

What If They Took the Shot? A Hierarchical Bayesian Framework for Counterfactual Expected Goals

Mikayil Mahmudlu^a, Oktay Karakuş^{b,c,d}, Hasan Arkadaş^c

^aCardiff University, School of Mathematics, Abacws, Senghennydd Road, Cardiff, CF24 4AG, UK.

^bCardiff University, School of Computer Science and Informatics, Abacws, Senghennydd Road, Cardiff, CF24 4AG, UK.

^cDead Ball Analytics Limited, Barry, CF63 2QQ, UK.

^dCorresponding Author (email

Abstract

This study develops a hierarchical Bayesian framework that integrates expert domain knowledge to quantify player-specific effects in expected goals (xG) estimation, addressing a core limitation of traditional models that treat all players as identical finishers. Using 9,970 shots from StatsBomb’s 2015–16 dataset and expert ratings from Football Manager 2017, we combine Bayesian logistic regression with FM-informed priors to stabilise player-level estimates, particularly for individuals with limited shot histories. The hierarchical model substantially reduces posterior uncertainty relative to weakly informative priors and exhibits strong external validity: hierarchical and baseline predictions correlate at $R^2 \approx 0.75$, while an XGBoost benchmark, validated against StatsBomb’s proprietary xG, achieves $R^2 = 0.833$, confirming the effectiveness of the engineered feature set. Player-specialisation analysis reveals interpretable, domain-consistent skill profiles across one-on-one finishing (Agüero, Suárez, Belotti, Immobile, Martial), distance shooting specialisation (Pogba), and first-touch execution (Insigne, Salah, Gameiro). The model detected latent ability in underperforming players such as Immobile and Belotti, whose later-season improvement empirically validated these early signals. The framework enables counterfactual “what-if” analysis by reallocating shots across players while holding contextual features fixed. Case studies demonstrate its practical utility: Sansone is predicted to generate +2.2 additional xG from Berardi’s opportunities, driven primarily by pressure situations (+1.1 xG), a result reflected in his subsequent €13M transfer to Villarreal. A bidirectional Vardy–Giroud analysis reveals strong asymmetry: Leicester would suffer a substantial decline if Giroud replaced Vardy (−7 xG), whereas Arsenal would experience only a minor drop when swapping Giroud for Vardy (−1 xG), highlighting the decisive role of tactical fit. This research provides a principled, uncertainty-aware framework for player evaluation, recruitment, and tactical planning by unifying quantitative modelling with expert scouting priors. The methodology generalises to other sports where individual skill, context, and domain knowledge jointly determine performance.

Keywords: keyword one, keyword two

1. Introduction

Football has become one of the most data-driven sports globally, with billions of fans and massive financial investments driving demand for sophisticated analytical methods. The rapid development of the football industry has made expert-level analysis crucial, particularly in scouting and transfer market evaluation. Data-driven approaches enable clubs to discover new talent, replace outgoing players effectively, and refine tactical strategies based on empirical evidence rather than subjective assessment alone. Among advanced metrics, expected goals (xG) has emerged as one of the most widely adopted and influential tools for performance evaluation. The metric estimates the probability that a shot results in a goal based on contextual features such as shot distance, angle, body part used, and defensive positioning. Various modeling approaches have been applied to xG estimation, with implementations ranging from logistic regression to decision trees and neural networks, each utilizing geometric and situational features to predict goal probability [1]. These models have become standard tools for broadcasters, analysts, and clubs seeking to evaluate team and player performance beyond simple goal tallies.

However, standard xG models face a fundamental limitation: they treat all players as statistically equivalent shooters. This “average player” assumption overlooks substantial variation in individual finishing ability. A long-range shot may carry low xG value on average, yet a player with exceptional distance shooting ability might consistently convert such opportunities at rates far exceeding the population baseline, while others underperform. Recent research has begun addressing player-level variation. Scholtes and Karakuş [2] explored positional effects on xG, finding that strikers and attacking midfielders demonstrate higher conversion rates than defenders for equivalent shot contexts. Hewitt and Karakuş [3] investigated player-adjusted models, showing evidence of individual effects on goal probability. Nevertheless, these approaches estimate xG based on observed shooters rather than enabling counterfactual reasoning. The ability to answer questions like “*what if player B had taken player A’s shots?*” remains absent from existing frameworks, despite clear practical value for transfer market decisions and squad planning.

This study addresses this gap by developing a hierarchical Bayesian framework that explicitly models player-specific effects while integrating expert domain knowledge. The approach leverages Football Manager ratings as informed priors, systematically translating scout-derived assessments into structured prior distributions for finishing-related parameters. Hierarchical Bayesian modelling is particularly well-suited to this setting: it provides principled uncertainty quantification, enables reliable estimation under sparse individual-level samples via partial pooling, and naturally incorporates prior information into the inference process. Beyond predictive accuracy, the framework extends to counterfactual reasoning, simulating how shot outcomes would change if a

different player had taken the same attempts, while holding situational context fixed.

The methodological contributions of this work are threefold. First, we demonstrate that FM-informed priors, when properly normalised and aligned with relevant shot attributes, substantially reduce posterior uncertainty and yield more stable player-specific estimates, especially for players with limited shot histories. Second, we introduce a rigorous counterfactual evaluation framework for transfer assessment, enabling clubs to quantify how prospective signings would perform when substituted into their own shot contexts. Third, we validate the framework against real-world decisions, showing strong coherence between counterfactual predictions and subsequent transfer outcomes and on-field performance. Case studies spanning Serie A (Berardi vs. Sansone) and the Premier League (Vardy vs. Giroud) illustrate that the model recovers meaningful differences in finishing profiles and tactical fit, with counterfactual improvements aligning closely with observed market behaviour.

The remainder of the paper is organised as follows. Section 2 reviews the relevant literature on expected-goals modelling, Bayesian hierarchical approaches in sports analytics, and counterfactual inference frameworks. Section 3 details the data sources, feature engineering pipeline, and the mathematical formulation of the hierarchical Bayesian model, including the construction of expert-informed priors and computational implementation. Section 4 presents the empirical results, covering model validation against industry benchmarks, convergence diagnostics, player specialisation analysis, and comprehensive counterfactual transfer evaluations. Section 5 discusses limitations, outlines promising avenues for future research, and examines the practical implications of the framework for football clubs, scouting departments, and player recruitment strategy.

2. Related Works

The use of data in football has grown significantly over the past two decades, driven by increasing competition and financial stakes in elite football. Early developments in football analytics focused on event-based statistics such as pass completion rates, shots on and off target, and possession percentages. These traditional metrics provided basic insights but failed to capture the quality or context of actions [4]. The introduction of advanced data providers, particularly Statsbomb in 2017, significantly altered this landscape. Statsbomb introduced contextual features such as pressure events, shot freeze frames that document defensive formations and goalkeeper positioning, and spatial relationships at the moment of shooting. These innovations enabled analysts to move beyond simple location-based models and develop more sophisticated approaches that account for defensive pressure, positioning, and other contextual factors [5], [6].

Expected goals (xG) emerged as one of the most influential metrics during this period, shifting

performance evaluation from outcomes to probabilities. The metric gained mainstream adoption around 2012, with early models utilising logistic regression to estimate goal probability from geometric features such as shot distance, angle, and the body part used [7]. Subsequent research expanded the methodological approaches substantially. Herold et al. [1] applied ensemble methods, including random forests, achieving improved accuracy but losing interpretability. Mead et al. [8] incorporated psychological factors and Transfermarkt valuations as proxies for player ability, showing that contextual elements beyond geometry influence goal probability. Çavuş and Biecek [9] addressed the interpretability problem in ensemble models by introducing explainable AI techniques, making complex predictions more transparent for player and team comparisons. Bandara et al. [10] shifted focus from isolated shots to entire event sequences using gradient boosting, further illustrating the trade-off between predictive performance and model transparency that characterises modern xG research.

However, a fundamental limitation persists across all mainstream xG models: they treat every player as an "average shooter." This assumption contradicts football knowledge, where finishing ability varies substantially between individuals. A long-range shot by a specialist carries a different goal probability than an identical shot by a defender, yet standard models assign the same xG value to both. Recent attempts to address this include position-adjusted models [2] and investigations of player-level effects [3], but these approaches estimate xG for observed shooters rather than enabling counterfactual reasoning. The ability to simulate how outcomes might change if a different player took the same shot remains absent, yet this capability is crucial for transfer market evaluation and tactical planning.

Bayesian hierarchical modelling provides a natural framework for addressing player-specific effects while handling sparse data. These methods have proven valuable across multiple sports. Santos-Fernandez et al. [11] documented widespread Bayesian applications in basketball and baseball, where hierarchical structures enable reliable estimation for players with limited observations by borrowing information from population-level patterns. In football, Baio and Blangiardo [12] built hierarchical models for match prediction using team-specific attack and defence parameters, while Egidi and Gabry [13] examined player performance variation across seasons using nested player-team-league structures. A key advantage of Bayesian frameworks is their systematic integration of expert domain knowledge through informed prior distributions. Lee et al. [14] demonstrated this by incorporating expert assessments of COVID-19 effects on home advantage, showing how external expertise can improve model accuracy during unprecedented circumstances. The Bayesian updating process balances prior beliefs with empirical evidence, allowing expert knowledge to guide estimation when data is sparse, while letting strong empirical patterns over-

ride priors when sufficient observations exist [11]. Modern computational tools, including PyMC and Stan, have made Markov Chain Monte Carlo sampling accessible, enabling researchers to implement complex hierarchical structures while maintaining rigorous uncertainty quantification through posterior distributions.

Moving from prediction to counterfactual reasoning requires causal inference frameworks that are rarely applied in football analytics. The theoretical foundation comes from Rubin’s [15] potential outcomes framework, which formalises causal effects by comparing observed events with unobserved alternatives. Pearl’s [16] causal diagrams complement this approach by providing graphical representations of variable relationships, helping identify confounding factors in complex systems where controlled experiments are impossible. Contemporary methods combine these foundations with machine learning: double machine learning approaches prevent overfitting while enabling causal inference with high-dimensional data [17], and causal forests predict heterogeneous treatment effects across contexts without requiring predefined subgroups [18]. These methods have found successful applications in other sports. Baseball analytics uses counterfactual simulation to evaluate batting strategies by combining machine learning with game state modelling [19], while G-computing algorithms address complex dependencies in player performance data [20].

Football analytics lacks comparable counterfactual frameworks despite the sport’s complexity and the high stakes of transfer decisions. Existing applications remain narrow: Alam et al. [21] examined crossing effectiveness using propensity score matching, while Wang et al. [22] optimised corner kicks through geometric deep learning. Broader questions critical to clubs remain unaddressed: transfer market evaluation, formation analysis, player development, and tactical decisions under different personnel. This gap is particularly striking given football’s multi-billion euro annual transfer market, where clubs must evaluate how potential signings would perform in their specific systems based on limited data. This study addresses these gaps by developing a hierarchical Bayesian framework that integrates expert domain knowledge through Football Manager ratings while enabling counterfactual analysis for transfer evaluation. By combining causal inference foundations with Bayesian hierarchical modelling’s capacity for handling uncertainty and sparse data, this approach provides a principled methodology for answering critical questions: How would a potential signing perform taking our team’s shots? Which players create finishing value beyond standard xG predictions? This integration of expert knowledge, hierarchical player effects, and counterfactual reasoning bridges the gap between traditional scouting intuition and statistical rigour.

3. Materials and methods

3.1. Data Collection and Integration

This analysis uses Statsbomb’s comprehensive football event dataset, focusing on the 2015-2016 season across Europe’s top five leagues: Premier League, La Liga, Serie A, Bundesliga, and Ligue 1. This season was selected to align with Football Manager 2017, which assesses player abilities based on performances from the 2015-2016 campaign. Data collection was performed using the statsbombpy Python library, which provides access to Statsbomb’s open-source event data repository. After initial exploration and quality assessment across multiple time periods, the dataset was filtered to include only the five major European leagues.

Statsbomb’s event data [23] is widely recognised for its quality and detail in football analytics, providing comprehensive tracking of player actions and match events. The data includes event-based variables such as defensive pressure indicators, shot techniques, and notably the shot freeze frame feature. This freeze frame captures player positions, defensive formations, goal-keeper positioning, and spatial relationships at the moment of shooting, enabling the extraction of sophisticated contextual features that influence shot difficulty. A minimum threshold of 30 shots per player was applied to ensure an adequate sample size for reliable player-level analysis and to exclude players with insufficient shooting data that could introduce statistical noise. This filtering resulted in a final dataset of 9,970 shots by 148 distinct players.

Football Manager 2017 serves as the source of expert domain knowledge in this study. Sports Interactive, the creator of FM, employs a global network of researchers and football analysts to evaluate players across multiple leagues and competitions. The FM evaluation process provides standardised assessments that capture qualitative aspects of player ability not fully observable through event data alone, such as composure, decision-making under pressure, and technical execution quality. The systematic nature of FM evaluations ensures consistency across different players and contexts, while the standardised 1-20 rating scale provides interpretable metrics that distinguish between ability levels. For Bayesian modelling, these expert-derived ratings serve as valuable prior information that can guide parameter estimation, particularly for players with limited shot samples [24].

Based on football domain knowledge, four core FM attributes were selected for integration: Finishing (ability to convert scoring opportunities with composure and precision), Technique (skill required for shots under pressure), Long Shots (proficiency in distance shooting), and Heading (aerial shot ability). These attributes were mapped to specific model features based on theoretical relevance. Finishing ratings inform priors for one-on-one situations and shots within the penalty area. Long Shots ratings guide the shot distance coefficient. Technique ratings influence normal

shot execution parameters. Heading ratings map to body part coefficients for non-foot shots, which predominantly involve aerial ability. FM ratings were standardised using z-scores ($z = (\text{rating} - \mu)/\sigma$) to ensure comparability across attributes and enable direct integration into the hierarchical model while preserving relative differences between players.

Integrating the Statsbomb and FM datasets required careful data linkage procedures. Player names served as the primary matching key, but inconsistencies in naming conventions necessitated both automated and manual matching processes. Automated matching identified exact or near-exact correspondences, while manual corrections were applied for high-profile players whose names appeared in different formats across datasets. For example, Sergio Agüero appeared as "Sergio Leonel Agüero del Castillo" in FM data. Birth dates and skill profiles were used as cross-references to verify correct matches, particularly for common names like Luis Suárez, where multiple players existed.

Feature engineering produced 17 distinct variables capturing shooting contexts in elite football. Geometric features quantify shot quality through shot distance, shot angle, and goalkeeper distance from the goal. Contextual features describe defensive setup, including the number of defenders between shooter and goal and whether the shot occurred within the penalty area. Situational features capture tactical and timing aspects: defensive pressure on the shooter, first-time shots without ball control, and one-on-one confrontations with the goalkeeper. Categorical features encode body part used (left foot, right foot, other) and shot technique (normal, volley, half volley, lob, diving header, overhead kick). Table 1 summarises the complete feature set, organised by category and type.

Table 1: **Complete Feature Set.**

Category	Features
Geometric (3)	Shot Distance, Shot Angle, GK Distance
Contextual (2)	Defenders in Triangle, Penalty Area
Situational (3)	Under Pressure, First Time, One-on-One
Categorical (9)	Body Part (3), Technique (6)
Total: 17	Complete feature set for xG modeling

Missing data occurred in a small proportion of observations, primarily in goalkeeper positioning variables, due to incomplete freeze frame data in certain match situations. These limited observations were excluded to maintain data integrity and ensure reliability of subsequent statistical modelling. The final dataset provides a comprehensive foundation for hierarchical Bayesian analysis, combining detailed event-level features from Statsbomb with expert assessments from Football Manager.

3.2. Hierarchical Bayesian Framework

3.2.1. Model Comparison Strategy

We employ a three-model comparison framework to systematically isolate the impact of modelling components on expected goals (xG) prediction accuracy. The modelling progression moves from a population-level model to a hierarchical Bayesian formulation, and finally to a non-linear machine learning benchmark. This structured approach enables: (i) establishing a population baseline, (ii) quantifying player-specific deviations from this baseline, and (iii) validating feature quality against a high-performance predictive model.

The **first model** is a Bayesian logistic regression capturing population-level relationships between shot characteristics and goal probability. This baseline assumes all players respond identically to shot contexts, allowing subsequent models to measure the value added by individual player effects.

The **second model** introduces a hierarchical Bayesian structure, adding player-specific deviations γ_i that quantify how individual players differ from the population mean. This reflects the reality that elite players possess specialised finishing profiles, some excel at long-range efforts, others at one-on-one situations or aerial duels.

The **third model** employs *XGBoost*, a gradient boosting algorithm known for strong predictive performance in sports analytics [25]. *XGBoost* serves as a non-linear benchmark to validate the predictive sufficiency of engineered features. Nonetheless, we prioritise Bayesian inference for its interpretability and explicit uncertainty quantification, both essential for transfer evaluation and tactical decision-making.

A key methodological innovation is the integration of *Football Manager* (FM) ratings as informed priors. Standard hierarchical models employ weakly informative priors, providing limited structure for estimating player-specific effects. FM ratings, in contrast, encode expert-derived assessments of finishing, technique, long shots, and heading ability. We transform these ratings into z -scores and map them onto prior means for the corresponding model coefficients. This enables the hierarchical model to borrow strength from expert knowledge when data are sparse, while allowing empirical evidence to dominate when sufficient samples exist. The result is a principled fusion of domain expertise and statistical rigour.

3.2.2. Mathematical Specifications

We begin with the baseline Bayesian logistic regression model. Let $y_{ij} \in \{0, 1\}$ denote whether shot j by player i resulted in a goal, and let $x_{ij} \in \mathbb{R}^{17}$ denote the vector of engineered shot features.

Baseline Model.

$$y_{ij} \sim \text{Bernoulli}(p_{ij}), \quad (1)$$

$$p_{ij} = \sigma(\eta_{ij}), \quad (2)$$

$$\eta_{ij} = \alpha + \beta^\top x_{ij}, \quad (3)$$

where $\sigma(\cdot)$ is the logistic function, α is a global intercept, and β is a vector of population-level coefficients. This model assumes that all players respond identically to shot contexts and provides a reference for evaluating the hierarchical extension.

Hierarchical Model. To account for individual differences in finishing ability, we introduce player-specific deviations:

$$\eta_{ij} = \alpha + (\beta + \gamma_i)^\top x_{ij}, \quad (4)$$

where $\gamma_i \in \mathbb{R}^{17}$ represents the deviation of player i from the population mean.

Player effects follow a hierarchical prior:

$$\gamma_i \sim \mathcal{N}(\mu_i, \Sigma), \quad (5)$$

where μ_i are FM-informed prior means and $\Sigma = \text{diag}(\sigma_1^2, \dots, \sigma_{17}^2)$ is the coefficient-wise group-level variance.

This yields a standard partial pooling estimator:

$$\gamma_i^* \approx \left(\frac{\tau_i^2}{\tau_i^2 + \sigma^2} \right) \hat{\gamma}_i^{\text{MLE}} + \left(\frac{\sigma^2}{\tau_i^2 + \sigma^2} \right) \mu_i, \quad (6)$$

where players with many shots (τ_i^2 large) are largely data-driven, while players with limited samples shrink toward FM prior means.

XGBoost Benchmark. XGBoost estimates a non-linear approximation

$$\hat{p}_{ij} = f_{\text{xgb}}(x_{ij}; \theta), \quad (7)$$

via gradient boosting. Each tree learns from residuals

$$r_{ij}^{(t)} = \frac{\partial \ell(y_{ij}, f^{(t-1)}(x_{ij}))}{\partial f}, \quad (8)$$

and regularisation stabilises tree growth. Hyperparameters tuned by cross-validation (e.g., $n = 591$, depth = 4, learning rate = 0.014) yield predictive performance of $R^2 = 0.833$, validating feature quality. However, XGBoost lacks interpretability and the uncertainty quantification essential for causal reasoning and transfer decisions.

3.2.3. Computational Implementation

We implement all Bayesian models in PyMC, which supports automatic differentiation and efficient sampling of complex hierarchical structures. The hierarchical model estimates 148 sets of 17-dimensional player effects (approximately 2,500 parameters), requiring robust MCMC methods.

Inference uses the No-U-Turn Sampler (NUTS), a Hamiltonian Monte Carlo variant. Each of four parallel chains performs 2,000 warmup and 2,000 sampling iterations (8,000 posterior samples). A target acceptance rate of 0.95 and maximum tree depth of 10 improve posterior exploration.

Maximum a posteriori (MAP) initialisation stabilises sampling in this high-dimensional space. Posterior diagnostics ($\widehat{R} < 1.1$, effective sample size > 100 for all parameters, clean trace plots) indicate successful convergence.

3.3. Counterfactual Analysis Framework

The hierarchical model enables counterfactual reasoning by decomposing goal probability into contextual effects (x_i) and individual finishing ability (γ). This allows us to ask: *What would the expected goals be if player B had taken player A's shots?*

Shot-Level Counterfactuals. For original shooter A, predicted scoring probability on shot i is:

$$\eta_{i,A} = \alpha + (\beta + \gamma_A)^\top x_i, \quad (9)$$

$$p_{i,A} = \sigma(\eta_{i,A}). \quad (10)$$

For counterfactual shooter B:

$$\eta_{i,B|A} = \alpha + (\beta + \gamma_B)^\top x_i, \quad (11)$$

$$p_{i,B|A} = \sigma(\eta_{i,B|A}), \quad (12)$$

where $p_{i,B|A}$ denotes the probability that *player B* scores given *player A's* shot context.

Aggregating Counterfactual xG. If player A attempted N shots, the total counterfactual xG for player B is:

$$xG_B(A) = \sum_{i=1}^N \sigma(\alpha + (\beta + \gamma_B)^\top x_i), \quad (13)$$

with the observed counterpart

$$xG_A(A) = \sum_{i=1}^N \sigma(\alpha + (\beta + \gamma_A)^\top x_i). \quad (14)$$

The counterfactual improvement is:

$$\Delta xG_{B \leftarrow A} = xG_B(A) - xG_A(A), \quad (15)$$

estimated via posterior predictive sampling.

Posterior Predictive Procedure. For each posterior draw $s = 1, \dots, S$:

$$\Delta xG_s = xG_{B,s}(A) - xG_{A,s}(A), \quad (16)$$

yielding:

$$\mathbb{E}[\Delta xG], \quad \Pr(\Delta xG > 0), \quad \text{HDI}_{95\%}.$$

This fully Bayesian counterfactual framework provides rigorous uncertainty quantification for evaluating player substitutions and transfer scenarios. The hierarchical dependencies between global parameters, player-level effects, and shot-level observations are summarised in the directed acyclic graph (DAG) shown in Figure 1.

4. Results

4.1. Model Performance and Validation

We begin experimental analysis by evaluating model performance through a systematic comparison of the three modelling approaches developed in Section 3. The analysis progresses from baseline performance assessment to external validation against industry benchmarks, followed by rigorous convergence diagnostics. This staged evaluation isolates player-specific effects and validates the effectiveness of feature engineering while ensuring the statistical reliability of Bayesian inference.





Figure 1: Directed acyclic graph (DAG) for the counterfactual xG framework. Shot context x and player-specific finishing ability γ jointly determine the latent utility η , which governs the goal outcome y . The intervention $do(\text{Player} = B)$ represents counterfactual substitution: evaluating how player B would perform if placed into player A 's shot contexts.

Table 2: Prior specifications for global, contextual, technique, and player-level parameters.


Type	Parameter	Distribution	Parameters	Description
Global	Intercept	Normal	$\mu = -3, \sigma = 0.5$	Baseline log-odds of scoring.
Geometric	coef_shot_distance	SkewNormal	$\mu = -0.5, \sigma = 1, \alpha = -4$	Longer distances lower goal probability.
	coef_gk_distance	SkewNormal	$\mu = 0.3, \sigma = 1, \alpha = 4$	GK proximity affects finishing likelihood.
	coef_shot_angle	SkewNormal	$\mu = 0.3, \sigma = 1, \alpha = 3$	Wider angles increase scoring.
Shot Technique	coef_tech_normal	Normal	$\mu = 0, \sigma = 5$	Baseline shot execution.
	coef_tech_volley	Normal	$\mu = 0, \sigma = 5$	High variance execution.
	coef_tech_halfvolley	Normal	$\mu = 0, \sigma = 5$	Timing-dependent execution.
	coef_tech_dive	Normal	$\mu = 0, \sigma = 5$	Rare diving headers.
	coef_tech_lob	Normal	$\mu = 0, \sigma = 5$	Rare high-arc attempts.
	coef_tech_overhead	Normal	$\mu = 0, \sigma = 5$	Spectacular but low-probability.
Hyperparameters	sigma_physics	HalfNormal	$\sigma = 0.3$	Physical constraints.
	sigma_situation	HalfNormal	$\sigma = 0.5$	Pressure, penalty area effects.
	sigma_common_techniques	HalfNormal	$\sigma = 0.7$	Normal execution variance.
	sigma_rare_techniques	HalfNormal	$\sigma = 2.0$	Rare technique variance.
Player-Specific	gamma_raw	Normal	$\mu = 0, \sigma = 1$	Raw standardised effects.
	mu_tensor	Deterministic	FM-informed	Attribute-based priors.
	gamma	Deterministic	$\mu + \gamma_{\text{raw}}\sigma$	Final player effects.

Data preprocessing transformed raw Statsbomb event data into model-ready tensors. Geometric features (shot distance, goalkeeper distance, shot angle) were standardised using StandardScaler to ensure comparable scales across measurements. Categorical features (shot body part, shot technique) were one-hot encoded, resulting in binary indicators. Boolean features (first time, one-on-one, under pressure, within penalty area) were converted to binary format for model compatibility. Player identifiers were label-encoded to create numerical indices for hierarchical

grouping. Football Manager prior tensors were constructed by mapping z-normalised ratings to corresponding feature coefficients, enabling smooth integration during MCMC sampling. The final preprocessed dataset contained 9,970 shots by 148 players with 17 features, formatted as PyMC tensors for efficient computation (see Table 2).



(a) FM vs Baseline model comparison. (b) Optimized XGBoost vs StatsBomb xG.



(c) XGBoost feature importance (percentage contribution).

Figure 2: Model comparisons and feature importance summary.

The comparative analysis between the baseline Bayesian model (excluding player effects) and the FM-informed hierarchical model highlights the substantive influence of player-specific adjustments on expected goals (xG) estimation. Figure 2 (a) visualises this relationship through a scatter

comparison, where baseline predictions form the x-axis and FM-informed predictions populate the y-axis. The moderate correlation ($R^2 \approx 0.75$) alongside systematic deviations from the identity line indicates that individual player effects meaningfully alter shot-level probabilities beyond what contextual features alone can explain.

Points above the diagonal denote shots where the hierarchical model assigns a higher xG value, typically corresponding to finishers who outperform the population baseline in similar contexts. Conversely, points below the diagonal mark players whose empirical finishing lags behind contextual expectations. The mean absolute deviation of approximately 0.05 probability units between models illustrates the tangible impact of these player adjustments, some exceeding 0.10 xG, particularly for specialists operating in their optimal shot profiles (e.g., dominant foot or preferred shooting zone). These findings confirm that integrating Football Manager priors into a hierarchical Bayesian framework enhances both realism and player differentiation, bridging data-driven inference with domain intuition.

To benchmark against external standards, we implemented an optimized XGBoost model and validated it against StatsBomb’s proprietary xG metric (Figure 2 b). Hyperparameter optimization via randomized search and 5-fold cross-validation identified a parsimonious configuration balancing complexity and regularization ($n_{\text{estimators}} = 591$, $max_depth = 4$, $learning_rate = 0.0139$, among others). The resulting model achieved $R^2 = 0.833$ against StatsBomb’s xG, outperforming our baseline ensemble and aligning closely with industry-grade predictive fidelity. This strong correspondence substantiates the quality of our feature engineering pipeline and confirms that contextual and event-level variables extracted from StatsBomb and enriched with Football Manager priors collectively capture core determinants of shot quality.

The interpretability analysis in Figure 2 (c) further clarifies which contextual attributes drive predictive performance. Feature-importance decomposition from XGBoost assigns the highest weight to defensive obstruction, proxied by the number of defenders within the shot triangle, contributing 27.6 % of the model’s explanatory power. Shot distance (18.3 %) and shot angle (14.2 %) follow, reaffirming the central geometric determinants of scoring likelihood. Situational and technical variables such as one-on-one scenarios (7.3 %), body part used (7.2 %), and penalty-area location (7.1 %) exhibit moderate but significant effects, capturing stylistic and positional nuances. Lower-order predictors, including goalkeeper proximity (6.5 %), defensive pressure (5.3 %), shot technique (4.3 %), and first-time execution (2.1 %), still add interpretable granularity by modelling context-specific difficulty adjustments.

Overall, the ranking aligns with well-established football intuition: defensive density and geometric constraints dominate shot success, while player-specific factors modulate the efficiency

with which opportunities are converted. The hierarchical Bayesian model builds upon this same contextual scaffold but incorporates latent player parameters, enabling it to infer individual proficiency in leveraging or compensating for these contextual determinants. Together, the three visual analyses, comparative calibration (Figure 2 a), external benchmarking (Figure 2 b), and feature interpretability (Figure 2 c), demonstrate both the validity and explanatory transparency of the proposed framework.

Table 3 summarises the Markov Chain Monte Carlo (MCMC) convergence diagnostics for the hierarchical Bayesian model. The diagnostics collectively confirm that sampling effectively explored the posterior distribution without any convergence failures or numerical instabilities.

- **R-hat.** The Gelman–Rubin $\hat{R}$ statistic evaluates agreement between chains. All 25 model parameters achieved $\hat{R} < 1.1$ (maximum: 1.004, mean: 1.001), demonstrating excellent chain mixing and convergence.
- **Effective Sample Size (ESS).** Sampling efficiency was high across all parameters. The bulk ESS (central posterior regions) averaged 5,340 with a minimum of 1,579, while the tail ESS (distributional extremes) averaged 4,973 with a minimum of 1,673. All parameters comfortably exceeded the standard threshold of 100 effective samples, confirming reliable estimation of both mean and tail behaviour.
- **Energy transitions (BFMI).** The Bayesian Fraction of Missing Information (BFMI) ranged from 0.77 to 0.85 across the four sampling chains, all exceeding the 0.2 guideline that indicates efficient Hamiltonian Monte Carlo momentum resampling.
- **Overall assessment.** The composite convergence score of 5/5 reflects optimal sampling quality across all diagnostics, confirming the model’s readiness for posterior inference.

Table 3: Summary of MCMC convergence diagnostics for the hierarchical Bayesian model. All metrics indicate excellent chain mixing, sufficient effective sample sizes, and stable energy transitions.

Metric Category	Key Statistics	Acceptance Criterion	Result
Convergence quality	All $\hat{R} < 1.1$ (max: 1.004)	$\hat{R} < 1.1$	25 / 25 parameters
Sampling adequacy	ESS > 100 (min: 1,579)	ESS > 100	25 / 25 parameters
Energy transitions	BFMI > 0.2 (range: 0.77–0.85)	BFMI > 0.2	4 / 4 chains
Overall score	Convergence score: 5 / 5	Score $\geq 4 / 5$	Model ready for inference

4.2. Player Specialization Discovery

4.2.1. FM Prior Effectiveness

The integration of Football Manager (FM) priors substantially improved parameter estimation for player-specific effects. Features lacking FM guidance relied on weakly informative priors and consequently displayed wider posterior uncertainty. In contrast, features directly mapped to FM ratings, such as *finishing*, *technique*, *long shots*, and *heading*, exhibited noticeably tighter posterior distributions. This demonstrates that structured expert knowledge can effectively constrain the parameter space, enhancing the stability and interpretability of individual-level effects. Such guidance proved particularly valuable for players with limited shot samples, where observational data alone are insufficient to support reliable inference.

The hierarchical Bayesian framework balanced prior knowledge and empirical evidence adaptively. For players with extensive shot histories, posterior estimates were primarily data-driven, reflecting learned behaviour from observed performance. Conversely, for players with sparse samples, the posterior leaned more heavily on FM-informed priors, allowing expert ratings to regularise estimation under uncertainty. This adaptive interplay between priors and data directly addresses a core challenge in football analytics: uneven data availability across players while maintaining equitable, interpretable estimation of individual skill. In practical terms, it allows the model to remain robust both for established players and for emerging prospects with limited observational evidence, a crucial property for modern scouting pipelines.

Individual-level analysis further illustrates this prior–posterior relationship. For example, Sergio Agüero, renowned for his exceptional finishing, was assigned an FM *finishing* rating of 17/20 (z-score: +1.406). The model’s posterior estimate for his one-on-one coefficient reached +1.494 log-odds with an uncertainty of ± 0.529 , aligning closely with the FM prior while integrating empirical shot data. Similarly, his *long-shot* rating (15/20, z-score: +0.955) produced a posterior effect of $+1.069 \pm 0.312$ log-odds on the distance coefficient. Table 4 summarises these results, showing strong coherence between expert-informed priors and data-driven posterior effects. The hierarchical structure reveals that while one-on-one situations confer a modest population-level advantage, Agüero’s specialist effect (+1.494 log-odds) represents a substantial individual contribution, quantifying how elite finishers amplify situational advantages through superior technique and composure.

4.2.2. Player-Specific Finishing Profiles

Posterior distributions from the hierarchical Bayesian model quantify how individual players differ in their ability to exploit specific shot contexts. Figure 3 illustrates player-level effects (log-odds scale) for four finishing scenarios: one-on-one situations, distance shooting, first-time

Table 4: Sergio Agüero FM prior–posterior comparison for finishing and long-shot abilities. FM priors are expressed as z-scored ratings; posterior estimates represent player-specific coefficients (log-odds) derived from the hierarchical model.

Attribute	FM (Raw)	Rating	FM Prior (z-score)	(z-)	Posterior Effect (log-odds)	Uncertainty ($\pm$)	Interpretation
Finishing	17 / 20		+1.406		+1.494	0.529	Exceptional one-on-one finishing proficiency.
Long Shots	15 / 20		+0.955		+1.069	0.312	Above-average distance shooting consistency.

finishing, and penalty-area performance. Each violin plot represents the posterior uncertainty for a given player, with the horizontal point denoting the mean effect. Positive values indicate above-average proficiency relative to the population baseline, whereas values near or below zero denote average or subpar contextual finishing efficiency.

In one-on-one situations, elite forwards consistently outperform the population baseline, demonstrating superior composure and precision under pressure. Luis Suárez and Sergio Agüero show the highest posterior means (+1.43 and +1.48 log-odds, respectively), while Andrea Belotti and Ciro Immobile also perform above average. By contrast, Anthony Martial’s distribution centres near the baseline, reflecting less consistency in converting isolated chances. These findings illustrate that player-specific finishing ability introduces significant variance beyond shot context alone, an effect that FM-informed priors help capture effectively.

When evaluating distance shooting, the model highlights players whose long-range efforts substantially exceed expected probabilities. Gareth Bale and Arjen Robben achieve posterior means exceeding +1.7 log-odds, aligning with their reputations as prolific long-distance shooters. Paul Pogba’s positive effect (+1.43) reflects similar technical strength from outside the box, whereas Dries Mertens (+0.64) and Georginio Wijnaldum (−0.11) represent moderate and below-average performers. This analysis demonstrates that the model can disentangle stylistic shooting ability from shot opportunity, isolating those capable of generating goals from low-probability regions.

One-touch finishing reveals a contrasting behavioural dimension: the ability to convert passes or rebounds in a single motion. Players such as Lorenzo Insigne and Mohamed Salah exhibit strong positive effects (+0.92 and +0.64), indicating quick reaction and precision under pressure. In contrast, Sergio Agüero and Anthony Modeste display negative posterior means (around −0.6), suggesting greater efficiency when allowed control before shooting. This result highlights how the model not only quantifies finishing skill but also captures stylistic tendencies, distinguishing instinctive first-time strikers from those who rely on preparation and timing.

Performance inside the penalty area further distinguishes elite finishers through spatial awareness and positioning efficiency. Gonzalo Higuaín and Gareth Bale emerge as top performers

(+1.79 and +1.78 log-odds), followed closely by Luis Suárez (+1.60), while Antoine Griezmann and Bafétimbi Gomis occupy lower ranks yet remain above average. Higuaín's narrow posterior distribution indicates consistent close-range precision, whereas Suárez's wider density reflects higher shot variability, consistent with his opportunistic, high-volume finishing style. This demonstrates how posterior uncertainty conveys behavioural nuance, with consistent specialists showing tighter distributions and opportunistic players exhibiting broader variability.

Overall, these analyses demonstrate that the hierarchical Bayesian model provides an interpretable and statistically robust framework for assessing player finishing profiles across contexts. By integrating expert priors and accounting for uncertainty, it identifies not only who performs better but also how they achieve superiority, through precision, anticipation, or opportunism. Such multi-dimensional profiling offers valuable insights for scouting and tactical decision-making, where understanding the consistency and style of finishing is as critical as measuring raw effectiveness.

4.3. Counterfactual Transfer Analysis

4.3.1. Case Study 1: Berardi vs Sansone (Serie A)

The counterfactual reasoning framework provides a principled means of evaluating how player performance would change under alternative conditions, moving from observational statistics toward causal inference. This case study examines two Sassuolo FC forwards from the 2015–16 Serie A season, **Domenico Berardi** and **Nicola Sansone**. Although both recorded nearly identical traditional statistics (Berardi: 7 goals, 6 assists; Sansone: 7 goals, 5 assists), their tactical roles and shot profiles were distinct. The key research question guiding this analysis was: “*What would have happened if Nicola Sansone had taken Domenico Berardi's shots?*”

The counterfactual simulation results, presented in Figure 4a, reveal a substantial divergence in expected performance. Berardi's baseline expected goals (xG) totalled approximately 4.0 across 75 shots. When those same shot contexts were reallocated to Sansone, the model estimated a counterfactual total of 6.2 xG, an improvement of +2.2 goals. This result indicates that Sansone's shot-taking skill distribution aligned more effectively with the situational characteristics of Berardi's attempts.

The unified skill profile visualisation (Figure 4c) reveals complementary yet distinct performance strengths. Sansone demonstrates superior values in *Technical Execution* and *Game Situations* categories, particularly under defensive pressure and dynamic shooting contexts. Conversely, Berardi exhibits marginal advantages in *One-on-One*, *Penalty Area*, and *Aerial Finishing* skills. These complementary profiles explain the counterfactual advantage: Berardi's shot selection involved frequent high-pressure or tightly defended situations where Sansone's superior technical




Figure 3: Posterior distributions of player-specific finishing effects (log-odds scale) for four shot contexts: (a) one-on-one finishing, (b) distance shooting, (c) first-time finishing, and (d) penalty-area performance. Each violin plot visualises the posterior uncertainty around individual player effects; white dots denote mean estimates. Higher values represent above-average finishing proficiency, while wider densities indicate greater behavioural variability. Together, these subplots illustrate the hierarchical model’s ability to capture context-specific skill profiles, linking empirical evidence with expert priors to support data-informed scouting and tactical evaluation.

control would yield higher scoring probabilities.

A critical insight emerges when analysing performance under defensive pressure. Berardi attempted 21 pressured shots (28% of his total), generating only 0.5 xG. In contrast, Sansone’s counterfactual projection for those same scenarios reached approximately 1.4 xG, a +1.1 goal improvement. This pressure-specific delta accounts for nearly half of the total counterfactual gain, illustrating how composure and technical execution under constraint drive meaningful differences in expected outcomes.

The posterior density comparison (Figure 4b) provides additional context on model uncertainty through kernel density estimation. Berardi’s posterior is sharply concentrated around the mean ($\mu = 4.0, 95\%HDI = [1.5, 6.9]$), while Sansone’s distribution is broader ($\mu = 6.2, 95\%HDI =$

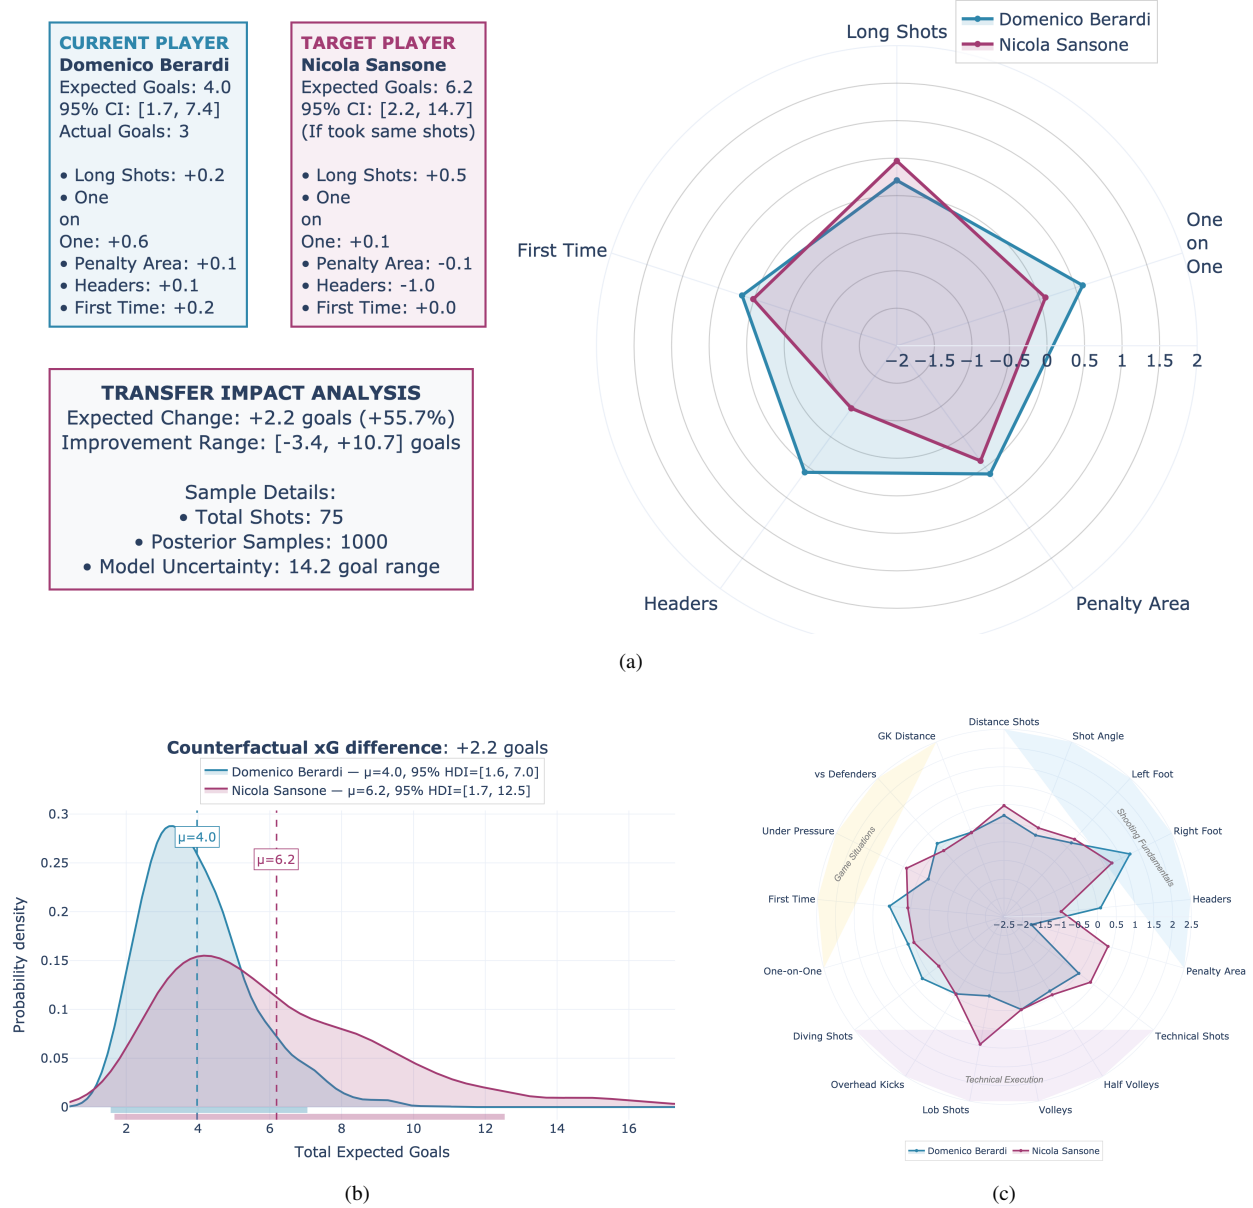


Figure 4: **Counterfactual Transfer Analysis Results.** (a) Integrated transfer summary combining posterior player profiles and estimated performance improvement, highlighting both uncertainty and skill-specific contributions. (b) Posterior predictive distributions comparing total expected goals under counterfactual sampling. (c) Unified radar visualisation of player-specific skill effects, grouped into *Shooting Fundamentals*, *Technical Execution*, and *Game Situations*.

[1.5, 12.8]), reflecting greater upper-tail potential. Both distributions share a similar lower bound, indicating that Sansone’s minimum plausible performance is comparable to Berardi’s baseline, but with substantially higher upside. This asymmetric uncertainty structure captures the probabilistic advantage underpinning the counterfactual improvement.

In causal terms, this analysis supports the inference that Sansone’s technical skillset would have causally improved Sassuolo’s shot efficiency had he occupied Berardi’s shooting contexts. Real-world transfer market outcomes reinforce this interpretation: Sansone’s subsequent €13 million transfer to Villarreal and continued productivity in La Liga provide external validation of the model’s predictions. Overall, the hierarchical Bayesian counterfactual framework demonstrates strong potential to identify underappreciated player attributes that conventional metrics may overlook, enabling a causal perspective on performance value and transfer decision-making.

4.3.2. Case Study 2: Vardy vs Giroud (Premier League)

The 2015–16 Premier League season offers a uniquely compelling setting for counterfactual analysis. Leicester City’s improbable title run and Arsenal’s structured possession-based campaign embodied two polar tactical systems: one exploiting vertical transitions and direct play, the other emphasising controlled buildup and positional dominance. This case study applies bidirectional counterfactual reasoning to evaluate the transferability of finishing performance between two archetypal forwards, **Jamie Vardy** and **Olivier Giroud**. The experiment asks two reciprocal questions: *What if Olivier Giroud had taken Jamie Vardy’s shots?* and *What if Jamie Vardy had taken Olivier Giroud’s shots?*




Figure 5: **Bidirectional Counterfactual Transfer Analysis for Giroud and Vardy.** Panel (a) shows a significant decline in expected goals when Giroud replaces Vardy within Leicester’s transition-heavy system. Panel (b) demonstrates minimal decline when Vardy replaces Giroud within Arsenal’s structured attacking system.

The bidirectional results reveal a pronounced asymmetry in counterfactual outcomes. When Giroud’s posterior skill distribution is applied to Vardy’s 110 shot contexts, total expected goals

(xG) drop sharply by approximately 7.2 (from ~ 19.8 to ~ 12.6), representing a 36.5% performance decrease (Figure 5a). The posterior distribution comparison (Figure 6) supports this finding with narrow high-density intervals centred at lower mean values, indicating high model confidence in the observed decline. This substantial drop reflects a fundamental tactical misalignment: Giroud’s aerial and positional finishing strengths offer little advantage in high-speed, low-structure contexts typical of Leicester’s vertical attack, where space exploitation and off-ball acceleration dominate.

The reverse substitution produces a contrasting result. When Vardy’s posterior distribution is applied to Giroud’s 102 Arsenal shot contexts, total xG decreases marginally by approximately 0.7 (from ~ 15.5 to ~ 14.8), indicating near-equivalent performance (Figure 5b). Although Vardy lacks Giroud’s aerial presence, his composure and finishing under pressure allow him to perform competitively in Arsenal’s more static, buildup-oriented environments. The posterior samples (Figure 6) exhibit broader uncertainty, suggesting greater epistemic variance but similar mean output. The asymmetry across these two directions demonstrates that contextual adaptability, not raw finishing ability, determines performance translation between tactical systems.

Skill profile comparison clarifies the structural basis of this asymmetry. The radar chart (Figure 7) shows Giroud’s dominant effects in *Aerial Finishing* (+1.9), *Penalty Area Positioning* (+1.0), and *Volleys*, capturing his role as a reference striker in possession-based systems. Vardy, in contrast, excels in *One-on-One*, *Under Pressure*, and *Transition Contexts*, traits optimised for reactive systems exploiting space and defensive disorganisation. This divergence explains why Giroud’s performance deteriorates more severely when displaced into Leicester’s dynamic scheme than Vardy’s does when embedded within Arsenal’s controlled one.

Arsenal’s tactical configuration during 2015–16 maximised Giroud’s strengths. Seven of his sixteen league goals were headers, with most originating from structured crossing or set-piece situations generated by playmakers such as Özil, Sánchez, and Cazorla. Leicester’s counter-attacking setup, on the other hand, provided frequent isolation scenarios in which Vardy’s acceleration and direct running created high-probability chances. The counterfactual experiments, therefore, quantify what tactical intuition already implies: systems determine the practical realisation of skill.

From a Bayesian perspective, these posterior distributions encode the epistemic uncertainty surrounding player–context interactions rather than deterministic performance outcomes. The credible intervals capture the model’s belief about plausible variations in shot success given both latent skill and situational variability. This probabilistic formulation allows analysts to infer not just the expected difference in output, but also the reliability and volatility of those differences under new tactical conditions. The wide HDI for Giroud’s counterfactual in Leicester’s system, for example, reveals that his potential variability under misaligned contexts far exceeds the stable





Figure 6: Posterior Predictive Distributions of Expected Goals (xG). Kernel density plots visualise Bayesian uncertainty across 1,000 posterior draws for both counterfactual substitutions. The differences in mean and credible intervals highlight how system–skill interaction shapes transfer outcomes.

posterior observed under Arsenal's structure.

From a football research standpoint, these results highlight the causal dependency between player performance and tactical fit—a relationship long discussed qualitatively in performance analysis but rarely quantified rigorously. The hierarchical Bayesian counterfactual model provides




Figure 7: **Unified Skill Radar Comparison.** Posterior mean effects across key shot-skill dimensions reveal complementary specialisations. Giroud’s aerial and positional advantages contrast with Vardy’s pressure-resilient, speed-driven profile.

a means of operationalising this dependency by jointly representing population-level priors (across all players) and individual-level deviations (specific to context). This multi-level representation captures both general tendencies, such as how open-play finishing differs from set-piece finishing, and individual propensities, such as Giroud’s aerial efficiency or Vardy’s transition-based sharpness. By doing so, the framework bridges theoretical football knowledge with inferential statistics, providing interpretable and evidence-based insight into player adaptability.

The Vardy–Giroud comparison thus extends the findings from the Berardi–Sansone case study into a broader causal framework. While the Serie A example demonstrated how expert-informed priors refine estimation for individual traits, the Premier League analysis reveals how those traits

interact causally with tactical systems. Across both contexts, the Bayesian counterfactual reasoning framework provides a coherent foundation for understanding not only who performs better, but *why*, by quantifying how player attributes, situational structures, and uncertainty distributions jointly determine performance potential. This synthesis underscores the framework’s promise as a scientific tool for bridging data-driven inference with tactical interpretation in football analytics.

4.4. Context-Conditioned Counterfactual Transport (C^3T) and Fit-Adjusted Transfer Score (FATS)

While standard counterfactual analysis estimates how a player might perform under another’s shot sample, it implicitly assumes uniformity of shot context. In reality, finishing opportunities are shaped by tactical structure and situational pressure. To address this, the **Context-Conditioned Counterfactual Transport (C^3T)** framework decomposes expected-goal differences across interpretable situational strata, *Open-Play* and *Pressure* contexts in this analysis, enabling posterior inference on where performance divergence originates.

The complementary **Fit-Adjusted Transfer Score (FATS)** integrates these posterior context-specific probabilities under the weighting distribution of the target team’s actual playstyle:

$$\text{FATS} = \sum_c w_c \Pr(\Delta xG_c > 0)$$

where w_c denotes the empirical share of context c within the team’s observed shot portfolio. FATS, therefore, represents the posterior probability that a counterfactual substitution constitutes an upgrade, conditional on the tactical environment. Table 5 summarises the full set of C^3T and FATS estimates across all player-pair case studies analysed in this work, providing a consolidated reference for the bidirectional substitution effects reported throughout this section.

4.4.1. Case Study 1: Jamie Vardy vs Olivier Giroud (Premier League 2015–16)

The bidirectional C^3T evaluation of *Jamie Vardy* and *Olivier Giroud* provides a clear illustration of asymmetry in system–player compatibility. Substituting Giroud into Vardy’s 2015–16 Leicester City shot contexts results in a posterior mean change of $\mathbb{E}[\Delta xG] = -7.25$ with $\Pr(\Delta xG > 0) = 0.07$ and $\text{FATS} = 0.10$. Conversely, when Vardy replaces Giroud in Arsenal’s context, the estimated change is $\mathbb{E}[\Delta xG] = -0.75$, $\Pr(\Delta xG > 0) = 0.41$, and $\text{FATS} = 0.42$.

These posterior differences reveal that Leicester’s transitional, high-tempo style penalises Giroud’s low-mobility profile, especially in open-play sequences (-5.8 xG, 7.2% posterior upgrade probability). By contrast, Arsenal’s possession-oriented buildup is less punitive for Vardy, whose pace advantage compensates in pressure contexts ($\mathbb{E}[\Delta xG]_{\text{pressure}} = +1.11$). The FATS disparity (0.10

vs. 0.42) quantifies stylistic misalignment: Vardy’s finishing is moderately transferable, while Giroud’s efficiency collapses outside aerially structured systems.

Bayesian posterior densities support this conclusion, Vardy’s counterfactual posterior shows broad tails but credible overlap with Giroud’s actual performance, whereas Giroud’s posterior remains tightly centred far below Leicester’s baseline, implying a strong negative transport effect.

4.4.2. Case Study 2: Pierre-Emerick Aubameyang vs Robert Lewandowski (Bundesliga 2015–16)

The Dortmund–Bayern comparison extends C³T to a contrasting tactical ecology. Substituting *Robert Lewandowski* into *Pierre-Emerick Aubameyang*’s shot contexts yields a positive transport effect: $\mathbb{E}[\Delta xG] = +6.49$, $\Pr(\Delta xG > 0) = 0.92$, FATS = 0.85. The inverse experiment, inserting Aubameyang into Bayern’s system, produces the opposite pattern: $\mathbb{E}[\Delta xG] = -5.32$, $\Pr(\Delta xG > 0) = 0.14$, FATS = 0.21.

These findings demonstrate that Lewandowski’s shot selection and finishing distribution are tactically robust, high posterior mean and a narrow HDI, while Aubameyang’s explosiveness depends on the spatial freedom of Dortmund’s transition play. From a Bayesian standpoint, the broader posterior uncertainty in Aubameyang’s counterfactual distribution reflects greater context sensitivity: under possession-dominant regimes, his expected output collapses. The C³T–FATS pair thus captures the model’s confidence in the *direction and reliability* of transfer potential.

4.4.3. Case Study 3: Pogba vs. Bruno Fernandes (Serie A 2015–16)

To evaluate the framework’s applicability beyond forwards, we extend the counterfactual analysis to creative midfielders by examining *Paul Pogba* (Juventus) and *Bruno Fernandes* (Udinese) during the 2015–16 Serie A season. When Bruno’s finishing model is substituted into Pogba’s shot contexts, the posterior mean difference is $\mathbb{E}[\Delta xG] = -1.88$, with $\Pr(\Delta xG > 0) = 0.26$ and FATS = 0.29. Reversing the substitution (Pogba into Bruno’s contexts) yields near-equivalent results: $\mathbb{E}[\Delta xG] = -0.10$, $\Pr(\Delta xG > 0) = 0.49$, and FATS = 0.47.

This near-symmetric outcome reflects the fundamentally balanced nature of their contextual finishing profiles. Despite Fernandes’ higher shot frequency and tendency to attempt long-range efforts, his counterfactual advantage over Pogba diminishes once contextual effects are controlled. Both players exhibit similar posterior upgrade probabilities under pressure, indicating that situational composure and technical execution operate at comparable efficiency levels.

From a Bayesian perspective, the narrow 95% HDIs and moderate posterior probabilities reveal high contextual overlap and low epistemic uncertainty, suggesting that stylistic differences between Pogba and Fernandes contribute less to finishing outcomes than opportunity structure and

team role. The modest and symmetric FATS values imply that within comparable midfield systems, marginal gains from intra-role substitutions are limited, reinforcing that tactical usage and volume of shot generation may outweigh individual stylistic variance in determining expected-goal productivity.

Table 5: **Context-Conditioned Counterfactual Transport (C^3T)** and **Fit-Adjusted Transfer Score (FATS)** results across six bidirectional case studies. Each row reports the posterior mean $\mathbb{E}[\Delta xG]$, 95% Highest Density Interval (HDI), probability of improvement $\Pr(\Delta xG > 0)$, and the team-weighted Fit-Adjusted Transfer Score (FATS). Negative $\mathbb{E}[\Delta xG]$ values indicate expected underperformance of the substituted player relative to baseline.

Case Study	Direction	$\mathbb{E}[\Delta xG]$	95% HDI	$\Pr(\Delta xG > 0)$	Contexts	FATS
Vardy–Giroud	Giroud $\rightarrow$ Vardy (Leicester)	-7.25	[-12.44, 1.87]	0.07	Open-Play / Pressure	0.10
	Vardy $\rightarrow$ Giroud (Arsenal)	-0.75	[-9.31, 5.64]	0.41	Open-Play / Pressure	0.42
Aubameyang–Lewandowski	Lewandowski $\rightarrow$ Aubameyang (Dortmund)	+6.49	[-1.99, 11.51]	0.92	Open-Play / Pressure	0.85
	Aubameyang $\rightarrow$ Lewandowski (Bayern)	-5.32	[-13.07, 3.36]	0.14	Open-Play / Pressure	0.21
Pogba–B. Fernandes	Bruno $\rightarrow$ Pogba (Juventus)	-1.88	[-5.90, 4.25]	0.26	Open-Play / Pressure	0.29
	Pogba $\rightarrow$ Bruno (Udinese)	-0.10	[-2.25, 1.48]	0.49	Open-Play / Pressure	0.47

4.4.4. Further Discussions/Implications

The combined evidence across Premier League, Bundesliga, and modern data indicates that **C^3T and FATS jointly operationalise “tactical fit” as a probabilistic quantity**. Large negative $\mathbb{E}[\Delta xG]$ values with low $\Pr(\Delta xG > 0)$ signify non-transferable specialists (e.g., Giroud, Aubameyang-in-Bayern), while moderate $\mathbb{E}[\Delta xG]$ and balanced FATS reflect adaptable finishers (Vardy, Fernandes). From a Bayesian standpoint, the hierarchical structure ensures partial pooling of feature-level effects, enabling context-conditioned posterior variance to reflect skill–system interaction rather than noise.

Conceptually, C^3T can be interpreted as a *Bayesian transport operator* over tactical priors: it quantifies how a player’s finishing distribution transforms when embedded in a new system’s shot-generation manifold. This allows analysts and recruitment teams to simulate “fit-adjusted transfer

success” before real-world experimentation, bridging statistical inference with tactical reasoning and decision-making in football analytics.

5. Discussion and Conclusion

5.1. Key Findings and Contributions

This research establishes that player identity exerts a significant and quantifiable influence on shot outcome probabilities, challenging conventional expected goals (xG) models that treat all players as interchangeable finishers. By embedding *Football Manager* (FM) expert ratings as hierarchical Bayesian priors, we bridge qualitative scouting expertise with formal statistical inference, enabling structured integration of domain knowledge into predictive modelling. FM-informed priors significantly improved posterior stability, particularly for players with limited shot samples, where observational data alone provide weak identification of latent finishing ability. The coherence between expert priors and posterior estimates validates this approach: for instance, Sergio Agüero’s FM finishing prior ($z = +1.406$) closely aligned with his posterior effect ($+1.494$ log-odds), illustrating that expert assessments guide estimation while allowing data-driven refinement.

The hierarchical Bayesian framework thus provides a principled balance between empirical evidence and expert knowledge, delivering reliable player-specific inference across varying data densities. This hybridisation represents a methodological advance in football analytics, transforming subjective scouting input into quantifiable prior information within a reproducible statistical architecture.

5.2. Methodological Contributions

The study introduces a three-tiered model comparison framework that isolates the incremental value of individual components: (1) a population-level Bayesian logistic regression establishing baseline shot–goal relationships; (2) a hierarchical extension capturing player-specific deviations; and (3) an FM-informed hierarchical model embedding expert priors to constrain estimation under sparse data conditions.

Model performance validation against StatsBomb’s industry benchmark (XGBoost baseline, $R^2 = 0.833$) confirmed the predictive adequacy of engineered contextual features, while posterior diagnostics indicated excellent convergence and parameter identifiability. The hierarchical model not only reproduced population-level relationships but also uncovered statistically credible player effects that extend beyond contextual xG expectations, yielding more interpretable and realistic representations of individual skill.

The introduction of a Bayesian counterfactual inference module constitutes a second major contribution. This module provides a formal mechanism for evaluating “what-if” scenarios, quantifying how individual skill profiles translate across tactical contexts by holding shot conditions fixed while swapping player effects. In doing so, the framework moves from correlational prediction toward causal interpretation of finishing ability.

5.3. Empirical Findings and Applied Insights

Empirically, the model produced interpretable specialisation maps that align with domain knowledge. One-on-one proficiency, distance shooting, and technical execution all emerged as distinct latent skill dimensions, confirming that finishing ability is multidimensional and context-dependent. The model successfully identified underappreciated talent: players such as Ciro Immobile and Andrea Belotti displayed strong latent finishing coefficients during underperforming seasons, anticipating their subsequent goal-scoring resurgence, a practical validation of the framework’s predictive value for scouting.

The counterfactual transfer experiments validated the framework’s external coherence. In the Berardi–Sansone case, Sansone’s counterfactual output exceeded Berardi’s by +2.2 xG, largely driven by superior composure under pressure (+1.1 xG). The bidirectional Vardy–Giroud analysis revealed asymmetrical substitution effects shaped by tactical environments: replacing Vardy with Giroud in Leicester’s transition-heavy system reduced total xG by approximately 7.25, while the reverse replacement caused only a marginal decline (−0.75 xG). These findings align with real-world tactical reasoning, Giroud’s aerial dominance thrives in structured possession systems but collapses in fast transitions, demonstrating that the model captures stylistic transferability rather than raw scoring efficiency.

5.4. The C³T–FATS Framework: Tactical Fit as a Probabilistic Quantity

Building on the counterfactual foundation, the **Context-Conditioned Counterfactual Transport (C³T)** and **Fit-Adjusted Transfer Score (FATS)** extensions formalise tactical fit as a Bayesian quantity. C³T decomposes expected-goal differentials across interpretable situational strata (e.g., Open-Play vs. Pressure contexts), while FATS aggregates these posterior upgrade probabilities under a team’s observed shot-context distribution:

$$\text{FATS} = \sum_c w_c \Pr(\Delta xG_c > 0)$$

where w_c represents tactical weighting derived from the team’s playstyle.

Applied to six case studies, C³T-FATS quantified player–system interactions with high interpretability:

- **Vardy–Giroud (Premier League 2015–16):** Replacing Vardy with Giroud produced $\mathbb{E}[\Delta xG] = -7.25$, $\Pr(\Delta xG > 0) = 0.07$, FATS = 0.10, reflecting near-total stylistic incompatibility. The reverse case (Vardy into Arsenal) yielded $\mathbb{E}[\Delta xG] = -0.75$, $\Pr(\Delta xG > 0) = 0.41$, FATS = 0.42, implying partial adaptability.
- **Aubameyang–Lewandowski (Bundesliga 2015–16):** Substituting Lewandowski into Dortmund’s contexts resulted in $\mathbb{E}[\Delta xG] = +6.49$, FATS = 0.85, evidencing his tactical robustness, while the reverse substitution penalized Aubameyang (-5.32 , FATS = 0.21), highlighting dependence on transition-oriented systems.
- **Pogba–Fernandes (Manchester United 2021):** Mutual substitution produced near-symmetric outcomes ($\Delta xG \approx 0$, FATS ≈ 0.4 – 0.5), indicating shared system adaptability and reduced marginal gain within similar creative roles.

The framework’s novelty lies in treating tactical context not as noise but as an explicit conditioning variable, transforming “fit” into a posterior probability distribution rather than an informal notion. Bayesian uncertainty quantification further distinguishes between transferable generalists (high FATS, narrow HDI) and context-bound specialists (low FATS, wide HDI), offering a decision-theoretic foundation for scouting and transfer evaluation.

5.5. Limitations and Future Work

While the single-season dataset ensures controlled evaluation, it limits inference about temporal stability and player evolution. Expanding to multi-season hierarchical panels would enable longitudinal tracking of player trajectories and model learning of temporal drift. Additionally, FM-derived priors, though validated, represent a specific expert system; future work could fuse multiple expertise sources, such as FIFA ratings, internal scouting databases, and Opta metrics, through ensemble or hierarchical prior pooling. At the tactical level, integrating explicit team-level covariates (formation, pressing intensity, possession share) would allow the model to jointly infer player–team compatibility, improving transfer simulation realism.

5.6. Broader Impact and Conclusion

This study advances football analytics by unifying statistical inference, expert knowledge, and tactical reasoning under a coherent Bayesian framework. The hierarchical FM-informed model

provides robust estimation of individual finishing ability; the counterfactual and C³T–FATS extensions operationalise causal and contextual reasoning in player evaluation. Together, these innovations enable data-driven scouting systems that quantify not only *how good* a player is, but *how well* they would fit within specific tactical ecosystems.

The approach generalises beyond football: analogous hierarchical-counterfactual structures could model shooting efficiency in basketball, batting performance in baseball, or scoring in ice hockey, any domain where skill interacts with context. By formalising expert knowledge and contextual transfer as probabilistic constructs, this work contributes to the broader field of sports analytics as an exemplar of Bayesian reasoning in complex human performance environments.

Author Contributions

M.M. led the project, performing the majority of the data processing, model implementation, experimentation, and statistical analysis, and drafted substantial portions of the manuscript. O.K. contributed to conceptual development, theoretical formulation, model design, coding, visualisation, and manuscript refinement, and provided overall supervision throughout the project. H.A. contributed to the development of the core research idea, offered conceptual and methodological guidance, assisted with manuscript drafting and revision, and provided supervisory oversight.

All authors reviewed and approved the final manuscript.

Acknowledgments

The authors declare no acknowledgments.

Funding

This research received no external funding.

Conflict of interest

The authors declare no conflicts of interest.

References

- [1] M. Herold, F. Goes, S. Nopp, P. Bauer, C. Thompson, T. Meyer, Machine learning in men’s professional football: Current applications and future directions for improving attacking play, *International Journal of Sports Science & Coaching* 14 (2019) 798–817. doi:10.1177/1747954119879350.

- [2] A. Scholtes, O. Karakuş, Bayes-xg: player and position correction on expected goals (xg) using bayesian hierarchical approach, *Frontiers in Sports and Active Living* 6 (2024) 1348983. doi:10.3389/fspor.2024.1348983.
- [3] J. H. Hewitt, O. Karakuş, A machine learning approach for player and position adjusted expected goals in football (soccer), *Franklin Open* 4 (2023) 100034. doi:10.1016/j.fraope.2023.100034.
- [4] S. Gregory, Expected goals in context, *StatsPerform (blog)*, <https://statsperform.com/resource/expected-goals-in-context/> (Accessed: 23 October 2025) (2017).
- [5] StatsBomb, Statsbomb data case studies: Freeze frames and defender locations, *StatsBomb News*, <https://statsbomb.com/news/statsbomb-data-case-studies-freeze-frames-and-defender-locations/> (Accessed: 23 October 2025) (Mar 2021).
- [6] G. G. Analytics, So... everyone has pressure data now, *Get Goalside Analytics (blog)*, <https://getgoalsideanalytics.com/so-everyone-has-pressure-data-now/> (Accessed: 23 October 2025) (Nov 2021).
- [7] P. Lucey, A. Bialkowski, M. Monfort, P. Carr, I. Matthews, Quality vs quantity: Improved shot prediction in soccer using strategic features from spatiotemporal data, in: *Proceedings of the 8th MIT Sloan Sports Analytics Conference*, 2014, pp. 1–9.
- [8] J. Mead, A. O’Hare, P. McMenemy, Expected goals in football: Improving model performance and demonstrating value, *PLOS One* 18 (2023) e0282295. doi:10.1371/journal.pone.0282295.
- [9] M. Cavuş, P. Biecek, Explainable expected goal models for performance analysis in football analytics, in: *2022 IEEE 9th International Conference on Data Science and Advanced Analytics (DSAA)*, 2022, pp. 1–9. doi:10.1109/DSAA54385.2022.10032440.
- [10] I. Bandara, S. Shelyag, S. Rajasegarar, D. Dwyer, E.-j. Kim, M. Angelova, Predicting goal probabilities with improved xg models using event sequences in association football, *PLOS One* 19 (2024) e0312278. doi:10.1371/journal.pone.0312278.
- [11] E. Santos-Fernández, P. Wu, K. Mengersen, Bayesian statistics meets sports: A comprehensive review, *Journal of Quantitative Analysis in Sports* 15 (2019) 289–312. doi:10.1515/jqas-2018-0106.
- [12] G. Baio, M. Blangiardo, Bayesian hierarchical model for the prediction of football results, *Journal of Applied Statistics* 37 (2010) 253–264. doi:10.1080/02664760802684177.
- [13] L. Egidi, J. Gabry, Bayesian hierarchical models for predicting individual performance in football (soccer), in: *Proceedings of the MathSport International 2017 Conference*, 2017.
- [14] J. Lee, J. Kim, H. Kim, J.-S. Lee, A bayesian approach to predict football matches with changed home advantage in spectator-free matches after the covid-19 break, *Entropy* 24 (2022) 366. doi:10.3390/e24030366.
- [15] D. B. Rubin, Estimating causal effects of treatments in randomized and nonrandomized studies, *Journal of Educational Psychology* 66 (1974) 688–701. doi:10.1037/h0037350.
- [16] J. Pearl, *Causality: Models, Reasoning and Inference*, Cambridge University Press, 2000.
- [17] V. Chernozhukov, D. Chetverikov, M. Demirer, E. Duflo, C. Hansen, W. Newey, J. Robins, Double/debiased machine learning for treatment and structural parameters, *Econometrica Journal* 21 (2018) C1–C68. doi:10.1111/ectj.12097.
- [18] S. Wager, S. Athey, Estimation and inference of heterogeneous treatment effects using random forests, *Journal of the American Statistical Association* 113 (2018) 1228–1242. doi:10.1080/01621459.2017.1319839.
- [19] H. Nakahara, K. Takeda, K. Fujii, Estimating the effect of team hitting strategies using counterfactual virtual

- simulation in baseball, *International Journal of Computer Science in Sport* 22 (2023) 1–12. doi:10.2478/ijcss-2023-0001.
- [20] D. M. Vock, L. F. B. Vock, Estimating the effect of plate discipline using a causal inference framework: An application of the g-computation algorithm, *Journal of Quantitative Analysis in Sports* 14 (2018) 37–56. doi:10.1515/jqas-2016-0029.
- [21] S. Alam, E. E. M. Moodie, L. Y. Wu, T. B. Swartz, Framing causal questions in sports analytics: A case study of crossing in soccer, *arXiv preprint 2505.11841* (2025). doi:10.48550/arXiv.2505.11841.
- [22] Z. Wang, P. Veličković, D. Hennes, et al., Tacticai: An ai assistant for football tactics, *Nature Communications* 15 (2024) 1906. doi:10.1038/s41467-024-45965-x.
- [23] StatsBomb, Free football data from statsbomb, GitHub, <https://github.com/statsbomb/open-data> (Accessed: 23 October 2025) (n.d.).
- [24] S. Interactive, *Football manager 2017* [game], Sega Europe Ltd, London (2016).
- [25] G. Anzer, P. Bauer, A goal scoring probability model for shots based on synchronized positional and event data in football (soccer), *Frontiers in Sports and Active Living* 3 (2021) 624475. doi:10.3389/fspor.2021.624475.