



Citation for published version:

Curran, T, Brady, A & Standage, M 2018, *The Luck Index*. University of Bath.

Publication date:

2018

Document Version

Publisher's PDF, also known as Version of record

[Link to publication](#)

University of Bath

Alternative formats

If you require this document in an alternative format, please contact:
openaccess@bath.ac.uk

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.



UNIVERSITY OF
BATH

Department for Health
Faculty of Humanities and Social Sciences

The Luck Index
Research Report

A report prepared for ESPN Sports Media Limited

by

Dr Thomas Curran, Aoibheann Brady, and Professor Martyn Standage
University of Bath

Correspondence:

Dr Thomas Curran
Department for Health
University of Bath
Claverton Down
Bath, UK
BA2 7AY

Tel: +44 (0)1225 383047
Email: t.curran@bath.ac.uk

CONTENTS

Acknowledgements	3
Executive Summary	4
1. Introduction	6
2. Methodology	6
3. Key Findings	13
4. References	15

ACKNOWLEDGEMENTS

There are two main parties that we wish to thank for supporting this work. First, we would like to thank the coders who took part in this project and the University of Bath for facilitating their involvement. Second, thanks are due to the research and marketing team at ESPN Sports Media Limited for initiating and supporting the project and to The Promotions Factory for providing much needed guidance on communications. Any errors found in this document are our own.

EXECUTIVE SUMMARY

The present report summaries the findings of a project commissioned by ESPN Sports Media Limited to inaugurate the Luck Index. Here, our schedule of work specifically determined, using a statistical method known as a Bayesian hierarchical model (Baio & Blangiardo, 2010), an index of luck for the Premier League season 2017/18.

The project had two objectives:

- 1) Consultation: Engage in an in-depth consultation period with ESPN to arrive at a number of agreed factors that can be deemed as good or bad luck (e.g., deflected goals, erroneous decisions, etc)
- 2) Data coding: Train student coders to view Premier League footage for the luck factors, interpret, and enter this into a database.
- 3) Data analysis: For each game identified as having a luck incident, model the expected outcome from certain parameters associated with the game and teams.
- 4) League table: Redraw the 2017/18 Premier League table and derive a luck index from this analysis.

A statistical construction of a Bayesian hierarchal model was employed to facilitate these objectives. This process comprises quantitative data collection and analyses. Data were collected via a combination of: (a) data coding for freely available Premier League footage; and (b) established constants for home advantage, team strength, red cards, and penalty conversion. The collected data were analysed using programming language in the R STAN package.

Key Findings

Huddersfield go down on goal difference by a single goal and Stoke stay up

Our model estimated that Stoke would have stayed up in the 2017/18 season had luck not been a factor. This would have come at the expense of Huddersfield who would instead have been relegated. Such was the closeness of this result, a difference of 1 goal on goal difference, we modelled several times with the finding remaining constant. Given the vast financial ramifications of relegation, our finding has quite substantial consequences for both teams, which runs into the millions of pounds.

Liverpool were the unluckiest team

Liverpool were denied 12 points in the 2017/18 Premier League season due to luck. This is a substantial points loss, one that would have seen them elevated to 2nd in the Premier League table had luck not been a factor. The resulting loss of 12 points also meant that they were unluckiest team according to our luck index.

Manchester United were the luckiest team

Manchester United, in contrast to Liverpool, were the luckiest team according to our luck index. They would have been 6 points worse off had luck not been a factor in the 2017/18 Premier League season. Without these points, Manchester United would have finished 4th in the Premier League behind Manchester City, Liverpool, and Tottenham.

Arsenal should have been 11 points better off away from home

Arsenal were the unluckiest team away from home. They would have been 11 points better off had luck not been a factor in the 2017/18 season. This tally was counteracted by some points that they picked up due to good luck at home but, regardless, Arsenal fan can be aggrieved that their team appears to have bad fortune on the road. This did not alter their position in terms of Champions League places, but did elevate them above Chelsea to 5th which would have afforded “bragging rights” within the capital.

Although Burnley had a great season, luck played a part

The surprise package for many this year were Burnley, who finished 7th. Undoubtedly a great season for the Clarets, yet luck also played a role in their success. Burnley were the second luckiest team in our analysis and should have been 4 points worse off according to our model. It should be pointed out that experiencing this good fortunate did not affect their 7th position finish.

1. INTRODUCTION

The present report summarizes the findings of a project commissioned by ESPN Sports Media Limited to inaugurate the Luck Index. Here, our schedule of work specifically determined, using a statistical method known as a Bayesian hierarchical model (Baio & Blangiardo, 2010), an index of luck for the Premier League season 2017/18.

The project had two objectives:

- 1) Consultation: Engage in an in-depth consultation period with ESPN to arrive at a number of factors that can be deemed as good or bad luck (e.g., deflected goals, erroneous decisions, etc)
- 2) Data coding: Train student coders to view Premier League footage for the luck factors, interpret, and enter this into a database.
- 3) Data analysis: For each game identified as having a luck incident, model the expected outcome from certain parameters associated with the game and teams. Then redraw the 2017/18 Premier League table based on these estimations.

2. METHODOLOGY

The project employed a period of consultation to derive at a set of indicators deemed to make up what we collectively considered to be lucky/unlucky. From here, we coded Premier League game highlights for these indicators and analysed the data. The analytical methodology is formally known as a Bayesian hierarchical model (Rovan, 2014). These models provide an extremely useful tool for the estimation of data that occur in the form of “events” such as scorelines because they can combine fixed (e.g., home advantage, penalty constants) and random effects (e.g., team-level defensive/attacking abilities) to arrive at a posterior distribution of likely outcomes that can be simulated hundreds of thousands of times (Baio & Blangiardo, 2010).

In being able to integrate a large amount of diverse information, Bayesian hierarchical models are a valuable analytic tool for our project that seeks to redraw the Premier League based on luck. This said, the construction of a Bayesian model for ‘in-game’ estimation is not straightforward, and can be subject to misinterpretation and/or manipulation. Naturally, then, questions of the accuracy, reliability, and appropriateness of our methods need to be addressed fully. In what

follows, we detail each step of our consultation, coding, and analysis to document the procedures that took place to establish ESPN's luck index.

Objective 1: Engage in an in-depth consultation period with ESPN to collectively arrive at factors that can be deemed as good or bad luck (e.g., deflected goals, erroneous decisions, etc).

In May 2018, the research team embarked on a period of consultation with ESPN to determine the incidents that would make up our coded luck factors. Here, an initial meeting was organised to discuss how we would define luck and what incidents we would code. This meeting took place with key staff involved in the project from the University of Bath and ESPN. Following this, in June 2015, the research team were in email communication to confirm a list of coded incidents, which included:

- 1) Incorrectly awarded goal (offside, foul in build-up, etc)
- 2) Incorrectly disallowed goal (offside, handball that was not, etc.)
- 3) Incorrectly awarded and converted penalty
- 4) Incorrectly awarded and converted free kick
- 5) Penalty not awarded that should have been
- 6) Incorrect red card
- 7) Red card that should have been
- 8) Goal scored outside of allotted time
- 9) Deflected goals

Objective 2: Train student coders to view Premier League footage for the luck factors, interpret, and enter this information into a database.

In June 2015 a research team including the first author trained 3 student coders to view and code Premier League footage (highlight reel) for the factors outlined in Objective 1. Each of these coders were instructed to code for these decisions in a diligent and systematic manner. If 2 of the 3 coders coded an incident, then the game that incident belonged to went forward for analysis. Our coders had training in game analysis via being trained referees or working in the industry. A random sub-sample of our coded incidents was sent to a professional Premier League referee for ratification. After this process, a total of 157 games (out of 380) had incidents and were carried forward for analysis. These incidents are contained in a spreadsheet that is available on

request from the first author.

Objective 3: For each game identified as having a luck incident, model the expected outcome from certain parameters associated with the game and teams. Then redraw the premier league table based on these estimations.

The aim of our analysis was to predict the scorelines for each of the games identified in our coding process. We first used results data for the 2017/18 season to work out scoring rates for a pair of teams (i.e. a match) given home advantage, and each team's attacking and defensive abilities — all of which will be estimated from observed data. We then adjusted the scoring rate based on type of incident and time of incident. These rates were used to project the scores from the time of incident, assuming no luck. A final score was then calculated. These scores were used to calculate the points and goal adjustments to the Premier League table, and a new table order was calculated based on these adjustments. Finally, the luck of the team was calculated with a Luck Index. Several steps are needed in this analysis.

Step 1: Bayesian hierarchical model to infer team scoring rates

We first needed to infer the latent parameters (each team's attacking and defensive strength) that were generating the observed data (i.e., the scorelines). From a measurement perspective, knowledge of scorelines provide quite crude assessments of team strength. Thus, we made use of a Bayesian model to help to quantify uncertainty about these parameters. Here, we followed (and amended and refined) the approach of Baio and Blangiardo (2010), which proposed a Bayesian hierarchical model for the prediction of football results in Serie A. This model suggests that elements of the vector of observed scores $y = (y_{g1}, y_{g2})$ are modelled as independent Poisson:

$$y_{gj} | \theta_{gj} \sim \text{Poisson}(\theta_{gj})$$

Conditionally on parameters θ_{gj} , which represent the scoring intensity in the g^{th} game for the team playing at home ($j = 1$) and visiting ($j = 2$), respectively. Parameters are then modelled according to a log-linear random effect model:

$$\log(\theta_{g1}) = \text{home} + \text{att}_{h(g)} + \text{def}_{a(g)} \quad \log(\theta_{g2}) = \text{att}_{a(g)} + \text{def}_{h(g)}$$

The home parameter represents home advantage, which we assume is constant for all teams and throughout the season. Additionally, the scoring intensity is determined jointly by the attack ("att") and defence ("def") abilities of the two teams involved. The nested indexes identify the team that is playing at home (visiting) in the g^{th} game of the season.

We fit the model through a tool called Stan, which carries out Bayesian inference using Monte Carlo techniques. The output of this provides posterior distributions for the home, attack and defence parameters, along with the home and away scoring rates θ .

Step 2: Initial scoring rates

Once we ran our Bayesian model, we had expected scoring rates for teams and games throughout the season. However, there are a number of games with luck incidents, and we wished to project the scoreline of these games from time of incident onwards. This means that we wanted to preserve the score at time of incident, then work out the scoring rate for each team for the remainder of the game and the corresponding goal additions. These were added to work out the final predicted match result had the incident(s) not happened. As a result, we were able to work out both the score and points differential for any given match, and for each team for the full season. Finally, we used this to re-draw the 2017-18 Premier League table to account for these adjustments.

Firstly, we extracted the simulated θ values. We ran the model for a total of 4 sets of 100,000 iterations. We then have a large sample for each θ , both home and away, for each game. The average of these was taken for the team's expected scoring rate for a given game.

Step 3.1: Scoring rate adjustments

We then needed to calculate scores for our incident match data. This consists of 157 games throughout the season, with incidents including penalties unfairly not being awarded, red cards that should have been awarded, etc. These all had their own individual adjustments that needed to be made to the scoring rate, which will be discussed in this section.

We first calculated a new rate based on time of incident, using the formula:

$$\theta_{new} = \frac{\theta_{old} * (97 - \text{TOI})}{97},$$

where TOI represents the time (in minutes) of incident. We then projected the score forward, adjusting according to type of incident. We assumed a 97 minute game as the maximum of all match lengths. This is a reasonable assumption with added time taken into account.

Step 3.2: Red card adjustments

According to Vecer and colleagues (2009), the team which drops to 10 men should expect their scoring intensity to drop approximately $2/3$ of the intensity prior to the red card, and the scoring intensity of the opposing team should increase to approximately $5/4$ of the prior intensity.

Therefore within this work, we simulated scores according to the Poisson distribution, as suggested above. This is calculated with rate θ_{new} as shown above. We generated large samples (100,000) to be sure of our estimates in this case.

Step 3.3: Penalty adjustments

The average penalty scoring rate across all games in the 2017-18 season was 0.8, thus we added on 1 goal to the team who should have been awarded a penalty with probability 0.8. This means that the team scores the penalty 0.8 of the time across these 100,000 simulations.

Step 3.4: Disallowed goals adjustment

Disallowed goals were immediately added back on to the projected score through a +1 adjustment if the incident was deemed unfair. This led to a final score adjustment.

Step 3.5: Deflected goals

If a team scored a deflected goal, we modelled from time of incident (i.e., just before they scored) with the rate only adjusted for time. This means that, on the whole, they are unlikely to score this goal again at that timepoint. This is worth consideration as it impacts 41 games out

of the 157.

Step 4: Final score calculations

Now we turn to constructing the final scores for each of these 157 games. Using the rates calculated above, we worked out the scores. As the rates originally came from our Bayesian hierarchical model, they take into account home advantage and each team's attacking and defensive capabilities. We calculated the final score 100,000 times for each of these 157 games, using the adjusted scoring rates calculated above, and adding on the score at time of incident for both home and away teams.

To get an idea of what the model says the end result should be in each case, we work out the probabilities of a home win, draw, or an away win. This can be seen in a Match Probabilities spreadsheet, which is available on request from the first author. As an example, see the histogram in Figure 1 below.




Figure 1: A plot of the histograms of expected scores for Burnley (red) and Chelsea (blue) over 100,000 simulations

This is a histogram of the density of scores (home and away) in the first incident game (Chelsea vs Burnley). We can see that Burnley (red in this case) have a much higher probability of scoring zero goals than any other result (probability of approximately 0.6). On the other hand, Chelsea (in blue) have a much higher probability of scoring greater than zero goals than scoring no goals at all.

We now worked out the *median* score for each game from the 100,000 simulations, as this represents the most likely outcome for that match. For each of the home and away sides respectively, this is the value that occurs most frequently in the 100,000 score simulations. We then made the goal adjustments to the table, by working out the difference in the model's final

score and the original final score. Using this, we worked out the type of original final result - home win/loss, away win/loss or draw - from the perspective of the home and away teams. We calculated the original points the team had for this game so we could work out the points adjustment to the table, and the corresponding adjusted results type. We calculated the original points that the team had for this game to make the adjustment.

Step 5: Redrawing the table

We then adjusted the points and goals for each of the 157 incident games. We needed to work out the home away goal and points adjustment by team. This just involved summing up all the home away goal adjustments, and home away points adjustments for each team. We read in the original Premier League table with positions, goal differences and points, and add the points and goal adjustments. Finally, we re-ordered the 2017/18 Premier League table according to the new points totals for each team. The new table can be seen in the Updated Table spreadsheet, available on request from the first author.

Step 6: Luck index calculations

The luck index is calculated as the weighted sum of the points a team missed out on due to unfair decisions and the goal adjustment as calculated by the model, divided by the total number of luck incidents the team was involved in:

$$\text{Luck index}_i = \frac{(0.8) \times (-\text{Points adjustment})_i + (0.2) \times (\text{Goal adjustment})_i}{(\text{Total incidents})_i}$$

3. KEY FINDINGS

The 'Luck Index' table

Original Position	Luck Index Position	Team	GoalDiff	Points	Total Incidents	Points Adjustment	Goal Adjustment	Adjusted GoalDiff	Adjusted Final Points	Adjusted Position	Luck Index
1	1	Man City	79	100	19	-3	-13	66	97	1	-0.01
4	2	Liverpool	46	75	14	12	-3	43	87	2	-0.73
3	3	Tottenham	38	77	21	0	-10	28	77	3	-0.10
2	4	Man United	40	81	16	-6	-4	36	75	4	0.25
6	5	Arsenal	23	63	15	8	-3	20	71	5	-0.47
5	6	Chelsea	24	70	13	0	-7	17	70	6	-0.11
7	7	Burnley	-3	54	13	-4	-2	-5	50	7	0.22
10	8	Newcastle	-8	44	18	4	0	-8	48	8	-0.18
15	9	Brighton	-20	40	15	6	0	-20	46	9	-0.32
8	10	Everton	-14	49	11	-5	-10	-24	44	10	0.18
11	11	Crystal Palace	-10	44	15	-2	-4	-14	42	11	0.05
13	12	West Ham	-20	42	14	-1	-6	-26	41	12	-0.03
14	13	Watford	-20	41	16	0	-7	-27	41	13	-0.09
9	14	Leicester	-4	47	21	-7	-7	-11	40	14	0.20
17	15	Southampton	-19	36	13	4	-5	-24	40	15	-0.32
12	16	Bournemouth	-16	44	17	-6	-6	-22	38	16	0.21
19	17	Stoke	-33	33	14	4	-3	-36	37	17	-0.27
16	18	Huddersfield	-30	37	17	0	-8	-38	37	18	-0.09
18	19	Swansea	-28	33	14	1	-3	-31	34	19	-0.10
20	20	West Brom	-25	31	18	2	-6	-31	33	20	-0.16

Figure 2. Luck index – the redrawn premier league table from our analyses

Huddersfield go down on goal difference by a single goal and Stoke stay up

Our model estimated that Stoke would have stayed up in the 2017/18 season had luck not been a factor. This would have come at the expense of Huddersfield who would instead have been relegated. Such was the closeness of this result, a difference of 1 goal on goal difference, we modelled several times with the finding remaining constant. Given the vast financial ramifications of relegation, our finding has quite substantial consequences for both teams, one that runs into millions of pounds.

Liverpool were the unluckiest team

Liverpool were denied 12 points in the 2017/18 Premier League season due to luck. This is a substantial points loss, which would have seen them elevated to 2nd in the Premier League table

had luck not been a factor. The result also meant that they were unluckiest team according to our luck index.

Manchester United were the luckiest team

Manchester United, in contrast to Liverpool, were the luckiest team according to our luck index. They would have been 6 points worse off had luck not been a factor in the 2017/18 premier League season. Without these points, Manchester United would have finished 4th in the premier league behind Manchester City, Liverpool, and Tottenham.

Arsenal should have been 11 points better off away from home

Arsenal were the unluckiest team away from home. They would have been 11 points better off had luck not been a factor in the 2017/18 season. This tally was counteracted by some points they picked up due to luck at home but, regardless, Arsenal fan can be aggrieved that their team appears to have bad fortune on the road. This did not alter their position in terms of Champions League places, but did elevate them above Chelsea to 5th.

Although Burnley had a great season, luck played a part

The surprise package for many this year were Burnley, who finished 7th. Undoubtedly a great season for the Clarets, yet luck also played a role in their success. Burnley were the second luckiest team in our analysis and should have been 4 points worse off according to our model. It should be pointed out that experiencing this good fortunate did not affect their 7th position finish.

Case study – Liverpool

We made sure that our findings are reasonable by looking at a case study. Liverpool gained the most from this approach, gaining a total of 12 points, which appears to be a large jump, so we look at Liverpool's contentious games. In three of their home games, there were score adjustments in going from a draw to a win in each case. In the first of these, Liverpool were drawing with Manchester United. The final actual result was a draw, but Liverpool were not awarded a penalty that should have been awarded. The model predicts a Liverpool win (+2 points). The same occurred with the Burnley game (+2 points). Finally, in the third of these games the opposition scored an outside of injury time goal to secure a draw. The model suggests a Liverpool victory (+2 points). In their away games, they were initially beating Watford who unfairly scored a goal. The final result was a draw, but the model suggests a Liverpool victory

(+2 points). In another game, Arsenal scored a deflected goal to draw, but projecting from time of incident gives a Liverpool win (+2 points). Finally, in the last of these, Liverpool should have been awarded a penalty, model predicts win not draw (+2 points). These add to a total score adjustment of +12 points for Liverpool, and a two-place jump in the Premier League table.

4. REFERENCES

Baio, G. & Blangiardo, M. (2010). Bayesian hierarchical model for the prediction of football results. *Journal of Applied Statistics*, 37, 253–264.

Carpenter, B., Gelman, A., Hoffman, M. D., Lee, D., Goodrich, B., Betancourt, M., ... & Riddell, A. (2017). Stan: A probabilistic programming language. *Journal of statistical software*, 76.

Vecer, J., Kopriva, F., & Ichiba, T. (2009). Estimating the effect of the red card in soccer: when to commit an offense in exchange for preventing a goal opportunity. *Journal of Quantitative Analysis in Sports*, 5