

The Manager Bounce: From Immediate Shock to Tactical Identity

Sean Hooker

Introduction

The demands placed on modern-day football managers are ever-growing. As recently as October 2025, Ange Postecoglou was dismissed from his role at Nottingham Forest after just 39 days in charge. This highlights the extreme pressures that managers and their staff face to deliver immediate results at the highest level.

Previous work has explored the short-term impact of managerial changes, both from a results perspective (1,2) and a physical performance standpoint (3). The impact on results tends to be short-lived (if present at all), with the so-called “manager bounce” often diminishing quickly—particularly once contextual factors such as fixture difficulty are accounted for. Similarly, studies assessing physical performance have found no clear improvements in player output following mid-season managerial changes.

Through our work, we aim to analyse this “bounce” effect through an *identity* lens. Specifically, we seek to evaluate the immediate tactical impact a manager can impose on a team, and how long it takes for an identifiable tactical identity to emerge. We aim to quantify these identities using what we term *micro metrics*—metrics that capture key tactical characteristics commonly associated with a manager’s style, such as pressing line height, directness, compactness, and speed in attacking play.

By constructing and combining these micro metrics into a vectorised standalone representation, we are able to capture a unique “fingerprint” of a manager for each game. Using these vectorised identities, we then measure the team’s tactical profile before and after a managerial change to analyse both the immediate “manager bounce” and the trajectory of tactical convergence under the new regime. We further assess how quickly different aspects of a manager’s identity take hold, and which contextual or structural factors influence this rate of convergence.

Main Contributions

Our research aims to provide the following key contributions:

1. **A context-adjusted framework for continuous performance metrics** – allowing for fairer evaluation of managerial impact across differing fixture difficulties and team contexts.
2. **A framework for constructing manager identity micro-vectors** – providing a method to quantify a manager’s tactical fingerprint.
3. **Identification of tactical convergence events** – analysing the rate and stability with which a manager’s identity is established over time.

1.1 Manager Bounce

The term “*Manager Bounce*” refers to the short-term improvement in team performance often observed after a new manager’s appointment, particularly during mid-season changes when the turnaround between fixtures is minimal.

Ramani (2) examined this phenomenon across the top five European leagues over six seasons, defining a *New Manager Bounce (NMB)* as a case where a new manager achieves better results than their predecessor. Their analysis showed a clear shift towards more wins and draws immediately after managerial changes, highlighting a temporary uplift in performance focused primarily on results.

Further work by Besters et al., building on Van Ours and van Tuijl (2016), found little lasting impact from mid-season appointments. Using the *cumulative surprise metric*—which compares actual points gained to expected points while controlling for opposition strength—they showed that most dismissals occurred after runs of poor form, creating the illusion of improvement once results reverted to expectation. Their Generalised Linear Model (GLM), which included controls for team strength and match location, confirmed that managerial changes had no significant effect on expected versus observed points.

Overall, while a new manager can spark an immediate response, the medium-term impact is typically limited. Explanations include shifts in morale, interpersonal dynamics, or tactical adjustments. Our work focuses on the latter—isolating the tactical component of the manager bounce to examine how quickly and to what extent a new manager can alter a team’s playing identity.

1.2 Tactical Identity

1.2.1 Team Identity

Numerous studies have sought to quantify *team identity* and *playing style* (4,5,6,22).

Defining a team's identity is inherently complex. Football produces constantly changing contexts, meaning any framework must integrate in- and out-of-possession actions, attacking and defensive metrics, and contextual variables such as match state and opposition strength. Less tangible aspects—like leadership or mentality—also shape identity but remain difficult to measure.

Fernandez-Navarro et al. (2016) used factor analysis on 19 performance indicators (14 attacking, 5 defensive) to develop team style profiles in La Liga and the Premier League (2006–07). Their work identified key stylistic metrics but, as Sharpe (2022) notes, such models often overlook context and rely on subjective variable choice.

Building on Franks et al. (2015) in the NBA, Sharpe (2022) modelled playing styles spatially—encoding open-play events and compressing pitch heat maps through *Non-Negative Matrix Factorisation (NMF)*. This approach learns stylistic traits directly from the data rather than relying on predefined metrics, offering a more objective description of team characteristics. Sharpe also proposed that these learned identities could extend to analysing *managerial identity*, given the strong link between team and coach.

Our work draws on this concept but focuses specifically on the *immediate tactical impact* of new managers. Instead of adopting Sharpe's NMF framework, we target "coachable" metrics—behaviours that can realistically shift early in a manager's tenure—to observe the *manager bounce* from an identity perspective. We also introduce a framework allowing greater control over *context adjustment*, whereby we felt this would have been difficult adapting their frequency based context adjustment.

1.2.2 Manager Identity

Monte (2024) and Sharpe (2022) each explore links between measurable metrics and a manager's tactical style. Monte, for example, uses passing and pressing data to assess influence, though he notes that such analyses often raise more questions than answers—especially regarding fair evaluation of a manager's ability or adaptability.

We adopt the view that the modern manager's role increasingly resembles that of a *coach* rather than a traditional overseer (8,9). This idea underpins our metric selection: while a new manager cannot immediately make a misfiring striker score, they can influence collective behaviours such as defensive compactness, pressing intensity, or passing versus carrying tendencies.

Accordingly, our analysis focuses on these *coachable tactical dimensions*—aspects most likely to shift quickly and therefore reveal a manager's emerging identity.

2. Data

2.1.1 Hudl Data

This research made use of Hudl's lineups, events, and 360-frame datasets.

- The *lineups data* provided insight into starting XIs, in-match tactical shifts, and key match events such as goals, substitutions, and dismissals.
- The *events data* was used to construct all performance metrics in this study, drawing on detailed information such as action types, locations, and the players involved.
- The *360-frame data* complemented this by capturing the spatial positioning of all players visible within each event frame, providing essential contextual information.

Hudl data formed the core of our dataset, from which all primary tactical and identity metrics were derived.

2.1.2 Manager Appointments

Managerial appointment data was sourced from Transfermarkt (10), which provides comprehensive records of managerial transitions, including predecessors, successors, match results, and appointment dates.

This dataset was later merged with the Hudl data to create a unified record linking managerial tenures, results, and team identities across the study period.

2.1.3 Premier League Data

All data were collected from the Premier League covering the 2020/2021 to 2024/2025 seasons. During this period, there were 85 managerial appointments (including caretaker roles).

Compared with other top-five European leagues, this represents a relatively moderate level of managerial turnover.

| League | Manager Changes |
|----------------|-----------------|
| Serie A | 111 |
| Ligue 1 | 90 |
| La Liga | 90 |
| Premier League | 85 |
| Bundesliga | 82 |

Figure 2.1. Manager changeover from 2020/21 to 2024/25 in the top five European leagues (including caretaker managers). Data taken from Transfermarkt.

The Premier League was selected to balance the availability of sufficient “bounce” events with the ability to observe longer-term *convergence* effects.

Leagues with higher managerial turnover risked fragmenting the dataset and reducing the reliability of convergence analysis. Additionally, the Premier League hosts many globally recognised managers, making identity associations more intuitive and well-documented.

2.1.4 Elo Rankings

Elo rankings were used as a historical, unbiased measure of team strength based on past performance.

The Elo system (12) begins with all teams at an equal baseline rating, adjusting points after each result depending on the relative strength of the opponent—more points are awarded for wins or draws against higher-rated teams.

Historical Elo data were obtained from Club Elo (13).

Benefits of Elo Rankings:

- Provide a consistent, quantitative measure of team strength over time.
- Enable context-adjusted comparisons across different fixtures and time periods.

Alternative approaches to modelling team strength exist, such as using bookmaker odds, as demonstrated by Skripnikov et al. (2025). However, we opted for Elo ratings to maintain simplicity and avoid short-term market biases—odds can fluctuate due to player availability or public perception rather than representing true historical team strength.

2.1.5 Pitch Dimensions

For all spatial calculations requiring pitch dimensions (in metres), measurements were obtained from 14 & 15. Any stadiums where we were not able to obtain a dimension were defaulted to 105x68m.

2.1.6 Meta Information

In total, 1,900 Premier League matches were analysed across five seasons, yielding 3,800 fixture-level data points (home and away).

From these, we focused on teams under managers appointed during the study period.

Managers with tenures shorter than 30 days—predominantly caretaker appointments—were excluded, as such brief spells would distort convergence analysis.

| Club | Permanent Coach In |
|-------------------------|--|
| AFC Bournemouth | {Andoni Iraola, Gary O'Neil} |
| Aston Villa | {Steven Gerrard, Unai Emery} |
| Brentford FC | {Keith Andrews} |
| Brighton & Hove Albion | {Fabian Hürzeler, Roberto De Zerbi} |
| Burnley FC | {Vincent Kompany, Scott Parker} |
| Chelsea FC | {Mauricio Pochettino, Enzo Maresca, Thomas Tuchel, Graham Potter} |
| Crystal Palace | {Patrick Vieira, Oliver Glasner, Roy Hodgson} |
| Everton FC | {Rafael Benítez, Frank Lampard, Sean Dyche, David Moyes} |
| Fulham FC | {Marco Silva} |
| Leeds United | {Sam Allardyce, Javi Gracia, Jesse Marsch, Daniel Farke} |
| Leicester City | {Dean Smith, Enzo Maresca, Martí Cifuentes, Ruud van Nistelrooy} |
| Liverpool FC | {Arne Slot} |
| Manchester United | {Rúben Amorim, Erik ten Hag} |
| Newcastle United | {Eddie Howe} |
| Norwich City | {Dean Smith} |
| Nottingham Forest | {Nuno Espírito Santo} |
| Sheffield United | {Chris Wilder, Slavisa Jokanovic} |
| Southampton FC | {Will Still, Ivan Juric, Russell Martin, Nathan Jones, Rubén Sellés} |
| Tottenham Hotspur | {Thomas Frank, Antonio Conte, Nuno Espírito Santo, Ange Postecoglou} |
| Watford FC | {Claudio Ranieri, Roy Hodgson, Rob Edwards} |
| West Bromwich Albion | {Sam Allardyce, Valérien Ismaël} |
| West Ham United | {Julen Lopetegui, Graham Potter} |
| Wolverhampton Wanderers | {Vítor Pereira, Bruno Lage, Julen Lopetegui, Gary O'Neil} |

Figure 2.2. Permanent Manager appointments in the Premier League from 2020/21 until 2024/25. Data from Transfermarkt.

Additionally, managers of newly promoted teams were excluded from the “new manager” category due to the absence of prior Championship-level data for comparison.

2.2 Metrics

We constructed tactical identities by creating *micro-vectors* — sets of metrics intrinsically linked to a manager’s tactical influence.

Sharpe (2022) highlights the subjective fragility of defining identities based solely on pre-selected metrics.

In this work, we deliberately selected variables that can realistically change within the immediate timeframe of a new managerial appointment.

Harder-to-coach characteristics (e.g., counter-pressing effectiveness, xG) were excluded, as these tend to stabilise only after a recruitment cycle and could distort convergence analysis.

We acknowledge that the chosen metrics do not capture all dimensions of managerial identity. Certain aspects — for example, pressing effectiveness — could be further refined through qualitative evaluation.

However, our intent is to isolate the most *coachable* and *immediate* stylistic signals. A manager can, for instance, instruct a team to press higher from their first session, but cannot instantly control the quality or success of those presses.

All metrics are calculated at the event level and aggregated per match using either the mean or median, depending on distribution.

2.2.1 Pressing Line Height

Morgan (2018) proposed a method to measure counter-pressing using “Pressure” events, defining a counter-press as a pressure occurring within five seconds of losing possession.

Similarly, we use pressure events to compute pressing line height.

All pressure events in the opponent’s half are identified, and the location of the pressure-applying player is recorded.

We then locate all teammates within ± 10 Hudl units on the x-axis (bounded by the halfway line) in the associated 360-frame and calculate their mean positions.

Intuition:

Managers can instruct teams either to press high in the opposition build-up or to sit deeper and absorb pressure.

Constraints:

- Minimum of 3 teammates captured in the frame at the time of pressure (to avoid single-player distortions).
- Players must be positioned in the opposition half.

Limitations:

- The success rate of presses is not captured.
- Possession-dominant teams may have fewer pressure events, introducing variance (accounted for later through event weighting).

Each pressure event is aggregated per match using the mean value.

Sense-checking for 2024/25, we see the following managers rank highest (min 10 games).

| Team Name | Manager | Mean PLH Line M |
|------------------------|------------------------|-----------------|
| Manchester City | Josep Guardiola i Sala | 72.091186 |
| Arsenal | Mikel Arteta Amatriain | 71.670026 |
| Tottenham Hotspur | Angelos Postecoglou | 71.233191 |
| Brentford | Thomas Frank | 70.507499 |
| Brighton & Hove Albion | Fabian Hürzeler | 70.470632 |

Figure 2.3. Managers with the largest Avg Pressing Line Heights aggregated across 2024/25 season (minimum 10 games).

2.2.2 Compactness

Rico-González et al. (2021) provide an overview of spatial-tactical metrics derived from positional data, including surface and coverage area calculations.

Building on their work and that of Guan et al. (18), we use convex hull geometry to measure a team's defensive compactness.

We calculate the convex hull area of the defending team within their own half for all opposition possession and defensive events.

Intuition:

Compactness is typically associated with defensively disciplined coaches *“If you have the ball you must make the field as big as possible, and if you don’t have the ball you must make it as small as possible”* (Cruyff, as cited in *Breaking the Lines*, 2023).

Constraints:

- Minimum of 8 players from the defending team present in the frame.
- All players are located within their own half.
- Goalkeepers excluded.

Limitations:

- Low sample sizes may occur for dominant teams, though this is mitigated by including all opposition attacking actions.

Compactness is expressed as the median normalised area (convex hull area divided by half-pitch area).

2.2.2.1 Convex Hull Area

The convex hull represents the smallest polygon encompassing all player locations. Its area is computed using Gauss’ shoelace formula (20):

$$A = \frac{1}{2} \sum_{i=1}^n x_i y_{i+1} - x_{i+1} y_i$$

The term “shoelace” refers to the criss-crossing pattern of summation resembling laced shoes.

Sense-checking the 2024/25 season, we find the following managers exhibiting the most compact defensive shapes, based on mean area (m²) across teams with a minimum of 10 games.

| Team Name | Manager | Mean Compactness M2 |
|-----------------|------------------------|---------------------|
| Brentford | Thomas Frank | 332.334718 |
| Arsenal | Mikel Arteta Amatriain | 354.450119 |
| Ipswich Town | Kieran McKenna | 360.743367 |
| Manchester City | Josep Guardiola i Sala | 361.586718 |
| Everton | Sean Dyche | 362.066506 |

Figure 2.4. Managers with the smallest average compactness m^2 aggregated across 2024/25 season (minimum 10 games).

2.2.3 Directness

This metric measures a team's forward progression ratio, defined as the proportion of total distance from "Pass" and "Carry" actions that move positively along the x-axis.

$$\text{Directness Ratio} = \frac{\sum \Delta x_i}{\sum \sqrt{(\Delta x_i)^2 + (\Delta y_i)^2}}$$

Backward or lateral movements are penalised to isolate direct attacking intent.

Intuition:

Captures teams that favour forward, vertical play over slower, lateral build-up.

Constraints:

- Actor events are limited to passes and carries.

Limitations:

- This simplified measure does not distinguish between defensive-third and attacking-third directness (Monte, 2024). Our framework intentionally begins with base metrics before extending into more granular ones in the future.

| Team Name | Manager | Average Forward Progression Ratio |
|-------------------|--------------------------------------|-----------------------------------|
| Nottingham Forest | Nuno Herlander Simões Espírito Santo | 0.417603 |
| Everton | Sean Dyche | 0.393560 |
| Crystal Palace | Oliver Glasner | 0.376673 |
| Everton | David Moyes | 0.367824 |
| Ipswich Town | Kieran McKenna | 0.365017 |

Figure 2.5. Managers with the highest average forward progression ratios across the 2024/25 season (minimum 10 games).

2.2.4 Progression Speed (m/s)

To assess how rapidly teams advance the ball, we calculate the average forward progression speed (m/s) across all positive actions.

Distance in Hudl units is converted to metres using known pitch dimensions, and time between sequential related events (e.g., "Shot", "Receive", "Dribble") is used to derive speed.

Event-level values are aggregated per match using the median, mitigating skew from long passes or counter-attacks. This is shown in Table 2.1, which highlights median better reflects central tendency.

| Mean-Median Gap (%) | P10 (%) | P50 (Median) (%) | P90 (%) | Within 10% of Median (%) | Within 20% of Median (%) |
|---------------------|---------|------------------|---------|--------------------------|--------------------------|
| 35.9 | 19.97 | 32.40 | 50.55 | 0.74 | 10.1 |

Table 2.1. Difference in mean and median values taken against all recorded progression speed metrics.

Figure 2.6 allows us to visually see an example of this distribution from a recorded match.




Figure 2.6. Distribution of progression speed metric counts for Manchester City vs Aston Villa (20 Jan 2021).

Intuition:

Represents the tempo of attacking play – slower speeds suggest methodical build-up; higher speeds indicate transitions, long balls or counter-attacking intent.

Constraints:

- Events must be "Pass" or "Carry".
- The ball must move forward (+x).
- Events must have positive On Ball Value (OBV).

Limitations:

- A broad approximation of attacking speed. More advanced models could separate structured build-up from transition play.

| Team Name | Manager | Median Forward M/S |
|-------------------|--------------------------------------|--------------------|
| Everton | Sean Dyche | 5.545409 |
| Everton | David Moyes | 5.464928 |
| Crystal Palace | Oliver Glasner | 5.383285 |
| Nottingham Forest | Nuno Herlander Simões Espírito Santo | 5.332410 |
| Leicester City | Ruud van Nistelrooy | 5.280296 |

Figure 2.7. Managers with the fastest forward median m/s aggregated across 2024/25 season (minimum 10 games).

2.2.5 Share of Carries

Finally, we compute the carry-to-pass ratio during positive progressions, using the same filters applied in the progression-speed metric:

$$\text{Share of Carries} = \frac{n_{\text{carry}}}{n_{\text{total}}}$$

Where we n_{total} is the total number of recorded positive progression events ("Pass" & "Carry").

Intuition:

Distinguishes teams preferring to advance via dribbling (carry-heavy) versus passing.

Constraints:

- Same as for progression speed.

Limitations:

- Carrying behaviour is largely player-dependent, influenced by player profiles and confidence levels. While a manager may encourage more ball-carrying, the extent to which it is executed depends on the technical quality of the players available. Figure 2.8 illustrates this to some extent, with Manchester City leading the way – supported by technically gifted ball carriers such as Doku and Haaland.

| Team Name | Manager | Carry to Pass Ratio |
|-------------------------|------------------------|---------------------|
| Manchester City | Josep Guardiola i Sala | 0.506677 |
| Arsenal | Mikel Arteta Amatriain | 0.482547 |
| Wolverhampton Wanderers | Gary Paul O'Neil | 0.473150 |
| Brighton & Hove Albion | Fabian Hürzeler | 0.473029 |
| Southampton | Russell Martin | 0.472158 |

Figure 2.8. Managers with the largest carry to pass ratio for positive progressive events aggregated across 2024/25 season (minimum 10 games).

3.1 Context Adjustment

When analysing managerial identity and the rate at which it converges, we sought to control for contextual factors to minimise noise. A team’s playing style is strongly influenced by elements such as opposition strength, venue, and game state.

To disentangle these external influences and obtain a purer reflection of the manager’s tactical signature, we context-adjusted all recorded metrics.

This process produced context-aware identity vectors, allowing us to track convergence more accurately and reduce bias from situational variation.

3.1.2 Modelling Context

Our approach extends the recent work of Skripnikov et al. (2025), who adjusted offensive metrics using event data to account for game context (e.g. winning, losing). They employed Generalised Additive Models (GAMs) to model count-based variables such as shots and corners, using a Negative Binomial distribution.

Their framework involved predicting a metric both with and without context:

- The *context-unaware* prediction assumed a level game, no red cards, and a home venue.
- The difference between these predictions represented the contextual effect, which was then applied to each observation.

Since our variables are continuous rather than count-based, we adapted their method. We introduced context adjustment for continuous metrics, reinforcing their findings that context-aware adjustments reduce misrepresentation of performance.

Our main extensions include:

1. Predicting context-unaware scenarios to derive delta adjustments for continuous variables.
2. Applying context at the match-aggregate level rather than event-by-event – a simpler yet robust approach given our research objectives.

3.1.3 Contextual Variables

3.1.3.1 Team Strength

Team strength was measured using Elo rating difference on the day of each match. Larger positive values indicate you are strong in comparison to your opposition.

As Elo differences exhibit wider variance than other contextual variables, we applied z-score normalisation (21) by season to stabilise sensitivity in our models.

3.1.3.2 Man Net

This variable captures the net player advantage a team experiences.

Using lineup data, we identified all red card events (“Second Yellow”, “Red Card”) and calculated the proportion of match time played with a numerical advantage or disadvantage.

$$man_{net} = \frac{1}{T} \int_0^T [P_i(t) - P_j(t)] dt$$

Where $P_i(t)$ is the number of on-field players for team i at time t, $P_j(t)$ is the number of on-field players for team j at time t and T is the total time of the match.

A quick example (90' match, T=5400 s):

- Opponent down 2 for 10 min (+2 for 600 s)
- Opponent down 1 for 5 min (+1 for 300 s)

man_net:

$$man_{net} = \frac{2 \cdot 600 + 1 \cdot 300}{5400} = \frac{1500}{5400} \approx 0.278$$

3.1.3.3 Man Abs

Similar to *Man Net*, but representing the absolute player differential, irrespective of whether the team was advantaged or disadvantaged.

$$man_{abs} = \frac{1}{T} \int_0^T |P_i(t) - P_j(t)| dt$$

3.1.3.4 Home Advantage

We included a binary variable `is_home` (0 = away, 1 = home) to model home advantage explicitly within all regression models.

3.1.3.5 Winning Share

Winning share quantifies the proportion of match time during which a team was leading. Goal event timestamps from the lineup data were used to determine when a team was ahead, and these durations were aggregated and divided by total match length (also derived from start/end timestamps).

3.1.3.6 Losing Share

Calculated identically to winning share, but capturing time spent trailing.

The remaining share of match time represents periods where the game was level.

3.1.3.7 Seasonal Fixed Effects

Because tactical trends and rules evolve across seasons, we included seasonal fixed effects, using 2020/2021 as the baseline.

This adjustment controls for variation outside a manager's control.

3.1.4 Model Evaluation

The objective of our models was to predict the expected metric value both with and without contextual adjustments.

Model performance and calibration were evaluated using the following diagnostics:

- Probability Integral Transform (PIT)(25). Assesses predictive uncertainty by transforming cumulative distribution function (CDF) values into uniform quantiles. Perfectly calibrated models exhibit uniform $U(0,1)$ distributions, while skewed "cone" or "U" shapes indicate bias.
- Kolmogorov-Smirnov (KS) Test (25). Tests whether the PIT distribution significantly deviates from uniformity. When $p < 0.05$, the null hypothesis of calibration is rejected.
- Residuals vs Fitted Values Plot. Used to check for heteroskedasticity visually – residuals should scatter uniformly around zero.
- Empirical coverage (pooled): Assesses model calibration across different sample sizes by measuring the proportion of observed values that fall within the 95% prediction interval for each decile of frame count per match. Consistent coverage near 95% across all deciles indicates stable model uncertainty regardless of the number of events or frames.

3.2 Model Selection

Model summaries and diagnostic plots are provided in the Appendix.

3.2.1 Ordinary Least Squares (OLS)

OLS regression is commonly applied in sports analytics, modelling a dependent variable Y as a linear combination of covariates X under the assumptions of independence, normality, and homoscedasticity.

$$y_i = \beta_1 x_{i1} + \beta_2 x_{i2} + \dots + \beta_p x_{ip} + \varepsilon_i$$

These assumptions were violated for several of our metrics, which are non-negative, bounded, and mildly heteroskedastic.

Applying OLS under such conditions risks biased coefficients and invalid predictions (e.g., negative expected values).

For example, an OLS model on Pressing Line Height yielded:

Breusch–Pagan LM: $p = 0.0166$; White test: $p = 0.0317$; PIT KS: $p = 0.0412$

indicating heteroskedasticity and miscalibration.

We also observed a nonlinear relationship (Figure 3.1) between Elo difference and the dependent variable, prompting us to include a quadratic term for Elo in all subsequent models.




Figure 3.1. Non-linear relationship between the z-score of Elo difference (explanatory variable) and Pressing Line Height (response variable).

While spline terms offered marginal improvements in KS scores, they introduced severe multicollinearity (condition numbers > 1000). We therefore retained the simpler quadratic specification.

3.2.2 Generalised Linear Models (GLMs)

GLMs (22) provided a more flexible and robust framework. Each GLM comprises:

1. A **distribution family** (e.g., Gamma, Beta, Poisson)
2. A **systematic component** ($X\beta$)
3. A **link function** $g(\mu)$ connecting predictors to the mean response.

These models naturally handle heteroskedasticity by allowing variance to scale with the mean.

We applied robust HC3 covariance estimators to stabilise inference against mild heteroskedasticity (27).

3.2.3 Gamma GLMs

Metrics such as Pressing Line Height and Progression Speed (m/s) were modelled using Gamma GLMs – suitable for continuous and positive variables.

$$y_i > 0 : \quad y_i \sim \text{Gamma}(k, \theta_i)$$

$$\text{Link function :} \quad \log(\mu_i) = \eta_i = x_i^\top \beta$$

$$\text{Mean :} \quad E[y_i] = \mu_i = k\theta_i$$

$$\text{Variance :} \quad \text{Var}(y_i) = \phi\mu_i^2$$

- The Gamma GLM assumes $y_i > 0$.
- μ_i is the conditional mean of the response.
- ϕ is the dispersion (scale) parameter.
- The log link ensures positive fitted means.

Diagnostic results for both models indicated good calibration (e.g., PIT KS $p > 0.18$; 95% empirical coverage ≈ 0.95) in Figure 3.2.

Minor heteroskedasticity, seen in Figure 6.2 & Figure 6.5 for the respective metrics, was observed at extreme fitted values but was deemed acceptable.

No multicollinearity was detected (VIF < 2.0 , condition numbers ≈ 2.7) in Figure 3.2.

All context variables were statistically significant except *Winning Share* in Pressing Line Height (Figure 6.1), and *Man Abs* in Progression Speed (Figure 6.4).

Given their theoretical relevance and minimal effect on model performance, both were retained for consistency across metrics.

3.2.4 Beta GLMs

Beta distributions are bounded between 0 and 1, making them appropriate for ratio-based metrics such as Compactness, Directness, and Carry Share.

$$y_i \in (0, 1) : \quad y_i \sim \text{Beta}(\mu_i \phi, (1 - \mu_i) \phi)$$

$$\text{Link function} : \quad \text{logit}(\mu_i) = \eta_i = x_i^\top \beta$$

$$\text{Mean} : \quad E[y_i] = \mu_i$$

$$\text{Variance} : \quad \text{Var}(y_i) = \frac{\mu_i(1 - \mu_i)}{1 + \phi}$$

- y_i is the dependent variable (bounded between 0 and 1).
- μ_i is the expected value of y_i .
- x_i is the vector of explanatory variables.
- β is the vector of regression coefficients.

- ϕ (precision parameter) controls the dispersion – higher ϕ means lower variance.
- The logit link ensures $\mu_i \in (0,1)$.

All Beta models showed strong calibration (e.g., KS $p > 0.08$; coverage ≈ 0.95) in Figure 3.2.

Residual patterns suggested mild curvature in Directness (Figure 6.9), likely due to variance scaling; this was addressed through weighted modelling in later stages.

All explanatory variables were significant except *Man Abs* (Figure 6.9 & 6.11), which was again retained for theoretical consistency.

| | Family | Link | PIT KS D | PIT KS p | Empirical 95% Coverage (pooled) | Condition Number (cond) | Max VIF |
|--------------------------------|--------|-------|----------|----------|---------------------------------|-------------------------|----------|
| Metric | | | | | | | |
| Pressing Line Height | Gamma | log | 0.017620 | 0.186792 | 0.947867 | 2.682900 | 1.963818 |
| Progression Speed (m/s) | Gamma | log | 0.014274 | 0.417201 | 0.956316 | 2.679606 | 1.961666 |
| Compactness | Beta | logit | 0.013522 | 0.486772 | 0.948394 | 2.682912 | 1.963972 |
| Share Carry | Beta | logit | 0.020504 | 0.080799 | 0.951053 | 2.679606 | 1.961666 |
| Directness | Beta | logit | 0.017800 | 0.179800 | 0.949000 | 2.679606 | 1.961666 |

Figure 3.2. Model diagnostics summary for fitted Gamma and Beta regressions across performance metrics. PIT-KS values and empirical 95% coverage indicate good calibration, while low condition numbers and variance inflation factors (VIF) suggest negligible multicollinearity.

3.3 Creating Context-Aware Metrics

With all models finalised, we adjusted each metric using its corresponding GLM's expected values.

Following Skripnikov et al. (2025), we defined a neutral context where all explanatory variables were set to zero.

Unlike their work, which assumed home conditions, we set `is_home = 0` (away). This avoids inflating performance for teams with disproportionate home advantages.

We also retained the team strength variable in the neutral scenario, but fixed it at a neutral value. This decision reflects our view that opponent strength substantially influences managerial tactics; by neutralising this factor, we can better assess how a manager might set up a team under ideal, context-independent conditions.

For Beta models, the context adjustment was defined as:

$$\text{logit}(\mu_i) = \eta_i = x_i^\top \beta$$

$$\text{Context adjustment : } \Delta_i = \eta_{i,obs} - \eta_{i,neutral}$$

$$\text{Adjusted(neutral)response : } \text{logit}(y_{i,adj}) = \text{logit}(y_i) - \Delta_i$$

$$\Rightarrow y_{i,adj} = \text{logit}^{-1}(\text{logit}(y_i) - (\eta_{i,obs} - \eta_{i,neutral}))$$

$$\text{where } \text{logit}^{-1}(z) = \frac{1}{1 + e^{-z}}$$

- $\eta_{i,obs}$: linear predictor under *observed* match context.
- $\eta_{i,neutral}$: predictor under *neutralised* context (e.g., equal strength, 11v11, neutral venue).
- The difference Δ_i quantifies the contextual bias.
- The adjusted value $y_{i,adj}$ rescales the observed performance to that neutral baseline.

For Gamma models, expected values were derived as:

$$\text{Link function : } \log(\mu_i) = \eta_i = x_i^\top \beta$$

$$\text{Context adjustment : } \Delta_i = \eta_{i,obs} - \eta_{i,neutral}$$

$$\text{Adjusted(neutral)response : } y_{i,adj} = y_i \cdot e^{-\Delta_i}$$

$$\text{Equivalently : } y_{i,adj} = y_i \cdot \exp(-(\eta_{i,obs} - \eta_{i,neutral}))$$

$$\text{where } \eta_{i,obs} = \log(\mu_{i,obs}), \quad \eta_{i,neutral} = \log(\mu_{i,neutral})$$

- $\mu_{i,obs}$: expected value under the *observed* context (e.g., actual home/away, manpower, score state).
- $\mu_{i,neutral}$: expected value under the *neutral* context (e.g., equal strength, 11v11, neutral venue).
- Δ_i : log-scale difference capturing the contextual effect.
- $y_{i,adj}$: adjusted value that removes contextual bias while remaining on the original (positive, continuous) scale.

These adjusted metrics form the foundation of our context-aware identity vectors, used to evaluate the temporal convergence of managerial identities.

4.1 Micro-Vector Identities

4.1.1 Identity Creation

To construct each team's tactical identity, we merged all context-adjusted metrics into a single *micro-vector*. The term *micro* reflects that this vector represents an aggregation of stylistic indicators rather than a fully fledged, definitive measure of identity.

$$\text{identity_vector} = \begin{bmatrix} \text{pressing_line_height_context_adj_z} \\ \text{progression_speed_context_adj_z} \\ \text{compactness_context_adj_z} \\ \text{share_carry_context_adj_z} \\ \text{directness_context_adj_z} \end{bmatrix}$$

4.1.2 Standerdisation

All stylistic metrics were standardised using robust global z-scores (28) to ensure comparability across dimensions:

$$z_i = \frac{x_i - \text{median}(x)}{1.4826 \times \text{MAD}(x)}$$

We adopted a *global* rather than season-specific scale to enable cross-season comparisons. This approach also ensures that each component contributes equally to the identity vector, preventing metrics with larger inherent variance (especially those derived from Gamma models) from dominating the analysis.

Metrics exhibiting (mild) skews (Figures 6.13 - 6.17) were first log-stabilised:

$$x'_i = \log(1 + x_i)$$

to reduce the influence of extreme values and approximate symmetry prior to scaling.

4.1.3 Calculating Distance

Similarity between consecutive identity vectors was measured using a weighted Mahalanobis-type distance (23) that accounts for uncertainty and reliability differences across metrics:

$$D_t^2 = \sum_{j=1}^k w_{t,j} (\Delta z_{t,j})^2, \quad \Delta z_{t,j} = z_{t,j} - z_{t-1,j}.$$

where $\Delta z_{t,j} = z_{t,j} - z_{t-1,j}$ represents the change in the standardised value of metric j between successive matches.

Although full covariance modelling was initially explored, it was limited by sample size and computational feasibility. Consequently, we adopted a diagonal approximation of the Mahalanobis distance, which in practice behaves similarly to a weighted Euclidean (z-score) distance.

This approach omits cross-metric covariances and therefore may not capture all sources of correlated noise between metrics. However, it retains sensitivity to relative within-metric variance and provides a stable, interpretable measure of tactical deviation over time. In this context, the simplification offers a pragmatic balance between model interpretability and computational tractability, while still isolating genuine tactical shifts once contextual effects are removed.

4.1.4 Modelling Uncertainty

Each metric was weighted inversely to its measurement uncertainty and data sparsity. Weights were defined as:

$$w_{t,j}^{-1} \propto \frac{1}{\tilde{w}_{t,j}} + \frac{1}{\tilde{w}_{t-1,j}} + \lambda (\tilde{R}_{t,j} + \tilde{R}_{t-1,j})$$

- $\tilde{w}_{t,j} \propto N_{frames,t}^{\alpha_j}$ reflects data reliability based on event coverage,
- $\tilde{R}_{t,j}$ is the normalised model variance.
- λ is a small regularisation constant (set to $\lambda = 0.2$ for our experiments) that controls the influence of model uncertainty. We explored values in the range 0.1–0.5 and selected 0.2 based on visual sense checks of stability and convergence behaviour; this was **not** a formal empirical optimisation, but a pragmatic choice that provided a meaningful—yet not dominant—weight on uncertainty.

The normalised model variance $\tilde{R}_{t,j}$ is obtained by dividing the raw model variance $R_{t,j}$ (derived from the GLM dispersion term) by its long-run standard deviation across matches:

$$\tilde{R}_{t,j} = \frac{R_{t,j}}{sd(R_{.,j})}, \quad sd(R_{.,j}) = \sqrt{\frac{1}{T_j - 1} \sum_{t=1}^{T_j} (R_{t,j} - \bar{R}_{.,j})^2}$$

This normalisation removes scale differences between metrics, ensuring that the reliability weighting operates on a comparable basis across all metrics.

4.1.5 Temporal Smoothing Over Time

To capture underlying stylistic stability, match-to-match distances were smoothed using an Exponentially Weighted Moving Average (EWMA):

$$EWMA_t = \alpha D_t + (1 - \alpha) EWMA_{t-1}$$

Where α was chosen to be 0.22, this was chosen as it provides a balance between short-term responsiveness and long-term stability.

This approach filters noise arising from transient conditions (e.g., injuries, match congestion) and highlights meaningful tactical adaptation.

4.2 Identity Convergence

Identity convergence is defined as the point where a manager's adjusted tactical profile stabilises around a consistent pattern.

For each manager, we tracked the EWMA of match-to-match identity distances and defined the threshold as the global median, computed only from observations recorded after each manager's first eight matches.

$$THR = \text{median} (D_t \mid t > 8)$$

This approach excluded early, potentially unstable performances at the start of a manager’s tenure. The resulting threshold (THR) was $THR=1.57$. We explored alternative cut-offs based on lower and upper quartiles, but adopted the median after visual checks indicated it delivered an optimal balance in the number of managers meeting the convergence criterion.

A convergence event occurs when a manager’s EWMA distance remains below this threshold for six consecutive matches.

Convergence occurs when $EWMA_t < THR$ for 6 consecutive matches.

This window approximates one month of competitive fixtures—long enough to capture sustained tactical consistency while filtering short-term noise.

4.3 Example of Convergence

Bringing all components together, Figure 4.1 illustrates Ange Postecoglou’s identity convergence trajectory. The plot visualises his match-to-match distances and marks the point of first sustained convergence.




Figure 4.1 Ange Postecoglou’s context-aware and context-unaware identity-vector distances over his tenure at Tottenham Hotspur. Convergence occurs after 12 games.

This example demonstrates how context-adjusted micro-vector identities substantially reduce noise (such as opposition strength or game state), isolating the tactical imprint attributable to the manager. We will explore these graphs further in the results section.

4.4 Summary of Identity Framework

This framework integrates context-adjusted performance metrics, multivariate distance modelling, and temporal smoothing to quantify how managerial identities evolve and stabilise over time.

By merging event-level data with contextual modelling, the approach captures both the immediacy of tactical influence and the long-term convergence of managerial style.

The resulting identity trajectories provide a foundation for subsequent analysis of managerial impact, enabling comparisons of adaptation rates, stylistic stability, and contextual resilience across managers and seasons.

5 Results

5.1 Convergence

Of the 85 Premier League managerial changes within the five-season window, 58 met our inclusion criterion of at least 30 days in post. Using our distance-and-convergence framework, we identified 31 convergence events among these 58 tenures—instances where a manager’s tactical identity stabilised according to the criteria defined earlier.

| Team | Prev Manager | New Manager | Prev Match Date | First Match Date | Gap Days | Matches until Convergence |
|-------------------------|--------------------------------------|--------------------------------------|-----------------|------------------|----------|---------------------------|
| Leicester City | Ben Dawson | Ruud van Nistelrooy | 2024-11-30 | 2024-12-03 | 3 | 6.0 |
| Crystal Palace | Patrick McCarthy | Oliver Glasner | 2024-02-19 | 2024-02-24 | 5 | 6.0 |
| Southampton FC | Ralph Hasenhüttl | Nathan Jones | 2022-11-06 | 2022-11-12 | 6 | 6.0 |
| Leeds United | Michael Skubala | Javier Gracia Carlos | 2023-02-18 | 2023-02-25 | 7 | 6.0 |
| Southampton FC | Nathan Jones | Rubén Sellés Salvador | 2023-02-11 | 2023-02-18 | 7 | 6.0 |
| Wolverhampton Wanderers | Bruno Miguel Silva do Nascimento | Steven Peter Davis | 2022-10-01 | 2022-10-08 | 7 | 6.0 |
| Tottenham Hotspur | Nuno Herlander Simões Espírito Santo | Antonio Conte | 2021-10-30 | 2021-11-07 | 8 | 6.0 |
| Everton FC | Frank Lampard | Sean Dyche | 2023-01-21 | 2023-02-04 | 14 | 6.0 |
| Chelsea FC | Thomas Tuchel | Graham Potter | 2022-09-03 | 2022-10-01 | 28 | 6.0 |
| Chelsea FC | Frank Lampard | Mauricio Roberto Pochettino Trossero | 2023-05-28 | 2023-08-13 | 77 | 6.0 |
| Manchester United | Ralf Rangnick | Erik ten Hag | 2022-05-22 | 2022-08-07 | 77 | 6.0 |
| Wolverhampton Wanderers | Nuno Herlander Simões Espírito Santo | Bruno Miguel Silva do Nascimento | 2021-05-23 | 2021-08-14 | 83 | 6.0 |
| Brighton & Hove Albion | Roberto De Zerbi | Fabian Hürzeler | 2024-05-19 | 2024-08-17 | 90 | 6.0 |
| Chelsea FC | Mauricio Roberto Pochettino Trossero | Enzo Maresca | 2024-05-19 | 2024-08-18 | 91 | 6.0 |
| Watford FC | Claudio Ranieri | Roy Hodgson | 2022-01-21 | 2022-02-05 | 15 | 9.0 |
| Wolverhampton Wanderers | Steven Peter Davis | Julen Lopetegui Argote | 2022-11-12 | 2022-12-26 | 44 | 9.0 |
| Liverpool FC | Jürgen Klopp | Arne Slot | 2024-05-19 | 2024-08-17 | 90 | 10.0 |
| Tottenham Hotspur | Ryan Mason | Angelos Postecoglou | 2023-05-28 | 2023-08-13 | 77 | 12.0 |
| Everton FC | Carlo Ancelotti | Rafael Benítez Maudes | 2021-05-23 | 2021-08-14 | 83 | 12.0 |
| Aston Villa | Aaron John Danks | Unai Emery Etxegoien | 2022-10-29 | 2022-11-06 | 8 | 14.0 |
| West Ham United | Julen Lopetegui Argote | Graham Potter | 2025-01-04 | 2025-01-14 | 10 | 14.0 |
| Crystal Palace | Patrick McCarthy | Roy Hodgson | 2023-03-19 | 2023-04-01 | 13 | 14.0 |
| Brighton & Hove Albion | Graham Potter | Roberto De Zerbi | 2022-09-04 | 2022-10-01 | 27 | 14.0 |
| Aston Villa | Dean Smith | Steven Gerrard | 2021-11-05 | 2021-11-20 | 15 | 17.0 |
| Crystal Palace | Patrick McCarthy | Roy Hodgson | 2023-09-16 | 2023-09-23 | 7 | 18.0 |
| Leeds United | Marcelo Alberto Bielsa Caldera | Jesse Marsch | 2022-02-26 | 2022-03-05 | 7 | 20.0 |
| Chelsea FC | Frank Lampard | Thomas Tuchel | 2021-01-19 | 2021-01-27 | 8 | 24.0 |
| Norwich City FC | Daniel Farke | Dean Smith | 2021-11-06 | 2021-11-20 | 14 | 27.0 |
| Nottingham Forest | Steve Cooper | Nuno Herlander Simões Espírito Santo | 2023-12-15 | 2023-12-23 | 8 | 34.0 |
| Crystal Palace | Roy Hodgson | Patrick Vieira | 2021-05-23 | 2021-08-14 | 83 | 50.0 |
| Newcastle United FC | Graeme Jones | Eddie Howe | 2021-11-06 | 2021-11-20 | 14 | 115.0 |

Figure 5.1. Games to convergence for new managers, with the number of days since their predecessor’s final match shown for context.

From Figure 5.1, it can be seen that the majority of managers converged within the first half of their debut season, suggesting that most establish a consistent tactical identity relatively quickly. However, notable exceptions include Patrick Vieira and Eddie Howe, both of whom exhibited longer convergence periods.




Figure 5.2 Patrick Vieira's context-aware and context-unaware identity-vector distances over his tenure at Crystal Palace. Convergence occurs after 50 games.

As shown in Figure 5.2, Vieira's context-adjusted identity distance begins to decline sharply from August to January of his first season, indicating the formation of a recognisable tactical identity. Although the curve narrowly misses the global convergence threshold around February–March, it is evident that his system becomes substantially more stable by this stage. This suggests that the defined threshold may be slightly conservative in certain cases.




Figure 5.3 Eddie Howe's context-aware and context-unaware identity-vector distances over his tenure at Newcastle United. Convergence occurs after 110 games.

In contrast, Figure 5.3 for Eddie Howe presents a considerably noisier trajectory. His identity distances fluctuate frequently, with full convergence not occurring until the fourth season of his tenure. This pattern could reflect external influences such as tactical experimentation, injury disruptions, or fixture congestion. We will explore general causations in later sections.

5.2 Distances

We begin to observe the largest pairwise shifts in distance between a manager's first match and the final match of their predecessor. Here, distance represents the reliability-weighted, diagonal Mahalanobis-type sum of squared changes in the context-adjusted z-scores.

Figure 5.4 presents the five largest observed changes in distance between incoming managers and their predecessors, calculated across all 58 managerial transitions, including those that did not meet the formal convergence criteria.

| Team | Prev Manager | New Manager | Prev Match Date | First Match Date | Gap Days | Distance From Last Managers Identity |
|------------------------|--------------------|-----------------------|-----------------|------------------|----------|--------------------------------------|
| Sheffield United | Paul Heckingbottom | Chris Wilder | 2023-12-02 | 2023-12-06 | 4 | 15.917431 |
| Crystal Palace | Roy Hodgson | Patrick Vieira | 2021-05-23 | 2021-08-14 | 83 | 13.400964 |
| Tottenham Hotspur | Ryan Mason | Angelos Postecoglou | 2023-05-28 | 2023-08-13 | 77 | 13.056680 |
| Brighton & Hove Albion | Roberto De Zerbi | Fabian Hürzeler | 2024-05-19 | 2024-08-17 | 90 | 10.340279 |
| Everton FC | Carlo Ancelotti | Rafael Benítez Maudes | 2021-05-23 | 2021-08-14 | 83 | 7.731229 |

Figure 5.4 Top five managerial transitions ranked by context-adjusted distance between the incoming manager's first match and the predecessor's final match.

With the exception of Chris Wilder, most managers who exhibit substantial stylistic differences from their predecessors had the benefit of a full pre-season to prepare their teams. This preparation window appears crucial in allowing new managers to embed tactical and structural changes before competitive play resumes.

| | Paul Heckingbottom (before) | Chris Wilder (after) | Diff (After - Before) | % Change |
|--|-----------------------------|----------------------|-----------------------|------------|
| Context Adjusted Carry Share | 0.405003 | 0.268752 | -0.136251 | -33.641973 |
| Context Adjusted Compactness | 0.098991 | 0.103867 | 0.004876 | 4.9257 |
| Context Adjusted Directness | 0.382133 | 0.623986 | 0.241853 | 63.290268 |
| Context Adjusted Pressing Line Height | 77.224122 | 72.032377 | -5.191745 | -6.722958 |
| Context Adjusted Progression Speed Median Ms | 4.798121 | 8.60922 | 3.811099 | 79.428989 |
| Points | 0.0 | 0.0 | 0.0 | NaN |
| Goal Difference | -5.0 | -2.0 | 3.0 | 60.0 |
| Opposition | Burnley | Liverpool | | |
| Side | away | home | | |

Figure 5.6 Context-adjusted metric comparison between Paul Heckingbottom's final match and Chris Wilder's first match for Sheffield United.

Figure 5.6 illustrates the context-adjusted differences between Chris Wilder's first match and Paul Heckingbottom's final match. A closer examination of Wilder's debut shows that this significant shift stems primarily from a deeper defensive setup and a greater emphasis on direct, vertical play, with reduced reliance on carrying compared to passing.

Further to this, Figure 5.7 illustrates that the shorter the time between a manager's appointment and their first match, the more similar their team appears to that of their predecessor. This finding is intuitive: it is unrealistic to expect a manager to implement substantial tactical or structural changes within such a limited time frame. Conversely, larger stylistic differences are typically associated with longer preparation periods, which provide managers with more opportunities to adjust training methods, refine tactical systems, and, in some cases, modify the playing squad during a transfer window.

Managerial Distance vs Gap Days




Figure 5.7 Mean and median pairwise distance (D^2) between incoming managers and their predecessors plotted against the number of days between their respective matches.

We next examine how stylistic similarity to a predecessor relates to short-term results. As shown in Figure 5.8, there appears to be an optimal balance between maintaining continuity and introducing tactical change. Managers whose teams are either *too similar* or *too dissimilar* to their predecessors tend to record weaker first-game outcomes. The strongest results occur when the new manager's setup diverges moderately from the previous style—suggesting that a measured level of change may yield the most effective early performances.

First-Game Points by D² Step Bin




Figure 5.8 Mean points achieved in a manager's first game, grouped by bins of Mahalanobis-type distance (D^2) between the new manager and their predecessor's tactical identity.

Building on the relationship between stylistic similarity and short-term results shown in Figure 5.8, Figure 5.14 illustrates how this similarity also affects the time it takes for managers to reach tactical convergence. The figure shows the average number of matches required for new managers to stabilise tactically, grouped by the stylistic distance between their first match and their predecessor's final match.

Although some variability is expected given the sample size, a clear trend emerges. Managers whose initial setups closely resembled their predecessors (distance 0-5) tended to converge most quickly, with a median of around six matches. Those introducing more drastic stylistic changes (distance > 10) typically required much longer—often more than twice as many matches on average.

The middle range (distance 5-10) appears to represent the most balanced outcome, combining moderate tactical change with relatively quick convergence. This range also corresponds with the higher short-term point returns observed in Figure 5.8, suggesting that a measured level of adaptation—rather than a complete overhaul—may deliver both quicker tactical stability and stronger early performance.

Overall, this pattern supports the idea that continuity helps early adjustment, but a moderate level of change is often needed to achieve genuine improvement.

Matches to Convergence by Distance Bin




Figure 5.14. Mean and median number of matches to convergence, grouped by the stylistic distance between a new manager's first match and their predecessor's final match. Although the sample size is limited (31), the pattern suggests that managers whose initial setups are most similar to their predecessors tend to reach convergence fastest, while those implementing larger tactical changes require more time to stabilise.

5.3 Similarities

Figure 5.9 examines how similar teams appear within their respective convergence bins. The analysis includes the 31 teams identified as having converged according to our predefined criteria. For each manager, context-adjusted identity vectors were averaged across their tenure and grouped by the number of matches required to reach convergence. Within each bin, pairwise similarities were computed using cosine similarity, and the mean similarity score was then derived.

The results show that managers who converged most rapidly tended to produce teams with higher overall similarity in their aggregated identities by the end of their tenure. Conversely, managers who required longer to converge exhibited greater dissimilarity, suggesting a higher degree of tactical flexibility, adaptation, or stylistic individuality.

Similarity within Convergence Bins




Figure 5.9 Mean cosine distance between teams within convergence bins, based on context-adjusted identity vectors aggregated over each manager's tenure.

From Table 5.10, we observe that teams converging most quickly tend to sit slightly deeper on the pitch, maintain a compact structure, and exhibit a style that favours quicker, more direct ball progression with a lower emphasis on carrying. This may indicate a preference for structured, transition-oriented play rather than prolonged possession. However, it is important to note that the sample size in this analysis is relatively small, and as such, the findings should be interpreted as informative rather than conclusive.

| Convergence Bin | Sample Size (n) | Carry Share (Adj.) | Compactness (Adj.) | Directness (Adj.) | Pressing Line Height | Progression Speed (ms) |
|-----------------|-----------------|--------------------|--------------------|-------------------|----------------------|------------------------|
| 0-10 | 17 | 0.431 | 0.109 | 0.344 | 77.889 | 5.144 |
| 11-20 | 9 | 0.429 | 0.108 | 0.346 | 78.346 | 5.250 |
| 21+ | 5 | 0.438 | 0.105 | 0.350 | 79.376 | 4.937 |

Figure 5.10 Mean context-adjusted tactical metrics across convergence bins ($n = 31$ managers).

5.4 Case Study and Possible Hypothesis of Tactical Shift

The following section provides a high-level exploration of potential factors influencing changes in tactical fingerprints over a manager's tenure. While a more in-depth causal analysis would have been valuable, time constraints limited the scope of this investigation. Nonetheless, the insights presented here highlight possible links between contextual factors and variations in managerial tactical identity.

5.4.1 Sean Dyche

From Figure 5.11, we observe that Sean Dyche converged relatively early to a stable tactical identity. This aligns with our broader findings that managers whose teams initially resemble their predecessors tend to establish a consistent identity more quickly. Such early convergence may reflect a pragmatic approach – focusing on stability, small tactical adjustments, and securing short-term results rather than major changes.

Figure 5.11 also shows periods where Dyche’s tactical identity diverged before settling again, particularly after the summer transfer window. This likely reflects a phase of tactical experimentation or adaptation following squad changes, with stability returning later in the season.




Figure 5.11 Sean Dyche’s context-aware and context-unaware identity-vector distances over his tenure at Everton. Convergence occurs after 6 games.



Figure 5.12 Sean Dyche’s 3-game rolling exponential weighted mean of points accumulated during his tenure at Everton. The smoothed trend highlights fluctuations in team performance across consecutive fixtures.

Everton — Dyche: 3 Game Rolling EMA Defensive and Midfield Tactical Shifts




Figure 5.13 Sean Dyche's 3-game rolling exponential weighted mean of defensive and midfield tactical shifts during his tenure at Everton. Shifts are derived from Hudl lineups data and represent in-game formation or structural adjustments.

Figures 5.12 and 5.13 provide further context through the 3-game rolling exponential weighted mean (EMA) of points earned and in-game tactical shifts. Tactical shifts were derived from Hudl lineups data, while the rolling EMA was used to smooth short-term variations and highlight longer-term patterns.

When comparing Figures 5.11 and 5.12, we can see that sharp changes in tactical identity often coincide with changes in points gained. This suggests that periods of tactical adjustment sometimes lead to positive short-term outcomes, while returns to more consistent tactical setups may be linked to steadier but less variable results. While convergence correlates with improved stability and results in some cases, the relationship we've provided is descriptive rather than causal.

In Figure 5.13, the link between tactical shifts and identity variation appears weaker but still notable. For instance, increases in midfield adjustments before October 2023 are followed by rises in tactical identity distance shortly afterwards. Similarly, a rise in defensive tactical shifts in late September 2024 coincides with a decline in identity variation. This may indicate a period of tactical fine-tuning — with Dyche adapting his system to player availability or form before settling on a more stable setup. Again, this relationship is purely descriptive rather than causal.

5.5. Summary of Results

Across the 58 managerial changes analysed, 31 showed clear tactical convergence, typically within the first 12 to 15 matches. Managers appointed with little preparation time tended to field teams tactically similar to their predecessors, while those afforded longer lead-in periods implemented greater stylistic changes before stabilising. Together, these findings suggest that tactical identity tends to evolve progressively rather than shifting immediately after appointment.

6. Discussion, Limitations and Future Work

This study examined the immediate tactical influence of newly appointed Premier League managers. Using a convergence framework based on context-adjusted performance metrics, we assessed how quickly, and how much, a manager's tactical identity settles after appointment. We found 31 convergence events among 58 eligible tenures, so about half of new managers reached a stable tactical profile within the study window.

The convergence rule depends on a chosen threshold, so it is partly subjective. Even so, the framework behaved in a sensible way: managers with little preparation time tended to set up teams that looked similar to their predecessors, while longer lead-in periods were linked to bigger stylistic changes. This fits with work on the "manager bounce", which suggests short-term improvements are often psychological or contextual rather than structural. We also saw that very large tactical changes right away were linked to weaker first results, which supports a more measured, step-by-step approach.

Eddie Howe is a clear outlier: his convergence came late (about 110 games). This shows the current framework does not suit every case, especially when a manager's approach develops slowly or in stages. One aim of this work was to make genuine tactical change easier to spot so that future studies can look at causes with more confidence; Howe's case suggests we may need extras such as phase-specific identities (defence, attack, structure) or a threshold that can adapt over time.

This study is limited by its sample size (58 appointments) and by using only five micro-metrics to define identity. Future work could split identity into defensive, attacking and structural fingerprints, and allow the weight of each metric to be adjusted to match different analysis needs. It should also look at what drives within-season changes—injuries, line-up rotation, transfers, and fixture congestion—using correlations and simple time-series models to check for delayed effects. Finally, context adjustment reduces noise but does not prove causation, so the patterns we report should be read as descriptive, not causal.

Overall, the framework offers a practical way to measure a manager's tactical identity and how it changes over time. With further tuning and testing, it could be useful for both researchers and practitioners in performance analysis.

References

- [1] Besters, L.M., van Ours, J.C. & van Tuijl, M.A. Effectiveness of In-Season Manager Changes in English Premier League Football. *De Economist* 164, 335–356 (2016).
<https://doi.org/10.1007/s10645-016-9277-0>
- [2] Ramani, H., 2023. Decoding football's new manager bounce: A journey into the unknown. *Sideline Strategists*. [online] 15 September. Available at:
<https://medium.com/sideline-strategists/decoding-footballs-new-manager-bounce-9f861aa1ae3b>
- [3] Guerrero-Calderón, B., Owen, A., Morcillo, J.A. and Castillo-Rodríguez, A., 2021. How does the mid-season coach change affect physical performance on top soccer players? *Physiology & Behavior*, 232, 113328. doi:10.1016/j.physbeh.2021.113328.
- [4] Fernandez-Navarro J, Fradua L, Zubillaga A, Ford PR, McRobert AP. Attacking and defensive styles of play in soccer: analysis of Spanish and English elite teams. *J Sports Sci*. 2016 Dec;34(24):2195–2204. doi: 10.1080/02640414.2016.1169309. Epub 2016 Apr 7. PMID: 27052355.
- [5] Sharpe, C. (2022, September 5). *Modelling team playing style*. StatsBomb Blog Archive.
<https://blogarchive.statsbomb.com/articles/soccer/modelling-team-playing-style/>
- [6] Monte, J. (2024, June 13). *Using data in the search for a new manager*. StatsBomb Blog Archive.
<https://blogarchive.statsbomb.com/articles/soccer/using-data-in-the-search-for-a-new-manager/>
- [7] Franks, A., Miller, A., Bornn, L., & Goldsberry, K. (2015). *Characterizing the spatial structure of defensive skill in professional basketball*. *The Annals of Applied Statistics*, 9(1), 94–121.
<https://doi.org/10.1214/14-AOAS799>
- [8] Pearce, J. (2024, June 10). *Liverpool 2.0: Who does what in FSG's new-look structure?* *The New York Times*.
<https://www.nytimes.com/athletic/5545133/2024/06/10/liverpool-fsg-explained-hierarchy/>
- [9] International Soccer Science & Performance Federation. (2024, April 30). *Coaching development in football: Requirements of the elite coach*.
<https://www.isspf.com/articles/coaching-development-in-football/>
- [10] Transfermarkt GmbH. (n.d.). Premier League – Manager changes per season. Retrieved September 16, 2025, from
<https://www.transfermarkt.co.uk/premier-league/trainerwechselprosaison/wettbewerb/GB1>
- [11] Skripnikov, A., Cemek, A., & Gillman, D. (2025). Leveraging minute-by-minute soccer match event data to adjust team's offensive production for game context (Preprint). arXiv.
<https://arxiv.org/abs/2508.04008>
- [12] Wikipedia contributors. (n.d.). *Elo rating system*. In *Wikipedia, The Free Encyclopedia*. Retrieved October 26, 2025, from https://en.wikipedia.org/wiki/Elo_rating_system
- [13] *Football Club Elo Ratings*. Retrieved September 16, 2025, from <https://www.clubelo.com/>

- [14] Hayward, B. (2024, May 19). *What are the pitch sizes for every Premier League stadium?* FourFourTwo. Retrieved September 16, 2025, from <https://www.fourfourtwo.com/features/premier-league-pitch-sizes>
- [15] Huck Nets. (2023, May 17). *The complete 2023 list of Premier League pitch sizes by club (22/23 season)*. Huck Nets. Retrieved September 16 2025, from <https://www.huck-net.co.uk/news/2023-05/premier-league-pitch-sizes/>
- [16] Morgan, W. (2018, September 5). *How StatsBomb data helps measure counter-pressing*. StatsBomb Blog Archive. Retrieved October 26, 2025, from <https://blogarchive.statsbomb.com/articles/soccer/how-statsbomb-data-helps-measure-counter-pressing/>
- [17] Rico-González, M., Pino-Ortega, J., Nakamura, F. Y., Moura, F. A., & Los Arcos, A. (2021). Identification, computational examination, critical assessment and future considerations of spatial tactical variables to assess the use of space in team sports by positional data: A systematic review. *Journal of Human Kinetics*, 77(1), 205-221. <https://doi.org/10.2478/hukin-2021-0021>
- [18] Guan, T., Cao, J., & Swartz, T. B. (2023). "Parking the bus?" : Defensive compactness and soccer outcomes. *Journal of Quantitative Analysis in Sports*, 19(4), 263-272. <https://doi.org/10.1515/jqas-2021-0059>
- [19] Breaking the Lines. (2023, March 12). *Tactical analysis: Compactness and coverage in possession*. <https://breakingthelines.com/tactical-analysis/tactical-analysis-compactness-and-coverage-in-possession/>
- [20] Wicklin, R. (2022, November 2). *The area and perimeter of a convex hull*. The DO Loop (SAS Blogs). Retrieved September 16 2025, from <https://blogs.sas.com/content/iml/2022/11/02/area-perimeter-convex-hull.html>
- [21] Wikipedia contributors. (n.d.). *Standard score*. In *Wikipedia, The Free Encyclopedia*. Retrieved October 26, 2025, from https://en.wikipedia.org/wiki/Standard_score
- [22] Wikipedia contributors. (n.d.). *Generalized linear model*. In *Wikipedia, The Free Encyclopedia*. Retrieved October 26, 2025, from https://en.wikipedia.org/wiki/Generalized_linear_model
- [23] Wikipedia contributors. (n.d.). *Mahalanobis distance*. In *Wikipedia, The Free Encyclopedia*. Retrieved October 27, 2025, from https://en.wikipedia.org/wiki/Mahalanobis_distance
- [24] Gleeson, M. (2024) *Modelling team play style using tracking data to evaluate the effectiveness of variance in tactical behaviour*. Paper presented at the StatsBomb Research Conference, University College Dublin. Available at: <https://blogarchive.statsbomb.com/uploads/2024/09/Michael-Gleeson-Statsbomb-Research-Conference-2024-Draft-2.docx-1.pdf>
- [25] Tichy, M. (2024) 'Demystifying the Probability Integral Transform', *Medium*, 22 August. Retrieved October 26, 2025, from: <https://medium.com/@maltetichy/demystifying-the-probability-integral-transform-77b7de3a3af9>
- [26] GeeksforGeeks. (n.d.). *Kolmogorov-Smirnov Test (KS Test)*. In *GeeksforGeeks*. Retrieved October 29, 2025, from <https://www.geeksforgeeks.org/machine-learning/kolmogorov-smirnov-test-ks-test/>

[27] Wikipedia contributors. (n.d.). Heteroskedasticity-consistent standard errors. In *Wikipedia, The Free Encyclopedia*. Retrieved October 29, 2025, from https://en.wikipedia.org/wiki/Heteroskedasticity-consistent_standard_errors

[28] Wikipedia contributors. (n.d.). *Median absolute deviation*. In *Wikipedia, The Free Encyclopedia*. Retrieved October 29, 2025, from https://en.wikipedia.org/wiki/Median_absolute_deviation

Appendix

6.1 Model Evaluation

6.1.1 Pressing Line Height

| Generalized Linear Model Regression Results | | | | | | |
|---|--------------------------|---------------------|-----------|-------|--------|--------|
| Dep. Variable: | Q('PLH_line_mean_units') | No. Observations: | 3798 | | | |
| Model: | GLM | Df Residuals: | 3786 | | | |
| Model Family: | Gamma | Df Model: | 11 | | | |
| Link Function: | log | Scale: | 0.0017868 | | | |
| Method: | IRLS | Log-Likelihood: | -9936.1 | | | |
| Date: | Sat, 25 Oct 2025 | Deviance: | 6.7531 | | | |
| Time: | 20:48:31 | Pearson chi2: | 6.76 | | | |
| No. Iterations: | 6 | Pseudo R-squ. (CS): | 0.1041 | | | |
| Covariance Type: | HC3 | | | | | |
| | coef | std err | z | P> z | [0.025 | 0.975] |
| Intercept | 4.3487 | 0.002 | 2121.265 | 0.000 | 4.345 | 4.353 |
| C(season, Treatment(reference='2020/2021')) [T.2021/2022] | 0.0020 | 0.002 | 0.883 | 0.377 | -0.002 | 0.006 |
| C(season, Treatment(reference='2020/2021')) [T.2022/2023] | 0.0050 | 0.002 | 2.263 | 0.024 | 0.001 | 0.009 |
| C(season, Treatment(reference='2020/2021')) [T.2023/2024] | 0.0148 | 0.002 | 6.575 | 0.000 | 0.010 | 0.019 |
| C(season, Treatment(reference='2020/2021')) [T.2024/2025] | 0.0225 | 0.002 | 10.384 | 0.000 | 0.018 | 0.027 |
| elo_diff_z | 0.0104 | 0.001 | 13.731 | 0.000 | 0.009 | 0.012 |
| I(elo_diff_z ** 2) | 0.0015 | 0.001 | 2.724 | 0.006 | 0.000 | 0.002 |
| is_home | 0.0052 | 0.001 | 3.779 | 0.000 | 0.003 | 0.008 |
| man_net | 0.0203 | 0.005 | 3.838 | 0.000 | 0.010 | 0.031 |
| man_abs | 0.0153 | 0.005 | 2.828 | 0.005 | 0.005 | 0.026 |
| winning_share | -0.0021 | 0.003 | -0.805 | 0.421 | -0.007 | 0.003 |
| losing_share | 0.0090 | 0.002 | 3.624 | 0.000 | 0.004 | 0.014 |

Figure 6.1. Coefficient summary – Gamma(GLM), PLH. Key effects align with expectations: strength (elo_diff_z), home advantage, and manpower terms are positively associated with higher PLH; season controls absorb a small upward trend. (HC3 SEs; log link.)





Figure 6.2. Residuals vs fitted. Residuals are centred with near-constant spread; only a slight uptick at high fitted values, suggesting at most mild nonlinearity.




Figure 6.3. Pressing Line Height Gamma GLM PIT Histogram. PIT is close to uniform, indicating good predictive calibration and interval coverage.

6.1.2 Progression Speed M/S

| Generalized Linear Model Regression Results | | | | | | |
|---|-----------------------|---------------------|----------|-------|--------|--------|
| Dep. Variable: | Q('median_speed_mps') | No. Observations: | 3800 | | | |
| Model: | GLM | Df Residuals: | 3788 | | | |
| Model Family: | Gamma | Df Model: | 11 | | | |
| Link Function: | log | Scale: | 0.019043 | | | |
| Method: | IRLS | Log-Likelihood: | -3962.3 | | | |
| Date: | Mon, 20 Oct 2025 | Deviance: | 71.143 | | | |
| Time: | 22:49:15 | Pearson chi2: | 72.1 | | | |
| No. Iterations: | 8 | Pseudo R-squ. (CS): | 0.3207 | | | |
| Covariance Type: | HC3 | | | | | |
| | coef | std err | z | P> z | [0.025 | 0.975] |
| Intercept | 1.6462 | 0.007 | 240.288 | 0.000 | 1.633 | 1.660 |
| C(season, Treatment(reference='2020/2021')) [T.2021/2022] | 0.0241 | 0.007 | 3.456 | 0.001 | 0.010 | 0.038 |
| C(season, Treatment(reference='2020/2021')) [T.2022/2023] | 0.0139 | 0.007 | 2.009 | 0.044 | 0.000 | 0.027 |
| C(season, Treatment(reference='2020/2021')) [T.2023/2024] | -0.0081 | 0.007 | -1.091 | 0.275 | -0.023 | 0.006 |
| C(season, Treatment(reference='2020/2021')) [T.2024/2025] | -0.0416 | 0.007 | -5.916 | 0.000 | -0.055 | -0.028 |
| elo_diff_z | -0.0812 | 0.002 | -33.927 | 0.000 | -0.086 | -0.077 |
| I(elo_diff_z ** 2) | -0.0073 | 0.002 | -4.146 | 0.000 | -0.011 | -0.004 |
| is_home | -0.0311 | 0.005 | -6.903 | 0.000 | -0.040 | -0.022 |
| man_net | -0.1556 | 0.016 | -9.783 | 0.000 | -0.187 | -0.124 |
| man_abs | 0.0018 | 0.016 | 0.107 | 0.915 | -0.030 | 0.034 |
| winning_share | 0.0752 | 0.008 | 9.350 | 0.000 | 0.059 | 0.091 |
| losing_share | -0.0775 | 0.008 | -9.547 | 0.000 | -0.093 | -0.062 |

Figure 6.4. Coefficient summary – Gamma(GLM), Progression Speed. Progression speed decreases with higher opponent strength (elo_diff_z), indicating stronger opponents slow build-up play. Home advantage and manpower advantages slightly reduce speed, while winning_share and losing_share have opposite, offsetting effects. Season effects show modest variation.





Figure 6.5. Residuals vs fitted. Deviance and Pearson residuals are well centred with limited patterning. A faint curvature at low fitted values suggests mild model nonlinearity but overall residual dispersion is stable, supporting model adequacy.




Figure 6.6. Progression Speed M/S Gamma GLM PIT Histogram. The PIT distribution is approximately uniform, indicating the Gamma model provides a well-calibrated fit with no major deviations in predictive reliability.

6.1.3 Compactness

| BetaModel Results | | | | | | |
|---|-------------------------|-----------------|------------|-------|--------|--------|
| Dep. Variable: | Q('compact_raw_median') | Log-Likelihood: | 10165. | | | |
| Model: | BetaModel | AIC: | -2.030e+04 | | | |
| Method: | Maximum Likelihood | BIC: | -2.022e+04 | | | |
| Date: | Sun, 26 Oct 2025 | | | | | |
| Time: | 21:51:54 | | | | | |
| No. Observations: | 3798 | | | | | |
| Df Residuals: | 3785 | | | | | |
| Df Model: | 11 | | | | | |
| | coef | std err | z | P> z | [0.025 | 0.975] |
| Intercept | -2.1420 | 0.008 | -259.416 | 0.000 | -2.158 | -2.126 |
| C(season, Treatment(reference='2020/2021')) [T.2021/2022] | -0.0013 | 0.009 | -0.153 | 0.879 | -0.018 | 0.016 |
| C(season, Treatment(reference='2020/2021')) [T.2022/2023] | 0.0198 | 0.009 | 2.226 | 0.026 | 0.002 | 0.037 |
| C(season, Treatment(reference='2020/2021')) [T.2023/2024] | -0.0156 | 0.009 | -1.700 | 0.089 | -0.034 | 0.002 |
| C(season, Treatment(reference='2020/2021')) [T.2024/2025] | -0.0203 | 0.009 | -2.251 | 0.024 | -0.038 | -0.003 |
| elo_diff_z | 0.0461 | 0.003 | 15.185 | 0.000 | 0.040 | 0.052 |
| I(elo_diff_z ** 2) | -0.0127 | 0.002 | -5.840 | 0.000 | -0.017 | -0.008 |
| is_home | 0.0288 | 0.006 | 4.956 | 0.000 | 0.017 | 0.040 |
| man_net | 0.1781 | 0.024 | 7.540 | 0.000 | 0.132 | 0.224 |
| man_abs | -0.1550 | 0.024 | -6.430 | 0.000 | -0.202 | -0.108 |
| winning_share | -0.0370 | 0.010 | -3.568 | 0.000 | -0.057 | -0.017 |
| losing_share | 0.0626 | 0.011 | 5.804 | 0.000 | 0.041 | 0.084 |
| precision | 5.8092 | 0.026 | 222.081 | 0.000 | 5.758 | 5.860 |

Figure 6.7. Coefficient summary – Beta(GLM), Compactness. Compactness increases modestly with opponent strength and home advantage, while greater manpower and losing states correspond to higher compactness. Seasonal effects are minor, and the model precision parameter indicates low residual variance.




Figure 6.8. Compactness Residual vs Fitted & PIT Histograms - Beta(mean). Residuals are evenly distributed around zero with no systematic trend, indicating a well-specified mean-variance structure. The LOWESS line

remains nearly flat, suggesting stable fit quality across the range of predicted values. The PIT distribution is close to uniform, supporting that the Beta model is well-calibrated for compactness.

6.1.4 Share of Carries

| BetaModel Results | | | | | | |
|---|--------------------|-----------------|------------|-------|--------|--------|
| Dep. Variable: | Q('share_carry') | Log-Likelihood: | 6915.3 | | | |
| Model: | BetaModel | AIC: | -1.380e+04 | | | |
| Method: | Maximum Likelihood | BIC: | -1.372e+04 | | | |
| Date: | Sat, 18 Oct 2025 | | | | | |
| Time: | 19:46:00 | | | | | |
| No. Observations: | 3800 | | | | | |
| Df Residuals: | 3787 | | | | | |
| Df Model: | 11 | | | | | |
| | coef | std err | z | P> z | [0.025 | 0.975] |
| Intercept | -0.2772 | 0.008 | -35.269 | 0.000 | -0.293 | -0.262 |
| C(season, Treatment(reference='2020/2021')) [T.2021/2022] | -0.0107 | 0.008 | -1.281 | 0.200 | -0.027 | 0.006 |
| C(season, Treatment(reference='2020/2021')) [T.2022/2023] | -0.0289 | 0.008 | -3.489 | 0.000 | -0.045 | -0.013 |
| C(season, Treatment(reference='2020/2021')) [T.2023/2024] | 0.0157 | 0.008 | 1.871 | 0.061 | -0.001 | 0.032 |
| C(season, Treatment(reference='2020/2021')) [T.2024/2025] | 0.0645 | 0.008 | 7.881 | 0.000 | 0.048 | 0.081 |
| elo_diff_z | 0.0621 | 0.003 | 22.379 | 0.000 | 0.057 | 0.067 |
| I(elo_diff_z ** 2) | 0.0072 | 0.002 | 3.500 | 0.000 | 0.003 | 0.011 |
| is_home | 0.0310 | 0.005 | 5.965 | 0.000 | 0.021 | 0.041 |
| man_net | 0.1284 | 0.019 | 6.810 | 0.000 | 0.091 | 0.165 |
| man_abs | 0.0161 | 0.020 | 0.827 | 0.408 | -0.022 | 0.054 |
| winning_share | -0.0629 | 0.009 | -6.656 | 0.000 | -0.081 | -0.044 |
| losing_share | 0.0698 | 0.009 | 7.396 | 0.000 | 0.051 | 0.088 |
| precision | 5.0685 | 0.024 | 212.864 | 0.000 | 5.022 | 5.115 |

Figure 6.9. Coefficient summary – Beta(GLM), Share of Carries. The share of carries increases with opponent strength (elo_diff_z), home advantage, and manpower advantage, while decreasing when teams are winning. Season effects are generally small, with a notable increase in 2024/2025. Overall, the model captures expected behavioural adjustments in possession style.

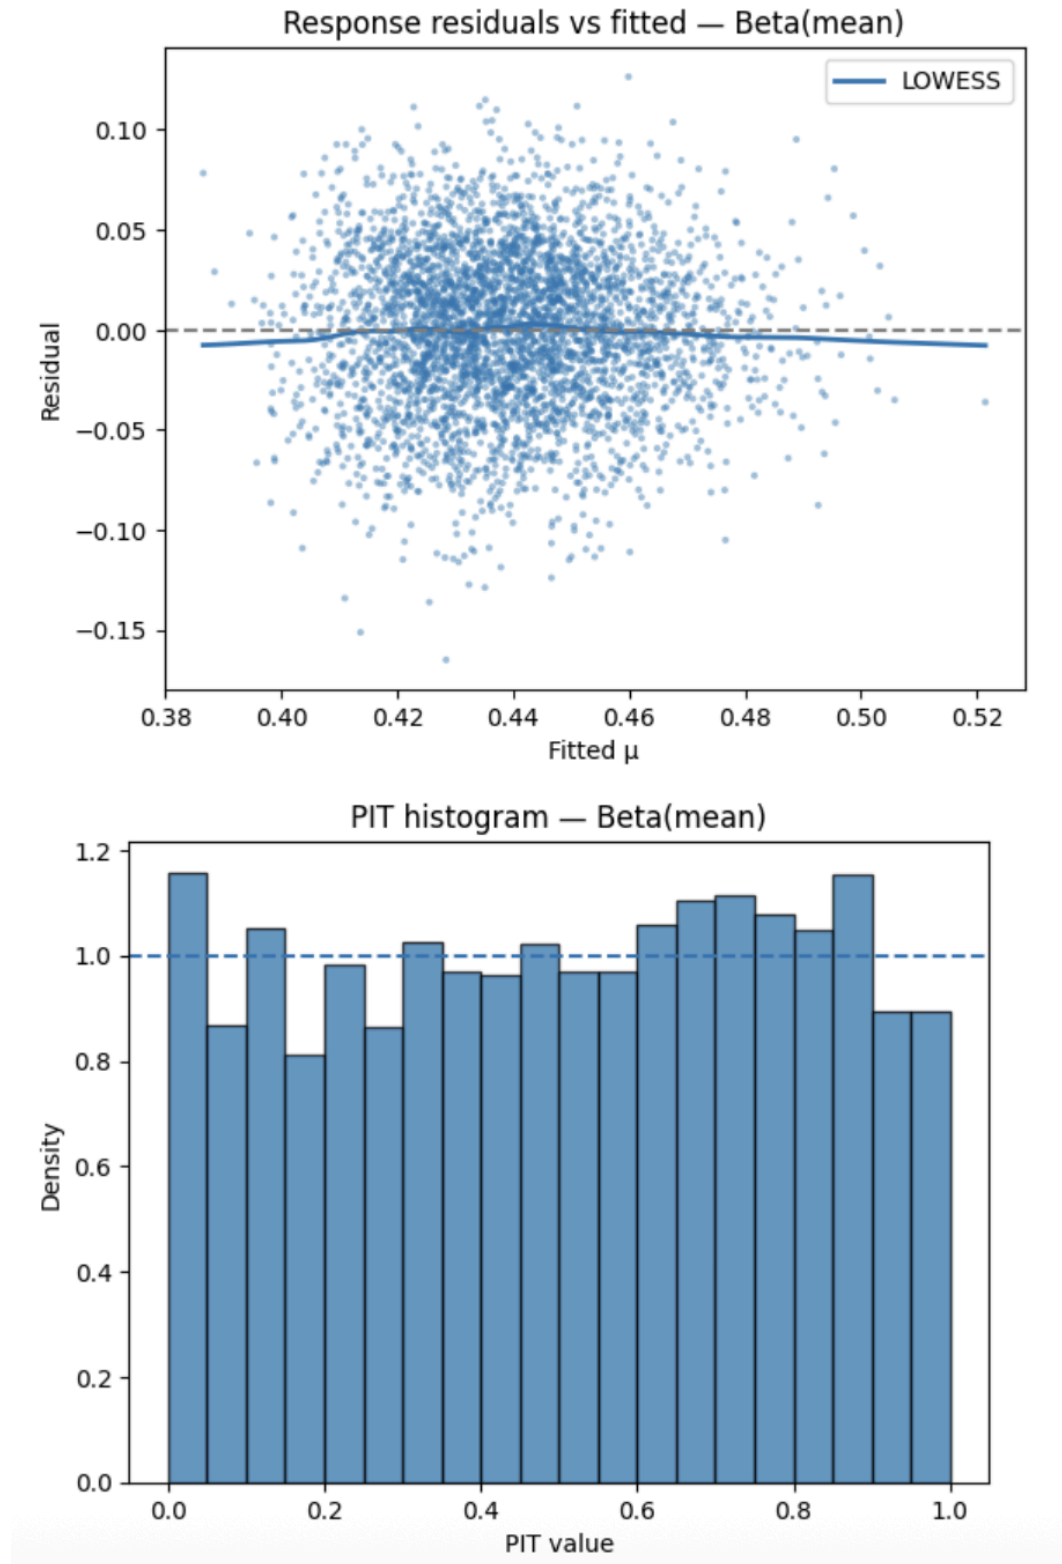


Figure 6.10. Share of Carries Residual vs Fitted & PIT Histograms - Beta(mean). Residuals are symmetrically scattered with no clear trend, and the LOWESS line remains flat across the fitted range, suggesting the

mean-variance structure of the Beta model is appropriate. The PIT values are approximately uniform, indicating that the model's predictive distribution is well calibrated.

6.1.5 Directness

| BetaModel Results | | | | | | |
|---|-----------------------|-----------------|---------|-------|--------|--------|
| Dep. Variable: | Q('global_ratio_raw') | Log-Likelihood: | 4673.6 | | | |
| Model: | BetaModel | AIC: | -9321. | | | |
| Method: | Maximum Likelihood | BIC: | -9240. | | | |
| Date: | Sat, 25 Oct 2025 | | | | | |
| Time: | 21:27:10 | | | | | |
| No. Observations: | 3800 | | | | | |
| Df Residuals: | 3787 | | | | | |
| Df Model: | 11 | | | | | |
| | coef | std err | z | P> z | [0.025 | 0.975] |
| Intercept | -0.5832 | 0.017 | -35.207 | 0.000 | -0.616 | -0.551 |
| C(season, Treatment(reference='2020/2021')) [T.2021/2022] | 0.0323 | 0.017 | 1.879 | 0.060 | -0.001 | 0.066 |
| C(season, Treatment(reference='2020/2021')) [T.2022/2023] | -0.0020 | 0.017 | -0.113 | 0.910 | -0.036 | 0.032 |
| C(season, Treatment(reference='2020/2021')) [T.2023/2024] | -0.0509 | 0.018 | -2.879 | 0.004 | -0.085 | -0.016 |
| C(season, Treatment(reference='2020/2021')) [T.2024/2025] | -0.0794 | 0.017 | -4.751 | 0.000 | -0.112 | -0.047 |
| elo_diff_z | -0.2185 | 0.006 | -37.581 | 0.000 | -0.230 | -0.207 |
| I(elo_diff_z ** 2) | -0.0302 | 0.004 | -6.876 | 0.000 | -0.039 | -0.022 |
| is_home | -0.0611 | 0.011 | -5.784 | 0.000 | -0.082 | -0.040 |
| man_net | -0.3976 | 0.037 | -10.677 | 0.000 | -0.471 | -0.325 |
| man_abs | 0.0372 | 0.038 | 0.974 | 0.330 | -0.038 | 0.112 |
| winning_share | 0.0604 | 0.020 | 3.017 | 0.003 | 0.021 | 0.100 |
| losing_share | -0.2343 | 0.019 | -12.400 | 0.000 | -0.271 | -0.197 |
| precision | 3.7447 | 0.024 | 154.233 | 0.000 | 3.697 | 3.792 |

Figure 6.11. Coefficient summary – Beta(GLM), Directness. Directness decreases significantly with higher opponent strength (elo_diff_z), suggesting stronger opposition forces teams to play less directly. Home advantage and manpower advantages are also associated with slightly lower directness, while losing states increase it—indicating more urgent, vertical play when behind. Seasonal effects show modest but consistent variability across years.




Figure 6.12. Directness Residual vs Fitted & PIT Histograms - Beta(mean). Residuals appear symmetrically distributed around zero with minimal curvature in the LOWESS line. The PIT values are approximately uniform,

confirming that the model provides a good probabilistic fit. Minor deviations from perfect uniformity are within expectation, supporting that the Beta distribution appropriately represents directness across the observed range.

6.2 Metric Distributions

6.2.1 Pressing Line Height




Figure 6.13. Distribution of Pressing Line Height (PLH_line_mean_units) across matches. The distribution is approximately normal with a slight right skew.

6.2.2 Progression Speed M/S




Figure 6.14. Distribution of Progression Speed (*median_speed_mps*) across matches. The distribution is slightly right-skewed.

6.2.3 Compactness




Figure 6.15. Distribution of Compactness (*compact_raw_median*) across matches. The distribution appears approximately normal, with only minor skewness.

6.2.4 Share of Carries




Figure 6.16. Distribution of Share Carry (share_carry) across matches. The distribution appears approximately normal with a slight right skew.

6.2.5 Directness




Figure 6.17. Distribution of Directness (*global_ratio_raw*) across matches. The distribution is moderately right-skewed.