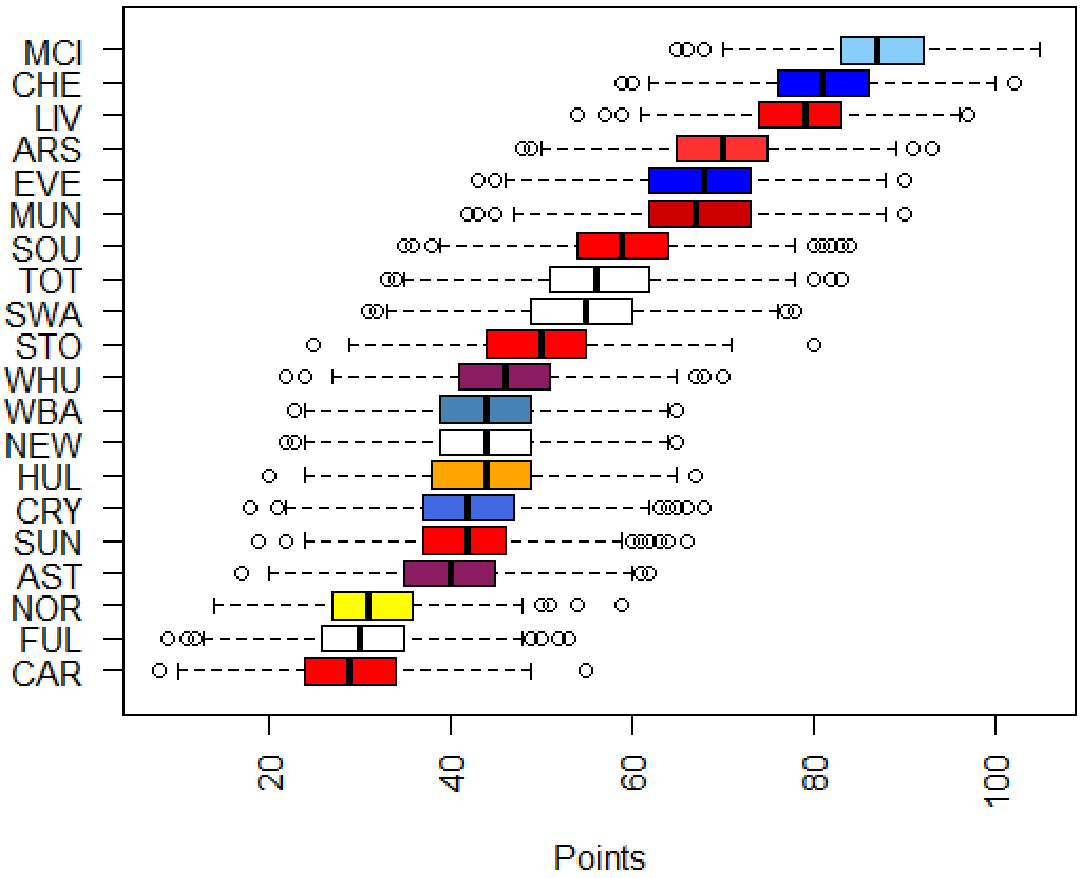

Figure 4.1: Boxplots of Points Scored by Teams over 1000 Simulated Seasons

have a worse goal difference than Cardiff. However, Fulham scored more goals than what Cardiff did, further supporting the idea that scoring goals is more important than not conceding them towards the foot of the table.

An interesting thing to note looking at Figure 4.1 is the amount of overlap between the number of points teams scored in at least one season. The only teams that seem to be dramatically far ahead are the top three teams (Manchester City, Chelsea and Liverpool), whilst the only teams lagging behind a large distance are the three relegated teams (Cardiff, Fulham and Norwich). There are not many instances where it would be impossible for a team to finish ahead of everyone else in a simulation, as unlikely as it would have been. The only impossibilities in the simulations were Manchester City finishing below the bottom four teams on the simulated averages and Chelsea finishing below the bottom two. This shows how tight in ability the midfield at least in the 2013/14 Premier League were, with a large number of teams being able to finish above the high ability teams, as unlikely as it was to have actually happened in one of the simulated seasons.

4.1.2 Dynamic Model

We can also look at how the parameters change over the course of the season, and potential reasons why parameters may have changed. For any given time point, a batch of 100 games are considered (meaning our first parameter estimates are after every team has played ten games), though there are some times where not every team has played an equal amount of times due to a various amount of postponements.

One thing we may want to consider is how a team's performance changes if they end up changing manager over the course of the season. There were 10 managerial changes over the course of the season, with Fulham changing their manager twice. However, not all of these changes will be considered. For example, Manchester United sacked their manager very late in the season, leaving not a lot of time for performances to change. On the other hand, Sunderland sacked and replaced their manager during our defined burn in period at the beginning of the season, meaning their time of change is not explicitly defined in our plots, though we can still somewhat infer the effects. Similarly, Crystal Palace sacked their manager in the burn in period, but replaced him when we start measuring our parameters, meaning we can still use this information. All other managerial replacements can be measured normally.

Another factor that may impact a team's performance is what transactions they make in the transfer window. Unlike replacing managers, all teams go through this process, though every team obviously works through the transfer window in different ways. If a team's performance drastically increases after the transfer window, we can often attribute that to them bringing in players that help the team. Likewise, if a team's performance tails off, it may be due to selling a key part of the team and not replacing them. The transfer window starts before time 20 and ends after time 23 in the forthcoming plots.

The first time progression we will consider is that of the three relegated teams; Norwich, Fulham and Cardiff. All three sacked and replaced their managers at least once over the course of the season, so we can draw parallels between them and see how their performances changed between managers. We can also look at their performances before and after the transfer window, to see if there were any dramatic changes in performance between those times. The transfer window period is defined as the time between the purple lines on each subsequent time plot. Additionally, this plot and all subsequent plots of this type will have a light blue line indicating the baseline attacking and defensive parameter of 1, for a visual representation of comparing a team's ability to an "average" ability. The progression of both the α and β parameters can be found in Figure 4.2.

We see that all three teams were worse than the baseline attacking power throughout the season, which may have contributed to their downfall. The same holds true for defensive power as well, with the exception of Norwich for a brief time after the transfer window closed. It is also worth considering the times where performances started to falter, if action was taken in light of these performances and whether this action worked.

Considering Cardiff, we see that they sacked their manager just before the transfer window started, which owners who are seeking a change will generally do since the new manager will often want to bring his own players in. Coincidentally, this was when

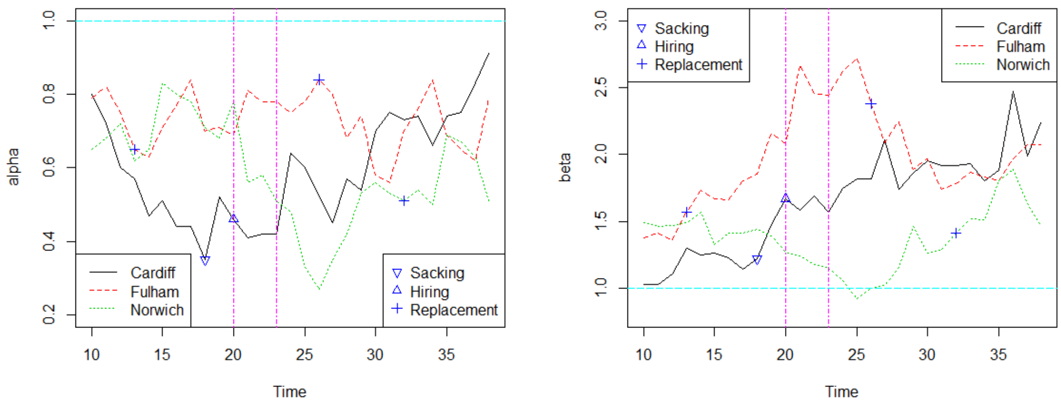

(a) α Parameters (b) β Parameters

Figure 4.2: Progression of α and β for Cardiff, Fulham and Norwich

Cardiff's attack was at its worst, meaning there was improvement from the new manager coming in, though this may also coincide with the transfer window and Cardiff strengthening their attack a little. However, the Welsh team start to become undone because of their defensive abilities towards the end of the season. After the transfer window, their defensive parameter never fell below 1.5, not a great value to begin with, and it only got worse from there, peaking just below 2.5.

One of Fulham's flaws appear to be their defensive abilities in the middle of the season, when they were on their second manager. Their defensive abilities were much worse than the other two teams relegated at this point of the season, whilst their attack was not improving to compensate for this, resulting in more defeats. After the second manager was replaced, Fulham did get better defensively, though their attack also got worse at this point, essentially making no difference to their woes. However, their attack did get back onto the levels they had been consistently at towards the end of the season, which may have saved them from the ignominy from finishing bottom of the table.

Norwich's nightmares over the season can mostly be attributed to their lack of attacking prowess. While it was only at a below average point at the start of the season, the attacking parameter dramatically tailed off between the start of the transfer window and some time afterwards. Curiously, action was not taken then, potentially because they were also defending well during this time period. The fact of the matter is that they were simply not scoring enough, and replacing their manager was perhaps an action that was taken too late to save them. While their attack did improve in the time the new manager came in, their defensive abilities also took a hit, showing they were still not winning games, and they ended up relegated.

Another comparison we may wish to make is between teams who finished at around the same positions but had quite differing forms over the season. Two teams who fit this are Crystal Palace and Newcastle. Their progression can be seen in Figure 4.3.

From the plots, we see that Crystal Palace improved as the season went on, while New-

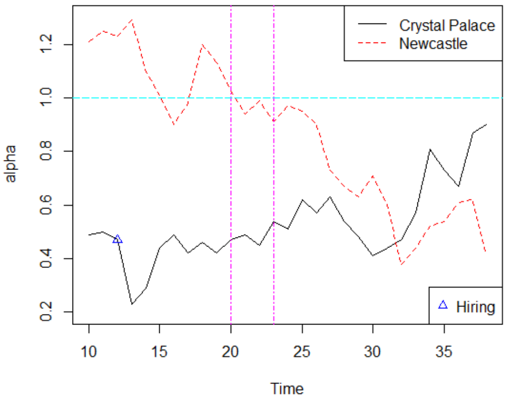

(a) α Parameters

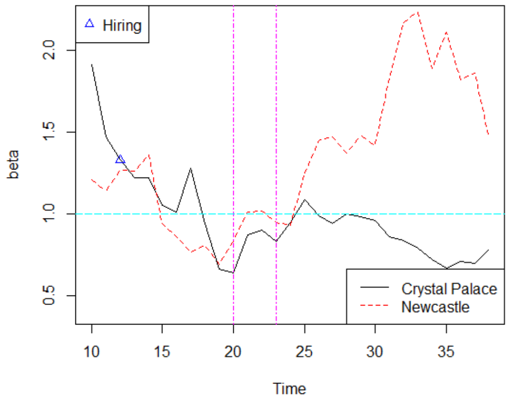

(b) β Parameters

Figure 4.3: Progression of α and β for Crystal Palace and Newcastle

castle regressed in ability over time. In Crystal Palace's case, this may be due to the effect of bringing in a new manager. Whilst their attack was never really considered as stellar, their defense showed dramatic improvement over the season, backed up by the fact that only 6 teams conceded fewer goals than them. Compare this to their position at the end of November (after 13 games), where they were bottom of the table, with the fewest goals and with only 3 teams conceding more than them.

In contrast, at this same point in the season, Newcastle were in 5th position, the highest they would obtain over the season. This can be attributed to their attacking ability at this point, with it being the highest it would get. Things start to noticeably go wrong for Newcastle around the transfer window period, with their attack faltering and their defense becoming weaker. A potential reason for this is that Newcastle sold one of their key players during the transfer window without replacing him, which may have disrupted the team spirit and hindered their ability to perform as well. Indeed, looking at Newcastle's record between the start of the transfer window and the end of the season (Time 20 to Time 38), they had a record of 5 wins, 1 draw and 13 losses, with the vast majority of losses being heavy losses where Newcastle failed to score. This would lead to a heavy impact to both their attacking and defensive parameters, as Figure 4.3 notes.

We now consider the teams who changed managers over the season that have not already been discussed. As previously mentioned, Manchester United sacked their manager very late into the season, leaving not enough time for performances to dramatically change, so will not be considered. Additionally, Sunderland will also not be considered due to changing their manager very early on. This leaves 3 teams to contemplate the effects of a manager change; Swansea, Tottenham and West Brom, whose progressions are seen in Figure 4.4.

Considering Swansea first, they made a managerial change just after the transfer window closed, at time 24. Of note is that, when this change was made, their attacking power was at their worst at that particular point of the season, which may have prompted the change. There was a subsequent improvement in their attacking prowess when the

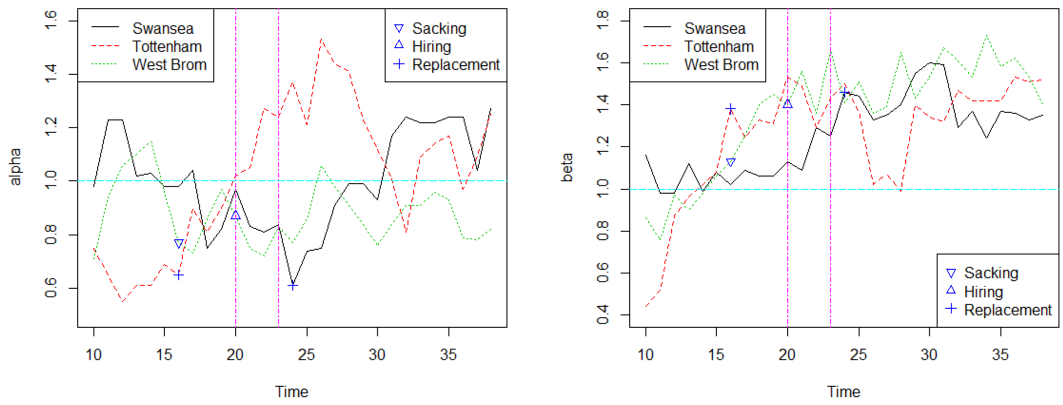

(a) α Parameters (b) β Parameters

Figure 4.4: Progression of α and β for Swansea, Tottenham and West Brom

change was made, which may have been down to the manager, but may also be due to other potential factors. On the other hand, Swansea's defensive abilities did seem to get marginally worse when the new manager came in, though it did get closer to their defensive abilities earlier on in the season towards the end.

Tottenham decided before the halfway point of the season that there needed to be a managerial replacement to turn around their fortunes. Like Swansea, the main benefit of this change was a significant improvement in attacking ability, particularly as Tottenham's attack was well below average at the start of the season. This changed after the beginning of the transfer window, where Tottenham's attack was always above average, apart from at one time point towards the end of the season. Another potential reason for the managerial change was that Tottenham's defensive abilities suddenly fell apart, which may have instigated a reaction. This didn't really improve after the new manager came in, though it is important to note that Tottenham were involved in some heavy defeats which may have skewed the parameters.

West Brom decided to sack their manager at around the same time that Tottenham did. Unlike Tottenham though, they hired a replacement a few weeks later rather than immediately. In terms of attacking, there didn't seem to be an effect of the new manager coming in, with their attack rising and falling by small amounts rather than a consistent improvement or regression. On the other hand, their defense did seem to become worse after the new manager came in, which may explain the relegation threat that West Brom were having. In the end, West Brom survived, but their manager did not survive in the job and was sacked at the end of the season.

One final comparison we can make is for the teams that occupied top of the table at some point over the course of the season. Five teams had this honour; Arsenal, Chelsea, Liverpool, Manchester City and Manchester United. Manchester United will be excluded from this comparison though, since they were only top of the table for 2 days at the very beginning of the season.

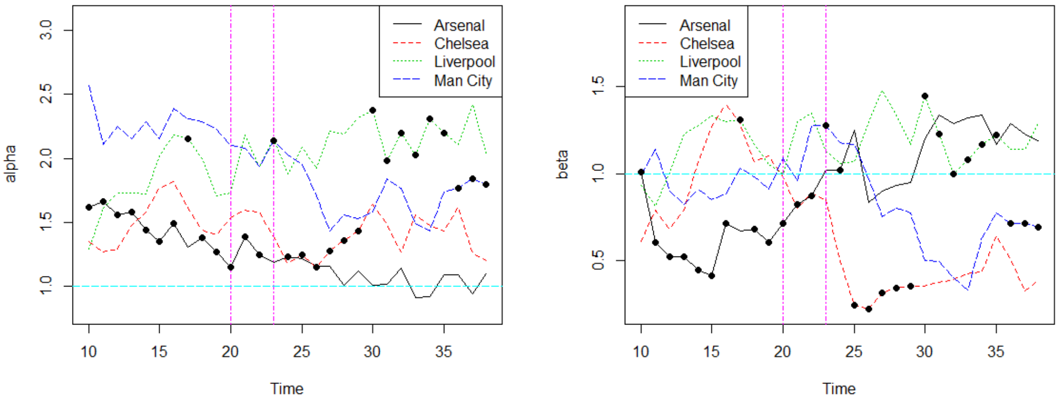

(a) α Parameters

(b) β Parameters

Figure 4.5: Progression of α and β for Arsenal, Chelsea, Liverpool and Manchester City

The black dots on each plot in Figure 4.5 indicate that particular team was top of the table at that time point. This suggests that the foundations of Arsenal being on top was based on their defense, with the best defensive parameter of the 4 teams for most of the first half of the season. Indeed, it appears that once this strong defensive foundation weakened, this resulted in Arsenal becoming undone in the race for the title, especially since their attack was, for the majority of the season, weaker than the other 3 teams at the top of the table.

Chelsea's main strength in the title race seemed to lie on having a superb defense in the second half of the season. Indeed, after the transfer window closed, no other team, besides Manchester City for a few weeks, had a better defensive ability than Chelsea. The problems Chelsea seemed to have are that their attack, while consistent, was never particularly outstanding, especially compared to Liverpool and Manchester City early in the season, and that their defensive ability in the first half of the season was not too great, at one point even being the worst of the 4 teams.

Liverpool's shot at the title was built on their offensive firepower. From time 11 onwards, their α parameter never fell below 1.5 and consistently hovered around 2, well above what the baseline attack power would be. Liverpool's main issue was their defensive abilities in comparison to the rest of the top 4. Rather alarmingly, Liverpool have the worst defensive parameter for the majority of time points over the season in comparison to the other teams. While this does get somewhat mitigated by their attacking power, looking at it in comparison to Manchester City towards the end of the season, where City's defense was a lot more efficient while still maintaining a high goal scoring threat, may explain why Liverpool's title challenge fell just short in the end.

Curiously, Manchester City spent the least amount of time at the top of the table compared to the other 3 teams, but they were up there when it mattered the most. Whilst City were very proficient in attack at the beginning of the season, they weren't on top, potentially due to their rather lacklustre defense. City's attacking power even suffered a blip directly after the transfer window, though it did improve slightly afterwards. However, while this blip was happening, City's defense was constantly improving, meaning they were letting in fewer goals while still scoring a good amount. This improvement, as well as getting their attack back on track, would culminate in Manchester City's second Premier League title.

4.2 Varying γ

4.2.1 Static Model

Similarly to when contemplating a common home effect across all teams, we first look at a static model implementing all of the results of the season with no time dependence, this time setting a parameter so that every team has their own home effect parameter.

Team	α	β	γ
Arsenal	1.41	0.94	1.09
Aston Villa	0.74	1.35	1.31
Cardiff	0.53	1.64	1.67
Chelsea	1.21	0.62	1.52
Crystal Palace	0.67	1.09	1.19
Everton	1.00	0.88	1.64
Fulham	0.74	1.93	1.49
Hull	0.89	1.19	1.08
Liverpool	2.12	1.21	1.11
Manchester City	1.68	0.87	1.64
Manchester United	1.52	0.98	0.85
Newcastle	0.90	1.34	1.13
Norwich	0.49	1.37	1.53
Southampton	0.96	1.06	1.46
Stoke	0.82	1.16	1.43
Sunderland	0.87	1.35	1.08
Swansea	0.95	1.23	1.54
Tottenham	1.10	1.18	1.22
West Brom	0.84	1.33	1.28
West Ham	0.65	1.12	1.69

Table 4.3: Parameter Estimates with Varying γ

Looking at Table 4.3, the β parameters remain relatively unchanged from the previous model since it covers the defensive abilities for both home and away games. Since we have now allowed the home effect γ to alter between teams, this impacts the attacking parameters for each team. The only team that is unchanged from before is Aston Villa, though this seems to only be by happenstance since their home effect parameter is near identical to the general home effect of the previous model. As to be expected, the vast majority of teams have a home effect greater than 1, since it is generally seen as advantageous to play at home. Interestingly, there is one team, Manchester United, whose home effect is less than 1, showing that they actually performed better away from home. It makes sense that this would be the case when looking at how they did in reality, since they lost seven games at home compared to five away and, of the teams they lost against, three of the teams that beat them at home finished below them compared to only one away, suggesting that the home parameter would take more of a hit because of this.

In general, the changes seem to be that the attacking parameter of each team will be increased from the previous model if their home effect is less than the general home effect parameter of the previous model. Similarly, if the attacking parameter has decreased from before, then the team's home effect will generally be greater than the previously defined home effect. This leads to, overall, not much change in positions when simulating 1000 seasons with this set of parameters, seen in Table 4.4.

Team	Simulated Position	Actual Position
Arsenal	4th	4th
Aston Villa	17th	15th
Cardiff	20th	20th
Chelsea	2nd	3rd
Crystal Palace	15th	11th
Everton	6th	5th
Fulham	19th	19th
Hull	12th	16th
Liverpool	3rd	2nd
Manchester City	1st	1st
Manchester United	5th	7th
Newcastle	14th	10th
Norwich	18th	18th
Southampton	7th	8th
Stoke	10th	9th
Sunderland	16th	14th
Swansea	9th	12th
Tottenham	8th	6th
West Brom	13th	17th
West Ham	11th	13th

Table 4.4: Simulated Positions for Varying γ

Looking at the simulated table, only 5 teams (Everton, Hull, Manchester United, Newcastle and West Brom) have changed position. Even then, the teams that have changed were very close in position for the previous model, which suggests differences may only be circumstantial.

Figure 4.6 shows a similar structure to Figure 4.1, besides the teams that have swapped their simulated positions. Again, there is a large amount of overlap in the amount of points teams acquired over the simulations, showing teams could potentially finish in a large variety of positions. A curious matter to note is that the outliers seem to have helped Hull leapfrog West Brom and Newcastle on the average position, since those three teams were very close in average points and Hull had a brilliant simulated season by their standards. By contrast, West Brom and Newcastle have not received this benefit, but

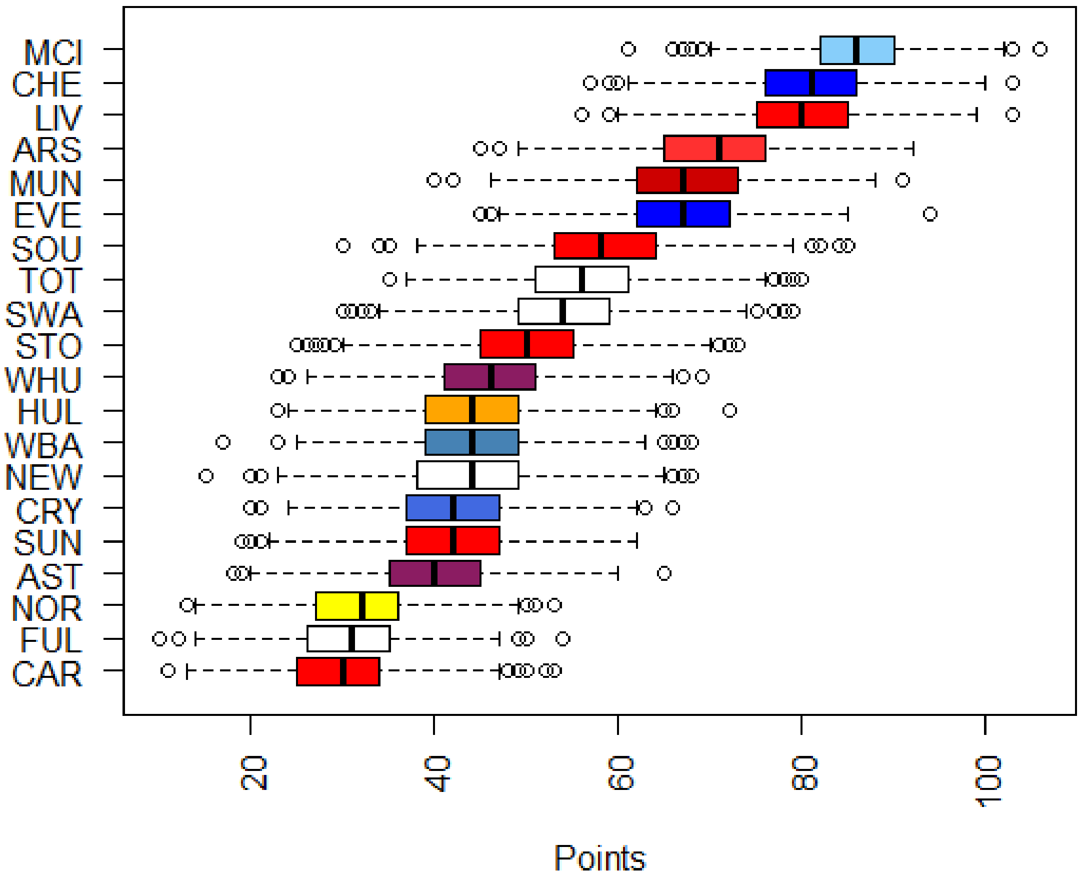

Figure 4.6: Boxplots of Points Scored by Teams over 1000 Simulated Seasons

rather they each had one simulated season where they performed much worse by their standards.

4.2.2 Dynamic Model

We can conduct slightly different investigations for this model. In addition to comparing attack and defense, since the β values do vary slightly between models over time, we can also compare home and away attacking power over the season. Again, we consider the relegated teams, the mid-table teams with drastically different forms and the teams at the top end of the table.

Viewing Figure 4.7, we look at how the relegated teams' attacking power both home and away evolve over the season. Cardiff's problems seem to stem from not being able to keep a consistent form going both home and away. Just before the transfer window, Cardiff's away attack was atrocious, far worse than what the other two relegated teams were managing at the time. While they were in the process of somewhat fixing this dur-

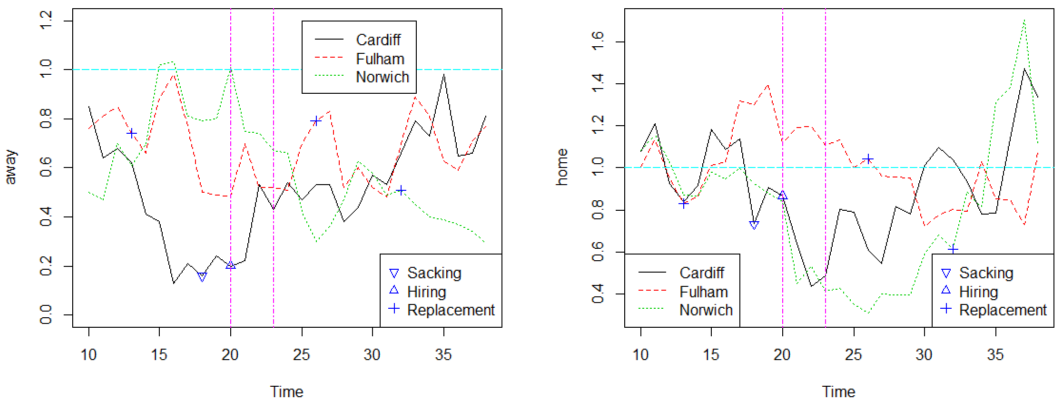

(a) Away Attack (b) Home Attack

Figure 4.7: Progression of Home and Away Form for Cardiff, Fulham and Norwich

ing and after the transfer window, their home attacking power started to suffer over this time, meaning they were still in just as much as trouble as they were beforehand. While both forms did improve towards the end of the season, they will still below par, leaving Cardiff with little chance of safety.

Looking at Fulham, their issue seemed to be scoring goals away, as well as the previously mentioned defensive mishaps. The α parameter for Fulham, the measure of away attack over the season, never exceeded the base value of 1, leading to consistent problems over the season. On the other hand, their home form was not as poor, apart from towards the end of the season where things seemed to be going wrong on both fronts. The main problem that Fulham appeared to have was touched upon in the previous model; that their defense was very poor in comparison to the other teams in the league.

Norwich seemed to have similar problems to Cardiff in that, at any point in the season, either their home or away attack was dreadful. Their capitulation looks to have began around the beginning of the transfer window, where their home attack seemed to be faltering and not really recovering until the tail end of the season. Unfortunately, this recovery seemed to coincide with their away form going awry, being much worse than the other 2 relegated teams and almost as bad as Cardiff's away woes earlier in the season. Norwich's inability to maintain consistency in their attack home and away may have led to their downfall.

As with the previous model, we make comparisons between two teams who had differing forms, but still finished close to each other; Crystal Palace and Newcastle. The plots showing how each teams' home and away attack progressed are shown in Figure 4.8.

As to be expected from our previous analysis, we see that Crystal Palace's attack is generally sub par for most of the season before improving towards the end, though it is important to note that Crystal Palace's success appeared to be built on their strong defensive abilities, shown in the analysis for the previous model. It should be noted that Crystal Palace seemed to maintain a consistent attacking ability both home and away

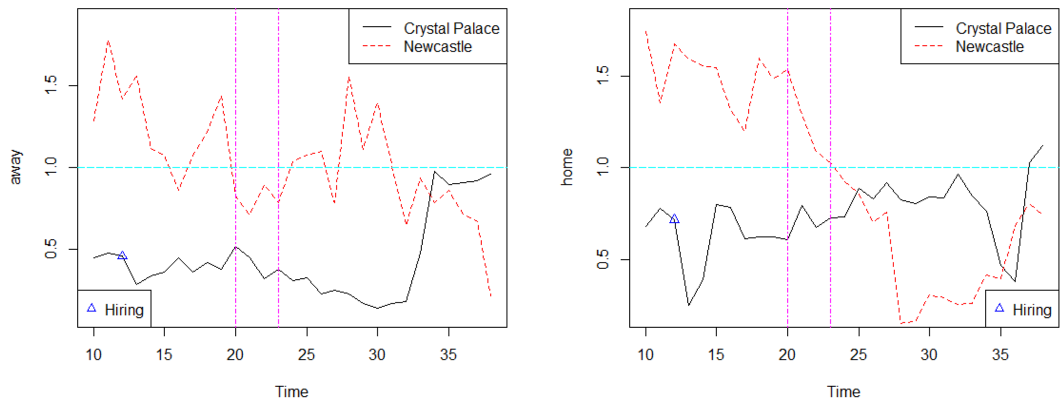

(a) Away Attack (b) Home Attack

Figure 4.8: Progression of Home and Away Form for Crystal Palace and Newcastle

over the season, though they were far better at home than away.

Newcastle's performance, generally, gets progressively worse as the season goes on. Curiously, between times 28 and 32, Newcastle's away parameter seems to be drastically better than their home parameter. Investigating, this is due to Newcastle having a couple of high scoring wins away from home in the games used at that time measurement, while in this same time period were on a barren run at home, having a mere 2 home goals in this time.

Again, we consider the teams who changed managers over the season who had a substantial amount of time with each manager; Swansea, Tottenham and West Brom. The home and away form over the season for each team are showcased in Figure 4.9.

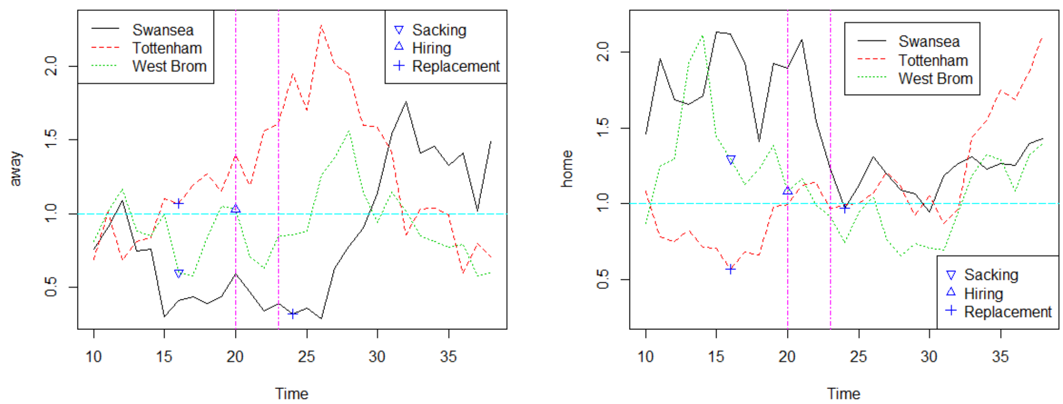

(a) Away Attack (b) Home Attack

Figure 4.9: Home and Away Form for Swansea, Tottenham and West Brom

Looking at the plots, Swansea's away attacking prowess was alarmingly lacking just before they decided to change manager. This did see an improvement after the managerial change after some time; it even exceeded what they had been managing at the start of the season and then stayed above that level. While Swansea's away form toiled, their home form was much better, though even this fell to a more average level towards the tail end of the first manager. While the home form didn't improve a great deal under the second manager, it at least stayed above average so, combined with an improving away attack, we see an overall improvement in attacking ability for Swansea.

Tottenham seemed to show an all around improvement in attacking both home and away after they changed their manager. Their away form in particular saw a huge improvement in the middle of the season, though it did fall just as fast as it rose towards the end of the season. This fall appears to be counteracted by their steady improvement in home attack after the new manager came in, reaching a very high value at the end of the season.

Investigating West Brom's home and away form shows a more worrying trend for those that are supporters of the club. West Brom's home form appears to drop off after appointing their new manager, though it does gradually pick up towards the end of the season. However, at this point, a new problem introduces itself in that their away form gets worse at this point after showing some sign of improvement after the transfer window. West Brom's lack of consistency to play well both home and away at the same time may have lead to their troubles of flirting with relegation.

Once more, we compare the teams who topped the table past the beginning of the season, namely Arsenal, Chelsea, Liverpool and Manchester City. The progression over time for the away and home attack are seen in Figure 4.10.

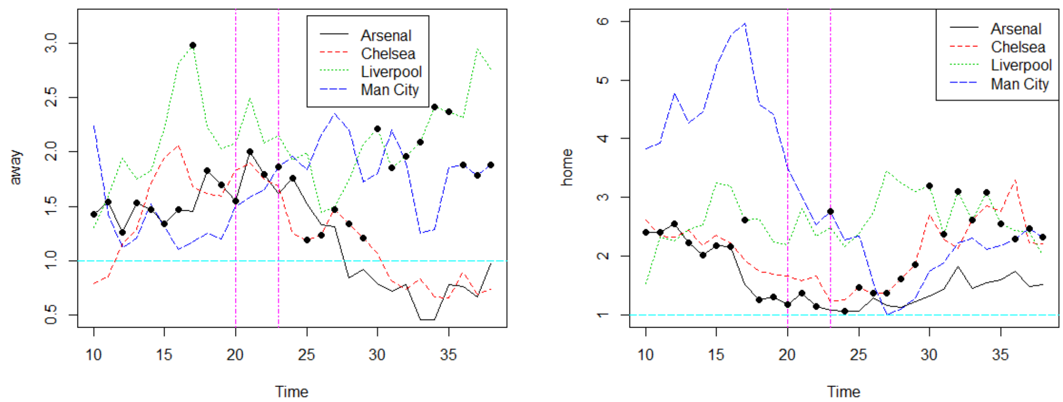

(a) Away Attack (b) Home Attack

Figure 4.10: Home and Away Form for Arsenal, Chelsea, Liverpool and Manchester City

An interesting thing to note is that none of the teams in this comparison ever had their home attack fall below 1, showing they were all above average in this department over the course of the season. Curiously, it was often not the case that the team with the best attacking parameters were top of the league, Liverpool providing most of the exceptions when they were top, plus one case where Manchester City were top with the best home attack of the 4 teams. This highlights the importance of having a good ability to stop goals being conceded as well as the ability to score goals.

Comparing the attacks show off how Arsenal's lack of attacking power compared to the other three teams towards the second half of the season derailed their campaign. Their home attack especially seemed to be an issue, with it being worst out of the other top 4 teams for the majority of the season. Arsenal's away attack also started to suffer after the transfer window and, combining this with their home form, this left them lagging behind the other teams in offensive firepower.

As noted when analysing the previous model, Chelsea's title challenge seemed to be mostly based off their defensive capabilities. Their away attacking abilities were never really prolific in comparison to the other top teams, especially towards the end of the season, whilst their home attacking abilities were only the benchmark for one very brief period of time towards the end of the season. Whilst their attacking abilities weren't bad, they never reached the heights that Liverpool and Manchester City did at points in the season.

Liverpool's challenge was built on their overwhelming attacking ability, as noted previously. They had the best away attack for most of the first half of the season, as well as towards the tail end of the season. Their home attack was also on top, for most of the second half of the season. Liverpool's main problem was not scoring goals, but defending them, as seen in the previous analysis.

Manchester City's attacking was outstanding at certain points in the season, especially their home attack in the first half of the season. At one point, their home attack parameter was at least double of every other team in the league. Despite this, they were still not on top, potentially to do with a combination of their away attack not being as good as the teams around them at this point and, seen in the previous analysis, their defense was not up to scratch either. Their home form also took a hit in the second half of the season just when they were improving their away form. When they obtained consistency in both home and away results, they took the top spot of the table and remained there when it counted.

4.3 Comparison Between Models

While the two models are similar in structure and produce similar output, there are some differences between the two models that we can consider. To begin with, we can work out the parameters for λ and μ for each potential match in the models that take into account all games over the season and make a note of any differences in the expected outcome there may be. It is also worth considering which teams seem to benefit from each model. For the purposes of this comparison, we define each λ and μ in the models as λ_C and μ_C for the model with common γ and as λ_V and μ_V for the varying γ model.

Running the comparison, we find that 14 out of the 380 matches have a difference in the expected outcome between models. These matches that show the differences can be found in Table 4.5.

Home Team	Away Team	λ_C	μ_C	λ_V	μ_V
Arsenal	Liverpool	2.036	1.814	1.860	1.993
Aston Villa	Stoke	1.128	1.161	1.125	1.107
Fulham	Aston Villa	1.402	1.428	1.489	1.428
Liverpool	Manchester City	2.207	2.311	2.047	2.033
Manchester United	Arsenal	1.495	1.254	1.214	1.382
Manchester United	Liverpool	1.925	1.891	1.563	2.078
Southampton	Everton	1.180	1.208	1.233	1.060
Southampton	Manchester United	1.314	1.283	1.374	1.611
Sunderland	Stoke	1.189	1.161	1.090	1.107
Swansea	Everton	1.203	1.391	1.287	1.230
Tottenham	Everton	1.226	1.345	1.181	1.180
West Brom	Swansea	1.331	1.383	1.322	1.264
West Ham	Southampton	1.045	1.142	1.164	1.075
West Ham	Tottenham	1.163	1.187	1.296	1.232

Table 4.5: Matches with Differences in Outcomes Between Models

An interesting thing to note is that only 6 teams out of the 20 are not affected by a difference in outcome between models, those teams being Cardiff, Chelsea, Crystal Palace, Hull, Newcastle and Norwich, though investigating each match shows that every match has a difference in values between models, and thus, a different sized advantage or disadvantage. One of the things this comparison seems to show is that Everton appear to benefit more when the model for a common γ between teams is used, being involved in 3 matches that they would otherwise lose on average in the varying γ model, though the difference is marginal. Comparing the two models and which matches they are, this makes sense since Everton's general attacking parameter drops to make room for their improved home effect parameter in the varying γ model and all 3 matches involved are when Everton is the away side. On the other hand, their Merseyside rivals Liverpool seem to benefit more from using the model where γ varies between teams, again being involved in 3 matches that would turn the tide in their favour if this model was used. Unlike Everton though, it is a mixture of home and away matches that Liverpool benefit from, though their sole home benefit may be more to do with Manchester City's slight deficiency in away attack compared to Liverpool's home attack, where their home effect parameter has regressed from the value for the common home effect in the other model. Noticeably, there has been a dramatic change between models for Manchester United against Liverpool, where a combination of Manchester United's well below par home effect and Liverpool's high attacking parameter leading to a convincing advantage for Liverpool in the varying γ model compared to the common γ model which gave a slight advantage to Manchester United.

We can consider the differences between the expected number of home goals for each game between the two models, along with a similar concept for the number of away goals. The plot for the home differences, with each home team segmented in alphabetical order, can be found in Figure 4.11.

A positive home difference corresponds to the home team benefiting from the varying

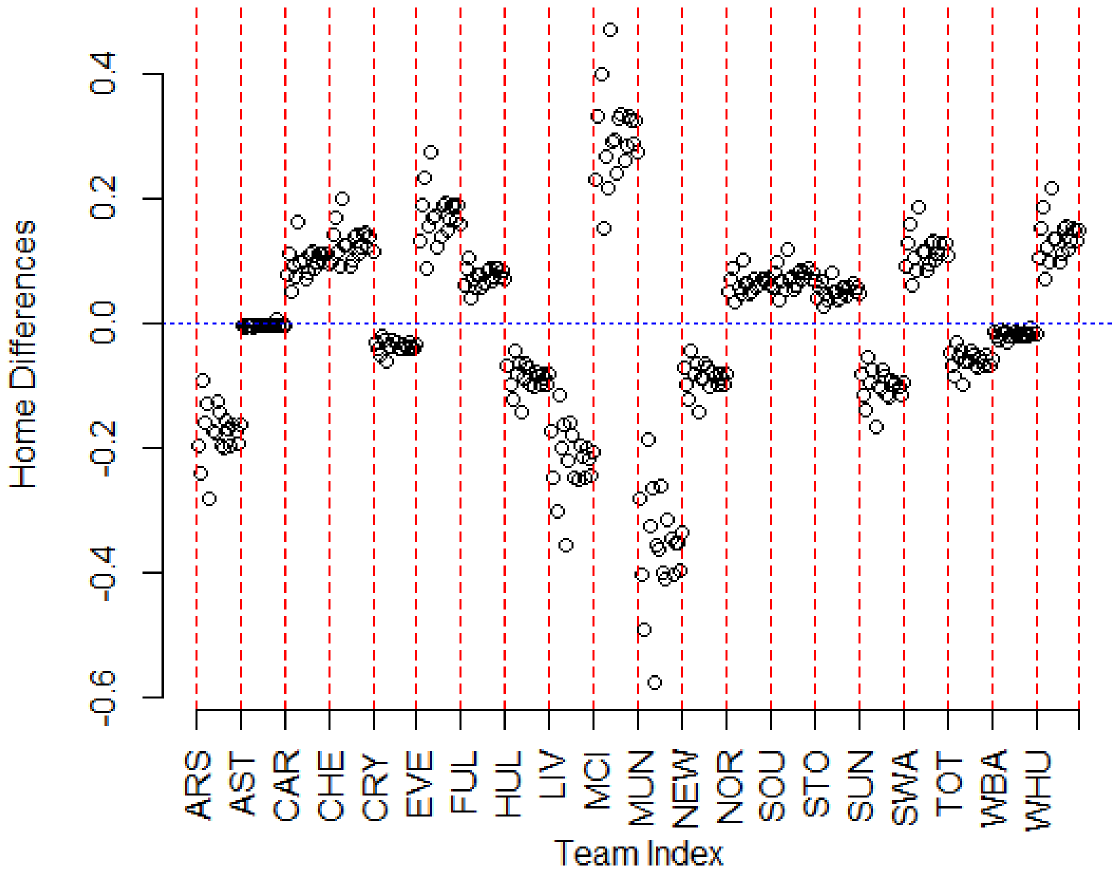

Figure 4.11: Home Differences Between Models

γ model, whilst a negative home difference means that the home team benefits more from the common γ model. The majority of teams have at least a slight benefit or hindrance at home from the use of the varying γ model, though teams like Aston Villa and West Brom see almost no change in home ability between models. The biggest beneficiaries of the varying γ model at home appear to be Manchester City, with only one match having a home difference between models of less than 0.2, a value which only 3 other teams (Chelsea, Everton and West Ham) manage to have at least one match above it. Teams who seem to benefit more from the common γ model when at home include Liverpool and Manchester United, with most of the differences in their home matches being negative.

We now consider the same concept, but for the number of away goals in a match. Again, we plot the differences, this time in alphabetical order of the away teams, in Figure 4.12.

Similarly to the home differences, a positive away difference shows a benefit for the away team for the varying γ model, whilst a negative away difference shows the team

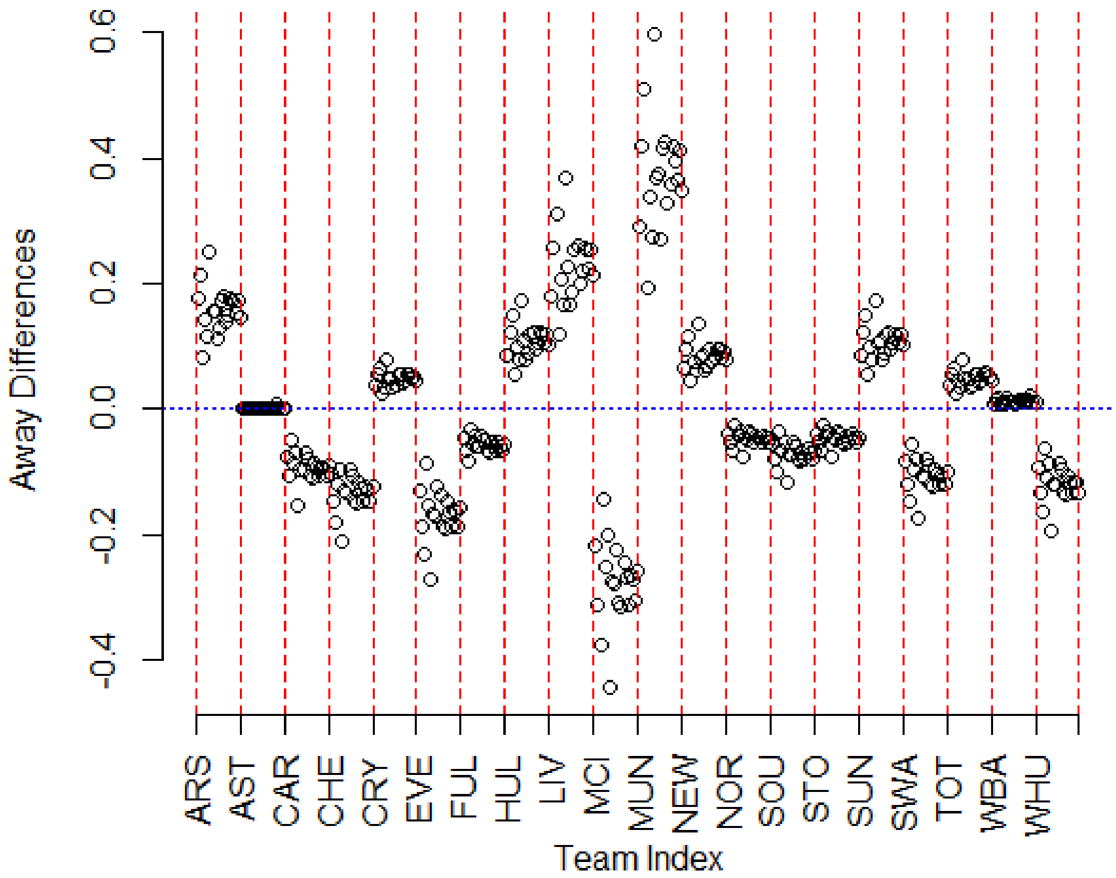

Figure 4.12: Away Differences Between Models

gains an advantage from the common γ model. Perhaps unsurprisingly, with the models giving similar end results, the teams that seem to benefit the most away with the varying γ model are the ones who receive the most detriment from the same model at home, such as Liverpool and Manchester United. Likewise, Manchester City seem to gain the most away when using the common γ model.

Another comparison we can make is a visual representation of the differences between the number of home and away goals, $\lambda_C - \mu_C$ and $\lambda_V - \mu_V$, for each match up in the respective models, seen in Figure 4.13.

A positive difference shows an expected home victory, whilst a negative difference corresponds to an expected away victory. The closer points are to the line, the less difference there is in that match numerically between models. The majority of games tend to have their differences between -1 and 2, showing the tight nature of the Premier League. Most points also tend to be close to the line, showcasing that a lot of matches have no significant difference numerically between models. Towards each of the tail ends, we see

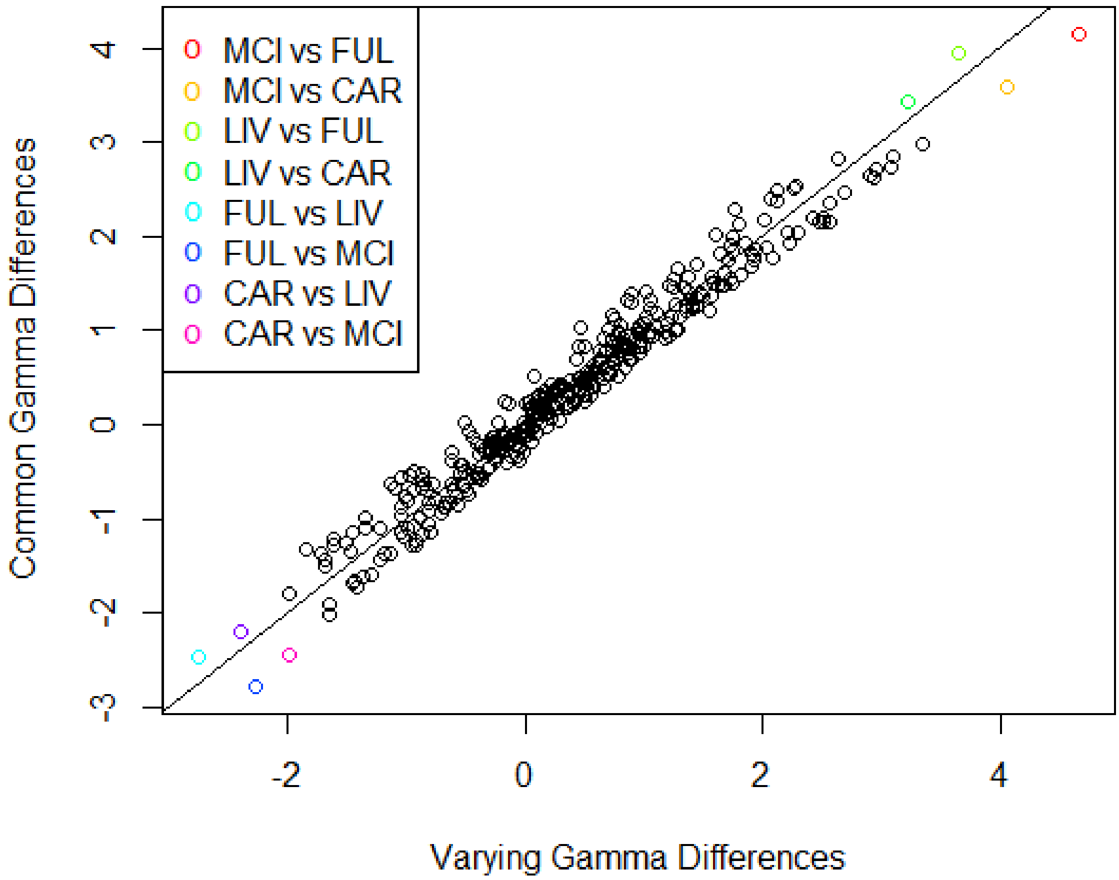

Figure 4.13: Comparison of Differences Between Models

games where there is a gulf in class between the two teams playing. Pulling out the 4 most extreme values at either end, we see that they only involve 4 different teams; Cardiff, Fulham, Liverpool and Manchester City. This makes sense considering that the match ups are when one of the top two teams play one of the bottom two, showcasing the contrast in abilities of the teams. Manchester City against Fulham is particularly noticeable, due to both models giving a difference between λ and μ of over 4, showcasing the high power of Manchester City's attack and the low defensive abilities of Fulham.

We can take another difference between the differences previously mentioned to see how much change is made between the models for each game, namely $(\lambda_V - \mu_V) - (\lambda_C - \mu_C)$. Taking the absolute values between these differences, we can visually compare them, seen in Figure 4.14.

Interestingly, the two matches with the highest differences between models involve the same teams; Liverpool and Manchester United, the match with Liverpool at home being the higher difference of the two. This may have happened due to the sharp differ-

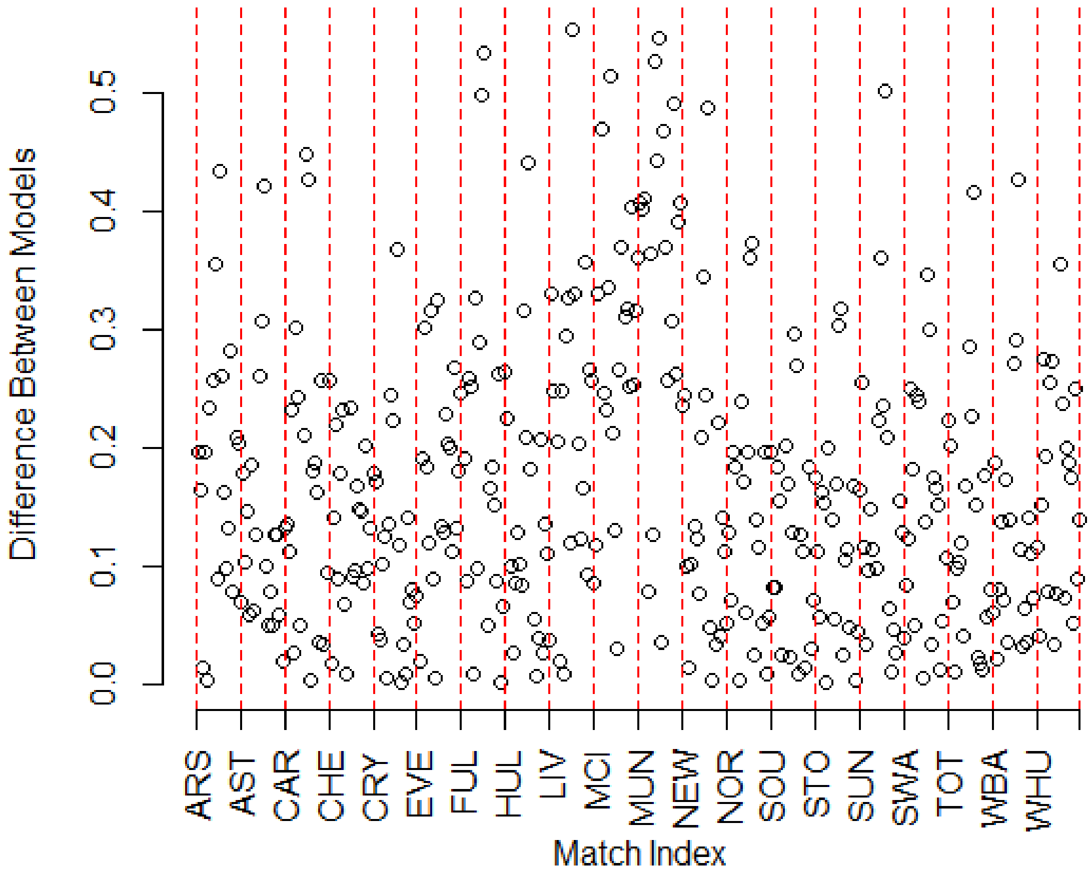

Figure 4.14: Differences in Parameters Between Models

ences made to each teams' attacking parameters between models, especially Manchester United's, who churn out a much better away performance in the varying γ model but a worse home performance in the same model. The match with the lowest change in difference was Stoke against Crystal Palace, with a change of only 0.000549, though the parameters in the varying γ model are noticeably higher. The majority of the changes in parameters between models are smaller than 0.3, suggesting that there isn't too much of a difference between models.

4.4 A Comparison Between Models and Expert Opinion

A comparison we can make to check how effective our models are is to use each model to predict results that have already happened and compare them to how a football expert predicted them at the time. To predict our results, we use the results of the ten preceding collection of matches prior to the collection of matches we wish to predict to estimate parameters and then calculate the probabilities of various scores occurring, the score with the highest probability being the predicted score of that game. We will be comparing our predictions with the predictions of football pundit Mark Lawrenson, who makes a prediction for each Premier League game every week for the BBC, looking at the results of the 2013/14 season [9]. Lawrenson also brings in a guest for all weeks of the season, so we will also compare our results to theirs.

We consider the same scoring system used on the BBC website for measuring how well each prediction has done, with 3 points given for the correct score, 1 is accrued for the correct result and no points are given for an incorrect result. It should be noted that Lawrenson's and the guest predictions are incomplete due to postponements of games arising from teams in the Premier League being involved in cup games, with the subsequent rearranged games not receiving new predictions. Some guests also did not make predictions for whatever reason for other games. We also start our predictions after every team has played 10 games so that the models can make predictions from each batch of games used for the parameters for each team.

The number of points earned for each medium of prediction can be seen in Table 4.6.

Varying γ	Common γ	Lawrenson	Guests
175 (280 games)	179 (280 games)	214 (274 games)	123 (262 games)

Table 4.6: Points Earned from Predictions

Both models have performed to roughly the same ability, which is to be expected since the 2 models give similar results for predicting the final league table. Curiously though, both models have performed vastly worse than what Lawrenson has. A potential explanation for this is that Lawrenson has a lot more information to go off when making his predictions, such as any injuries and suspensions a team has or indeed the opposite effect of a star player coming back. Our models on the other hand can only measure based on the number of goals each team scores and concedes over the previous ten fixtures. In spite of this, our models did perform better than the guests over the course of the season. It is important to note however that the guests have a wide range of footballing knowledge. Some guests may not follow football at all, whilst others may be former players who are now pundits.

It is also worth noting that, for a number of games, the difference between the highest predicted probabilities of scores were very small for quite different results for the two models. For example, in the game between Liverpool and Arsenal using the common γ model, the score with the highest predicted probability was 1-1, with a probability of 0.09643. The next highest probability however was only 0.004 smaller, with the predicted score corresponding to this probability being 2-1 to Liverpool. In reality, this game finished 5-1 to Liverpool, showing a very marginal difference between the model gaining a point or not gaining a point at all.

It is also worth investigating matches where the models and Lawrenson have predicted very different results, with one turning out to be correct in the end. One such example where the models were wrong but Lawrenson was correct was the game between Newcastle and Arsenal, with Arsenal winning 1-0. For this game, Lawrenson predicted a 2-0 Arsenal win, whereas both models predicted a 1-1 draw. A possible reason the models predicted a 1-1 draw may be due to rather skewed results that Newcastle had in their favour at the time, with a 5-1 victory against Stoke and a 3-0 victory against Crystal Palace standing out, whilst Arsenal had a more consistent form of constantly winning games, but by a small amount. While the models can only make their predictions based on goals scored, Lawrenson can view which games teams have won and make other inferences about the teams beyond the number of goals teams scored.

A match where the opposite situation occurred, in that the models predicted the correct result but Lawrenson did not, was the game between Crystal Palace and Manchester City late in the season, with Manchester City running out as 2-0 winners. While both models predicted a 1-0 Manchester City win, Lawrenson went for a 1-1 draw. This prediction may have been made due to Palace's recent form suggesting that they were very difficult to score against, while not necessarily winning games by large margins. Additionally, Palace had recently beat Chelsea, one of the other top teams in the league, so Lawrenson perhaps thought they could pull off a minor shock against Manchester City. This was not to be the case though, with Manchester City's prolific goalscoring ability shining through in the end.

Chapter 5

Applications to Betting

One potential use of the previously defined models is to try and predict the outcome of any given match between two teams. Using this knowledge, we can place bets on each match, whether it is simply the outcome of the match or the exact score of the match. We showcase this by making predictions on a recent weekend's Premier League games and placing a £10 bet on each game, for the scenarios of predicting outcome and predicting score. Each score prediction will be based on whatever the score with the highest probability is. All odds calculated are based on the odds Sky Bet offer and are correct as of the morning before the first match of the weekend.

5.1 Score Predictions

We consider the 2014/15 English Premier League games of the weekend commencing March 21. Using the time dependence of both models, where we only consider the results of the last ten games, our predictions for each of the weekend's games can be seen in Table 5.1.

Home Team	Away Team	Common γ	Odds	Varying γ	Odds
Manchester City	West Brom	1-0	13/2	1-0	13/2
Aston Villa	Swansea	1-1	9/2	1-1	9/2
Newcastle	Arsenal	0-2	7/1	0-2	7/1
Southampton	Burnley	1-0	5/1	0-0	10/1
Stoke	Crystal Palace	1-1	9/2	1-1	9/2
Tottenham	Leicester	2-1	7/1	2-2	14/1
West Ham	Sunderland	1-0	6/1	1-1	11/2
Liverpool	Manchester United	1-0	13/2	0-0	10/1
Hull	Chelsea	0-1	4/1	0-2	9/2
QPR	Everton	0-0	17/2	0-0	17/2

Table 5.1: Predictions for Games of Weekend of March 21

It is also worth considering the calculated probabilities that each of the scores will occur, to determine how "safe" of a bet the predicted scores are, seen in Table 5.2.

Generally, the calculated probabilities of the predicted scores happening fall between

Home Team	Away Team	Common γ Probabilities	Varying γ Probabilities
Manchester City	West Brom	0.173	0.157
Aston Villa	Swansea	0.133	0.148
Newcastle	Arsenal	0.123	0.136
Southampton	Burnley	0.204	0.216
Stoke	Crystal Palace	0.133	0.109
Tottenham	Leicester	0.085	0.064
West Ham	Sunderland	0.187	0.133
Liverpool	Manchester United	0.207	0.238
Hull	Chelsea	0.170	0.107
QPR	Everton	0.175	0.441

Table 5.2: Probabilities of Predicted Scores of Weekend of March 21

one tenth and two tenths, showing how difficult it can be to predict the exact score of a match. There are some particularly noteworthy probabilities to draw attention to though. For example, neither model seems confident in being able to predict the score of Tottenham vs Leicester, with probabilities of less than one tenth of that score occurring in both instances. Another probability that sticks out is the probability of QPR vs Everton finishing as a goalless draw being an unusually high value of 0.441 when the γ value is allowed to vary between teams. This is in stark contrast to the probability of this same result using a common γ between teams. One possible reason that this may have occurred is that, in the ten games being used for the data, Everton had scored a lot fewer goals away than at home (1 goal away compared to 6 at home), something which is not completely captured by the model using a common γ. Likewise, QPR had only scored 3 goals at home in the last ten games compared to 6 goals away in the same period. This would logically raise the probability of a goalless draw when considering home and away performance separately.

5.2 Outcome Predictions

While predicting scores can net a good return, they also have a greater risk due to the greater accuracy of prediction required. For this reason, we may only want to predict the outcome of a match rather than the exact score. This will lead to a lower potential return due to the odds being worse, though the risk of losing decreases. The odds for each match ending in either a home win, draw or away win can be seen in Table 5.3.

We can use these odds to infer Sky Bet's calculated probabilities for each outcome in each game. We do this by taking the inverse of each of the odds in a game, sum them and then normalise each of these inverted odds so we get a sum to 1. These calculated probabilities can be found in Table 5.4.

The probabilities of each outcome for each match according to both models are noted in Tables 5.5 and 5.6.

For the most part, the two tables predict the same winners for each game, with the

Home Team	Away Team	Home Win Odds	Draw Odds	Away Win Odds
Manchester City	West Brom	1/3	9/2	8/1
Aston Villa	Swansea	7/5	11/5	2/1
Newcastle	Arsenal	9/2	16/5	4/7
Southampton	Burnley	4/9	10/3	6/1
Stoke	Crystal Palace	23/20	11/5	5/2
Tottenham	Leicester	4/7	10/3	4/1
West Ham	Sunderland	17/20	5/2	18/5
Liverpool	Manchester United	11/10	12/5	12/5
Hull	Chelsea	13/2	10/3	4/9
QPR	Everton	23/10	23/10	6/5

Table 5.3: Odds for Matches of Weekend of March 21

Home Team	Away Team	Home Win	Draw	Away Win
Manchester City	West Brom	0.896	0.066	0.037
Aston Villa	Swansea	0.428	0.272	0.300
Newcastle	Arsenal	0.097	0.137	0.766
Southampton	Burnley	0.828	0.110	0.061
Stoke	Crystal Palace	0.504	0.263	0.232
Tottenham	Leicester	0.761	0.130	0.109
West Ham	Sunderland	0.634	0.216	0.150
Liverpool	Manchester United	0.522	0.239	0.239
Hull	Chelsea	0.057	0.111	0.832
QPR	Everton	0.255	0.255	0.489

Table 5.4: Sky Bet Probabilities for Matches of Weekend of March 21

occasions that the models predict a different winner being in cases where the differences between the probabilities were marginal. The only exception to this is the game between QPR and Everton, where the model with varying γ is more confident the game will be a draw compared to the other outcomes and the model with common γ, which is only marginally suggesting an Everton win. As previously alluded to, this may be due to the differing performances home and away of QPR and Everton. Another interesting case to point out is how different the probabilities are between the models for the game between Tottenham and Leicester. While both models do predict a Tottenham win based on outcome alone, the model for a common γ appears to have much more confidence for this outcome compared to the model for varying γ, which is stating a higher probability for Leicester winning. A potential reason for this can be found by considering the home and away form of Leicester in particular. The amount of goals Leicester scored over the ten match weeks was 7, though only one of them was at home. Under the varying γ model, this would leave Leicester's general attacking parameter at a decent amount, but leave their home parameter quite low. Indeed, the attacking parameter was 1.77, but their home effect was only 0.11. The home effect, naturally, would have no effect on Leicester's performance for their match since they were playing away. Since the model with common γ would combine Leicester's home and away performances, this has the effect of dulling

Home Team	Away Team	Home Win	Draw	Away Win
Manchester City	West Brom	0.573	0.264	0.163
Aston Villa	Swansea	0.364	0.286	0.350
Newcastle	Arsenal	0.061	0.123	0.816
Southampton	Burnley	0.613	0.263	0.124
Stoke	Crystal Palace	0.349	0.287	0.364
Tottenham	Leicester	0.733	0.147	0.120
West Ham	Sunderland	0.604	0.258	0.138
Liverpool	Manchester United	0.480	0.330	0.190
Hull	Chelsea	0.149	0.250	0.601
QPR	Everton	0.284	0.334	0.382

Table 5.5: Probabilities of Outcomes of Weekend of March 21 for Common γ

Home Team	Away Team	Home Win	Draw	Away Win
Manchester City	West Brom	0.469	0.271	0.260
Aston Villa	Swansea	0.330	0.316	0.354
Newcastle	Arsenal	0.101	0.211	0.688
Southampton	Burnley	0.454	0.365	0.181
Stoke	Crystal Palace	0.400	0.250	0.350
Tottenham	Leicester	0.456	0.190	0.354
West Ham	Sunderland	0.511	0.281	0.208
Liverpool	Manchester United	0.512	0.360	0.128
Hull	Chelsea	0.100	0.177	0.723
QPR	Everton	0.295	0.525	0.180

Table 5.6: Probabilities of Outcomes of Weekend of March 21 for Varying γ

Leicester's performance away, and hence, leaving their probability of winning as much lower.

An interesting point to note is that the highest probability of a certain score occurring is not necessarily going to coincide with the highest probability of a certain outcome. For example, in the Aston Villa and Swansea game, both models suggest that the score with the highest probability of occurring is 1-1. However, in terms of just the outcome, the common γ model suggests that an Aston Villa win is the most likely outcome, whilst the varying γ model infers the most likely outcome as a Swansea win. A similar situation also occurs for the game between Stoke and Crystal Palace.

A noticeable trait comparing the probabilities of the models with the probabilities calculated via the Sky Bet odds are that they, for the majority of the matches, put a tremendous weight on a team winning the game if they perceive that team to have a sizable advantage. This is especially noticeable for the game between Manchester City and West Brom, where the odds calculated a probability of of 0.896 for a Manchester City win compared to 0.573 and 0.469 for the same outcome in the two models. A similar trend is also noticeable for Southampton against Burnley, where the overwhelming probability goes for a Southampton win. Curiously, there is one match that goes against the trend, at least compared to the common γ model, that match being Newcastle against Arsenal, with the probability for an Arsenal win being 0.766 using the odds compared to 0.816 from the common γ model.

5.3 Results

The results for this set of fixtures are tabulated in Table 5.7.

Home Team	Away Team	Score
Manchester City	West Brom	3-0
Aston Villa	Swansea	0-1
Newcastle	Arsenal	1-2
Southampton	Burnley	2-0
Stoke	Crystal Palace	1-2
Tottenham	Leicester	4-3
West Ham	Sunderland	1-0
Liverpool	Manchester United	1-2
Hull	Chelsea	2-3
QPR	Everton	1-2

Table 5.7: Results for Weekend of March 21

Both models have performed rather poorly in predicting the scores of each game. Only one score was predicted correctly (West Ham against Sunderland for the common γ model), which showcases the difficulty in predicting correct scores since there are a lot of possible permutations. In both cases, we would have made an overall loss had we bet on all games, with a £30 loss for the common γ model and a £100 loss for the varying γ model.

In terms of predicting outcomes, both models have performed much better overall. If we were to bet on the outcome with the highest probability for each game, the common γ model would have 8 correct results, whilst the varying γ model would have 7 correct results. In both cases, we would have an overall profit of £49.13 and £22.13 respectively. Given the outcomes, our maximum possible profit would be £113.13 if we were able to predict all outcomes correctly. In this case, the common γ model has performed better than the varying γ model. We also see that Sky Bet predicted 7 correct results for the weekend's fixtures if we take the outcome with the highest probability for each game as their predicted outcome, a similar performance to the two models. If we were to bet £10 on each game on the outcome that Sky Bet deemed most likely, we would have made a profit of £14.13, a lower amount than both of the models. It is important to note however that there will be occasions where a lot of results don't follow what statistical models or the bookmakers predict, due to the nature of football and there being a wide variety of factors that can influence the outcome of one match.

Chapter 6

Conclusions

The preferential model to use depends on how much information we wish to extract, since both models have their advantages and disadvantages. While we get more information using the varying γ model, namely the differences in a team's attacking abilities home and away, we ultimately end up with similar outcomes from both models. Due to this, should we just want an end result and consider the basic properties of each team, the common γ would be the preferred choice since we have less parameters to deal with. If we did want to consider home and away form separately though, the varying γ model would be viewed as the better choice.

There are some improvements that could be made to both models. For example, as seen in the simulations, a team's positioning is based mostly on the number of goals scored and conceded. While this does give a decent structure to a team's ability, it is limited by the fact that, no matter how many goals a team scores, the points awarded don't change within the same outcome. A 7-0 win is worth the same in points as a 1-0 win, so if a team only seems to win games by a large margin, but gets beat more often by small margins, this would overestimate their abilities. Likewise, if a team only wins games by a narrow margin but lose games quite heavily, this would end up underestimating their abilities. For this, we may have to define another parameter, especially if a time dependence is being used, of the number of points a team has at that point in the season to get a slightly better idea of the team's ability to win games and their ability to not lose games.

We may also wish to extend our modelling to include multiple divisions, for the purposes of comparing the abilities of each division so we can have consistent parameters factoring in promotion and relegation should we wish to consider results across multiple seasons. These may be used to help predict outcomes for any cross-divisional match that may happen, such as a cup match.

In terms of betting, due to the uncertain nature of it, it is a lot more difficult to make a decision on which model is the better choice to use. A much larger sample of games would need to be taken before a more accurate judgment could be made. It is also impractical to base predictions purely on modelling attack and defense alone, as there are a high number of factors outside of this to consider, which we may wish to implement into future models. A few of these factors are qualitative, such as the appointment of a new manager and the injury or suspension of a key player, whilst others are quantitative, like the number of points the team has at a certain time in the season previously mentioned.

Bibliography

[1] http://www.sportingintelligence.com/2014/07/21/man-united-and-liverpool-remain-top-tv-draws-despite-2013-14-without-trophies-210703/ Retrieved 01/04/2015

[2] C. Reep and B. Benjamin, *Skill and Chance in Association Football*, 1968

[3] D. Karlis and I. Ntzoufras, *Analysis of sports data by using bivariate Poisson models*, 2003

[4] G. Baio and M. A. Blangiardo, *Bayesian hierarchical model for the prediction of football results*, 2010

[5] L. M. Hvattum and H. Arntzen, *Using ELO ratings for match result prediction in association football*, 2010

[6] M. J. Maher, *Modelling association football scores*, 1982

[7] M. J. Dixon and S. G. Coles, *Modelling Association Football Scores and Inefficiencies in the Football Betting Market*, 1997

[8] http://www.rsssf.com/tablese/eng2014.html Retrieved 20/04/2015

[9] http://www.bbc.co.uk/sport/0/football/27272476 Retrieved 22/04/2015