



OPEN ACCESS

EDITED BY

Korkut Ulucan,
Marmara University, Türkiye

REVIEWED BY

Diego Nascimento,
Universidad de Atacama, Chile
Paulo Henrique Ferreira,
Federal University of Bahia (UFBA), Brazil
Keisuke Fujii,
Nagoya University, Japan

*CORRESPONDENCE

Oktay Karakuş
✉ karakuso@cardiff.ac.uk

RECEIVED 03 December 2023

ACCEPTED 16 May 2024

PUBLISHED 14 June 2024

CITATION

Scholtes A and Karakuş O (2024) Bayes-xG: player and position correction on expected goals (xG) using Bayesian hierarchical approach.
Front. Sports Act. Living 6:1348983.
doi: 10.3389/fspor.2024.1348983

COPYRIGHT

© 2024 Scholtes and Karakuş. This is an open-access article distributed under the terms of the [Creative Commons Attribution License \(CC BY\)](https://creativecommons.org/licenses/by/4.0/). The use, distribution or reproduction in other forums is permitted, provided the original author(s) and the copyright owner(s) are credited and that the original publication in this journal is cited, in accordance with accepted academic practice. No use, distribution or reproduction is permitted which does not comply with these terms.

Bayes-xG: player and position correction on expected goals (xG) using Bayesian hierarchical approach

Alexander Scholtes and Oktay Karakuş*

School of Computer Science and Informatics, Cardiff University, Cardiff, United Kingdom

This study employs Bayesian methodologies to explore the influence of player or positional factors in predicting the probability of a shot resulting in a goal, measured by the expected goals (xG) metric. Utilising publicly available data from StatsBomb, Bayesian hierarchical logistic regressions are constructed, analysing approximately 10,000 shots from the English Premier League (for the years of 2003 and 2015) to ascertain whether positional or player-level effects impact xG. The findings reveal positional effects in a basic model that includes only distance to goal and shot angle as predictors, highlighting that strikers and attacking midfielders exhibit a higher likelihood of scoring. However, these effects diminish when more informative predictors are introduced. Nevertheless, even with additional predictors, player-level effects persist, indicating that certain players possess notable positive or negative xG adjustments, influencing their likelihood of scoring a given chance. The study extends its analysis to data from Spain's La Liga (≈20 K shots from 1973 and 2004 to 2020) and Germany's Bundesliga (≈7.5 K shots from 2015), yielding comparable results. Additionally, the paper assesses the impact of prior distribution choices on outcomes, concluding that the priors employed in the models provide sound results but could be refined to enhance sampling efficiency for constructing more complex and extensive models feasibly.

KEYWORDS

expected goals, football, Bayesian hierarchical models, player adjustment, position adjustment, prior effects, soccer

1 Introduction

One of the most common advanced football analytics metrics is the idea of expected goals (xG), which estimates the probability of a given shot resulting in a goal based on several features about the shot such as distance from the shooter to the goal or the body part used by the shooter. However, none of the mainstream xG models take into account any player-specific features when estimating these values. To illustrate this, imagine that you have two players taking the same shot from the same position, with defenders in the same place and everything else being the same. Still, one player is Lionel Messi and the other is a random player from the National League (English 5th tier). Obviously, players who play in the National League are good, but it is not unreasonable to assume that Lionel Messi would be more likely to score. However, xG metrics would assign the same value for both of these changes.

The objective of this paper is to investigate if there are position or player effects on xG, meaning that certain positions or players have higher or lower goal probabilities for a given chance than others. This will be achieved using a Bayesian hierarchical model, where the

hierarchies will be the position of the player or the player. The results of this method will initially be compared to a more traditional frequentist xG model to evaluate baseline results without any group effects. Then the hierarchical models will be compared to the non-hierarchical Bayesian models, to assess the impact of having hierarchies in the data on the results. If the rationale described above of the two players taking the same shot is valid, then it is expected that the xG predictions of the hierarchical models will differ significantly from the non-hierarchical models, supporting the idea that there is a position and/or player effect on the xG of a given shot.

This paper will begin with a review of the relevant literature around this topic, looking at the development of football analytics and xG, as well as any attempts to use Bayesian modelling in football analytics. Thereafter, the methodology will be described by going through the frequentist and Bayesian techniques used. Then, the data will be introduced and described with any changes made before the choice of Bayesian prior distributions for predictors is discussed. Next, the results of the modelling will be presented before validating the results of the Bayesian models on additional data. The aforementioned choice of prior distributions will then be evaluated. Finally, a discussion section will deliberate on the significance of the results before concluding the paper.

2 Related works

The use of data in football is often not fully embraced, with many decision-makers arguing the sport is too complex for data to be used effectively to improve results and performance (1). However, with its successful use in other sports, there was sufficient interest for some clubs, companies, and individuals to pursue using data to derive conclusions and make suggestions in football. With a growing demand for data, companies that specialise in sports data collection have grown too, along with their ability to track data. The result is that there is now an enormous amount of football data to use for several purposes, such as player/club performance, scouting, and player fitness and injury risk, to name a few (2).

At the heart of the idea of using data in football was the potential to gain a competitive advantage. As a result, clubs that use data tend to be secretive about their operations and procedures (2). Despite this, there is plenty of publicly available literature and sources showing how data can be used in football. Moreover, sports broadcasters have long used data when giving an overview of a match, such as possession statistics. Still, these have only recently moved away from simple counts and percentages to more complex metrics. The Bundesliga, for example, provides a goal probability value after each goal is scored, giving the chance of that given opportunity resulting in a goal (3).

This goal probability, also commonly called expected goals (xG), has been a central topic in the development of more advanced statistics using football data (1). Crucially, it moves away from the idea of things that did happen and focuses on things that could have happened. With football being such a complex and chaotic sport, outcomes often do not reflect expectations as matches are often decided by fine margins or decisions out of the

players' and coaches' control. Nevertheless, using expectations gives decision-makers an idea of the underlying performance of their team and allows them to see if their team is over or underperforming according to expectations (4).

There have been many versions of xG models created since the idea was founded, using a variety of machine-learning techniques and data sources. Herold et al. (5) provide a summary of many applications of machine learning in football, including xG models. The most common methods of estimating xG in their paper are logistic regressions, decision trees, ensemble methods (e.g., random forest), and neural networks. Lucey et al. (6) use player and ball tracking data from the 10 s leading up to a shot to estimate goal probabilities across an entire season and found that "defender proximity, interaction of surrounding players, speed of play, coupled with shot location play an impact on determining the likelihood of a team scoring a goal".

Furthermore, Anzer and Bauer (7) evaluate several machine-learning approaches with hand-crafted features from synchronised positional and event data of 105,627 shots in the German Bundesliga and conclude that their extreme gradient boosting-based xG model reaches the best performance. Madrero Pardo (8) uses qualitative data from the popular video game FIFA to account for player effects on xG using a logistic regression and an XGBoost model. They found that an adjusted model can better predict goals over a season for individual players and teams than an overall xG model. Fairchild et al. (9) built an xG model again using logistic regression and used it to estimate MLS teams' offensive efficiency in scoring. They also discuss evaluation metrics for expected goals models and suggest the use of the Brier score to compare predicted probability to ground-truth binary outcomes. Cavus and Biecek (10) apply a variety of ensemble and boosting methods to calculate xG values and find that a random forest model performs best, even compared to models from other papers using other techniques and data.

The closest study to this paper to date is that of Hewitt and Karakuş (11), which investigates position and player-adjusted xG models. They find evidence of positional adjustments with forwards having a positive adjustment, midfielders having a slightly negative adjustment, and defenders having a large negative adjustment. Moreover, they also find evidence of player effects on xG by fitting their model with only data from Lionel Messi and find a large positive adjustment in this case.

One of the features a lot of these models have in common is their frequentist approach, as opposed to using Bayesian methods. Spearman (12) uses a Bayesian approach to estimate the maximum a *posteriori* effects of parameters in a model for predicting future scoring of teams in games. Joseph et al. (13) used Bayesian networks to predict the results of matches played by Tottenham Hotspur and compared the results to ML techniques such as Naïve Bayes, K-nearest neighbour and decision tree. They reiterate one of the benefits of Bayesian modelling which is comparably accurate predictions in the absence of a large amount of data. Zamboni-Ferraresi et al. (14) use Bayesian methods to analyse team performance in Europe's top leagues to determine which features tend to be most significant in predicting team performance. They find that the most important features include the number of assists, the

number of shots conceded, saves made by the goalkeeper, passing accuracy, and number of shots on target.

One area of Bayesian modelling which is also often not considered in football analytics is using multi-level, or hierarchical, models. Tureen and Olthof (15) construct a multi-level model for player-adjusted expected goals but do not use a Bayesian approach to do so. Still, they use their model to calculate estimated player impact values on xG. On the other hand, Baio and Blangiardo (16) construct a Bayesian hierarchical model but use it to predict match results as opposed to xG directly and group their data by the team as opposed to by the player. Blanco et al. (17) in their study employ a Bayesian methodology to probabilistically estimate Chilean Premier League team positions, revealing insights into attack and defence strengths. While the model accurately predicts the top five positions, it identifies weak defensive capabilities across all teams, showcasing its competitiveness for soccer championship prediction. Still, the use of Bayesian hierarchical modelling in football is a relatively unexplored area.

By using Bayesian hierarchical modelling, group-level effects can be reliably estimated even with small group sizes. Therefore, the effect of a player's position or even the player themselves on the chance of a given shot resulting in a goal can be reliably measured. The result is that certain players could be identified as being more likely to score than others for given chances, which is a result that can be used for player selection or scouting purposes. This idea is present in the work of Hewitt and Karakuş (11), where Messi is found to be an extremely efficient goal scorer. This conclusion may be obvious to football fans, but the fact that the efficiency can be reliably measured is extremely interesting for potentially comparing the goalscoring efficiency of footballers. Tureen and Olthof (15) construct a metric they refer to as "estimated player impact", which is another calculation of a player's individual effect on the probability of scoring. The estimation of a player's impact on xG can potentially be another tool in evaluating player performance for team selection or scouting purposes.

After a general look at the literature on the xG metric and its evaluation throughout the years, in the following section, we explain the details of the methodology this paper is proposing to study positional and player-related corrections to generic xG approaches. We propose the utilisation of the Bayes formula by taking player and position information into a conditional probability formulation which is then evaluated under Bayesian hierarchical modelling.

3 Methodology

3.1 Preliminaries on xG calculation

Before showing how Bayesian methods can be applied to xG modelling, we now show how xG models are typically created. This involves using a frequentist approach and, in this case, a logistic regression appears as the natural choice to obtain goal probabilities. This paper will first attempt to build a generic xG model with comparable results to an xG model built by StatsBomb, an industry leader in data collection and analysis. To

do so, we gradually increase the number of predictors (both given and engineered) used in the logistic regression that generally follows the formulation given below

$$\text{logit}(p_i) = \log\left(\frac{p_i}{1-p_i}\right) = \beta_0 + \sum_{j=1}^N \beta_j \cdot X_{ji} \quad (1)$$

where p_i is the probability of the shot i resulting in a goal, X_{ji} is the value of predictor j for the shot i . Moreover, β_j for $j = 0, 1, \dots, N$ is the logistic model coefficients, and N is the number of parameters of the logistic regression. Traditional logistic regression can be seen as a method used to model the relationship between a binary dependent variable (in this context Goal or No-Goal) and one or more independent variables (given and engineered features). It uses the logistic function in Equation (1) to transform a linear combination of features into a probability of the dependent variable being one of the two classes.

3.2 On the Bayesian predictive modelling of expected goals (xG) models

In predictive modelling of xG, accurate prediction of xG is crucial for assessing team performance and predicting match outcomes. The posterior predictive distribution plays a pivotal role in this process as it provides a probabilistic framework for estimating the likelihood of different goal-scoring scenarios given observed data and model parameters.

The posterior predictive distribution, which can be denoted as $p(y_{xG}|\mathbf{X}, \mathbf{y})$, represents the distribution of expected goals y_{xG} for future events, conditioned on the observed features \mathbf{X} and corresponding outcomes \mathbf{y} . Mathematically, it is expressed as in Equation (2):

$$p(y_{xG}|\mathbf{X}, \mathbf{y}) = \int p(y_{xG}|\beta)p(\beta|\mathbf{X}, \mathbf{y})d\beta \quad (2)$$

where $p(\beta|\mathbf{X}, \mathbf{y})$ is the target posterior distribution of model parameters β given the observed data, and $p(y_{xG}|\beta)$ is the likelihood function representing the probability of observing expected goals y_{xG} given the parameters β .

In simpler terms, approaching modelling from a Bayesian perspective means we can incorporate prior knowledge about model parameters, resulting in a posterior density of these parameters ($p(\beta|\mathbf{X}, \mathbf{y})$). This posterior is then used in the inference process to predict future xG values by sampling from the posterior predictive distribution of $p(y_{xG}|\mathbf{X}, \mathbf{y})$. Essentially, estimating this distribution helps us measure the uncertainty in xG predictions by considering both uncertainty in parameters and variability in observed data.

3.3 Bayesian logistic regression

Bayesian logistic regression extends logistic regression by introducing a Bayesian framework for modelling. Instead of

estimating fixed model parameters as in traditional logistic regression, Bayesian logistic regression defines these parameters as unknowns and treats the uncertainty around these unknowns as random variables with associated probability distributions. This means that we get the single population value as the estimate for the model parameters as well as full probability distributions, allowing us to quantify uncertainty.

In Bayesian logistic regression, we specify prior distributions for the model parameters, representing our beliefs about their values before observing any data. The likelihood function represents the probability of observing the data given the model parameters. Using Bayes' theorem, we update the prior distribution with the likelihood distribution to obtain the posterior distribution, which reflects our updated beliefs after considering the data. To compute the posterior distribution in Bayesian logistic regression, various techniques can be used, including Markov Chain Monte Carlo (MCMC) methods and variational inference. These methods sample from the posterior distribution of the parameters to estimate their values and uncertainties.

Once we have the posterior distribution of the model parameters, we can perform various tasks such as parameter estimation, uncertainty quantification, and prediction. Bayesian logistic regression provides a natural way to make probabilistic predictions, as it generates a distribution of predicted probabilities for each class, rather than just point estimates.

For this paper, the Bayesian methods used will involve fitting Bayesian logistic regressions as both single-level models (baseline) without group-level effects, and multi-level, or hierarchical, models with group-level effects. The predictions from both models will then be compared to determine if there is evidence of group-level effects on xG. To formulate such Bayesian models, this paper specifies expressions in Equations (3) to (9) below:

$$Y_i = \text{binary outcome of shot } i \quad (1 = \text{goal}, 0 = \text{no goal}), \quad (3)$$

$$p_i = \text{probability of shot } i \text{ resulting in a goal}, \quad (4)$$

$$X_{j,i} = \text{the value of predictor } j \text{ for shot } i, \quad (5)$$

$$\beta_j = \text{model coefficient for predictor } j, \quad (6)$$

where the likelihood distribution is

$$Y_i \sim \text{Bernoulli}(p_i) \quad (7)$$

Hence, the Baseline and Hierarchical models are like

$$\text{Baseline Model} \quad \text{logit}(p_i) = \beta_0 + \sum_{j=1}^N \beta_j \cdot X_{j,i}, \quad (8)$$

Hierarchical Model

$$\text{logit}(p_{ik}) = \beta_{0,k} + \sum_{j=1}^N \beta_{j,k} \cdot X_{j,i} + \beta_{N+1,k} \cdot X_{N+1,i}, \quad (9)$$

where $N + 1$ is the index of the grouping variable in the data, p_{ik} refers to the probability of shot i for k th grouping effect resulting in a goal, and $\beta_{j,k}$ is the model coefficient for predictor j and k th grouping effect.

3.4 Data

The data used for this project is all freely available event data from StatsBomb, obtained using their Python package *StatsBombPy* (please see the Statsbomb GitHub page via <https://github.com/statsbomb> for details). From their database, only men's competitions were used because it could be that there is a difference in given goal probabilities in men's and women's football, and we have more data from men's competitions. Then, all open-play shots were extracted with all relevant information for each shot. Set-pieces were excluded because again goal probabilities could vary for set-pieces, and we are not interested in modelling this effect.

The resulting data has 63,309 shots from a variety of competitions and years, with 42 columns of information for each shot. Table 1 gives a breakdown of leagues and seasons that the full data consists of whilst Table 2 gives some summary statistics for the most relevant variables in the data. As well as the information given in the columns already in the data, there are several features which were not included by StatsBomb which could be useful for predicting goal probability. Many sources cite distance to goal and shot angle to be two of the most important predictors of goal probability (11, 18–20).

The data has the location of the shot, and the StatsBomb data specification (21) provides information about the coordinates of the goalposts. Distance to goal is therefore calculated as the Euclidean distance from the shot to the centre of the goal.

TABLE 1 Full dataset details including leagues and seasons where the total number of shots taken is 63,309.

Leagues	Season (# of samples)	Total
Argentina - Liga Profesional	1981 (16), 1997 (32)	48
England - Premier League	2003 (832), 2015 (9,410)	10,242
Europe - Champions League	1970 (33), 1971 (30), 1972 (41), 2003 (13), 2004 (39), 2006 (18), 2008 (19), 2009 (28), 2010 (26), 2011 (45), 2012 (29), 2013 (27), 2014 (31), 2015 (41), 2016 (25), 2017 (28), 2018 (28)	501
Europe - UEFA Euro	2020 (1,187)	1,187
Europe - UEFA Europa League	1988 (84)	84
France - Ligue 1	2015 (8,261)	8,261
Germany - 1. Bundesliga	2015 (7,360)	7,360
India - Indian Super League	2021 (2,877)	2,877
International - FIFA U20 World Cup	1979 (27)	27
International - FIFA World Cup	1958 (82), 1962 (39), 1970 (252), 1974 (159), 1986 (75), 1990 (20), 2018 (1,554), 2022 (1,377)	3,558
Italy - Serie A	1986 (32), 2015 (9,395)	9,427
North and Central America - North American League	1977 (48)	48
Spain - Copa del Rey	1977 (23), 1982 (20), 1983 (21)	64
Spain - La Liga	1973 (27), 2004 (139), 2005 (379), 2006 (601), 2007 (645), 2008 (693), 2009 (776), 2010 (688), 2011 (791), 2012 (656), 2013 (750), 2014 (861), 2015 (8,658), 2016 (805), 2017 (892), 2018 (813), 2019 (680), 2020 (771)	19,625

TABLE 2 Summary statistics for most relevant features for xG in the dataset.

Variable	Description	Summary	
distance_to_goal	Distance between the shooter and the middle of the goal line, normalised to StatsBomb pitch size of 120 × 80 due to varying pitch sizes.	N: 63,309, Min: 0.63 Mean: 18.96, Max: 88.83, SD: 8.58	
shot_angle	By creating a triangle between the shooter and the two goalposts, the shot angle is the angle that is by the shooter.	N: 63,309, Min: 0.66 Mean: 25.39, Max: 168.61, SD: 15.60	
gk_distance_to_goal	Same as distance_to_goal but for the goalkeeper instead of the shooter.	N: 63,309, Min: 0.00 Mean: 3.56, Max: 118.00, SD: 2.61	
players_in_shot_triangle	The number of players in the shot triangle, created by the shooter and the two goalposts.	N(0): 1,786, N(1): 30,262, N(2): 19,481 N(3): 6,802, N(4): 2,918, N(5): 1,217 N(6): 513, N(7): 211, N(8): 77 N(9): 29, N(10): 11, N(11): 2	
opponents_in_radius	The number of opposition players in a 1m radius of the shooter.	N(0): 55,536, N(1): 7,030 N(2): 662, N(3): 71, N(4): 10	
shot_body_part	The body part used by the shooter to hit the ball.	N(Preferred Foot): 30,738, N(Head): 10,647 N(Other Foot): 11,733, N(Other): 191	
shot_first_time	Whether the shot was a first-time shot, meaning the shooter took no additional touches of the ball before shooting.	N(True): 20,946	N(False): 42,363
gk_in_shot_triangle	Whether the goalkeeper was in the shot triangle created by the shooter and the two goalposts when the shot was taken.	N(True): 60,570	N(False): 2,739
shot_one_on_one	Whether the shooter was one-on-one with the goalkeeper when shooting.	N(True): 3,546	N(False): 59,763
shot_open_goal	Whether the shooter was shooting at an open goal.	N(True): 736	N(False): 62,573
shot_technique	The technique the shooter used.	N(Normal): 47,854, N(Overhead Kick): 385 N(Half Volley): 9,371, N(Diving Header): 284 N(Volley): 4,483, N(Backheel): 244 N(Lob): 688	
under_pressure	Whether the shooter was under pressure when shooting.	N(True): 16,149	N(False): 47,160
goal	Whether the shot resulted in a goal.	N(True): 6,559	N(False): 56,750
shot_statsbomb_xg	StatsBomb's own estimated xG value for each shot.	N: 63,309, Min: 0.00 Mean: 0.10, Max: 1.00, SD: 0.13	
general_position	The general position of the shooter (striker, attacking midfielder, other midfielder, or defender).	N(ST): 17,073 N(AM): 20,065	N(M): 15,858 N(D): 10,313
player	The name of the shooter.	N(Messi): 1,907 N(L. Suárez): 596 N(Iniesta): 374 N(Neymar): 352	N(T. Henry): 325 ... N(V. Migas): 1 N(R. Tricella): 1

N(x):y means that for y number of samples the variable has the value of x.

For shot angle, the cosine rule is used to calculate the angle from the shooter to the two goalposts. To calculate this reliably, any shots that are taken from the same x-coordinate as the goal line are excluded since this would create a straight line instead of a triangle with no shot angle.

Statsbomb's open-access datasets offer a valuable "freeze frame" feature, providing detailed insights into specific moments within a match. This freeze frame information captures crucial events such as passes, shots, and tackles, along with the positional coordinates of players involved, enabling thorough analysis of on-field actions. By offering this level of granularity, we are able to delve deeper into detailed analysis for the shot taken and to engineer new useful features that can affect the xG predictions. Thanks to this feature, the freeze-frame feature in the data is utilised, we engineered several features: the goalkeeper's distance to the goal, whether the goalkeeper is present in the shot triangle formed by the shot and two goalposts, the number of players present in the shot triangle, and the number of opponents within a 1m radius of the shooter. Opponents are used for the 1m radius as opposed to all players because the only time a non-opponent in the radius of a shooter would impact goal probability is when they are in the shot

triangle, which is already being accounted for. Otherwise, only opponents will try to put pressure or tackle the shooter outside of the shot triangle.

The initial dataset categorises each player into 25 distinct positions, such as Left Centre Forward, Second Striker, Right Wing, and Right Centre Back. Due to players' flexibility across positions and potential overlap in samples, clustering them into positional groups poses a challenge. To address this issue, we compute mode statistics for each player's position, condensing them into a single positional variable. However, given the analytical focus of this study and the abundance of players in certain positions, particularly those with low sample sizes, we consolidate these positions into four main categories: strikers (ST, encompassing L/R/C Centre Forward), attacking midfielders (AM, including Secondary Striker, L/R Wing, and L/R/C Attacking Midfield), non-attacking midfielders (M, comprising L/R/C Centre Midfield, L/R Midfield, and L/R/C Defensive Midfield), and defenders (D, covering L/R/C Centre Back, L/R Back, L/R Wing Back, and Goalkeeper). This process results in the engineering of the "general_position" feature, consisting of these four positions, with an expected average decline in goal probability for each respective group.

Finally, we modify the body part used by changing “left foot” and “right foot” to “preferred foot” and “other foot” according to the player’s apparent preference. To assign these, we go back through the open data and instead look at the passes for each player, we then assign whichever foot the player made the most passes with as their preferred foot because they often have more time to choose which foot to pass with and will then tend to go safely with their preferred foot, which is less feasible when taking a shot as players tend to be under more pressure and have less time.

3.5 Reference models

As outlined in the preceding sections, the primary aims of this paper are to support the adoption of the Bayesian hierarchical modelling approach in investigating the xG metric, specifically concerning player-specific attributes and position groups of players. To fulfil these objectives, we commence by providing definitions for the benchmark and baseline models in the subsequent sections.

3.5.1 Statsbomb xG model

This model, employed for comparison in this study, relies on Statsbomb’s open data set to provide corresponding xG calculations for each shot discussed in the dataset. The specifics of Statsbomb’s xG model remain undisclosed, and it serves as the benchmark model in this paper.

3.5.2 Baseline xG model

The initial frequentist model introduced is a fundamental model primarily utilising the distance between the shooter and the goal as a predictor, recognised as a key factor in goal probability determination. Another crucial feature considered is the angle of the shot, formed by the shooter and the two goalposts. Recognising the inherent link between distance and shot angle, an interaction term is incorporated to capture the combined influence of both features. This foundational parameterisation constitutes the formulation of the Baseline xG model under a logistic regression model as in Equation (10)

$$\text{logit}(p_i) = \beta_0 + \beta_1 \cdot (\text{distancetogoal})_i + \beta_2 \cdot (\text{shotangle})_i + \beta_3 \cdot (\text{distance}_i \cdot \text{angle}_i) \quad (10)$$

It is logical to infer that, under reasonable assumptions, the likelihood of scoring decreases, on average, as the distance to the goal increases. This implies that the coefficient β_1 is likely to be negative. Conversely, β_2 is expected to be positive, reflecting the observation that as the shot angle decreases, the shot’s position is likely to be farther away or from a wider position, both scenarios leading to a lower goal probability. Figure 1 illustrates the associations between shot angle, distance, and the proportion of goals scored to shots, providing a clearer insight into these relationships.

3.5.3 Extended model

While the primary factors considered for calculating the goal probability are typically assumed to be the distance and angle of the shot utilised in the Baseline Model, it is important to acknowledge that various other elements can influence the shot’s outcome as shown clearly in the literature. Section 3.4 has previously addressed each of these factors, and we now introduce our frequentist Extended model, outlined below in Equation (11)

$$\begin{aligned} \text{logit}(p_i) = & \beta_0 + \beta_1 \cdot \text{distancetogoal}_i + \beta_2 \cdot \text{shotangle}_i \\ & + \beta_3 \cdot (\text{distance}_i \cdot \text{angle}_i) + \beta_4 \cdot \text{gkdistancetogoal}_i \\ & + \beta_5 \cdot \text{playersinshottriangle}_i + \beta_6 \cdot \text{bodypart}_i \\ & + \beta_7 \cdot \text{firsttimeshot}_i + \beta_8 \cdot \text{gkinshottriangle}_i \\ & + \beta_9 \cdot \text{oneononeshot}_i + \beta_{10} \cdot \text{opengoal}_i \\ & + \beta_{11} \cdot \text{technique}_i + \beta_{12} \cdot \text{underpressure}_i \end{aligned} \quad (11)$$

3.6 Bayes-xG models

Bayesian logistic regression hierarchical modelling is a powerful statistical approach used to analyse and model complex relationships within data. In this context, an additional grouping parameter is employed to incorporate the hierarchical structure of the data, acknowledging potential dependencies or variations across groups. Unlike frequentist logistic regression, Bayesian hierarchical modelling allows for the inclusion of prior information, facilitating a more robust estimation of parameters and uncertainties. Below, we initiate by formulating our *Bayes-xG* models to account for player and position grouping effects. Subsequently, we provide a comprehensive explanation of the rationale behind selecting specific prior distributions for the model parameters.

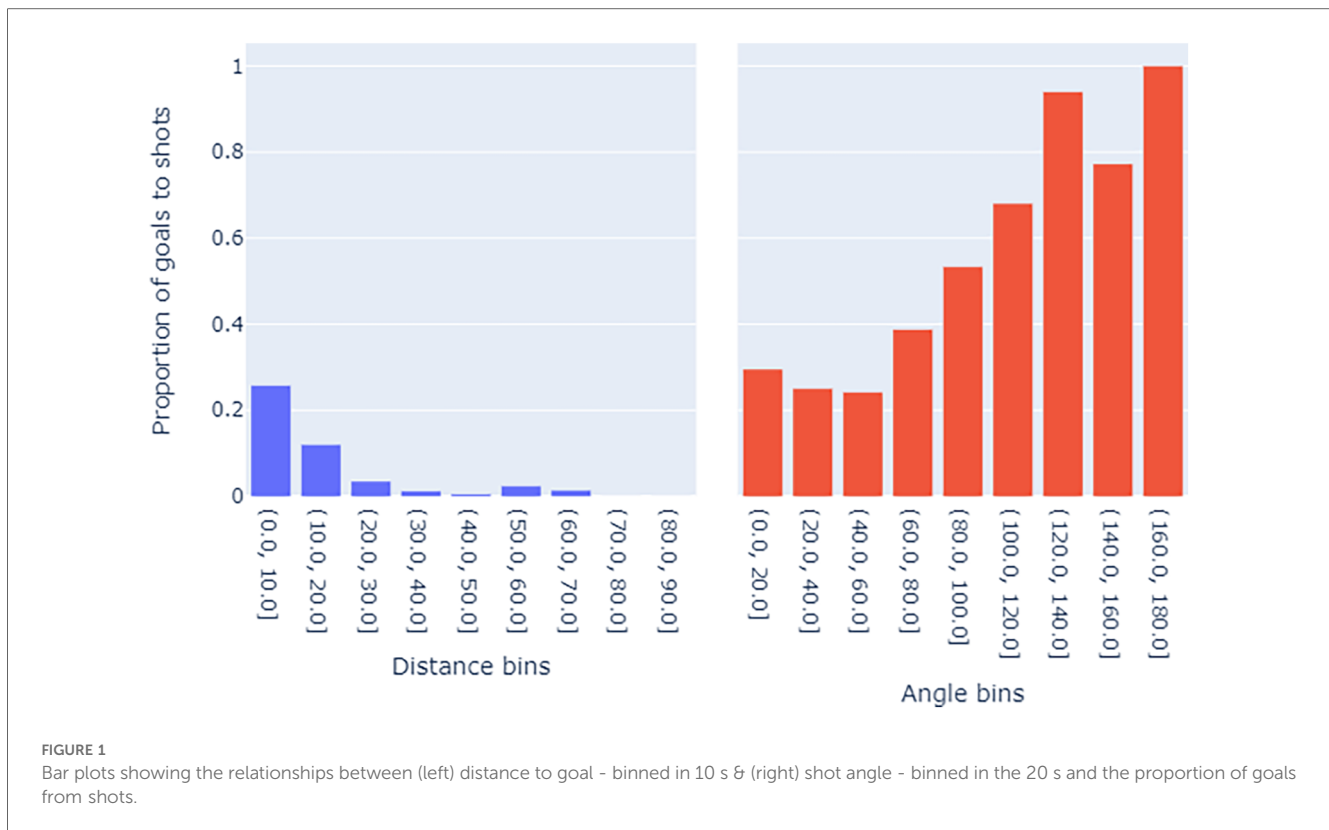
3.6.1 Model definitions

The foundational Bayesian logistic regression model extends the traditional logistic regression equation by adding a grouping effect as in Equation (12)

$$\text{logit}(p_{ik}) = \beta_{0,k} + \sum_{j=1}^N \beta_{j,k} \cdot X_{ji} + \underbrace{\beta_{N+1,k} \cdot X_{N+1,i}}_{\text{Grouping effect}}, \quad (12)$$

where $k = 1, \dots, \mathcal{K}$ with \mathcal{K} is the number of elements in the group. It is also clear that for this formulation, we extended logistic regression model coefficients $\beta_j \in \mathcal{R}^N$ into a complex form for grouping effect as $\beta_{j,k} \in \mathcal{R}^{\mathcal{K} \times (N+1)}$. This leads to fitting different logistic regression coefficients for each specific group where the same shot predictors will result in different xG calculations for each element in the group.

Specifically, $\beta_{0,k}$ is a *group-specific intercept* for k th element of the group which accounts for variations in the baseline success probability across different groups. On the other hand, $\beta_{\{1, \dots, N\}, k}$



refer to the group-specific slopes which capture variations in the effect of the covariate across groups.

Following the technical details above, for this paper, we decided to define three versions of *Bayes-xG* models which can be expressed as

- *Bayes-xG*₁ → uses the Baseline model with grouping parameter of *position*,
- *Bayes-xG*₂ → uses the Extended model with grouping parameter of *position*,
- *Bayes-xG*₃ → uses the Extended model with grouping parameter of *player*.

3.6.2 Choice of priors

When using Bayesian modelling methods, one important consideration is the choice of prior distributions for the predictors in the model. Table 3 lists the predictors in each of the models and the prior distributions used for them.

Variables lacking a clear rationale for a positive or negative value are assigned a prior distribution modelled on the normal distribution. In contrast, variables, where the direction of the effect can be reasonably predicted, are assigned a prior distribution skewed towards that direction. This choice is made to favour values in the predicted direction while still allowing for the possibility of values in the opposite direction, acknowledging the potential for an incorrect prediction. The justifications for prior distribution choices are given below:

- *distance_to_goal* ($\mu = -1$ and $\alpha = -1$): Scoring becomes progressively more challenging as the distance from the goal

increases, primarily because the goalkeeper gains additional time to react to the shot.

- *shot_angle* ($\mu = 1$ and $\alpha = 1$): This is because an increase in the shot angle implies either moving closer to the goal or positioning oneself more centrally, both of which result in a higher average scoring probability.
- *players_in_shot_triangle* ($\alpha = \{5, \dots, -5\}$): The presence of an increasing number of players within the shot triangle makes it increasingly challenging for a player to execute a shot in a manner that avoids hitting any of the players while still managing to score a goal.
- *opponents_in_radius* ($\alpha = \{1, \dots, -2\}$): In a manner akin to the previously mentioned feature, an increase in the number of opponents within the shooter's proximity corresponds to an increased challenge for the shooter. With more players in the vicinity, there is an augmented effort from opponents to disrupt and block the shot, consequently making it more difficult for the shooter to successfully score.
- *gk_in_shot_triangle* ($\alpha = -2$): If the goalkeeper is positioned within the shot triangle, their chances of successfully saving a shot are higher compared to when they are located outside of the shot triangle.
- *shot_one_on_one* ($\alpha = 2$): When a player finds themselves in a one-on-one situation with the goalkeeper, their sole task is to outplay the goalkeeper with their shot, without the need to navigate or consider other players. This circumstance makes scoring comparatively more straightforward.
- *shot_open_goal* ($\alpha = 4$): Similar to the previous feature, in this case, there is no goalkeeper present. Consequently, the

TABLE 3 Listing which predictors are used in each of the Bayesian models, and what prior distribution is given to the coefficient of the given predictor.

Predictor	Bayes-xG			Prior
	1	2	3	
Intercept	✓	✓	✓	$\mathcal{N}(\mu = 0, \sigma = 5)$
distance_to_goal	✓	✓	✓	$\mathcal{SN}(\mu = -1, \sigma = 5, \alpha = -1)$
shot_angle	✓	✓	✓	$\mathcal{SN}(\mu = 1, \sigma = 5, \alpha = 1)$
distance_angle_interaction	✓	✓	✓	$\mathcal{N}(\mu = 0, \sigma = 5)$
gk_distance_to_goal		✓	✓	$\mathcal{N}(\mu = 0, \sigma = 5)$
players_in_shot_triangle		✓	✓	$\mathcal{SN}(\mu = -1, \sigma = 5, \alpha = \{5, 4, \dots, -5\})$ where α is determined by the value of this feature (0 players=5, 1 player=4, etc.)
opponents_in_radius		✓	✓	$\mathcal{SN}(\mu = -1, \sigma = 5, \alpha = \{1, 0, \dots, -2\})$ where α is determined by the value of this feature (0 players=1, 1 player=0, etc.)
shot_body_part		✓	✓	$\mathcal{N}(\mu = 0, \sigma = 5)$
shot_first_time		✓	✓	$\mathcal{N}(\mu = 0, \sigma = 5)$
gk_in_shot_triangle		✓	✓	$\mathcal{SN}(\mu = 0, \sigma = 5, \alpha = -2)$
shot_one_on_one		✓	✓	$\mathcal{SN}(\mu = 0, \sigma = 5, \alpha = 2)$
shot_open_goal		✓	✓	$\mathcal{SN}(\mu = 0, \sigma = 5, \alpha = 4)$
shot_technique		✓	✓	$\mathcal{N}(\mu = 0, \sigma = 5)$
under_pressure		✓	✓	$\mathcal{SN}(\mu = 0, \sigma = 5, \alpha = -2)$
general_position	✓	✓		$\mathcal{SN}(\mu = 0, \sigma, \alpha = \{2, 1, 0, -2\})$ where scale parameter $\sigma \sim \mathcal{HN}(\gamma = 5)$ and α is {ST, AM, M, D}, respectively.
player			✓	$\mathcal{SN}(\mu = 0, \sigma, \alpha = \{2, 0\})$ where scale parameter $\sigma \sim \mathcal{HN}(\gamma = 5)$ and α is assigned depending on prior beliefs about a player ({2: good finisher, 0: not good finisher}).

\mathcal{N} , normal distribution; \mathcal{SN} , skew-normal distribution; \mathcal{HN} , half-normal distribution.

- shooter's sole objective is to direct the shot accurately towards the target, making the likelihood of scoring very high.
- *under_pressure* ($\alpha = -2$): When a player is under pressure, their ability to concentrate on accurate and well-targeted shooting diminishes, resulting in a decreased likelihood of scoring on average.
 - *general_position* ($\alpha = \{ST: 2, AM: 1, M: 0, D: -2\}$): The α values signify the expectation that strikers will exhibit the highest proficiency in finishing, followed by attacking midfielders, other midfielders, and, finally, a decrease for defenders.
 - *player* ($\alpha = \{2, 0\}$): If a player is anticipated to excel in finishing skills based on their name and reputation, they are given a value of 2 for α ; otherwise, a value of 0 is assigned.
 - A standard value of $\sigma = 5$ was chosen for the priors to strike a balance, ensuring sufficient variability. This choice aims to prevent the priors from becoming overly narrow in case the underlying prior knowledge is incorrect. Simultaneously, it avoids excessive largeness that could prolong convergence and necessitate numerous rounds of sampling.

3.7 Model development

Following the exposition of technical aspects related to the models employed in this paper, we proceed to elucidate the practical details of their implementation. The entire computational implementation was carried out using the Python programming language, specifically version 3.8 and above. In non-Bayesian modelling phases, logistic regression was executed using the Python *sklearn* module along with its associated methods.

Bayesian modelling stages were implemented by using the *bambi* module which is an open-source Python package and

purposefully crafted to simplify the fitting of generalised linear multilevel models (GLMMs) using a Bayesian framework. This encompasses a broad category of techniques widely employed in various research domains, including linear regression, analysis of variance (ANOVA), logistic and Poisson regression, as well as multilevel and crossed-group-specific effects models. *Bambi* facilitates the specification of intricate generalised linear hierarchical models through a formula notation reminiscent of R, providing a balance between user-friendly syntax for novices and direct access to internal objects for advanced users. This design allows beginners to swiftly articulate complex models with default priors, akin to popular R packages while providing seasoned users with the flexibility to directly manipulate internal objects for a more advanced and nuanced modelling approach (22).

Specifically, we employed the Python module *bambi* to construct Markov Chain Monte Carlo (MCMC) models, generating posterior distributions. Monitoring instantaneous parameter estimates from the chains confirmed their convergence, leading us to select 1,500 draws, with the initial 250 draws designated as the burning period. In total, 4 Markov chain sampling is developed resulting in 6,000 total samples for each *Bayes-xG* model. We set a target acceptance ratio of 95% in the model whilst using the prior distributions given in Table 3. Considering the target feature in the models is a binary variable of goal status, we decided to utilise a Bernoulli likelihood for all *Bayes-xG* models.

Furthermore, in the context of this paper, Bayesian implementation of the models was executed using data from a specific league. This decision was driven by the fact that a substantial portion of the entire dataset consisted primarily of Barcelona matches, as this information was publicly provided by StatsBomb. The concern was that including such a dominant dataset might introduce bias and influence the results. By

focusing on a single league, the analysis can encompass a diverse range of players and team matchups, ensuring a more balanced and comprehensive examination.

4 Experimental analysis

The experimental analysis of this paper was studied under six cases:

- (1) Frequentist model comparison to benchmark model for the whole 60 K+ shots data set,
- (2) Bayes-xG model-based “positional” analysis and comparisons for English Premier League data set (10 K+ shots),
- (3) Bayes-xG model-based “player-specific” analysis and comparisons for English Premier League data set (10 K+ shots),
- (4) Extending the developed Bayes-xG model evaluations into different countries, e.g., Spain (La Liga - 19 K shots) and Germany (Bundesliga - 7.5 K shots),
- (5) Investigating the choice of priors on Bayes-xG model outputs, and
- (6) Uncertainty quantification via discussing some posterior predictive analysis on developed Bayesian xG models.

4.1 Frequentist/non-Bayesian models

In the initial series of experiments, we assessed the modelling performance of frequentist models based on logistic regression, namely the Baseline xG and Extended xG, utilising two distinct sets of features. The analysis involved the complete data set comprising 63,309 instances. The objective was to observe and

compare how these models deviate from the predictions of the benchmark StatsBomb xG model.

The distributions of the predicted xG values for each model are shown in Figure 2, along with the StatsBomb xG values for comparison. As expected, the Extended model performs better than the Baseline model (with shot angle, distance to goal, and the interaction between the two) by also predicting much more extreme values and by decreasing the interquartile and whisker ranges. Furthermore, the extended model has a distribution very close to that of the StatsBomb model.

Concerning the evaluation criteria, Table 4 displays how each model performs compared to the StatsBomb xG model regarding R^2 , mean absolute error (MAE), and root mean square error (RMSE). As expected, the extended model outperformed the Baseline model due to its enriched data for predicting the goal likelihood. Particularly noteworthy is that the extended model achieved a Brier score significantly close to the StatsBomb model’s Brier score of 0.075, suggesting a similar performance to an industry-leading xG model. Please also note that the Brier score is a version of the DeFinetti measure (23) which is used to measure Euclidean distance between predictions and outcome probabilities of the event (24).

In the final analysis of the initial set of experiments, we explored the impact of incorporating engineered advanced features on the performance of the frequentist logistic regression model. Figure 3 illustrates the trends in performance evaluation metrics—Brier score, R^2 , MAE, and RMSE—relative to the number of features integrated into the model. We initiated the analysis with a single-parameter model, utilizing only the distance to the goal, and systematically added features one by one. Typically, there are 16 model parameters (as detailed in Table 3), but this count increases to a maximum of 33 after one-hot encoding categorical features. Examining Figure 3 reveals a

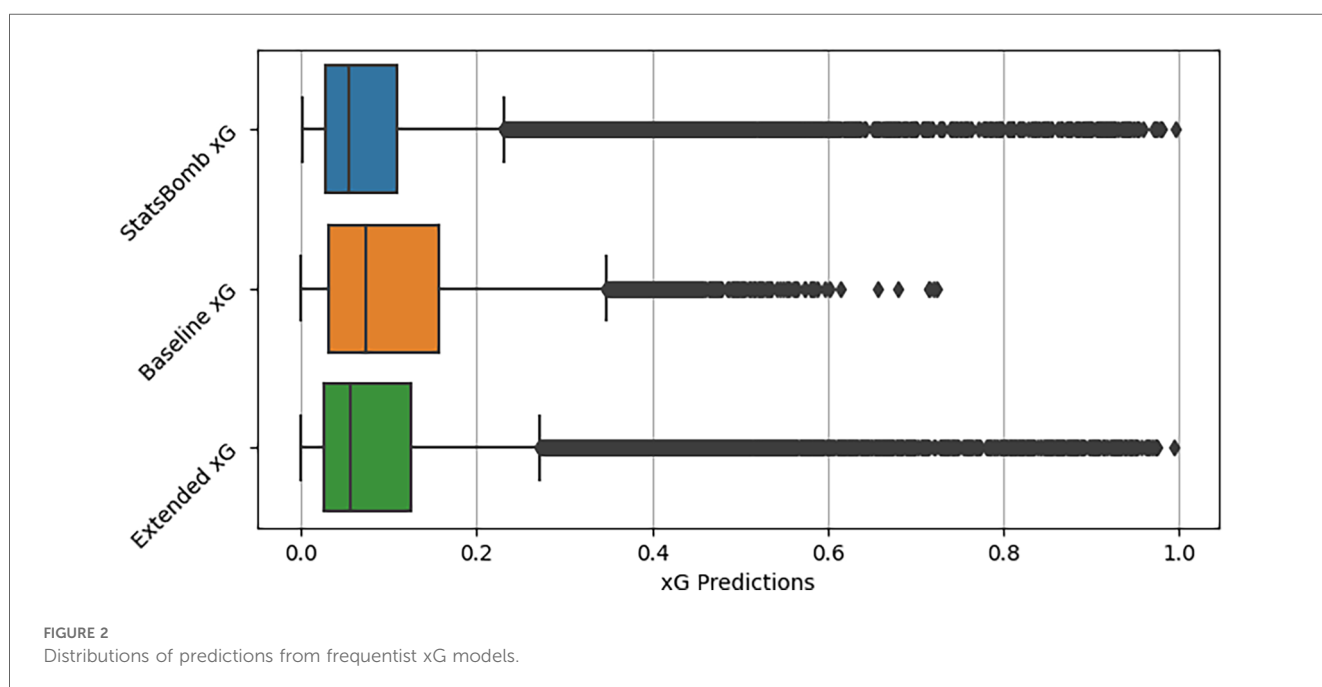


TABLE 4 Comparison of outcomes from non-Bayesian models, including the Baseline xG model integrating distance to the goal, shot angle, and their interaction, and the Extended xG model incorporating additional features, with reference to the StatsBomb xG model.

Metric	Expression	Baseline xG	Extended xG
RMSE	$\sqrt{\frac{1}{N} \sum_i^N (y_i^{SB} - \hat{y}_i)^2}$	0.095	0.055
MAE	$\frac{1}{N} \sum_i^N y_i^{SB} - \hat{y}_i $	0.058	0.029
R ² Score	$1 - \frac{\sum_i^N (y_i^{SB} - \hat{y}_i)^2}{\sum_i^N (y_i^{SB} - \bar{y}^{SB})^2}$	0.428	0.826
Brier Score	$\frac{1}{N} \sum_i^N (\hat{y}_i - Y_i^{(0,1)})^2$	0.086	0.076

y_i^{SB} : Statsbomb xG for shot i .
 \hat{y}_i : Model prediction xG for shot i .
 \bar{y}^{SB} : Sample mean of Statsbomb xG.
 $Y_i^{(0,1)}$: Binary correct outcomes (Goal, or No-Goal).
 N : number of shots.

substantial influence on model performance resulting from the introduction of these created parameters. Notably, there are certain plateaus in the trends, particularly associated with one-hot encoded parameters representing categories with a minimal number of samples in the dataset (e.g., 10 players in the shot triangle).

4.2 Positional analysis via Bayesian models

For the second experimental case in this paper, we investigate the effects of the general player positions on the pitch to the xG values by performing a Bayesian hierarchical modelling approach. Baseline xG and Extended xG models are redeveloped by using the positional grouping effect and we obtained *Bayes-xG₁* and *Bayes-xG₂* models, respectively.

We commence with a more straightforward model that examines only a few features within the logistic regression framework. As mentioned in the above sections, the baseline xG model incorporates distance to the goal, shot angle, and their interaction as predictors. To assess the influence of the grouping variable, “general position” on xG, we calculate the differences in xG predictions between the single-level frequentist model and its hierarchical Bayesian counterpart, referred to as the *Bayes-xG₁* model, for each shot in the dataset.

Figure 4 illustrates the distributions of the xG adjustment, considering the general position within the hierarchical model for each position category. The observed distributions align well with the prior beliefs regarding the impact of the general position. Defenders exhibit a substantial number of negative xG adjustments in comparison to the baseline model’s xG predictions, with some adjustments reaching as low as 0.1. As anticipated, non-attacking midfielders display smaller xG adjustments, encompassing both positive and negative values. Contrary to expectations, attacking midfielders exhibit larger positive xG adjustments on average compared to strikers, who also generally have positive adjustments. This unexpected finding may stem from the fact that strikers, by shooting more frequently from high xG scoring positions, often possess sufficiently high xG values without requiring a significant positional adjustment. On the other hand, attacking midfielders frequently take shots from areas around the goal, where xG chances are lower, yet they excel in scoring due to their above-average attacking and scoring capabilities.

Figure 5 illustrates the shot locations categorised by general playing positions, normalised for each position. It highlights the notable concentration of chances for defenders positioned

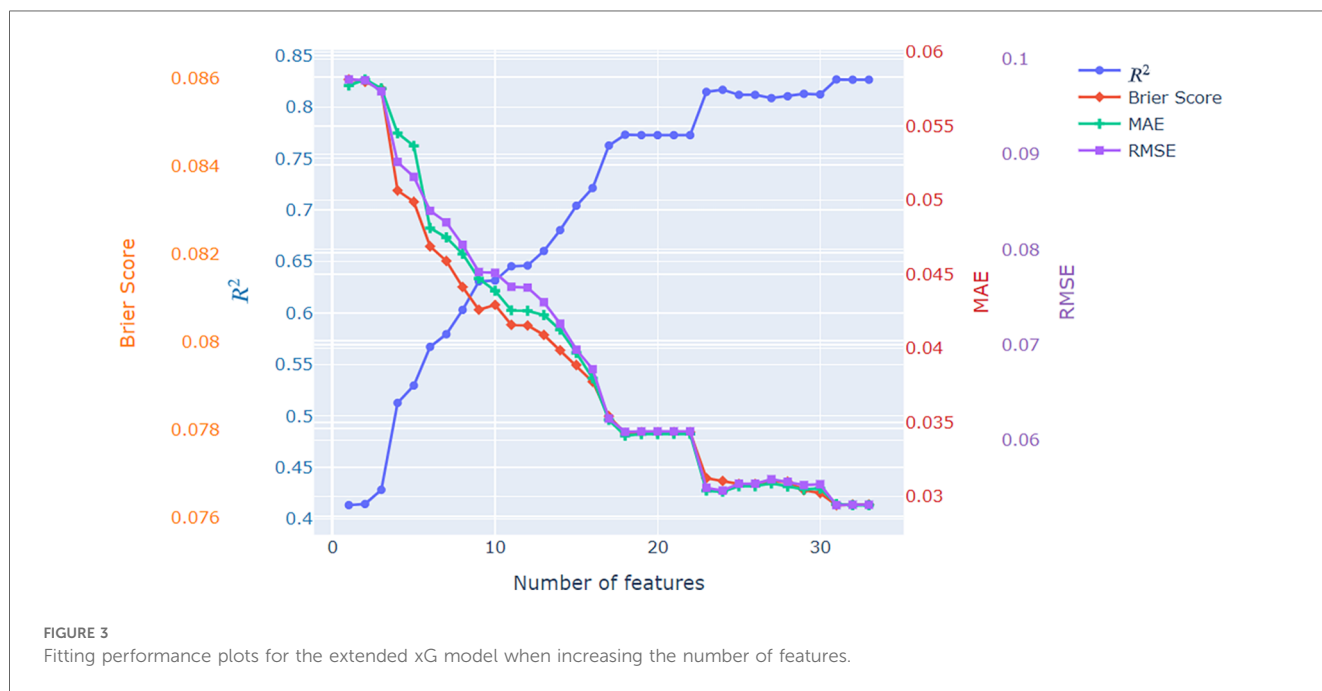
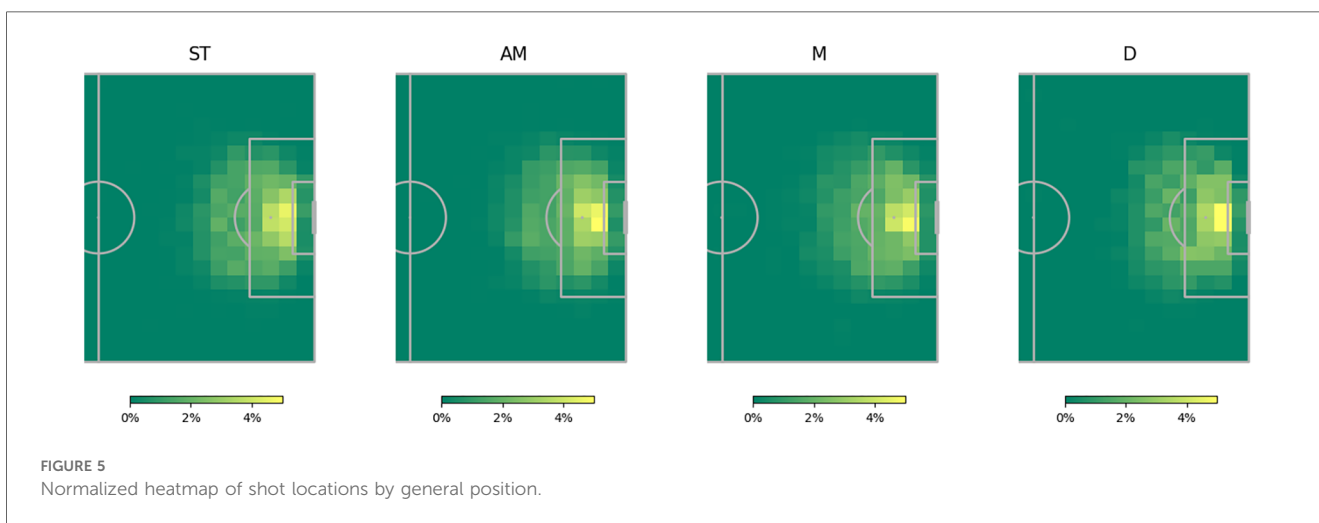
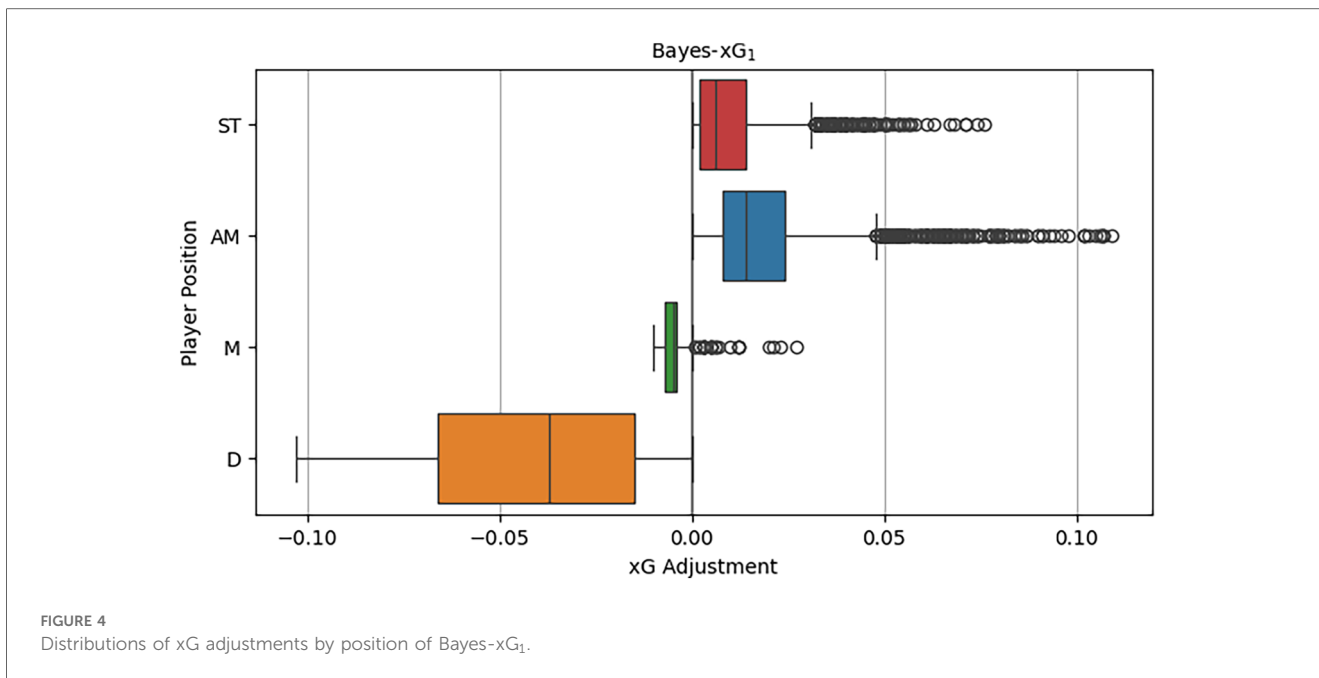


FIGURE 3 Fitting performance plots for the extended xG model when increasing the number of features.



directly in front of the goal, providing a potential explanation for the observed reversals in Figure 4. Contrary to the theorised expectation from the analysis in Figure 4, attacking midfielders are not inclined to take shots from distant or challenging angles. In fact, on average, strikers exhibit a higher tendency for such shots. This discovery, coupled with the observation that attacking midfielders, on average, have larger positive xG adjustments than strikers, suggests that attacking midfielders may have a superior ability, on average, to convert high xG chances situated right in front of the goal compared to their striking counterparts.

Before delving into the more intricate model analysis for Bayes-xG₂, we aim to showcase a validation step to demonstrate the accuracy of the MCMC-based sampling technique employed in developing Bayesian models in this paper. To achieve this, we

replicated the Bayesian analysis, this time utilising Bayes' Formula in Equation (13)

$$P(\text{goal}|\text{position}_i) = \frac{P(\text{position}_i|\text{goal}) \cdot P(\text{goal})}{P(\text{position}_i)} \quad (13)$$

This allowed us to conduct an analysis where the results of the baseline model could be adjusted using Bayes' Theorem, and these adjusted outcomes were then compared to the results of the hierarchical model. The comparison aimed to assess the proximity between theoretical adjustments and model adjustments. The outcomes of this process are detailed in Table 5. Notably, the mean model adjustments closely align with

TABLE 5 Mean xG adjustment for each general position from Bayes-xG₁ vs. theoretical adjustment of Baseline xG prediction using Bayes' Theorem.

Position	Mean model adjustment	Mean theoretical adjustment
ST	0.009	0.010
AM	0.019	0.020
M	-0.006	-0.005
D	-0.042	-0.044

the theoretical adjustments for each position, affirming that the model has effectively estimated the positional impact.

We proceed with our analysis by delving into the upgraded iteration of the Baseline model, referred to as the Extended model. Employing Bayes-xG₂, this advanced model involves a positional analysis akin to its predecessor, Bayes-xG₁. However, it incorporates numerous additional predictors, including factors such as *opponents_in_radius* and *gk_distance_to_goal*, aiming to enhance the baseline xG predictions. Upon inspecting the distributions of xG adjustments by position in Figure 6, it is evident that the values are notably smaller when compared to those derived from Bayes-xG₁. Few adjustments now exceed 0.01 away from the baseline xG values. This observation implies that the supplementary predictors effectively contribute to position prediction. This indicates that xG advantages stem less from the player's position and more from the specific play situation during shooting. For instance, while attackers generally enjoy better scoring chances on average, leading to positive xG adjustments in the basic model, the Extended model, by accounting for various features defining these improved scoring chances (e.g., one-on-ones), mitigates the impact of player position on xG adjustments.

It is intriguing to examine the xG adjustments across the predictor ranges of “*distance_to_goal*” and “*shot_angle*.” Figure 7 illustrates these xG adjustments grouped by the position for both Bayes-xG models. In Figure 7A for Bayes-xG₁, the convergence of each position towards an adjustment of 0 signifies that, at a certain distance, the xG value tends to be close to zero, regardless of other shot-related factors. Notably, defenders exhibit a slight dip in xG adjustments at the lowest distance before increasing. This suggests scenarios where defenders find themselves in goal-scoring positions, perhaps following set pieces or during a team's pursuit in a match. The order of adjustments in Figure 4 aligns consistently across positions, with attacking midfielders having the highest positive adjustments, followed by strikers, and other midfielders showing minimal adjustment from the baseline model. Examining the “*shot_angle*” for Bayes-xG₁ in Figure 7B, smaller values correspond to more challenging scoring opportunities, either due to being far from the goal or from tight angles. Defenders display gradually larger negative adjustments as the shot angle increases, reversing after a certain point for very high shot angles, likely corresponding to very close distances to the goal. Similar patterns emerge for other positions, with attacking midfielders having the largest positive xG adjustments, followed by strikers and other midfielders. For Bayes-xG₂ model outputs in Figures 7C,D, the compensation of positional effects through additional predictors is evident. Despite a clear distinction in the effects of distance to the goal and shot angle for each position in Bayes-xG₁, no significant differences between positions are observed in Bayes-xG₂. The diminishing impact of position on xG adjustment is attributed to the diverse player abilities and roles within each position category. A distinct observation evident in both Figures 7C,D is that defenders exhibit consistent trends for all angles just below 0, whereas

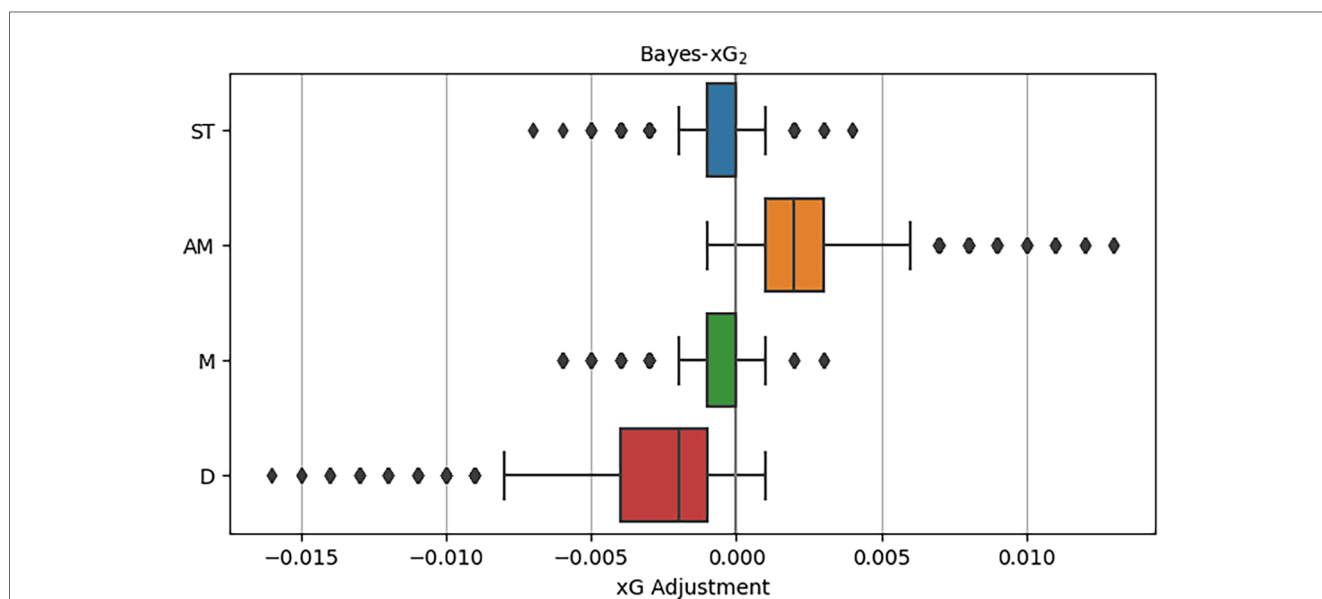
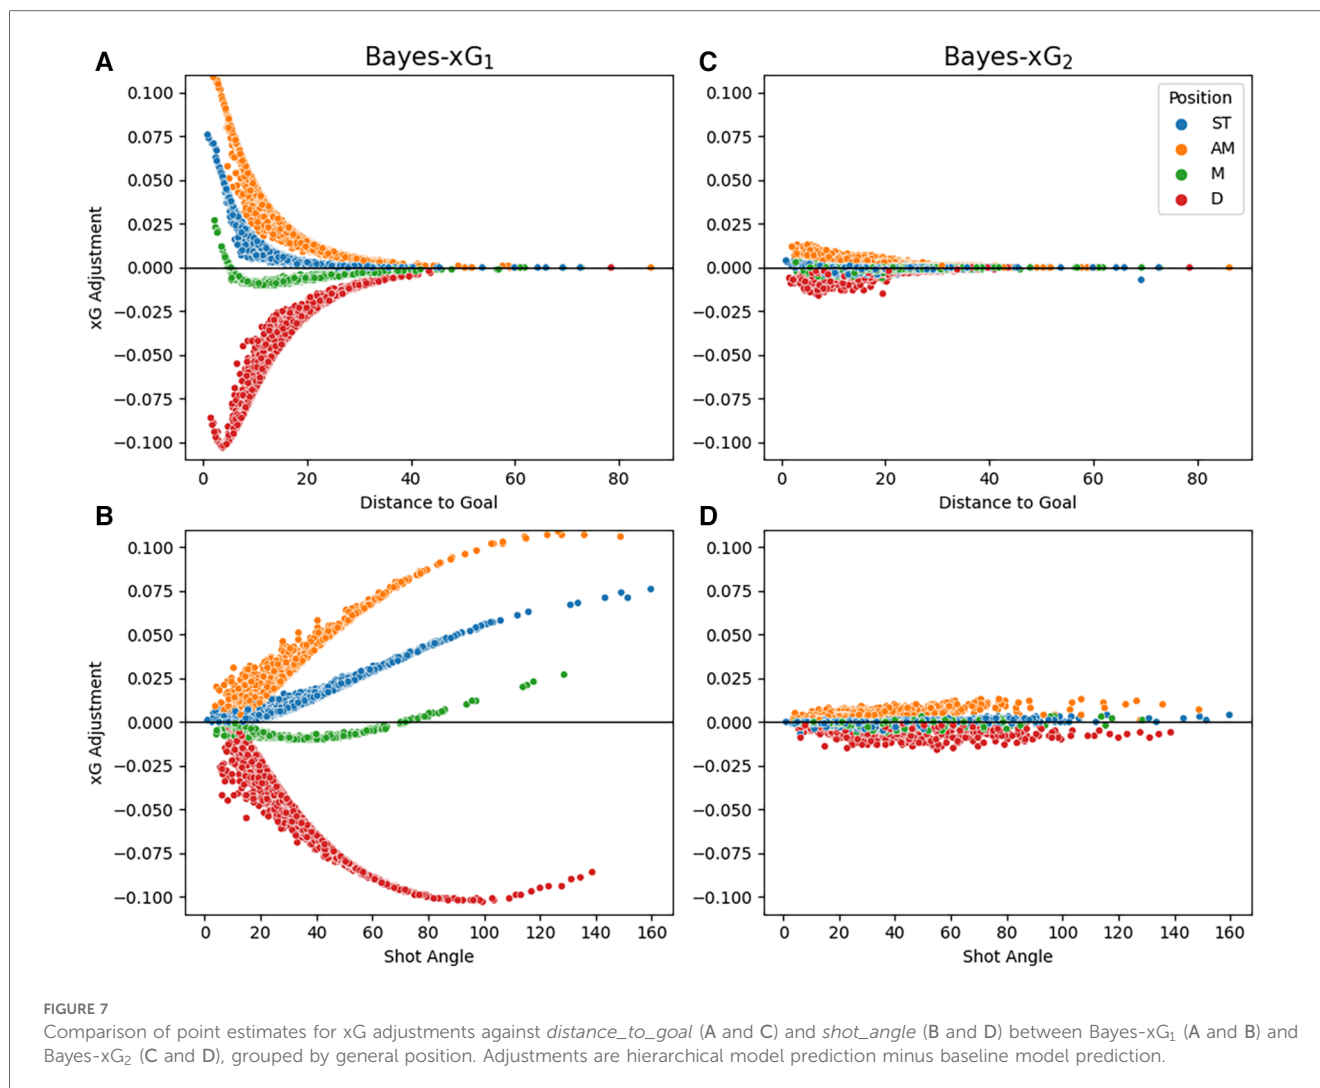


FIGURE 6 Distributions of xG adjustments for Bayes-xG₂, where adjustment is hierarchical model prediction minus baseline model prediction - grouped by general position.



attacking midfielders mirror the same pattern but above the 0 line. Adjustments made to strikers' and midfielders' xG values are barely discernible, with values hovering around 0 across all shot angles.

To further explore the aforementioned phenomena, the next experimental case involves a player-specific analysis with Bayes-xG₃, grouping data based on the player taking the shot rather than their general position.

4.3 Player-specific analysis via Bayesian models

In the context of the third experimental case explored in this paper, we introduced Bayes-xG₃, derived from the Extended model, with players designated as the grouping variable in Bayesian hierarchical modelling. The focus here is to assess whether the models necessitate player-specific adjustments. Unlike the previous experimental case that centred on positional analysis, conducting a player-specific analysis poses increased complexity due to the considerably larger pool of candidates within the group, making the analysis more challenging and

computationally intensive. Instead of individually representing each player in our group, a selective approach is employed, categorising the majority as "other" and opting for a few players with expected positive or negative xG adjustments. Player selection is based on the "conversion rate," i.e., the percentage of shots scored. To ensure relevance, only players with a minimum of 50 shots are considered, and the chosen players, along with their statistics, are detailed in Table 6. Players like R. Pirès, S. Agüero, and J. Vardy, recognized for their prolific goal-scoring, are expected to have positive xG adjustments. Pirès, notably, exhibits an exceptional conversion rate in the data subset. Conversely, players like P. Coutinho and R. Barkley, with below-average conversion rates, might have slight negative xG adjustments, while J. Shelvey, who failed to convert any of his 51 shots in the data, is likely to have a more substantial negative xG adjustment.

As explained in the preceding methodology section, the impact of each player will be characterised by a prior distribution, specifically a skewed normal distribution. The choice of distribution parameters is dependent upon the prior beliefs regarding a player's proficiency as a goal scorer, with the

TABLE 6 Selected players for Bayes-xG₃ and their goal-scoring statistics in the data set.

Player	Shots	Goals	Conversion rate
Robert Pirès	56	14	25.00%
Sergio Agüero	112	20	17.90%
Jamie Vardy	111	19	17.10%
Phillippe Coutinho	105	8	7.60%
Ross Barkley	82	6	7.30%
Jonjo Shelvey	51	0	0.00%

parameter α taking values of either 2 or 0. This selection is guided by qualitative beliefs about the players rather than direct utilisation of the data for informing the priors. Players such as Pirès, Agüero, Vardy, and Coutinho, acknowledged as talented attacking players, are attributed $\alpha = 2$. In contrast, Barkley and Shelvey, who are not commonly associated with being top-tier attackers but possess other defining qualities in their game, are assigned $\alpha = 0$.

Figure 8 illustrates the distributions of xG adjustments for individual players and the collective “other” players group. A prominent observation is the substantial positive xG adjustments for Robert Pirès, some reaching as high as 0.3 above the baseline xG. Notably, these adjustments persist even after incorporating additional predictors in Bayes-xG₂ that were intended to eliminate group effects in the previous experimental analysis. Pirès also exhibits a wide spread of adjustments, ranging close to 0, indicating a diverse array of shot types. Some were high xG chances, requiring minimal adjustment, while others were more challenging but consistently converted by Pirès, resulting in significant positive adjustments. Agüero displays consistently positive xG adjustments, albeit smaller on average and with a narrower spread compared to Pirès. Intriguingly, Vardy and

Coutinho exhibit minimal positive xG adjustments, not significantly greater than those of Barkley. Shelvey aligns with expectations, displaying substantial negative xG adjustments based on his conversion rate in the data. Lastly, the “other” group centres around 0 for xG adjustment, as anticipated, given its diverse player composition with no discernible group effect to capture.

Figure 9 displays the shot locations and outcomes for the selected players, offering insights into the findings presented in Figure 8. Beginning with Pirès, notable for his substantial positive xG adjustments, the observation centres on his efficiency in goal scoring despite a relatively low number of shots. His ability to score from challenging positions, such as both corners of the box and outside the area, contributes to the positive xG adjustments, indicating his prowess as a goal scorer even in demanding scenarios. Agüero and Vardy exhibit similar shot patterns, but the model assigns significantly higher positive xG adjustments to Agüero. This discrepancy may stem from the nature of Vardy’s shots being inherently high xG chances, like one-on-one opportunities, whereas Agüero manages to convert more challenging shots, resulting in larger adjustments. Comparing Vardy with Coutinho and Barkley, who exhibit similar xG adjustments in Figure 8, suggests that their goal-scoring patterns align with baseline xG values without substantial player adjustments. Lastly, Shelvey’s shot map lacks goals from various positions. While difficult-to-score shots receive minor adjustments, centrally located missed chances likely contribute to the notable negative adjustments, reflecting Shelvey’s poor conversion rates in this dataset.

Figure 10 displays the cumulative data for goals scored, baseline expected goals (xG) from the single-level model, and adjusted xG from the player-corrected model for the selected

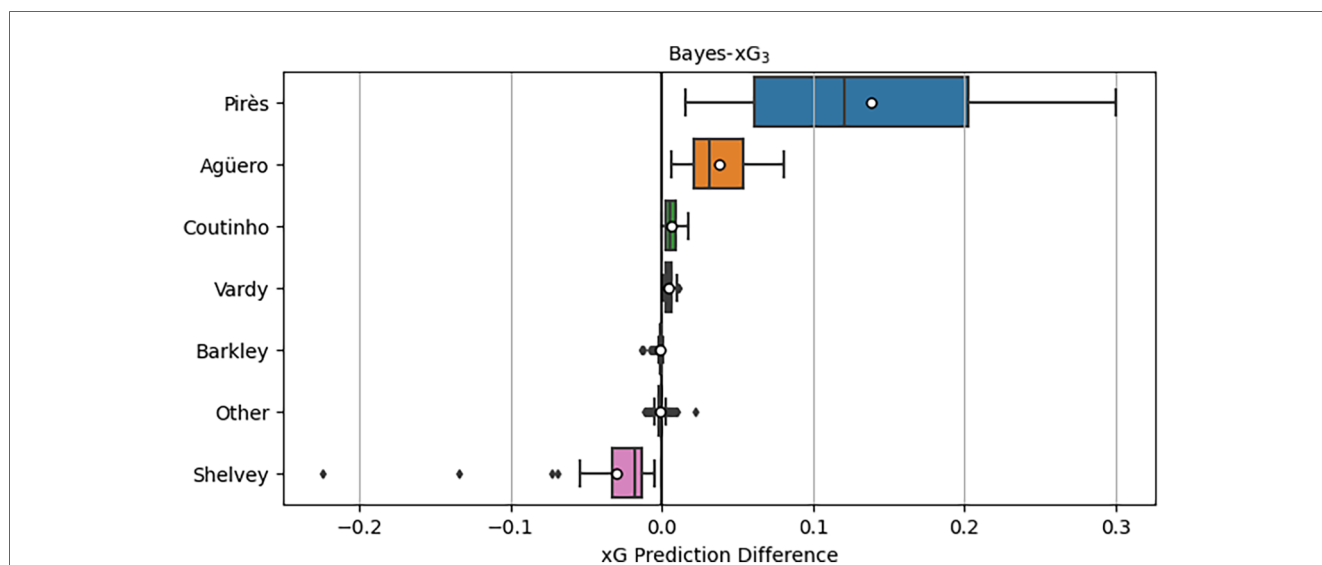
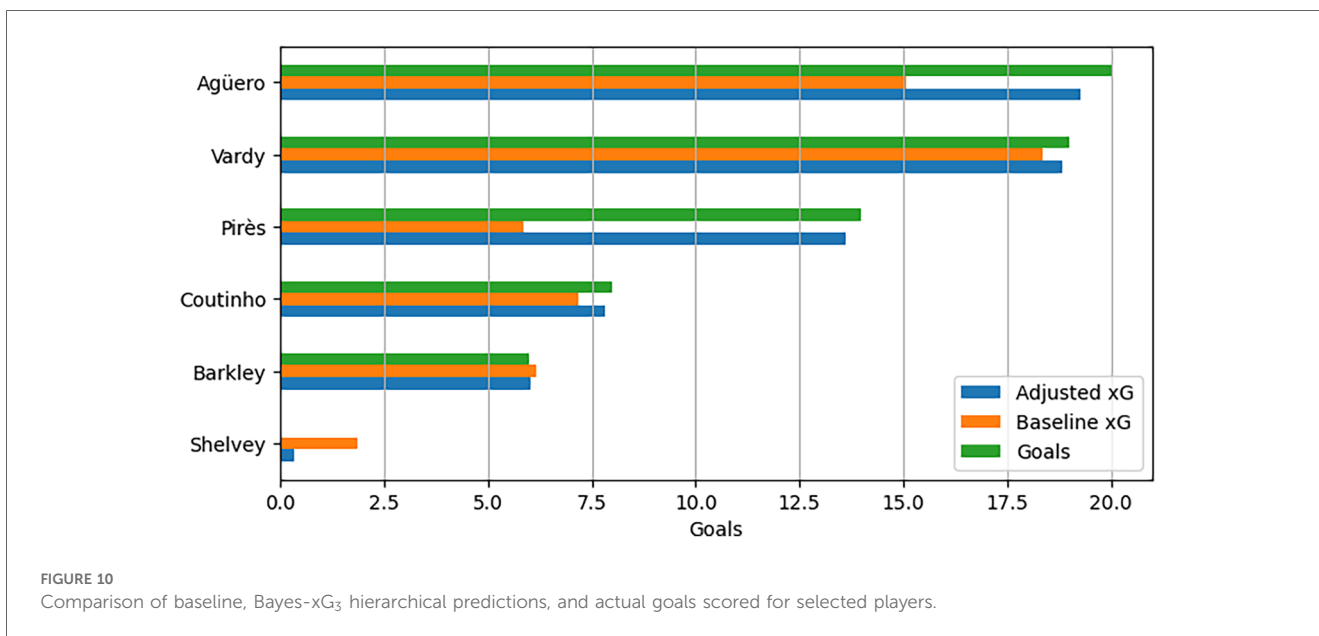
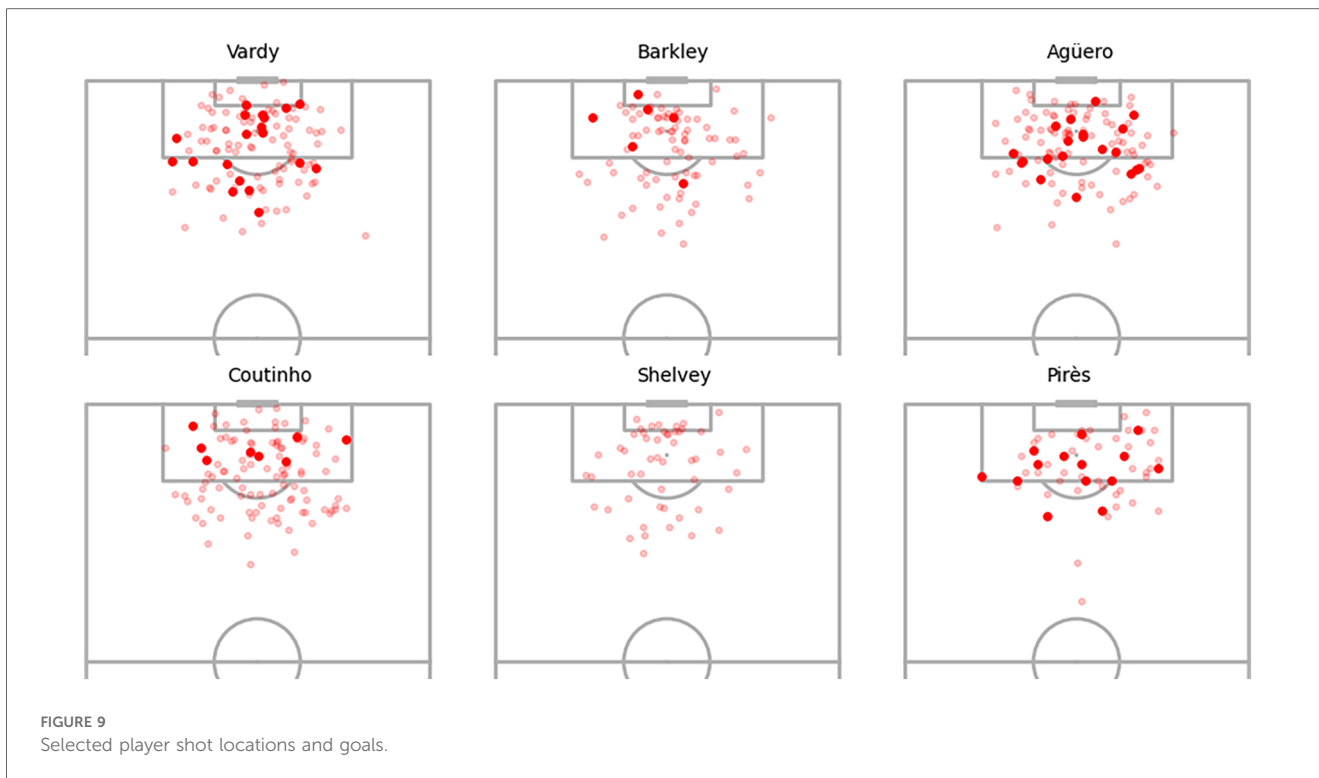


FIGURE 8 Distributions of xG adjustments for Bayes-xG₃, where adjustment is hierarchical model prediction minus baseline model prediction - grouped by player. Mean xG for each player shown as “o” in boxplots.



players in this analysis. The visual representation illustrates that, in comparison to the single-level model, the player-corrected model provides more accurate estimates of total goals scored by each player. Notably, players like Pirès and Agüero, who outperformed their baseline xG by scoring difficult chances, exhibit adjusted xG totals much closer to their actual goals scored. Conversely, Shelvey’s adjusted xG total is more aligned with the zero goals he scored, although it is crucial to emphasise that it is not precisely zero.

4.4 Extension into other leagues

To ascertain the generalizability of the conclusions drawn in this study beyond the Premier League dataset examined earlier and their applicability to football on a broader scale, as our fourth experimental case, a parallel analysis was conducted using data from Spain’s La Liga and Germany’s Bundesliga. The datasets for these leagues encompass approximately 19,000 and 7,500 shots, respectively. It is noteworthy that the La Liga dataset

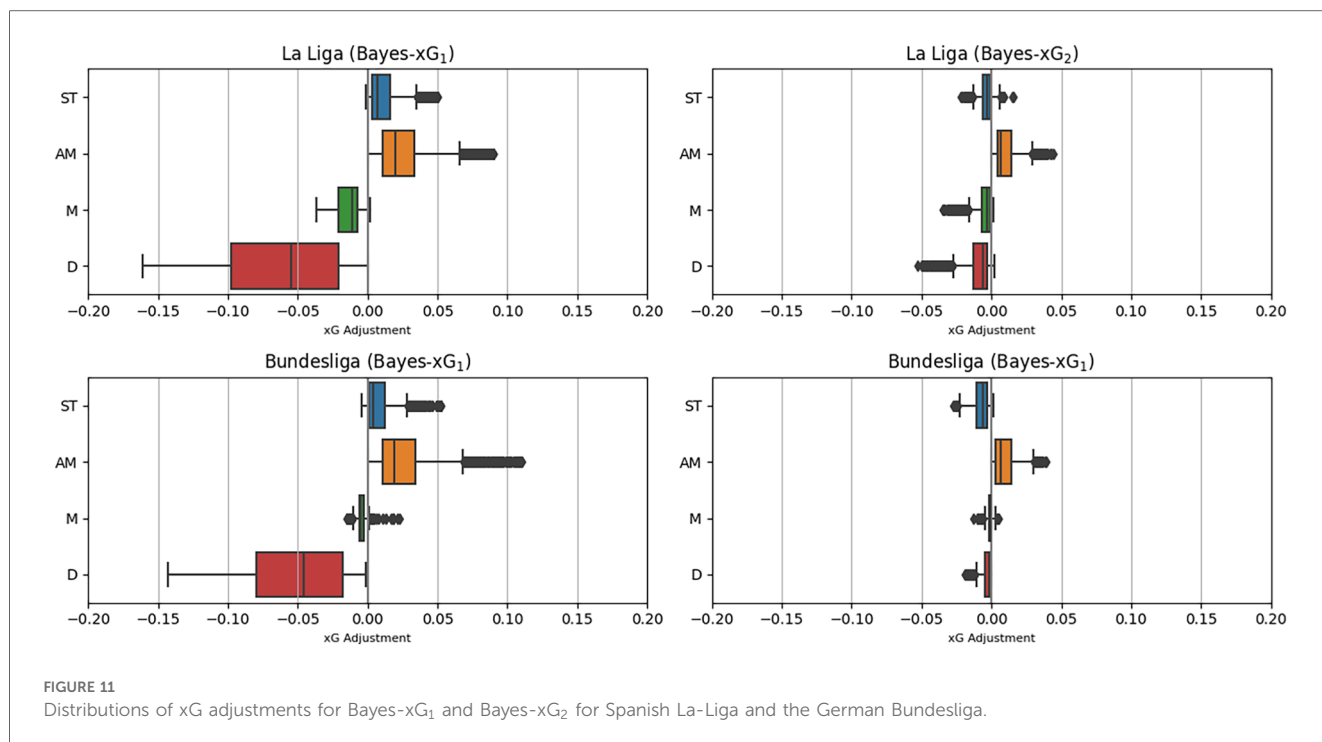


FIGURE 11 Distributions of xG adjustments for Bayes-xG₁ and Bayes-xG₂ for Spanish La-Liga and the German Bundesliga.

is notably influenced by Barcelona, primarily due to the fact that StatsBomb predominantly released data from games in which Lionel Messi, a prominent Barcelona player, participated.

Figure 11 displays the distributions of xG adjustments for the Bundesliga and La Liga, which serve as the data sets for the Bayes-xG₁ and Bayes-xG₂ models. The outcomes for both leagues closely resemble those of the Premier League in the Bayes-xG₁ model. The average adjustments follow a consistent order across all leagues, with attacking midfielders exhibiting the most positive adjustments, followed by strikers, other midfielders, and then a substantial drop to defenders. Additionally, the magnitudes of these adjustments exhibit similar patterns. The Bayes-xG₂ model reveals a noteworthy reduction in the magnitude of xG adjustments, a trend observed consistently across all leagues. Although the Bayes-xG₂ model indicates slightly larger adjustment magnitudes for the additional leagues, the variation is not significant enough to alter the model results significantly when compared to those of the Premier League.

To conduct a player-specific examination of Bayes-xG₃, the inclusion of new players from both leagues was necessary. The selection process mirrored that of the Premier League, where players were listed based on their conversion rates in the data sets, encompassing both the best and worst performers. Tables 7 and 8 display the chosen players and their respective statistics for La Liga and the Bundesliga, providing a comprehensive overview of the selected players' performance in each league.

Given these conversion rates, it is reasonable to infer that the top three performers in each table would, on average, experience positive adjustments in expected goals (xG), while the bottom three would, on average, encounter negative or negligible xG adjustments. The visual representation in Figure 12 illustrates these adjustments for players in both the Spanish La Liga (A)

TABLE 7 Selected players for Bayes-xG₃ and their scoring statistics from La Liga data set.

Player	Shots	Goals	Conversion rate
Gareth Bale	89	20	22.50%
Lionel Messi	1,862	375	20.10%
Samuel Eto'o	295	62	21%
Bebé	74	2	2.70%
Rafael Márquez	53	2	3.80%
Andrés Iniesta	362	25	6.90%

TABLE 8 Selected players for Bayes-xG₃ and their scoring statistics from Bundesliga data set.

Player	Shots	Goals	Conversion rate
Javier Hernández	63	16	25.40%
Pierre-Emerick Aubameyang	107	22	20.60%
Robert Lewandowski	147	28	19%
Pascal Groß	56	1	1.80%
Hakan Çalhanoğlu	51	1	2%
Timo Werner	64	6	9.40%

and the German Bundesliga (B). As anticipated, players like Bale, Messi, and Eto'o from La Liga predominantly exhibit positive xG adjustments, aligning with expectations. Conversely, other selected players either display minimal adjustments or predominantly negative adjustments. In the Bundesliga context, Figure 12B reaffirms many anticipated outcomes. An intriguing finding is that Aubameyang, despite maintaining a high conversion rate, exhibits predominantly negative xG adjustments. This suggests that the goals he scores tend to be from chances with already high expected goals, contrary to some expectations.

Figure 13 illustrates the shot locations of Bundesliga players, validating that the majority of Aubameyang's shots originate from

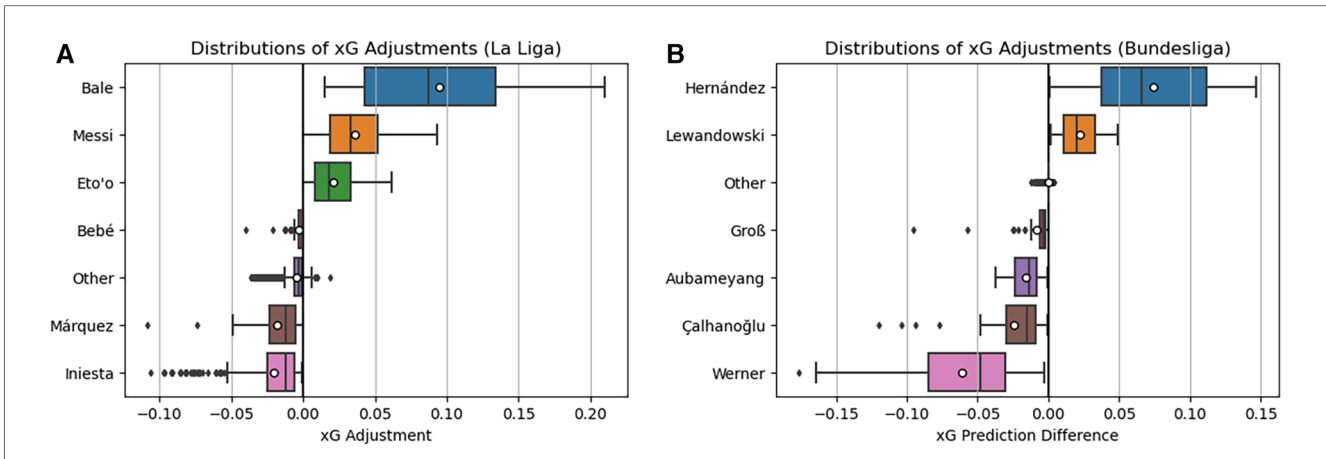


FIGURE 12 Distributions of xG adjustments for Bayes-xG₃ for Spanish La-Liga (A) and the German Bundesliga (B). Mean xG for each player is shown as “o” in boxplots.

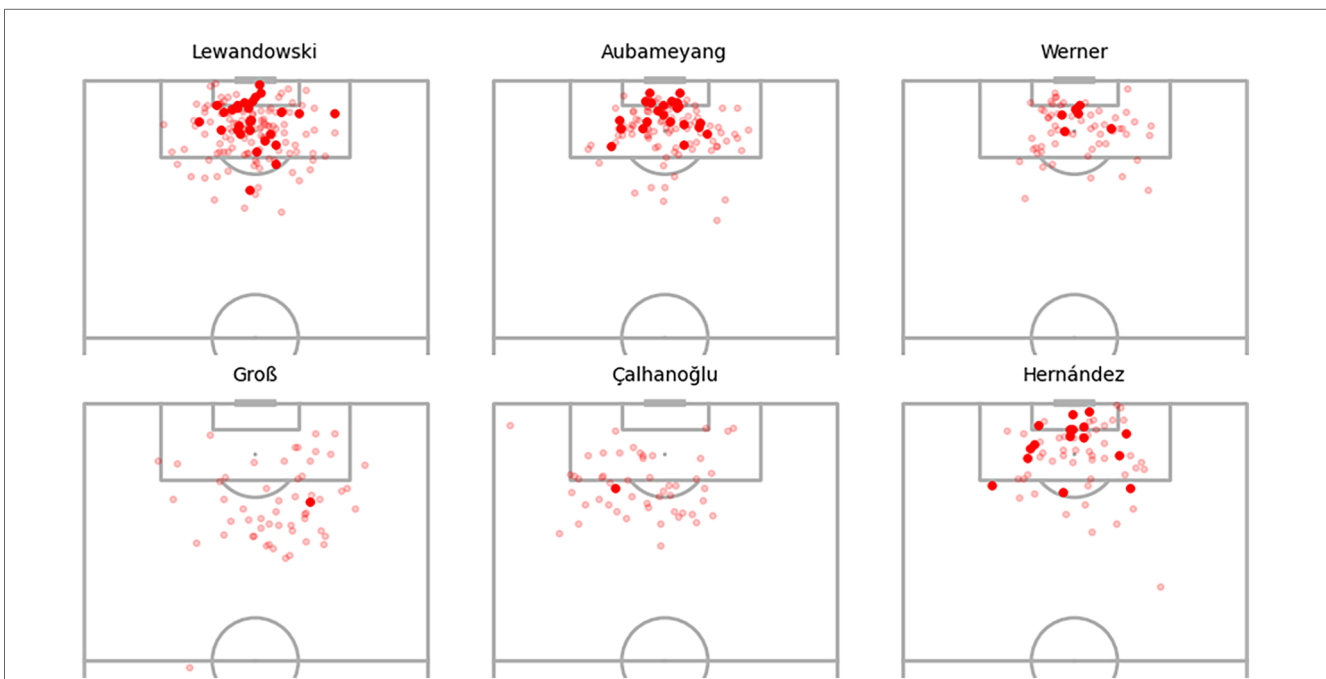


FIGURE 13 Selected player shot locations and goals for the German Bundesliga.

within the penalty area. This observation implies that these shots likely possess additional characteristics, such as one-on-one opportunities, making them high expected goals (xG) chances. In contrast, Timo Werner exhibits minimal goals despite a comparable shot map, and the substantial negative xG adjustments suggest that he should have scored more from these shot positions. On a different note, Çalhanoğlu records relatively few goals from shots that present higher difficulty due to their distance from the goal, resulting in slightly fewer negative xG adjustments.

In conclusion, the findings from both La Liga and the Bundesliga verify the results obtained in the Premier League sections above. Notably, there is an indication of a positional

impact on xG in a fundamental xG model (Bayes-xG₁), but such effects markedly diminish with the adoption of a more intricate model (Bayes-xG₂). Nevertheless, even with the extended model, there remains evidence of player-specific effects on xG, providing a quantitative measure of how certain players excel or lag behind others in scoring.

4.5 Analysing the choice of priors

The impact of prior choices is particularly evident when assessing the efficiency and accuracy of Bayesian models.

TABLE 9 Choices of prior distributions for analysing the impact of prior choice on results.

Predictor	2-Wide	3-Tight	4-Wide	5-Tight	6-Ill-suited
	Uniform	Uniform	Normal	Normal	Prior
Intercept	$\mathcal{U}(-100, 100)$	$\mathcal{U}(-1, 1)$	$\mathcal{N}(0, 10)$	$\mathcal{N}(0, 0.25)$	$\mathcal{SN}(0, 0.25, 2)$
distance_to_goal					$\mathcal{SN}(0, 0.25, 2)$
shot_angle					$\mathcal{N}(0, 0.25)$
distance*angle					$\mathcal{SN}(0, 0.25, -2)$
gk_distance_to_goal					$\mathcal{N}(0, 0.25)$
players_in_shot_triangle					$\mathcal{SN}(0, 0.25, \{-5, \dots, 5\})$
opponents_in_radius					$\mathcal{SN}(0, 0.25, \{1, \dots, -2\})$
shot_body_part					$\mathcal{N}(0, 0.25)$
shot_first_time					$\mathcal{N}(0, 0.25)$
gk_in_shot_triangle					$\mathcal{SN}(0, 0.25, 2)$
shot_one_on_one					$\mathcal{SN}(0, 0.25, -2)$
shot_open_goal					$\mathcal{SN}(0, 0.25, -4)$
shot_technique					$\mathcal{N}(0, 0.25)$
under_pressure					$\mathcal{SN}(0, 0.25, 2)$

Sensitivity to the choice of priors is crucial, and careful consideration is needed to strike a balance between the informativeness of priors and their impact on computational efficiency. In Bayesian hierarchical modelling, this emphasizes the need for a thoughtful approach to prior selection, ensuring that the chosen priors align with the characteristics of the data and contribute to the model's robustness and reliability.

The selection of prior distributions holds significant importance in Bayesian hierarchical modelling for several reasons. Firstly, the choice of prior widths plays a pivotal role. Employing overly wide priors, encompassing a large range of values, may necessitate an extensive number of samples for the model to converge to the true values of the variables, leading to prolonged computation times. Conversely, if the prior distributions are excessively narrow and the true values are unlikely to be sampled, the model's outcomes may be biased and inaccurate.

To evaluate the appropriateness of the chosen prior distributions, a reassessment will be conducted by refitting the extended baseline model (single-level) with different priors, followed by a comparison of predictions. The sets of priors under consideration include (1) the existing priors given in Table 3, (2–3) a pair of wide-tight uniform priors, (4–5) a pair of wide-tight normal priors, and (6) a deliberately ill-suited set of priors. The choices of priors for the cases of (2–6) are presented in Table 9. For the uniform priors, each variable has been given a wide and tight uniform distribution pair. Similarly, for normal priors, zero mean priors are chosen with two different σ values to represent a wide and tight value support. On the other hand, the ill-suited priors have been given very narrow distributions by using a small value for σ . Moreover, some of the skews in the distributions have been flipped such that the prior belief about the effect is the reverse of what was actually used.

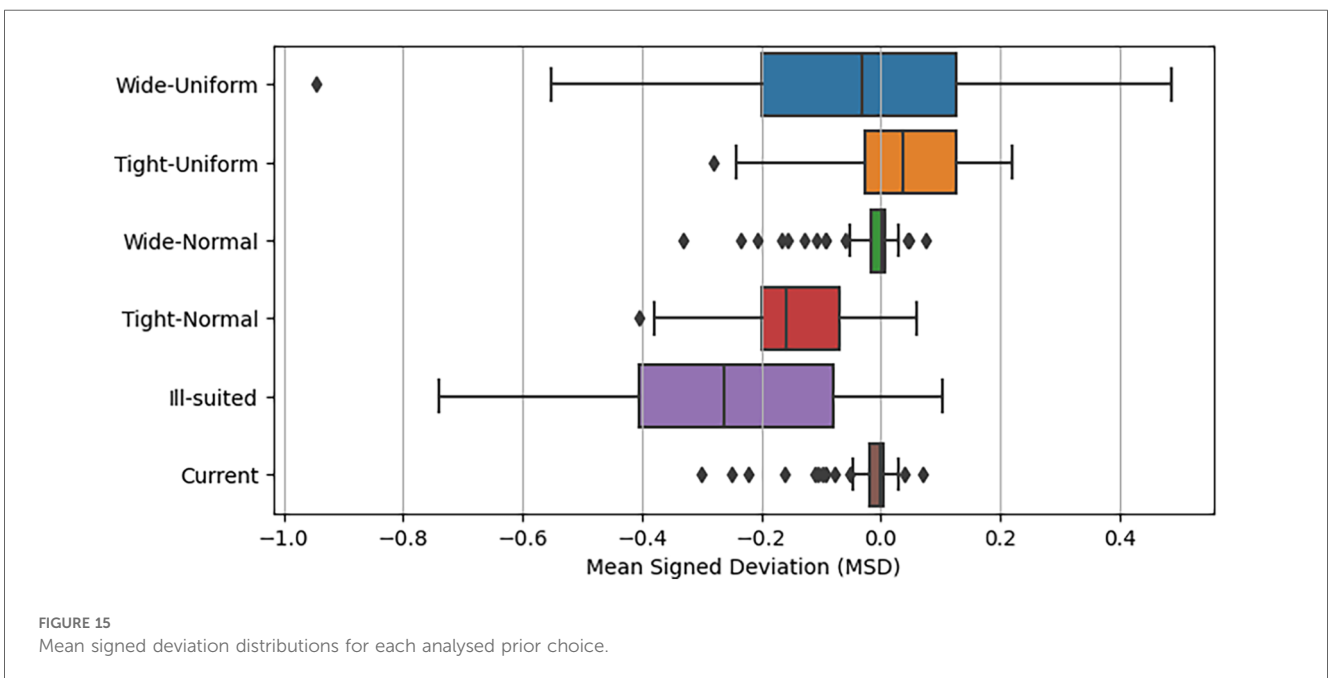
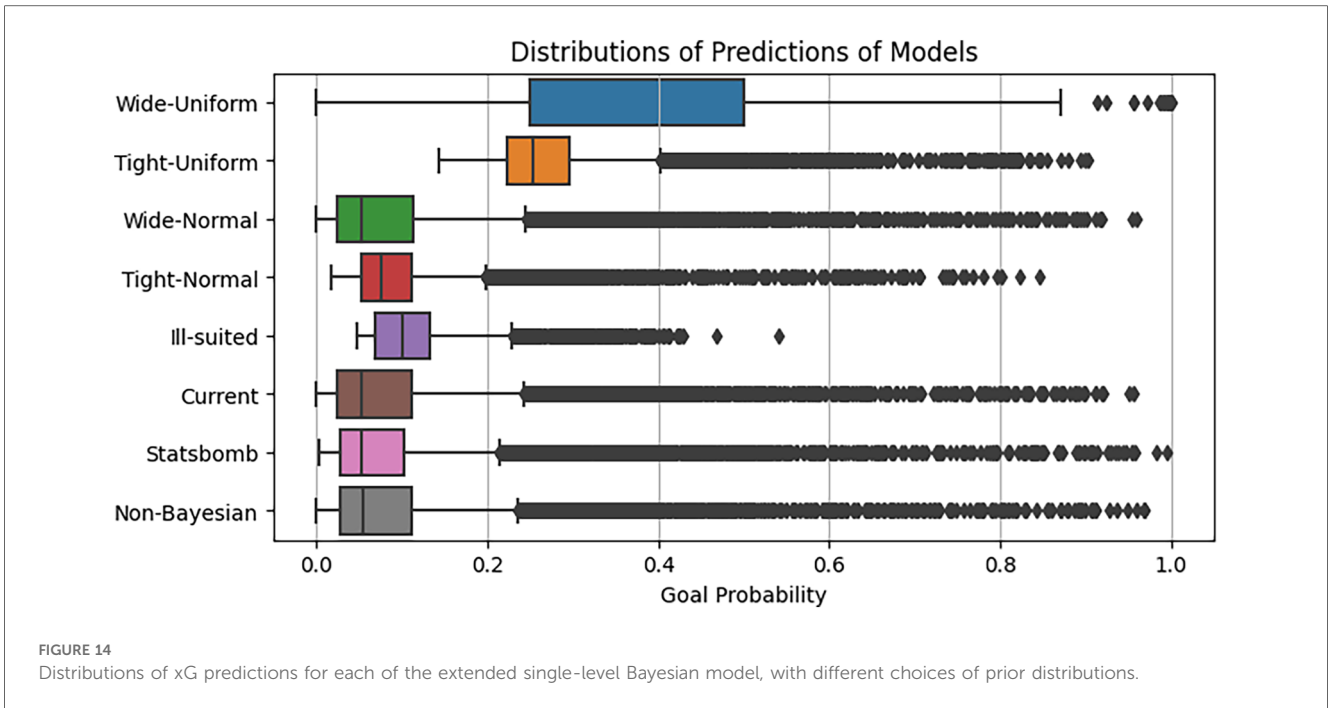
Furthermore, the predictions generated by these models will be juxtaposed with those of the extended non-Bayesian model on the same data, providing a baseline for assessing the efficacy of the selected priors in yielding accurate predictions.

Figure 14 illustrates the distributions of predictions generated by each of the prior models. The model utilising wide uniform

prior distributions exhibits notably poor performance when compared to both the non-Bayesian baseline model and the Statsbomb benchmark, displaying a considerable spread of predictions. On the other hand, employing tight uniform priors leads to more restricted predictions, although the average expected goals (xG) predictions tend to be relatively smaller. The utilisation of a tight normal prior yields similar performance, primarily underestimating xG values, particularly with the highest xG values concentrated around 0.8. In contrast, adopting wide normal priors results in enhanced performance compared to the tight setting, with predictions following a similar trend to the existing prior configuration. It is noteworthy that the mean and interquartile ranges of both normal priors depicted in Figure 14 exhibit a favourable correspondence with the baseline and benchmark models.

On the other hand, as depicted in Figure 14, ill-suited priors result in a notably narrow spread, with scarce xG predictions exceeding 0.5. Despite this, the ill-suited priors exhibit improved performance compared to uniform priors in terms of aligning the average with the baseline and benchmark predictions and maintaining a similar-sized interquartile range. This improved performance is likely attributed to the greater number of samples utilised for parameter estimation. In the case of the model with uniform priors, the same number of samples, however, prevented the model from converging to optimal parameter values, resulting in poor predictions. While, given more samples, this model could eventually yield accurate results, the computational time required is uncertain and could be extensive. The model with the existing prior distributions outperforms the others significantly, closely resembling the distribution of baseline and benchmark models' predictions while also offering valuable insights based on player positions.

In the 5th experimental case's final analysis, we explore the mean signed deviation (MSD) values between each prior case and the predictions of the Non-Bayesian extended model. The selection of MSD aims to emphasize instances of over or underprediction based on the prior choice, using the mean spread of MSD values as a performance metric. Figure 15



illustrates the distributions of MSD values for each prior group through various box plots. Notably, Figure 15 highlights the considerable performance of the Wide Normal prior choice, exhibiting results akin to the current prior approaches. Its interquartile range closely aligns with the current prior case, albeit with a few more instances of over-predicted outliers. It is crucial to observe that both uniform priors' MSD values are distributed around zero, despite with a broader spread. In contrast, the Tight-Normal and Ill-suited prior cases exhibit notably poor performance, marked by a higher frequency

of overestimated xG values compared to the Non-Bayesian extended model.

4.6 Uncertainty quantification

The final experimental case of this study involves examining and reflecting on the uncertainty inherent in the models and priors proposed and employed. While motivating for a Bayesian approach in earlier sections, we highlighted how Bayesian models

treat parameters as unknowns and assign random variables to represent the uncertainty surrounding these parameters. Each of these random variables is linked to a probability distribution, ultimately yielding a posterior distribution of estimations as a result of the inference process. In this subsection, unlike the

previous experimental cases, we visualize and discuss the posterior predictive densities of model outcomes to more accurately quantify the uncertainty offered by Bayesian xG models.

We initiate by examining the positional Bayesian models of Bayes-xG₁ and Bayes-xG₂ in this section. For each instance

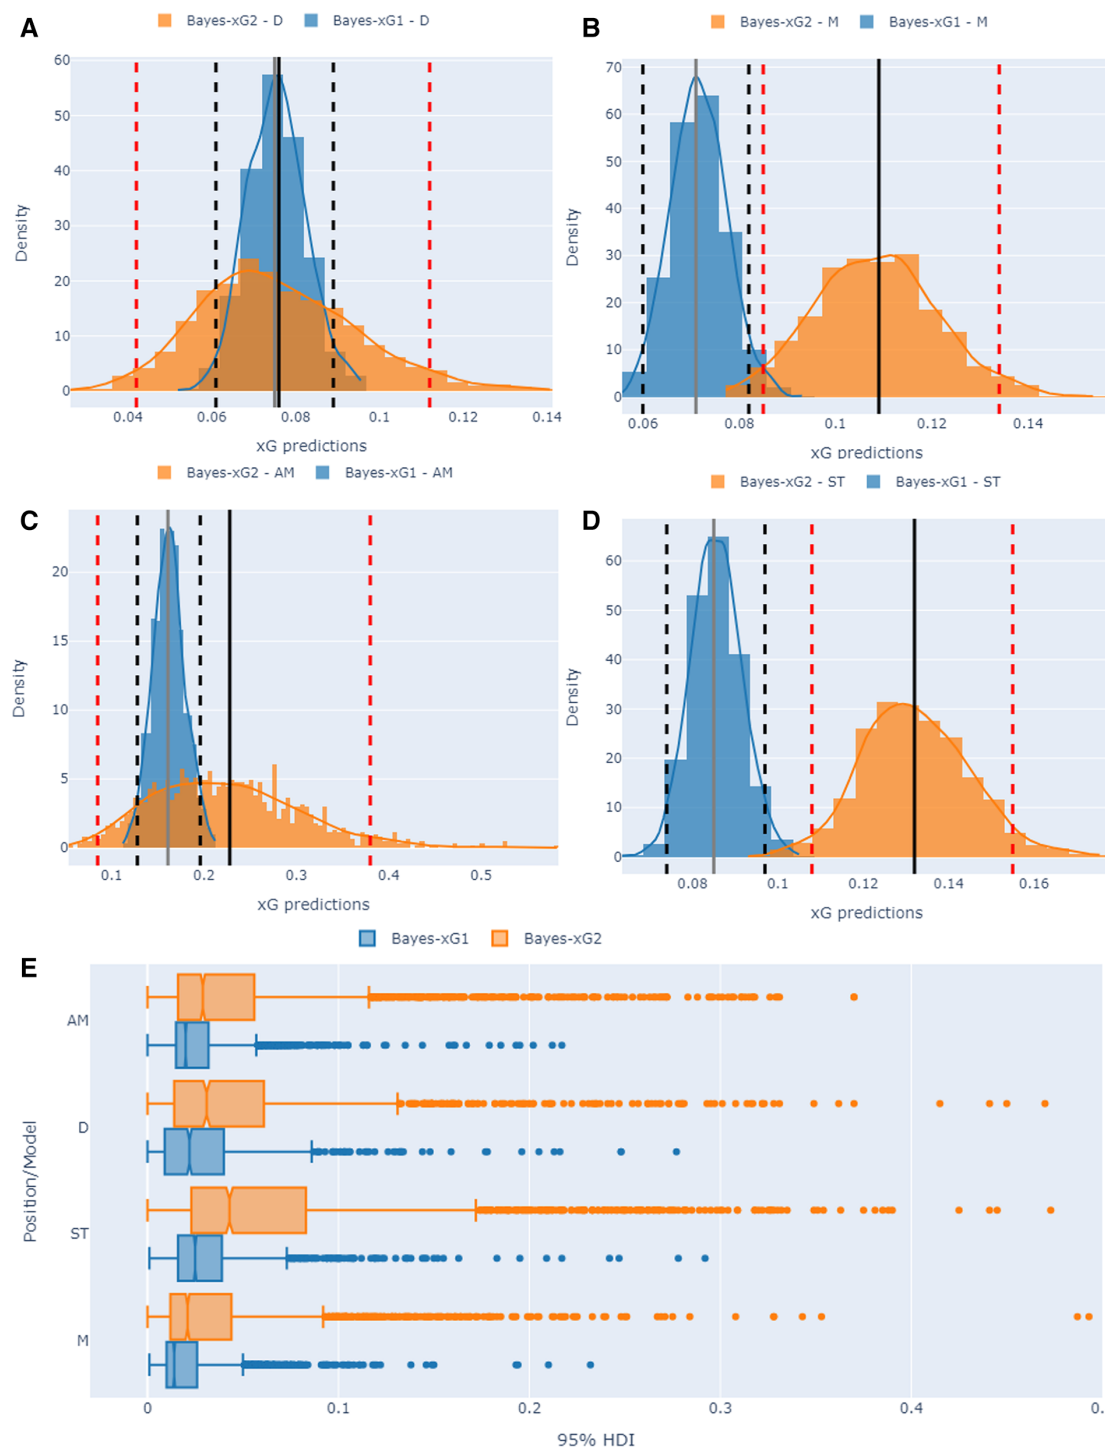


FIGURE 16 Uncertainty quantification visuals for positional Bayesian models (Bayes-xG₁ and Bayes-xG₂). Histograms in (A) to (D) refer to comparing two Bayesian models for the same shot with a Statsbomb xG of 0.15. Dashed vertical lines refer to 2.5% and 97.5% HDI whilst solid vertical lines refer to mean predicted xG values. (E) shows 95% HDI box plots for each position/model pair.

specific to a position in the English Premiership dataset, we selected a shot with a Statsbomb xG of approximately 0.15, resulting in a goal. Subsequently, we illustrated the posterior predictive densities of these two Bayesian positional models alongside their corresponding 95% high-density intervals (HDI) (2.5% from each tail). Figures 16A–D illustrate the outcomes of this uncertainty analysis for the D, M, AM, and ST positions, respectively. Additionally, Figure 16E presents a boxplot distributional analysis of the 95% HDI for each position obtained from both Bayesian models.

Analyzing Figures 16A–D, it becomes evident that the Bayes-xG₂ model consistently exhibits higher uncertainty in its xG

predictions across all positional models. This is indicated by its broader shape and lower probability peak, as observed in the analysis, where the 95% highest density interval (HDI) values consistently surpass those of the baseline positional model, Bayes-xG₁. Conversely, Bayes-xG₁ appears to consistently underestimate xG values compared to Bayes-xG₂, despite its narrower HDI ranges. This discrepancy can be understood from a probabilistic standpoint, where the enhanced positional model of Bayes-xG₂, incorporating more comprehensive engineered features, challenges the inference process with a more nuanced perspective. This is reflected in its wider 95% HDI values, which effectively account for positional effects.

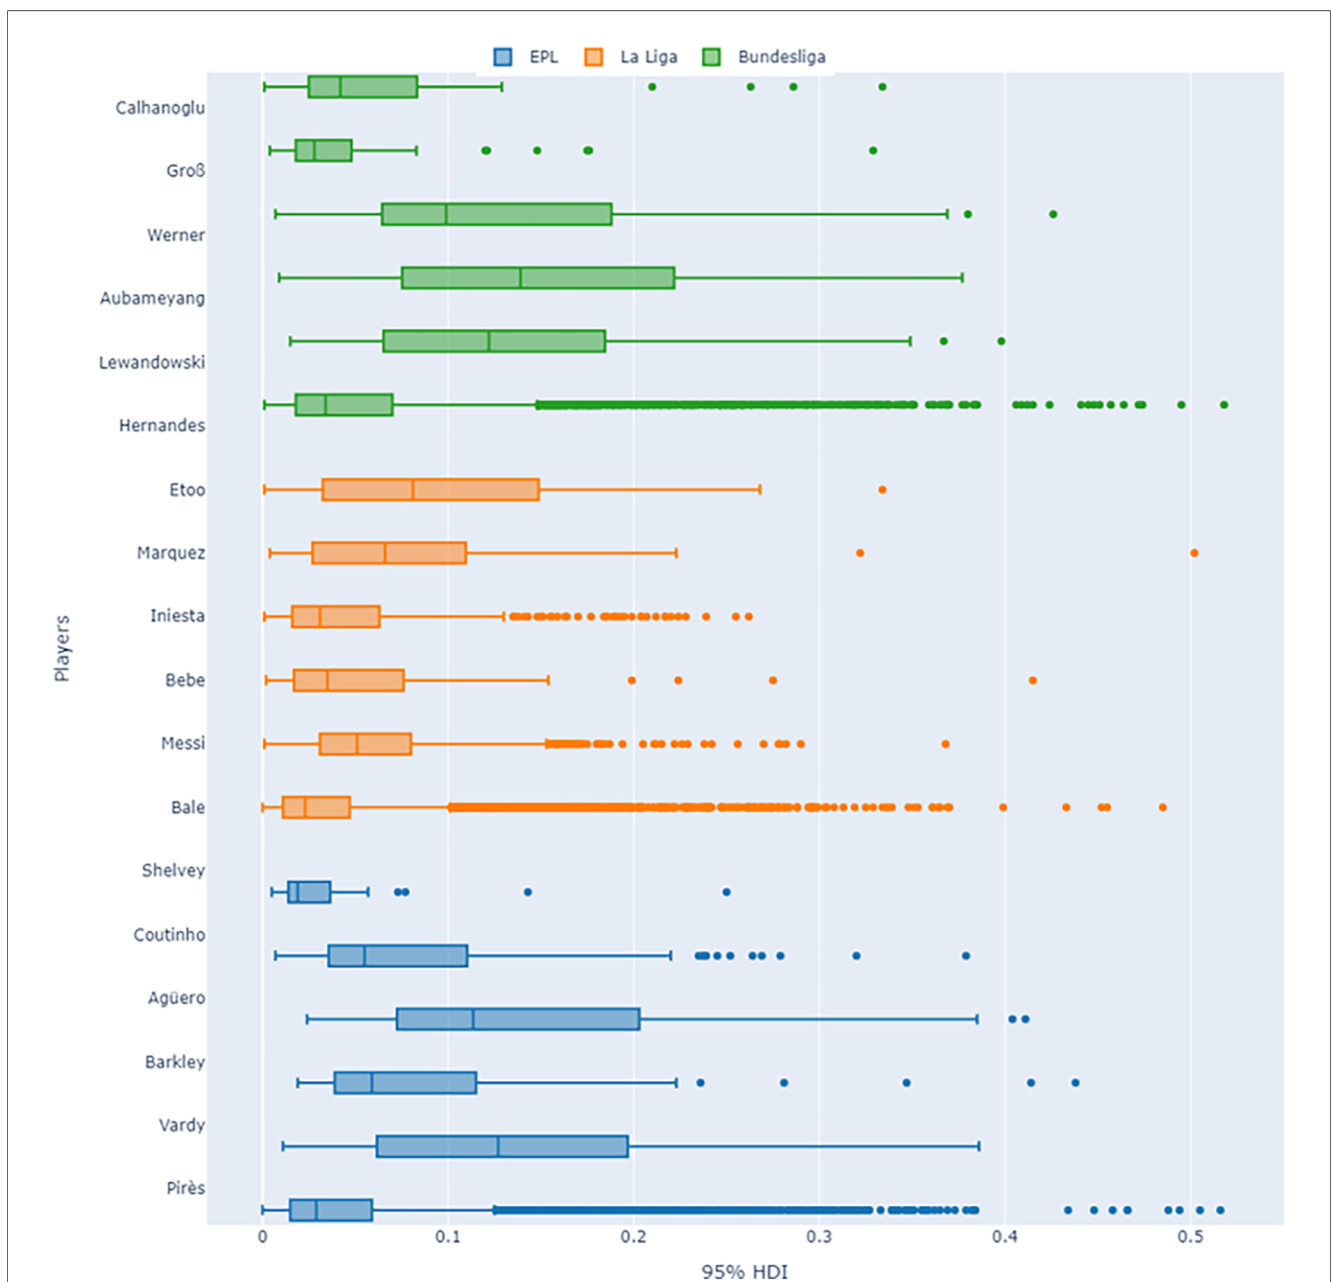


FIGURE 17
Uncertainty quantification visuals for player-specific Bayesian model (Bayes-xG₃).

This observation is further supported by Figure 16E, where all Bayes-xG₂ models demonstrate broader 95% HDI values across all positional parameters. Additionally, it is noteworthy that HDI intervals for defensive (D) and striker (ST) positions exhibit greater width compared to midfielder (M) and attacking midfielder (AM) positions. This implies that the positional corrections applied to xG predictions are more pronounced for defensive and striker positions, encompassing both positive and negative adjustments.

The higher uncertainty outcomes of Bayes-xG₂ model may initially seem concerning, but it actually signifies a more robust

and adaptable model that performs well in capturing the intricacies of football dynamics. It reflects the model's ability to account for the complexity and variability inherent in shots in football matches, where numerous factors can influence goal-scoring opportunities. The richer set of engineered features in Bayes-xG₂ enables it to capture more nuanced patterns and positional effects, thus providing a more comprehensive assessment of xG probabilities with wider HDIs.

For the player-specific analysis in our study, we examined the posterior predictive results derived from the extended player model of the Bayes-xG₃ model. Figure 17 illustrates a boxplot

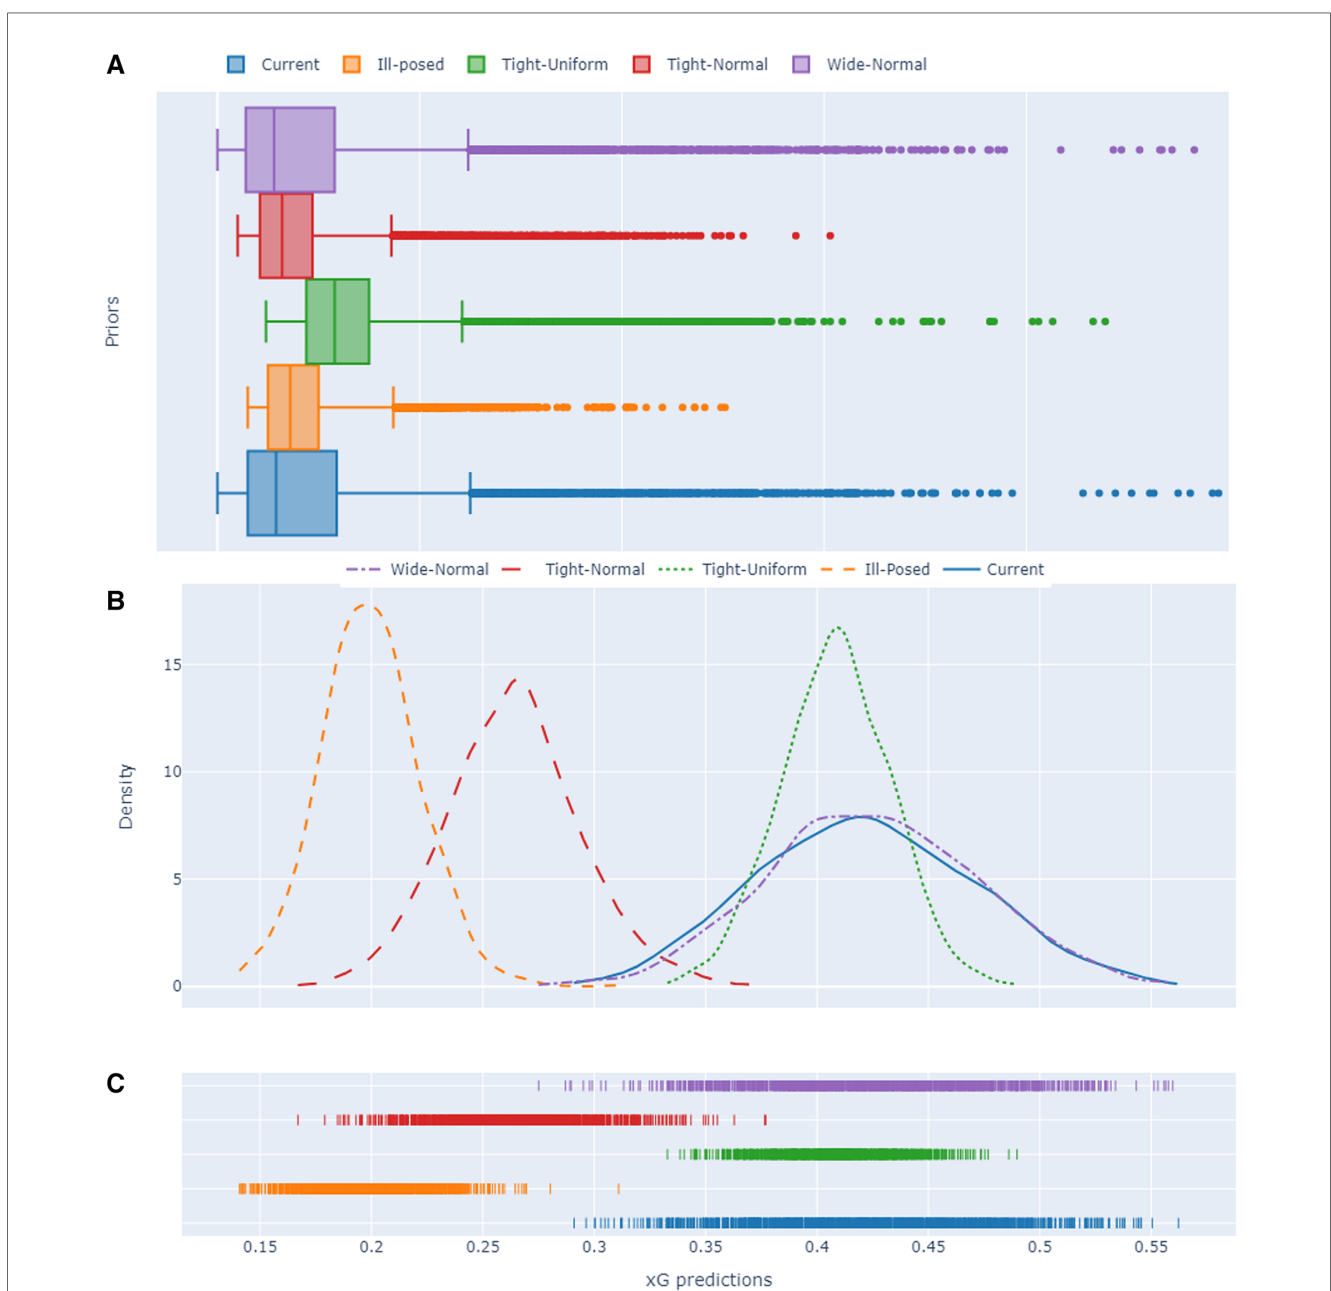


FIGURE 18 Uncertainty quantification visuals for different prior choices associated with the extended single-level Bayesian model. (A) 95% HDI box plots, (B) Posterior predictive estimate distribution plots and (C) rug plots for a randomly selected shot from the EPL data set with a Statsbomb xG value of approximately 0.5.

distributional analysis of the 95% HDIs for six players across three leagues. Upon reviewing the obtained 95% HDI values for our player set from the English Premier League (EPL), La Liga, and Bundesliga using the Bayes-xG₃ model, several observations emerge. Notably, players with significant positive or negative xG corrections, such as Hernandez (+), Werner (-), Bale/Messi (+), Shelvey (-), and Pires (+), exhibit higher HDI values. Particularly, among the top positively corrected players from each league (Hernandez, Bale, and Pires), we observe markedly low median HDIs with numerous outliers. This suggests that these players' shooting abilities yield nuanced and detailed posterior predictive estimates, indicating their scoring capabilities differ significantly from the average player in the dataset. Consequently, this heightened uncertainty in their model outcomes serves as a meaningful indicator within the context of our study.

Additionally, higher uncertainty in Bayes-xG₃ model outcomes provides more insightful perspectives for the performance analysis of players. When a player's predicted outcomes exhibit greater variability, it implies that their performance is not easily characterised by conventional metrics alone. Instead, it suggests a more complex and multifaceted contribution to their team's performance. By acknowledging and delving into this uncertainty, we believe one can uncover hidden strengths or weaknesses that may not be apparent when considering only deterministic and generalised predictions. Thus, embracing uncertainty opens avenues for a deeper understanding of player performance dynamics and facilitates more informed decision-making in player evaluation and team strategy development.

The final experimental set of this paper addresses the assessment of uncertainty arising from the selection of priors for the Bayesian models proposed. Alongside the hierarchical modelling of model parameters, the choice of priors represents a crucial aspect in the application of Bayesian models. While we previously examined the impact of various priors in the preceding subsection, our focus now shifts to evaluating the uncertainty inherent in the single-level extended Bayesian model resulting from the selection of prior distributions. **Figure 18** illustrates three subplots depicting this analysis. The first subplot (A) showcases a boxplot distributional examination of the 95% HDIs for each prior. Following this, a specific shot from the EPL dataset with a Statsbomb xG value of approximately 0.5 is utilized to generate posterior predictive estimates for each prior. These estimations are displayed in the form of a distribution plot (B) and a rug plot (C).

The issue of high uncertainty persists in the analysis of the best-performing models. As depicted in **Figure 18A**, both the proposed priors and Wide-Normal priors demonstrate comparable performance, with their HDI values notably larger than those of the less informative priors, Ill-posed and Tight Uniform. Upon examining the xG predictions under different prior selections in **Figures 18B,C**, it becomes evident that although less informative priors yield lower uncertainty, their xG predictions tend to be significantly underestimated, approximately around 0.2 and 0.27 for Ill-posed and Tight-Normal priors respectively. Conversely, the remaining three prior

choices yield similar xG predictions, with the Tight-Normal model exhibiting less uncertainty. Notably, the proposed and Wide-Normal priors, with their posterior predictive outcomes characterised by high uncertainty, offer a broader range of variability in xG predictions.

5 Final remarks & conclusion

The objective of this study was to explore whether there exist distinct effects based on player position on expected goal (xG) values, suggesting that different positions or players may exhibit specific xG adjustments for given scoring opportunities. The hypothesis was that proficient attacking players, such as strikers, attacking midfielders, or those recognised for their offensive prowess, would demonstrate positive xG adjustments. In contrast, players less renowned for their attacking contributions, such as defenders or those with defensive roles, were expected to show negative xG adjustments.

To reach the objective mentioned above, this study has developed several Bayesian models to evaluate the influence of a player's position and individual player effects on xG predictions. Initially, a basic xG model (Baseline xG), incorporating only distance to the goal, shot angle, and their interaction, indicated positional effects on xG (Bayes-xG₁). Strikers and attacking midfielders exhibited positive xG adjustments, midfielders displayed minimal adjustments, while defenders had notably negative xG adjustments on average. However, the introduction of additional predictors in the models diminished the positional effects to the extent that they became almost insignificant (Extended xG), suggesting that player position had minimal impact on xG when considering more shot-related factors (Bayes-xG₂). Subsequently, player effects were explored using the extended model employed for the second positional-effects model (Bayes-xG₃), grouping the data based on the player's shooting rather than the shooter's position. The model was illustrated using six players from each dataset from three of the European Top 5 leagues, revealing significant player effects on xG even when controlling for various shot-related factors. These effects were diverse in direction, notably positive for R. Pirés (as well as for G. Bale and J. Hernandez) and negative for J. Shelvey (as well as for A. Iniesta and T. Werner).

The indication that there exist player-specific effects in determining goal probability could prove beneficial in football scouting and player selection. By computing adjusted xG values for various players and comparing these adjusted values to their non-adjusted counterparts, as demonstrated in this analysis, it becomes possible to distinguish players who excel at converting challenging opportunities from those who consistently find themselves in advantageous positions. Examining the results for the English Premier League dataset, particularly for J. Vardy, reveals that his total adjusted xG is not significantly different from his baseline xG (see **Figure 10**). This suggests that, given the quality of chances Vardy receives, he scores at a relatively average rate. On the other hand, S. Agüero demonstrates a more consistent ability to score from more challenging positions,

evident in his larger adjusted xG. It is important to note that this observation does not imply that Agüero is a superior player or attacker compared to Vardy. Instead, it suggests that, on average, Agüero is more adept at converting chances with lower xG values than Vardy, indicating proficiency in scoring from less favourable situations.

It is important to acknowledge that there might also be team-related influences at play in this context. To further compare Vardy and Agüero, Vardy is part of a Leicester team known for its high-tempo and direct attacking style. This approach likely leads to shooting scenarios where the ball is played behind the defence, creating situations with fewer defenders to obstruct or impede a shot. This dynamic often results in one-on-one opportunities with the goalkeeper, contributing to Vardy consistently receiving numerous high xG chances. Conversely, Agüero played for one of the top teams in the league, causing opponents to adopt a more conservative approach. Teams facing Manchester City tend to minimise space, making chances more challenging with multiple players in the shot triangle and other complicating factors. We also recognise a potential limitation of the player-specific analysis proposed in this study concerning average-performing players. Our analysis emphasises players who either exceed their expected goals (xG) significantly or fall short, potentially overlooking consistent players who consistently score high xG but also miss smaller xG opportunities. This does not imply that these players are unworthy of analysis; rather, they represent player profiles that are less complex and easier to predict. However, effectively addressing the nuances of consistent players of this nature falls outside the scope of our current paper but remains an area for future research exploration.

Lastly, the examination of uncertainty analysis, particularly in the context of position-specific and player-specific corrections on xG, illuminates crucial insights into the intricacies of performance assessment in sports analytics. Through methodologies like Bayesian regression and posterior predictive distributions, researchers can effectively quantify and interpret the variability in xG predictions, taking into account factors such as prior knowledge, model parameters, and observed data. The observation of higher uncertainty in certain player positions or individual player corrections underscores the complexity and variability inherent in player performance dynamics. This increased uncertainty, while concerning and challenging, presents valuable opportunities for deeper insights and more precise evaluations. By acknowledging and incorporating uncertainty into the analytical framework, future studies can enhance the accuracy and reliability of xG predictions, ultimately leading to more informed decision-making processes and strategic interventions in sports management, scouting and coaching.

The Bayesian modelling results were collocated with non-Bayesian, or frequentist, modelling. Particularly in the case of player correction, Bayesian modelling offers a significant advantage by capturing uncertainty through posterior distributions rather than relying solely on point estimates. This

becomes particularly advantageous when dealing with younger players with limited match experience, as Bayesian hierarchical modelling effectively addresses data groups with few observations. Regardless, Bayesian modelling is rarely used in the literature of football analytics. Beyond its applications in scouting and player selection, Bayesian hierarchical modelling holds promise for various metrics in football. For instance, the assessment of injury risk among players is a common practice in large football clubs, where certain players may have a higher overall susceptibility to injuries. Hierarchical modelling, in this context, has the potential to provide more accurate assessments of individual players' injury risks by considering their specific injury history.

Data availability statement

Publicly available datasets were analyzed in this study. This data can be found here: <https://github.com/statsbomb/open-data>.

Author contributions

AS: Conceptualization, Data curation, Formal Analysis, Investigation, Methodology, Software, Validation, Visualization, Writing – original draft, Writing – review & editing; OK: Conceptualization, Methodology, Supervision, Validation, Visualization, Writing – original draft, Writing – review & editing.

Funding

The author(s) declare financial support was received for the research, authorship, and/or publication of this article.

Funding was provided by Cardiff University.

Conflict of interest

The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

Publisher's note

All claims expressed in this article are solely those of the authors and do not necessarily represent those of their affiliated organizations, or those of the publisher, the editors and the reviewers. Any product that may be evaluated in this article, or claim that may be made by its manufacturer, is not guaranteed or endorsed by the publisher.

References

- Smith R. *Expected Goals: The Story of How Data Conquered Football and Changed the Game Forever*. London: Mudlark - HarperCollins Publishers (2022).
- Tippett J, *The Expected Goals Philosophy: A Game-Changing Way of Analysing Football*. London: Wydawca Nieznany (2019).
- Aberle M, Figdor L, Mongrand L, Janetzke M. *The Tech Behind the Bundesliga Match Facts Xgoals: How Machine Learning is Driving Data-Driven Insights in Soccer*. AWS Machine Learning Blog (2020). Available online at: <https://aws.amazon.com/blogs/machine-learning/the-tech-behind-the-bundesliga-match-facts-xgoals-how-machine-learning-is-driving-data-driven-insights-in-soccer/> (accessed September 1, 2024).
- Brechot M, Flepp R. Dealing with randomness in match outcomes: how to rethink performance evaluation in European club football using expected goals. *J Sports Econ.* (2020) 21:335–62. doi: 10.1177/1527002519897962
- Herold M, Goes F, Nopp S, Bauer P, Thompson C, Meyer T. Machine learning in men's professional football: current applications and future directions for improving attacking play. *Int J Sports Sci Coaching.* (2019) 14:798–817. doi: 10.1177/1747954119879350
- Lucey P, Bialkowski A, Monfort M, Carr P, Matthews I. Quality vs quantity: improved shot prediction in soccer using strategic features from spatiotemporal data. In *8th MIT Sloan Sports Analytics Conference; 2005 Feb 27–28; Boston Convention and Exhibition Center, Boston, MA, United States.* (2015).
- Anzer G, Bauer P. A goal scoring probability model for shots based on synchronized positional and event data in football (soccer). *Front Sports Active Living.* (2021) 3:1–15. doi: 10.3389/fspor.2021.624475
- Madrero Pardo P. *Creating a model for expected goals in football using qualitative player information* (Master's thesis). Universitat Politècnica de Catalunya.
- Fairchild A, Pelechrinis K, Kokkodis M. Spatial analysis of shots in MLS: a model for expected goals and fractal dimensionality. *J Sports Anal.* (2018) 4:165–74. doi: 10.3233/JSA-170207
- Cavus M, Biecek P. Explainable expected goal models for performance analysis in football analytics. In: *2022 IEEE 9th International Conference on Data Science and Advanced Analytics (DSAA)*. IEEE (2022). p. 1–9.
- Hewitt JH, Karakuş O. A machine learning approach for player and position adjusted expected goals in football (soccer). *Frankl Open.* (2023) 4:100034. doi: 10.1016/j.fraope.2023.100034
- Spearman W. Beyond expected goals. In: *Proceedings of the 12th MIT Sloan Sports Analytics Conference* (2018). p. 1–17.
- Joseph A, Fenton NE, Neil M. Predicting football results using Bayesian nets and other machine learning techniques. *Knowl Based Syst.* (2006) 19:544–53. doi: 10.1016/j.knosys.2006.04.011
- Zambom-Ferraresi F, Rios V, Lera-López F. Determinants of sport performance in European football: what can we learn from the data? *Decis Support Syst.* (2018) 114:18–28. doi: 10.1016/j.dss.2018.08.006
- Tureen T, Olthof S. “Estimated Player Impact” (EPI): quantifying the effects of individual players on football (soccer) actions using hierarchical statistical models. In: *StatsBomb Conference Proceedings*. StatsBomb (2022).
- Baio G, Blangiardo M. Bayesian hierarchical model for the prediction of football results. *J Appl Stat.* (2010) 37:253–64. doi: 10.1080/02664760802684177
- Blanco LB, Ferreira PH, Louzada F, Nascimento DC. Is football/soccer purely stochastic, made out of luck, or maybe predictable? how does Bayesian reasoning assess sports? *Axioms.* (2021) 10:1–17. doi: 10.3390/axioms10040276
- Eggels H, van Elk R, Pechenizkiy M. Expected goals in soccer: explaining match results using predictive analytics. In: *The Machine Learning and Data Mining for Sports Analytics Workshop*. Vol. 16 (2016).
- Mead J, O'Hare A, McMenemy P. Expected goals in football: improving model performance and demonstrating value. *PLoS One.* (2023) 18:e0282295. doi: 10.1371/journal.pone.0282295
- Umami I, Gautama DH, Hatta HR. Implementing the expected goal (xg) model to predict scores in soccer matches. *Int J Inform Inf Syst.* (2021) 4:38–54. doi: 47738/ijis.v4i1.76
- Statsbomb. *Statsbomb open data specification* (2019).
- Capretto T, Piho C, Kumar R, Westfall J, Yarkoni T, Martin OA. Bambi: a simple interface for fitting Bayesian linear models in python. *J Stat Softw.* (2022) 103:1–29. doi: 10.18637/jss.v103.i15
- De Finetti B. *Probability, Induction and Statistics*. Hoboken, NJ: John Wiley and Sons (1972).
- Bastos LS, da Rosa JMC. Predicting probabilities for the 2010 FIFA World Cup games using a Poisson-Gamma model. *J Appl Stat.* (2013) 40:1533–44. doi: 10.1080/02664763.2013.788619