

This is a preprint of the following published article:

Florez, M., Guindani, M. Vannucci, M. (2024). Bayesian bivariate Conway–Maxwell–Poisson regression model for correlated count data in sports. *Journal of Quantitative Analysis in Sports*. <https://doi.org/10.1515/jqas-2024-0072>

Mauro Florez*, Michele Guindani, and Marina Vannucci

Bayesian Bivariate Conway–Maxwell–Poisson Regression Model for Correlated Count Data in Sports

<https://doi.org/10.1515/jqas-2024-0072>

Received Month DD, YYYY; revised Month DD, YYYY; accepted Month DD, YYYY

Abstract: Count data play a crucial role in sports analytics, providing valuable insights into various aspects of the game. Models that accurately capture the characteristics of count data are essential for making reliable inferences. In this paper, we propose the use of the Conway–Maxwell–Poisson (CMP) model for analyzing count data in sports. The CMP model offers flexibility in modeling data with different levels of dispersion. Here we consider a bivariate CMP model that models the potential correlation between home and away scores by incorporating a random effect specification. We illustrate the advantages of the CMP model through simulations. We then analyze data from baseball and soccer games before, during, and after the COVID-19 pandemic. The performance of our proposed CMP model matches or outperforms standard Poisson and Negative Binomial models, providing a good fit and an accurate estimation of the observed effects in count data with any level of dispersion. The results highlight the robustness and flexibility of the CMP model in analyzing count data in sports, making it a suitable default choice for modeling a diverse range of count data types in sports, where the data dispersion may vary.

Keywords: Conway–Maxwell–Poisson distribution, Intractable likelihood, Random effects, Markov chain Monte Carlo, Home Advantage, COVID-19.

1 Introduction

Count data play a ubiquitous role in sports analytics, providing valuable insights across various aspects of the game. They are used to record the frequency of events or occurrences, including the number of goals scored in a soccer game, the number of strikes or balls thrown in baseball, the number of points scored in basketball, or the number of strokes taken in golf. This type of data is essential for calculating key statistics such as batting averages, shooting percentages, or save percentages, as well as to analyze and compare players' performance, team performance, and overall game strategy. For example, count data can be used to evaluate a player's ability in scoring goals or making assists, or to assess a team's defensive tactics in preventing their opponents from scoring. Count data can also be used to develop predictive models that help coaches and managers make informed decisions about game strategy and player selection. For instance, count data may be used to determine which players are most likely to score in a given situation or which plays are most likely to lead to a successful outcome.

When analyzing such types of data, it becomes crucial to utilize models that accurately capture the distinctive characteristics of the data in order to make accurate and reliable inferences. Conventional models

*Corresponding author: Mauro Florez, Marina Vannucci, Rice University, Statistics Department, Houston, TX, e-mail: mf53@rice.edu, marina@rice.edu

Michele Guindani, UCLA, Department of Biostatistics, Los Angeles, CA, e-mail: mguindani@g.ucla.edu

used to describe count data include the Poisson and the Negative Binomial distributions. For example, Maher (1982), Dixon and Coles (1997), Lee (1997), Karlis and Ntzoufras (2000) were among the first to model the number of goals scored by opposing teams in soccer using bivariate Poisson likelihoods. Since the two teams interact during the game, the game outcomes might be correlated. Hence, models for correlated count data have also been proposed (Karlis and Ntzoufras, 2003; McHale and Scarf, 2011), as well models using a Poisson difference distribution to model the differences in goals Karlis and Ntzoufras (2009) and hierarchical Bayesian models (Baio and Blangiardo, 2010). The Negative Binomial distribution was originally proposed by Reep et al. (1971) to account for the overdispersion commonly observed in baseball scores, and it has since found applications in various other sports. However, whether to employ the Poisson or Negative Binomial distribution depends on the specific characteristics of the dataset being analyzed. For example, Karlis and Ntzoufras (2000) found no significant difference between the Poisson and the Negative binomial distributions when modeling soccer data since there is relatively little over-dispersion. In some cases, the data may even exhibit some degree of under-dispersion, as demonstrated in the analysis of data from the UEFA EURO 2020 by Fedrizzi et al. (2022) and the English Premier League seasons 2015-16 to 2017-18 by Backlund and Johdet (2018).

In this paper, we propose the use of a Conway-Maxwell-Poisson (CMP) model for the analysis of count data in sports, focusing in particular on the analysis of outcomes in baseball and soccer games. The CMP distribution extends the traditional Poisson distribution by introducing an additional *dispersion* parameter – in addition to the mean – that controls the degree of over- or under-dispersion. Some advantages of the CMP distribution for modeling count data, over other distributions, such as the Negative Binomial distribution, are that it is easier to interpret and more flexible in modeling data. Here, we illustrate these advantages in sports analytics through a series of simulations and an analysis of the home advantage (HA) in baseball and soccer games, pre-, during, and post-COVID pandemic. Karlis and Ntzoufras (2003) utilize a Poisson model to assess the HA effect, assuming equidispersion. In order to relax this assumption, Boshnakov et al. (2017) propose the use of a Weibull distribution and Higgs and Stavness (2021) employ the negative binomial distribution in the context of North American sports. However, when there is minimal or no overdispersion present, the goodness of fit is either slightly inferior or comparable to that of the Poisson distribution. In our modeling framework we introduce a bivariate CMP specification using a set of random effects to model within-game correlation. The approach is robust and capable of accurately fitting the data, irrespective of whether it exhibits overdispersion, equidispersion, or underdispersion. It is worth noting that Piancastelli et al. (2023) also employ a multivariate Conway-Maxwell-Poisson distribution; however, their parametrization may not lend itself to a straightforward interpretation of the parameters. In addition, in our model formulation we also investigate whether the HA changed due to the COVID-19 pandemic, when crowds were prohibited or limited. Studies have shown that the HA effect was significantly reduced in soccer during this period (McCarrick et al., 2021; Tilp and Thaller, 2020; Benz and Lopez, 2021; Piancastelli et al., 2023). However, a study by Higgs and Stavness (2021) found that there was little to no change in the HA in baseball during the COVID-19 season. Our findings also indicate a relatively smaller impact of the COVID-19 pandemic on the HA in baseball compared to soccer.

The paper is structured as follows: in Section 2.1, we briefly describe the Conway-Maxwell-Poisson distribution, our model, and the parametrization used in the paper. In Section 2.2, we describe how we obtain posterior inference, including the MCMC algorithm. In Section 3, we present a simulation study where we evaluate the performance of the model and the estimation of the HA effect. In Sections 4 and 5, we present our data analysis on Soccer English Premier-League and Major League of Baseball games, respectively, in seasons 2019-2022. In Section 6, we provide some concluding remarks.

2 Methods

2.1 A Bayesian Conway-Maxwell-Poisson regression model for sports data

We propose to model the number of points or the goals scored by the home and away teams in each game as Conway-Maxwell-Poisson (CMP or COM-Poisson) random variables. The CMP distribution, first introduced by Conway and Maxwell (1962), is a two-parameter generalization of the Poisson distribution that allows modeling count data with over-, under-, or equi-dispersion. In this paper, we consider the reparametrization proposed by Guikema and Goffelt (2008) where the probability mass function (pmf) is expressed as

$$P(Y = y | \mu, \nu) = \left(\frac{\mu^y}{y!} \right)^\nu \frac{1}{Z(\mu, \nu)}, \quad \mu > 0, \nu \geq 0, \quad (1)$$

with $q(y | \mu, \nu) = \left(\frac{\mu^y}{y!} \right)^\nu$ indicating the unnormalized component of the pmf and $Z(\mu, \nu) = \sum_{y=0}^{\infty} q(y | \mu, \nu)$ the intractable normalising constant. This parametrization relies on a centering parameter μ , which facilitates straightforward interpretation when fitting a generalized linear model (GLM) and a shape parameter ν , which enables the modeling of various dispersion scenarios across measurements. Specifically, $\nu < 1$ accounts for overdispersion, $\nu > 1$ accounts for underdispersion, and $\nu = 1$ corresponds to equidispersion, representing the Poisson distribution. The mode of the CMP distribution is given by the flooring function, $\lfloor \mu \rfloor$, which rounds down the value of μ to the nearest integer. However, the moments of the distribution cannot be directly computed due to the analytically intractable normalizing constant Z . Instead, they can be approximated using asymptotic results. The expected value of the distribution can be approximated as $E(Y) \approx \mu + \frac{1}{2\nu} - \frac{1}{2}$, while the variance can be approximated as $Var(Y) \approx \frac{\mu}{\nu}$ (Shmueli et al., 2005). These approximations are quite accurate, except for small μ or ν in $E(Y)$, and for $\nu > 1$ in the $Var(Y)$ (Benson and Friel, 2021).

Let $\mathbf{y}_i = (y_{i1}, y_{i2})$ denote the observed number of goals or points in game i , $i = 1, \dots, n$, for the home team ($j = 1$) and the away team ($j = 2$). We propose a CMP generalized linear model framework to predict y_{ij} , where, similarly as in Karlis and Ntzoufras (2003), we include a bivariate game-specific random effect specification to adjust for a team's relative strength offense, the opponent's relative strength defense, as well as to capture a COVID-19 effect and the correlation between the home and away scores. More specifically, we model the μ and ν parameters of the CMP distribution as a function of the available covariates, say \mathbf{x}_{ij} , and the latent correlated random effects:

$$y_{ij} | \mu_{ij}, \nu_{ij} \sim CMP(\mu_{ij}, \nu_{ij}), \quad i = 1, \dots, n, j = 1, 2,$$

where the centering vector $\boldsymbol{\mu} = (\mu_{i1}, \mu_{i2})^T$ represents the expected number of goals or points scored by the home team and the away team playing in game i , respectively, and it is specified as

$$\begin{aligned} \log(\mu_{i1}) &= \beta_H + \beta_{H_i}^\omega + \beta_{A_i}^\delta + c_H + c'_H + b_{i1}, \\ \log(\mu_{i2}) &= \beta_A + \beta_{A_i}^\omega + \beta_{H_i}^\delta + c_A + c'_A + b_{i2}. \end{aligned}$$

Here, the random effects $\beta_{H_i}^\omega$ and $\beta_{H_i}^\delta$ capture the offensive and defensive strength, respectively, of the home team playing in game i . Similarly, $\beta_{A_i}^\omega$ and $\beta_{A_i}^\delta$ capture the offensive and defensive strengths of the away team playing at game i . The model also incorporates random-effects to identify a COVID-19 effect, c_H , and a post-COVID effect, c'_H . The intercepts, β_H and β_A , represent the effects of playing at home and away, respectively. The HA effect is defined in the log scale as the rate of the average home team goals and the away goals $HA = \log\left(\frac{\mu_1}{\mu_2}\right) = \beta_H - \beta_A$. In addition, we let the dispersion in the scores of both teams depend on the specific teams that are competing, and also on whether they play at home or away, by modeling the shape parameters as

$$\begin{aligned} \log(\nu_{i1}) &= \gamma_H + \gamma_{H_i}^\omega + \gamma_{A_i}^\delta + k_H + k'_H \\ \log(\nu_{i2}) &= \gamma_A + \gamma_{A_i}^\omega + \gamma_{H_i}^\delta + k_A + k'_A, \end{aligned}$$

where the parameters $\gamma_{H_i}^\omega, \gamma_{A_i}^\delta, \gamma_{A_i}^\omega, \gamma_{H_i}^\delta, k_H, k'_H, k_A, k'_A$ have similar interpretation as above. We argue for this choice in the data applications below (see Figures 3 and 11).

The vector of random effects, $\mathbf{b}_i = (b_{i1}, b_{i2})$, is assumed to follow a multivariate normal centered at 0 with an unrestricted covariance matrix D ,

$$\mathbf{b}_i | D \sim N_2(\mathbf{0}, D), \quad i = 1, \dots, n, \quad (2)$$

and models the within-game correlation between the two teams. More specifically, we can prove that, under the model above, the covariance between the two team's observed counts, y_{i1} and y_{i2} , is approximately

$$\text{cov}(y_{i1}, y_{i2}) \approx \lambda_{i1} e^{0.5d_{11}} (e^{d_{12}} - 1) \lambda_{i2} e^{0.5d_{22}}, \quad (3)$$

with $\lambda_{ij} = \exp(\mathbf{x}_{ij}^T \boldsymbol{\beta}_j)$. The proof can be found in the Appendix. We complete the model by assuming standard independent priors on the vectors $(\boldsymbol{\beta}, \boldsymbol{\gamma}, D^{-1})$ as:

$$\begin{aligned} \boldsymbol{\beta} &\sim N(\boldsymbol{\beta}_0, B_0^{-1}), \\ \boldsymbol{\gamma} &\sim N(\boldsymbol{\gamma}_0, G_0^{-1}), \\ D^{-1} &\sim \text{Wishart}(v_0, R_0), \end{aligned}$$

where $(\boldsymbol{\beta}_0, B_0, \boldsymbol{\gamma}_0, G_0, v_0, R_0)$ are fixed hyperparameters and $\text{Wishart}(v_0, R_0)$ denotes the Wishart distribution with v_0 degrees of freedom and scale matrix R_0 .

2.2 Exchange algorithm and rejection sampler for the CMP

The posterior density of the proposed model is proportional to

$$\Phi_2(\boldsymbol{\beta} | \boldsymbol{\beta}_0, B_0^{-1}) \Phi_2(\boldsymbol{\gamma} | \boldsymbol{\gamma}_0, G_0^{-1}) f_W(D^{-1} | v_0, R_0) \prod_{i=1}^n p(\mathbf{y}_i | \boldsymbol{\beta}, \boldsymbol{\gamma}, b_i) \Phi_2(\mathbf{b}_i | \mathbf{0}, D), \quad (4)$$

which is not available in closed form. Most importantly, in the CMP regression, Bayesian inference for the parameters $\{\mathbf{b}_i\}$, $\boldsymbol{\beta}$, and $\boldsymbol{\gamma}$ is a *doubly-intractable problem* as highlighted by Murray et al. (2012), and Benson and Friel (2021). The likelihood is intractable, since it contains an intractable normalizing constant, $Z(\mu, \nu)$, which is parameter-dependent; thus, the posterior cannot be normalized either. Furthermore, since it is impossible to evaluate the likelihood point-wise, it is not possible to use standard MCMC methods for posterior inference. However, it is possible to use the Exchange algorithm devised by Murray et al. (2012) to augment the doubly-intractable posterior with auxiliary latent variables, since it is possible to draw samples exactly from the intractable likelihood (Chanialidis et al., 2018; Benson and Friel, 2021). Indeed, Chanialidis et al. (2018) showed that the CMP distribution is amenable to a rejection sampler. They proposed a first rejection sampler for the CMP without evaluating its normalizing constant, establishing an upper bound based on a piece-wise geometric distribution.

In this paper, we implement a more efficient rejection algorithm proposed recently by Benson and Friel (2021) where they use two envelope distributions based on the value of the shape parameter ν . More specifically, we construct a Markov chain procedure based on the Exchange algorithm to simulate from (4) using blocks of parameters $\{\mathbf{b}_i\}$, $\boldsymbol{\beta}$, $\boldsymbol{\gamma}$ and D and full conditional distributions,

$$[\mathbf{b} | \mathbf{y}, \boldsymbol{\beta}, \boldsymbol{\gamma}, D]; [\boldsymbol{\beta} | \mathbf{y}, \boldsymbol{\gamma}, \mathbf{b}]; [\boldsymbol{\gamma} | \mathbf{y}, \boldsymbol{\beta}, \mathbf{b}]; [D^{-1} | \mathbf{b}], \quad (5)$$

where $\mathbf{b} = \{\mathbf{b}_1, \dots, \mathbf{b}_n\}$. The simulation output is obtained by recursively simulating these distributions using the updated values of the conditional variables at each step. More in detail, we obtain samples from these distributions as follows:

- **Sampling \mathbf{b} :** The Exchange algorithm proposes to update the parameter vector \mathbf{b}_i to \mathbf{b}_i^* , using a Gaussian proposal density centered at the current value, $h(\mathbf{b}_i, \mathbf{b}_i^*)$, and augmenting the posterior with

n auxiliary draws, $\mathbf{y}_i^* = (y_{i1}^*, y_{i2}^*)$ generated from the likelihood at the proposed parameter value, \mathbf{b}_i^* . The augmented posterior is

$$p(\mathbf{b}_i, \mathbf{b}_i^*, \mathbf{y}_i^* | \mathbf{y}_i) \propto p(\mathbf{y}_i | \mathbf{b}_i, \boldsymbol{\beta}, \boldsymbol{\gamma}) p(\mathbf{b}_i | D) h(\mathbf{b}_i, \mathbf{b}_i^*) p(\mathbf{y}_i^* | \mathbf{b}_i^*, \boldsymbol{\beta}, \boldsymbol{\gamma}). \quad (6)$$

Therefore the acceptance ratio for the augmented posterior is

$$\alpha_{EA}(\mathbf{b}_i, \mathbf{b}_i^*) = \min \left\{ 1, \frac{\prod_{j=1}^2 \frac{q(y_{ij} | \boldsymbol{\beta}_j, \boldsymbol{\gamma}_j, \mathbf{b}_{ij}^*)}{Z(\boldsymbol{\mu}_{ij}^*, \boldsymbol{\nu}_{ij}^*)} h(\mathbf{b}_i^*, \mathbf{b}_i) \Phi_2(\mathbf{b}_i^* | \mathbf{0}, D) \prod_{j=1}^2 \frac{q(y_{ij}^* | \boldsymbol{\beta}_j, \boldsymbol{\gamma}_j, \mathbf{b}_{ij})}{Z(\boldsymbol{\mu}_{ij}, \boldsymbol{\nu}_{ij})}}{\prod_{j=1}^2 \frac{q(y_{ij} | \boldsymbol{\beta}_j, \boldsymbol{\gamma}_j, \mathbf{b}_{ij})}{Z(\boldsymbol{\mu}_{ij}, \boldsymbol{\nu}_{ij})} h(\mathbf{b}_i, \mathbf{b}_i^*) \Phi_2(\mathbf{b}_i | \mathbf{0}, D) \prod_{j=1}^2 \frac{q(y_{ij}^* | \boldsymbol{\beta}_j, \boldsymbol{\gamma}_j, \mathbf{b}_{ij}^*)}{Z(\boldsymbol{\mu}_{ij}^*, \boldsymbol{\nu}_{ij}^*)}} \right\}. \quad (7)$$

Since we can cancel out the normalizing constants, the acceptance ratio is

$$\alpha_{EA}(\mathbf{b}_i, \mathbf{b}_i^*) = \min \left\{ 1, \frac{\prod_{j=1}^2 q(y_{ij} | \boldsymbol{\beta}_j, \boldsymbol{\gamma}_j, \mathbf{b}_{ij}^*) \Phi_2(\mathbf{b}_i^* | \mathbf{0}, D) \prod_{j=1}^2 q(y_{ij}^* | \boldsymbol{\beta}_j, \boldsymbol{\gamma}_j, \mathbf{b}_{ij})}{\prod_{j=1}^2 q(y_{ij} | \boldsymbol{\beta}_j, \boldsymbol{\gamma}_j, \mathbf{b}_{ij}) \Phi_2(\mathbf{b}_i | \mathbf{0}, D) \prod_{j=1}^2 q(y_{ij}^* | \boldsymbol{\beta}_j, \boldsymbol{\gamma}_j, \mathbf{b}_{ij}^*)} \right\}. \quad (8)$$

- **Sampling $\boldsymbol{\beta}$ and $\boldsymbol{\gamma}$:** we need to sample $\boldsymbol{\beta}$ given $(\mathbf{b}, \boldsymbol{\gamma}, D)$ and $\boldsymbol{\gamma}$ given $(\mathbf{b}, \boldsymbol{\beta}, D)$. We can see that the posterior of $\boldsymbol{\beta}$ and $\boldsymbol{\gamma}$ is proportional to

$$p(\boldsymbol{\beta}, \boldsymbol{\gamma} | \mathbf{y}, \mathbf{b}, D) \propto \phi_k(\boldsymbol{\beta} | \boldsymbol{\beta}_0, B_0^{-1}) \phi_k(\boldsymbol{\gamma} | \boldsymbol{\gamma}_0, G_0^{-1}) \prod_{j=1}^2 p(\mathbf{y}_{\cdot j} | \mathbf{b}_{\cdot j}, \boldsymbol{\beta}_j, \boldsymbol{\gamma}_j),$$

where $p(\mathbf{y}_{\cdot j} | \boldsymbol{\beta}_j, \boldsymbol{\gamma}_j, \mathbf{b}_{ij}) = \prod_{i=1}^n p(y_{ij} | \boldsymbol{\beta}_j, \boldsymbol{\gamma}_j, \mathbf{b}_{ij})$. We sample each of the $\boldsymbol{\beta}_j$'s and $\boldsymbol{\gamma}_j$'s by implementing a sequence of Exchange algorithms for each component of the parameter vectors. More specifically, the proposal $\boldsymbol{\beta}_j^*$ is given as proposed in the Robust adaptive Metropolis algorithm (RAM) Vihola (2012), using a Gaussian proposal density. Then we draw n auxiliary data $\mathbf{y}_{\cdot j}^*$, and the acceptance rates for $\boldsymbol{\beta}$ are

$$\alpha_{EA_j}(\boldsymbol{\beta}_j, \boldsymbol{\beta}_j^*) = \min \left\{ 1, \frac{\phi_{k_j}(\boldsymbol{\beta}_j^* | \boldsymbol{\beta}_{0j}, B_{0j}^{-1}) q(\mathbf{y}_{\cdot j} | \boldsymbol{\beta}_j^*, \boldsymbol{\gamma}_j, \mathbf{b}_{\cdot j}) q(\mathbf{y}_{\cdot j}^* | \boldsymbol{\beta}_j, \boldsymbol{\gamma}_j, \mathbf{b}_{\cdot j})}{\phi_{k_j}(\boldsymbol{\beta}_j | \boldsymbol{\beta}_{0j}, B_{0j}^{-1}) q(\mathbf{y}_{\cdot j} | \boldsymbol{\beta}_j, \boldsymbol{\gamma}_j, \mathbf{b}_{\cdot j}) q(\mathbf{y}_{\cdot j}^* | \boldsymbol{\beta}_j^*, \boldsymbol{\gamma}_j, \mathbf{b}_{\cdot j})} \right\}, j \leq 2. \quad (9)$$

Similarly, we have the sequence of acceptance ratios for $\boldsymbol{\gamma}$ as

$$\alpha_{EA_j}(\boldsymbol{\gamma}_j, \boldsymbol{\gamma}_j^*) = \min \left\{ 1, \frac{\phi_{k_j}(\boldsymbol{\gamma}_j^* | \boldsymbol{\gamma}_{0j}, G_{0j}^{-1}) q(\mathbf{y}_{\cdot j} | \boldsymbol{\beta}_j, \boldsymbol{\gamma}_j^*, \mathbf{b}_{\cdot j}) q(\mathbf{y}_{\cdot j}^* | \boldsymbol{\beta}_j, \boldsymbol{\gamma}_j, \mathbf{b}_{\cdot j})}{\phi_{k_j}(\boldsymbol{\gamma}_j | \boldsymbol{\gamma}_{0j}, G_{0j}^{-1}) q(\mathbf{y}_{\cdot j} | \boldsymbol{\beta}_j, \boldsymbol{\gamma}_j, \mathbf{b}_{\cdot j}) q(\mathbf{y}_{\cdot j}^* | \boldsymbol{\beta}_j, \boldsymbol{\gamma}_j^*, \mathbf{b}_{\cdot j})} \right\}, j \leq 2. \quad (10)$$

- **Sampling D :** the full conditional for D is a Wishart distribution $p(D | \mathbf{b}) \propto p(\mathbf{b} | D) p(D) = \prod_{i=1}^n \phi_2(\mathbf{b}_i | \mathbf{0}, D) f_w(D^{-1} | v_0, R_0^{-1})$. Thus $D^{-1} | \mathbf{b} \sim \text{Wishart} \left(n + v_0, [R_0^{-1} + \sum_{i=1}^n (\mathbf{b}_i \mathbf{b}_i^T)]^{-1} \right)$, and is sampled directly.

3 Simulation Study

3.1 Data generation and parameter settings

This section aims at illustrating and evaluating the performance of the proposed CMP model under simulated data on different scenarios. More specifically, we assess how the model fits score data and compare performances with the widely employed independent Poisson and Negative Binomial models. Additionally, we test the capability of the proposed model to accurately capture the true HA effect in scenarios with general over-, equi-, and under-dispersed data.

Based on the model presented in Section 2.1, we fix the home advantage (HA) effect by setting $\beta_H = 0.6$ and $\beta_A = 0.1$, resulting in a true effect of $HA = 0.5$. We simulate scores within a framework comprising 20 teams, labeled as Team 1 through Team 20, across various seasons. Each season consists of a total of 380 games, where every team competes against each other, both at home and away. In this simulation

setup, we did not consider the effect of the COVID-19 season. We assign to every team specific strength attacking (β^ω) and defending (β^δ), distinguishing between their performance at home (H) and away (A). At home, attacking strengths range from 0.5 for Team 1 to -0.5 for Team 20. The attacking strengths playing away range from -0.42 to 0.42, with a notable exception for Teams 6 through 15, where the sign is inverted. The list of the effects for every team can be found in Table 1. This approach categorizes the teams into four groups: the first group (Teams 1:5) is characterized by strong performance both at home and away. The second group (Teams 6:10) demonstrates above-average attacking strength at home but below-average performance away. Conversely, the third group (Teams 11:15) displays below-average strength when playing at home but stronger away attacking capabilities. Finally, the last group (Teams 16:20) possesses the lowest strengths compared to the other teams and is, therefore, composed of the worst performing teams. As for the strengths defending, they were assigned from 0.5 to -0.5 at Home and from 0.6 to -0.6 playing Away.

Tab. 1: Simulation study: Team's strengths at Attack and Defense playing Home or Away

Team	Good attack at home and away					Good at home, bad away					Bad at home, good away					Bad at home, bad away				
	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20
β_H^ω	0.50	0.45	0.40	0.34	0.29	0.24	0.18	0.13	0.08	0.03	-0.03	-0.08	-0.13	-0.18	-0.24	-0.29	-0.34	-0.40	-0.45	-0.5
β_H^δ	0.50	0.45	0.40	0.34	0.29	0.24	0.18	0.13	0.08	0.03	-0.03	-0.08	-0.13	-0.18	-0.24	-0.29	-0.34	-0.40	-0.45	-0.5
β_A^ω	0.42	0.38	0.34	0.29	0.25	-0.20	-0.16	-0.11	-0.07	-0.02	0.02	0.07	0.11	0.16	0.20	-0.25	-0.29	-0.34	-0.38	-0.42
β_A^δ	0.60	0.54	0.47	0.41	0.35	0.28	0.22	0.16	0.10	0.03	-0.03	-0.10	-0.16	-0.22	-0.28	-0.35	-0.41	-0.47	-0.54	-0.60

As for the shape parameters γ , we consider that strongest teams have a larger dispersion in the number of goals, so their shape parameters will be larger than weakest teams, then we set them sequentially from 0.35 for Team 1 to -0.35 for Team 20, both at home and away. We consider three scenarios where data tend to be over-, equi- or under-dispersed based on the value of the intercept parameters γ_H and γ_A . In the first scenario, we let $\gamma_H = \gamma_A = 0$, in which we observe that the shape parameters ν_1, ν_2 are centered around 1, corresponding to equi-dispersed data. In the second scenario, we set $\gamma_H = \gamma_A = 0.5$, resulting in shape parameters mainly greater than 1, indicating an under-dispersed scenario. In the last scenario, we consider $\gamma_H = \gamma_A = -0.5$, which corresponds to data mostly over-dispersed.

We fit the CMP model assuming non-informative priors by setting the hyper-parameters to $\beta_0 = \gamma_0 = 0$, $B_0 = G_0 = 0.1I$, $\nu_0 = 50$, $R_0 = I$. A motivation for this choice of values and a prior sensitivity analysis can be found in the Appendix. We run MCMC chains with 30K samples in total, where 10K are used for burn-in. We tune the MCMC algorithm to achieve an acceptance ratio of approximately 40%. For each scenario, we simulate 40 replicated datasets, fit the three models, estimate the HA effect, and calculate the Deviance Information Criterion (DIC) (Spiegelhalter et al., 2002) to assess the goodness of fit, as also done by Chaniavidis et al. (2018). In the case of the Conway-Maxwell-Poisson model, direct computation of the DIC is not possible due to the intractable constant in the likelihood function. Thus, we adopt an approach suggested by Benson and Friel (2021) to utilize an unbiased likelihood estimator that relies on the number of acceptances (r) obtained from the rejection sampling algorithm used in the Exchange algorithm,

$$\hat{f}^{(r)}(y|\theta) = \prod_{i=1}^n \frac{q_f(y_i|\theta_i)}{Z_g(\phi_i)} \frac{\widehat{M}_{f/g,i}^{(r)}}{B_{f/g,i}},$$

where Z_g is the normalizing constant of the envelope distribution, $B_{f/g}$ is a tractable enveloping bound, and $\widehat{M}^{(r)}$ is the number of draws required for r acceptances. For more detailed information, please refer to Benson and Friel (2021). They recommend a high value for the number of acceptances (r). In our case, we selected $r = 1000$. For computational simplicity we calculate the DIC using 100 random samples from the MCMC algorithm.

3.2 Results

Figure 1 shows the estimation and the 95% credible interval of the HA effect under the models considered across 1, 3 and 5 seasons. Looking at those results, we observe that, by modeling the dispersion parameter, we indeed capture the true effect on almost all the scenarios, with improved precision for cases with larger sample sizes. In contrast, under the Negative Binomial and the Poisson models, the effect is not captured if the scores tend to be over or under-dispersed. These results highlight the fact that our model is robust in fitting data exhibiting any type of dispersion.




Fig. 1: Simulation study: comparison of HA effect estimation and their 95% credible intervals. Dashed line corresponds to the true value

Deviance Information Criterion (DIC) values for the different models averaged across replicated datasets are shown in Table 2. Both the Conway-Maxwell-Poisson and Poisson models demonstrate reasonable performance in capturing equi-dispersion, with no significant difference observed between the DIC values for the two response variables in cases with small sample size. In the over-dispersed scenarios, the Negative Binomial model generally outperforms the Poisson model but shows comparable or worse performance compared to the Conway-Maxwell-Poisson model. As the sample size increases to $n = 1900$, the Conway-Maxwell-Poisson model exhibits the lowest DIC values, and significant differences among the models are observed in both over- and under-dispersed scenarios. These results suggest that the Conway-Maxwell-Poisson model outperforms the Poisson and Negative Binomial models when dealing with large sample sizes. Furthermore, the Conway-Maxwell-Poisson model adapts very well to various dispersion scenarios.

4 Analysis of Soccer English Premier League Data

4.1 Data description

We illustrate the performance of our model on sports count data, by using game outcomes from the English Premier League over the span of three seasons: 2019-2020, 2020-2021, and a portion of 2021-2022.¹ The dataset comprises a total of 1237 games. On average, we observe 1.49 goals per game for the home team, with a corresponding variance of 1.79. Similarly, the away team records an average of 1.27 goals per game,


¹ Data extracted from: <https://www.football-data.co.uk/>

Tab. 2: Simulation study: DIC (Mean, SD) across 40 replicates. Bolded numbers represent the lowest values for each response in the bivariate outcome model.

Data	Model	1 Season ($n = 380$)		3 Season ($n = 1140$)		5 Seasons ($n = 1900$)	
		y_1	y_2	y_1	y_2	y_1	y_2
Equi	CMP	1355.91 (35.02)	1091.58 (31.28)	3922.93 (60.83)	3220.48 (55.66)	6487.45 (80.34)	5343.95 (73.92)
	Poisson	1333.79 (31.86)	1104.65 (32.39)	3932.85 (60.14)	3245.29 (52.67)	6542.70 (63.86)	5390.63 (69.82)
	NB	1349.06 (28.36)	1118.44 (30.92)	3937 (53.56)	3252.55 (50.09)	6538.25 (58.19)	5392.93 (67.41)
Over	CMP	1490.37 (30.92)	1251.37 (28.85)	4391.69 (54.74)	3724.38 (50.12)	7269.45 (73.87)	6194.46 (73.13)
	Poisson	1509.55 (37.25)	1279.62 (34.64)	4493.19 (75.45)	3805.90 (56.38)	7470.95 (99.78)	6340.21 (81.99)
	NB	1498.13 (30.25)	1280.05 (31.38)	4407.10 (58.35)	3754.65 (45.27)	7314.27 (77.73)	6238.71 (64.75)
Under	CMP	1236.08 (24.67)	962.93 (26.79)	3523.70 (46.89)	2810.50 (55.05)	5722.46 (62.80)	4596.15 (72.28)
	Poisson	1203.57 (18.45)	966.65 (28.49)	3522.18 (38.23)	2817.38 (40.78)	5841.81 (39.09)	4665.59 (54.50)
	NB	1228.44 (17.73)	985.75 (28.58)	3555.02 (36.92)	2840.20 (40.07)	5880.47 (37.91)	4692.64 (53.74)

accompanied by a variance of 1.49. These statistics indicate that both the home and away goals exhibit a slightly higher variance than their respective means. To assess the relationship between the mean and the variance, we utilize the dispersion statistic, denoted as σ_p . This measure takes into account the sample size and model complexity while quantifying the difference between the variance and the mean. Notably, for both the home and away goals, σ_p is found to be less than 1.04, indicating no significant deviation between the mean and the variance. It is worth noting that a threshold of $\sigma_p > 1.2$ is typically employed to identify over-dispersion (Payne et al., 2018). Hence, the results suggest that the Poisson model may be a reliable choice for modeling this data.

The negative Spearman's correlation coefficient of -0.143 suggests a slight monotonic relationship between the goals scored by the home and away teams. This finding is further supported by the distribution of goals depicted in Figure 2. Specifically, the relative frequency of the home team scoring more goals than the away team ($Fr(\text{Home Team Goals} > \text{Away Team Goals})$) is found to be 0.429, while the relative frequency of the home team scoring fewer goals than the away team ($Fr(\text{Home Team Goals} < \text{Away Team Goals})$) is 0.34. These findings underscore the presence of a correlation effect between the goals scored by the home and away teams. Therefore, it is important to consider this relationship when analyzing the data.

Fig. 2: Analysis of Premier League data: Distribution of the scores of the 1237 games in the English Premier League

To complement the analysis of dispersion, in Figure 3 we display the Index of dispersion (ID) for the Premier League data. The ID or Variance to Mean Ratio (VMR) is defined as the ratio between the sample variance and the sample mean, $\frac{s^2}{\bar{x}}$. If data tend to be equi-dispersed this value is around to 1, while for

over-dispersed data will be greater than 1. The heat map on the left of Figure 3 corresponds to the ID of the goals scored by the Home Team, while the heat map on the right is the ID of the goals scored by the Away team. We can observe that the index of dispersion is generally close to 1 for specific games with over-dispersion and others with under-dispersion, depending on the teams playing. Additionally, the two heatmaps are different, suggesting that the effect on the dispersion of the home goals and the away goals made by the teams is different when they play at home than when they play on the road. For instance, when Arsenal plays at home, the ID of the goals they scored (first column, left heatmap) tends to be greater than 1, but the goals they conceded (first column, right heatmap) tend to be more under-dispersed. Similarly we can interpret the values for Everton. When Everton plays away, the ID of the number of goals scored by the opponent (Everton row, left heatmap) varies depending on the opponent; 7 out of them are over-dispersed. However, the ID of the goals they score tends to be more equi-dispersed (Everton row, right heatmap).

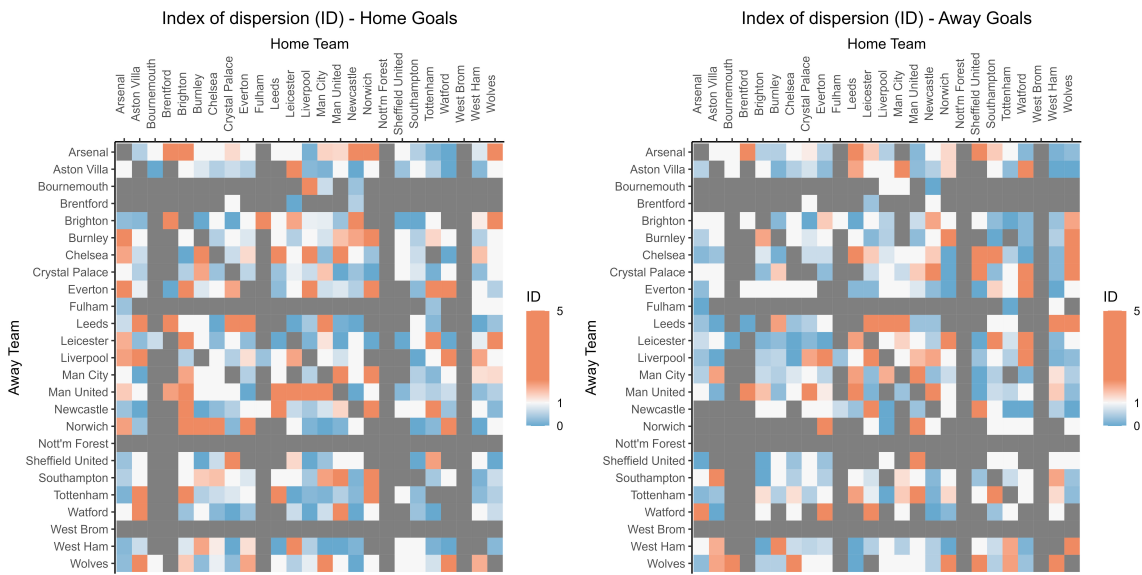


Fig. 3: Analysis of Premier League data: Indices of dispersion (ID)

Figure 4 shows the average team’s goals per game during the COVID-19 restrictions. We observe a decrease in the number of goals scored per game. Prior to the pandemic, the average rate of goals scored by home teams per game was 1.51. However, this figure decreased to 1.38 during the pandemic and subsequently increased to 1.57 after the restrictions were lifted. A detailed breakdown by teams can be found in Figure 4(a). From this plot, we can discern that nearly half of the teams experienced a reduction in the average number of goals scored per game at the home field during the pandemic, followed by an increase afterwards. Conversely, the average number of goals scored by away teams per game exhibited a small increase during the COVID-19 pandemic, rising from 1.22 before the pandemic to 1.31 during the restrictions. However, this trend reversed after the COVID-19 restrictions were lifted, resulting in a decrease to 1.27 goals per game. Figure 4(b) provides a visual representation of these observations.

Furthermore, in Figure 5 we can examine the difference between the number of goals scored by the home team and the away team and its distribution across the COVID-19 seasons under consideration. A negative difference indicates that the away team scored more goals than the home team, signifying an away team victory, while a positive difference implies that the home team scored more goals, representing a home team victory. We see that the distribution of the home goals difference remains similar before and after the COVID-19 restrictions. However, during the pandemic restrictions, there is a slight increase in the proportion of games with a negative difference. From these findings, it is apparent that the percentage of home team wins significantly decreased during the COVID-19 restrictions (% Home Wins: Before - 45%,

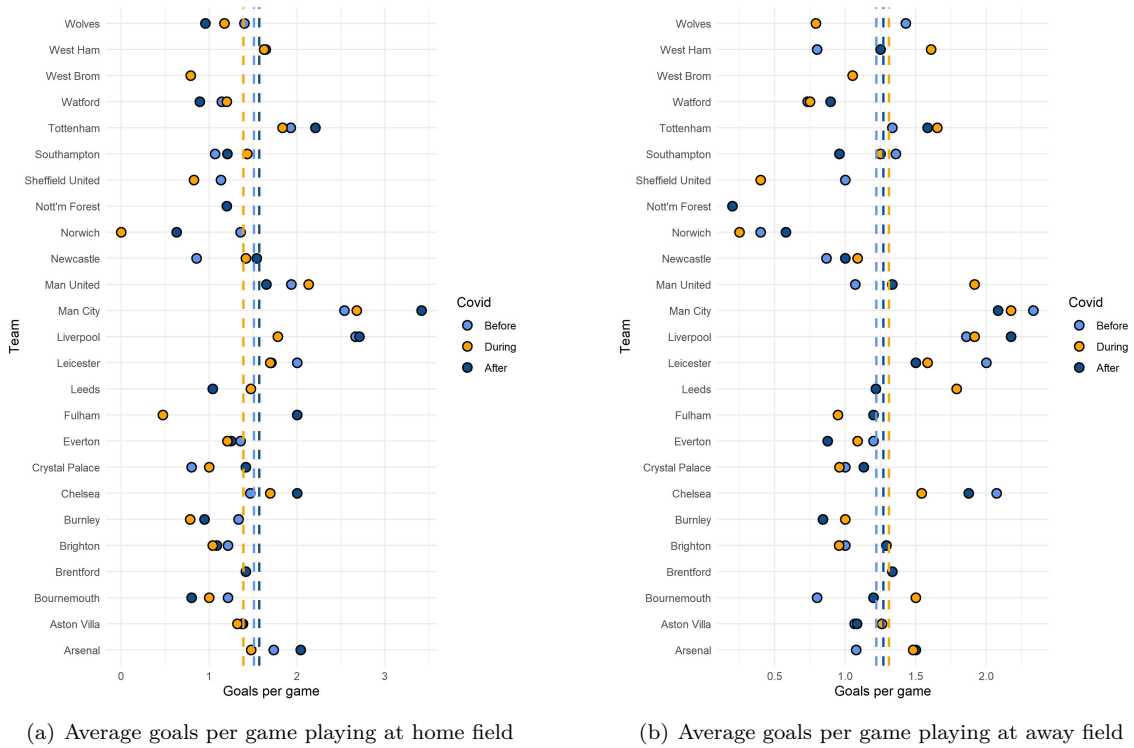


Fig. 4: Analysis of Premier League data: Average team goals per game by COVID-19 stage

After - 44%, During - 39%), while the percentage of home team losses increased (% Home Losses: Before - 31%, After - 31%, During - 39%). This observation supports the hypothesis that the HA diminished during the pandemic and was recovered afterwards. However, an ANOVA test to ascertain whether there were significant differences in the home goals difference across the stages of COVID-19 (before, during, and after the restrictions) yielded a p-value of 0.154, indicating no significant evidence to reject the null hypothesis that the average home goals difference per game did not change during the various stages of COVID-19 considered.




Fig. 5: Analysis of Premier League data: distribution of the difference between Home goals and Away goals

4.2 Model fitting

We fit the proposed CMP model on the Premier League data, and then compare the performance with those of the Negative Binomial and Poisson models. Given the negative correlation between the two responses and the apparent equi-dispersion of the data, our model appears to be a suitable approach for modeling the Premier League data and identifying any variations in the HA during the crowd restrictions imposed amid the COVID-19 pandemic.

For our analysis, we assume the following prior hyperparameters: $\beta_0 = \gamma_0 = 0$, $B_0 = G_0 = 0.1 I$, $\nu_0 = 50$, and $R_0 = I$, to reflect minimal prior knowledge. In this case, we used 180K MCMC iterations with a burn-in of 50K iterations, which achieved an effective sample size greater than 500 for the estimated parameters. Figure 6 displays a comparison between the empirical distribution of the data and the posterior predictive distribution obtained from our model. It is evident that our model effectively captures the distribution of the data, suggesting a good fit between the observed and the predicted values.

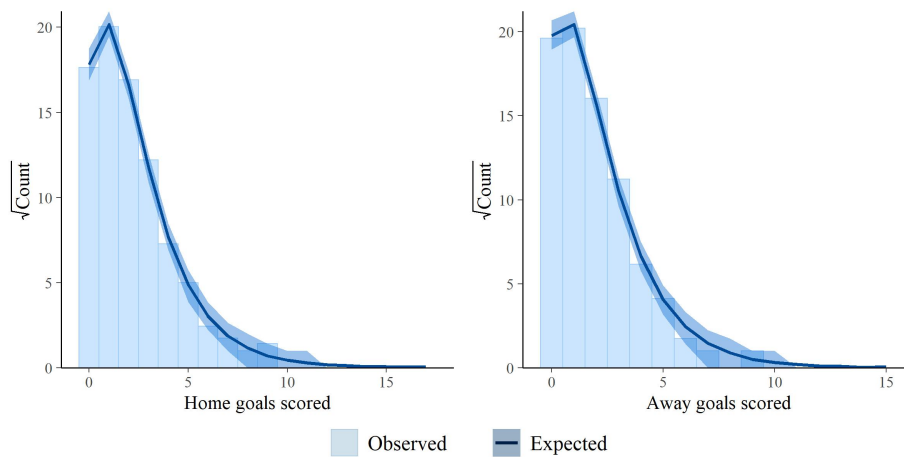


Fig. 6: Analysis of Premier League data: empirical distribution and the posterior predictive distribution of the home goals and away goals

Figure 7 presents a comparison of the attack and defense strengths of each team based on the posterior estimations. Higher values in the defense parameters of any team β^δ indicate that it is more likely to concede goals. We plot the exponential of the negative parameter such that teams positioned above the vertical and horizontal lines at 1 exhibit better performance than the team average. Notably, Manchester City and Liverpool, the champions in recent seasons, demonstrate the strongest attack and defense, whether playing at home or away. Interestingly, there are teams that seem to benefit from playing away, as indicated by their superior attack strength compared to other teams. West Ham, Brighton, and Leeds are notable examples. This is an intriguing finding because these teams were actually battling to avoid relegation during the seasons considered in the study. In an interview in *The Athletic*, Shaka Hislop, former West Ham goalkeeper, suggested a possible reason for this phenomenon. He points out that teams facing difficulties often find it easier to play away from home due to reduced pressure. Hislop notes that the intense scrutiny and expectations of home crowds can inhibit players' performance, whereas playing away allows for more freedom and different strategies, that could contribute to improved results (Thomas, 2019). On the other hand, several teams demonstrate a lower number of goals conceded when playing at home compared to playing away. Noteworthy teams in this regard include Everton, Crystal Palace, and Brighton. These findings provide valuable insights into the varying performance and strategies of different teams in terms of attack and defense strengths, shedding light on the dynamics of the Premier League. Similarly, we can interpret the estimated shape parameters. A detailed description and this estimation can be found in the Appendix.




Fig. 7: Analysis of Premier League data: Attack and defense strength of each team

In Figure 8, we examine the distribution of the HA before the COVID-19 pandemic (HA_B), during the pandemic (HA_D), and after the pandemic restrictions (HA_A). We define the HA effect in the log scale as the rate of the average home team goals and the away goals during the corresponding COVID-19 stage, i.e. $HA_B = \log(\frac{\mu_1}{\mu_2}) = \beta_H + c_H - (\beta_A + c_A)$, $HA_D = \beta_H - \beta_A$, and $HA_A = \beta_H + c'_H - (\beta_A + c'_A)$. The estimated effects of HA_B , HA_D , and HA_A are approximately 0.238, 0.0916, and 0.288, respectively. This suggests that we initially expected the Home Team to score 26.8% more goals than the Away Team before the pandemic. However, during the pandemic, this advantage reduced to 9.6%, and after the COVID-19 restrictions, it appeared to recover to 33.4%. Furthermore, we find that the HA was significantly diminished during the COVID-19 pandemic, as evidenced by the posterior probability $P(HA_D < HA_B | \text{data}) = 0.8525$. However, after the COVID-19 restrictions were lifted, there seems to be a recovery in the HA, as indicated by the posterior probability $P(HA_D < HA_A | \text{data}) = 0.9527$. Among the reasons as to why the HA was negatively affected during the COVID-19 pandemic and subsequently recovered as restrictions were lifted, some authors have provided evidence suggesting that the absence of crowds influenced the behavior of referees in soccer, resulting in fewer penalties being awarded to away teams (McCarrick et al., 2021; Tulp and Thaller, 2020). In recent work by Price et al. (2022), they evaluated the HA in soccer using a causal inference approach. They found that the HA is primarily driven by offensive-based statistics, where the home team gains an advantage by dominating the field, resulting in more shots, passes, and opportunities in the danger zone. Conversely, other authors suggest that crowd influence plays a significant role in the HA effect (Pettersson-Lidbom and Priks, 2010; Lopez, 2016; Reade et al., 2022). These perspectives highlight the multifaceted nature of the HA effect and the various factors that can contribute to its manifestation in different sports.

Finally, in Figure 9, we compare the fitting of our proposed model with the Negative Binomial and Poisson models. The posterior predictive rootograms (Kleiber and Zeileis, 2016) for the home and away team goals are displayed for each distribution. The rootograms provide a visual illustration of the goodness of fit, showing the alignment between the observed counts and the expected counts for each distribution. In the plot, the observed counts' bars are shown hanging from the curve that represents the expected counts. A perfect fit occurs when the observed bars coincide with the horizontal line at 0. Upon examination, it is evident that both the CMP and Poisson models exhibit similarly good fits to the data. The CMP distribution appears to effectively capture the pattern of the observed home team goals, especially for zero and lower counts. Additionally, in Table 3 we display the deviance information criterion (DIC) (Spiegelhalter et al., 2002). A smaller DIC indicates a better fit to the data set. As it can be seen, the lowest values correspond to the proposed CMP model.




Fig. 8: Analysis of Premier League data: distribution of the HA for pre, during, and post-COVID-19 season in the Premier League

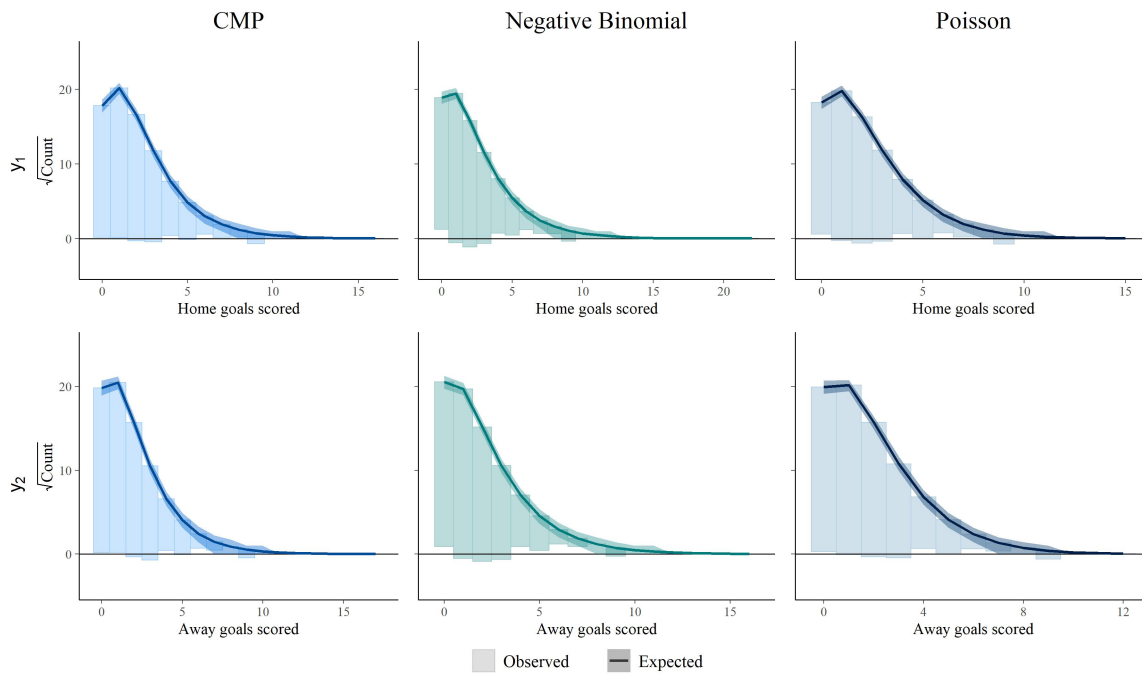


Fig. 9: Analysis of Premier League data: posterior predictive Rootogram of home team goals y_1 and away team goals y_2 for the CMP, Negative Binomial and Poisson models

Tab. 3: Analysis of Premier League data: Deviance Information Criterion (DIC) for the CMP, Negative Binomial (NB) and Poisson models. Bolded numbers represent the lowest DIC for each response

Model	y_1	y_2
CMP	3635.611	3475.992
NB	3759.975	3562.146
Poisson	3738.665	3543.986

5 Analysis of Major League of Baseball data

5.1 Data description

To further elucidate the performance of the CMP model on an additional dataset, we conducted an analysis of 5,756 Major League Baseball (MLB) games played during the regular seasons of 2019, 2020, and 2021.² The average number of runs scored by the home teams was 4.72 per game, with a variance of 10.24. Conversely, the away teams averaged 4.63 runs per game, with a variance of 11.01. Notably, both sets of data exhibit a dispersion statistic, σ_p , greater than 2.1, indicating the data is significantly over-dispersed Payne et al. (2018). In contrast to the soccer data, baseball scores appear to exhibit no correlation, as evidenced by the Spearman correlation coefficient of 0.004. The relative frequency of the home team scoring more runs than the away team, denoted as $Fr(\text{Home Team runs} > \text{Away Team runs})$, is 0.534, while the frequency of the home team scoring fewer runs or an equal number of runs, denoted as $Fr(\text{Home Team runs} \leq \text{Away Team runs})$, is 0.46. These statistics suggest a general positive HA effect. The distribution is visualized in Figure 10. It is important to note that in baseball, ties do not occur. If a game is tied after nine innings, extra innings are played until one team scores and the other does not.

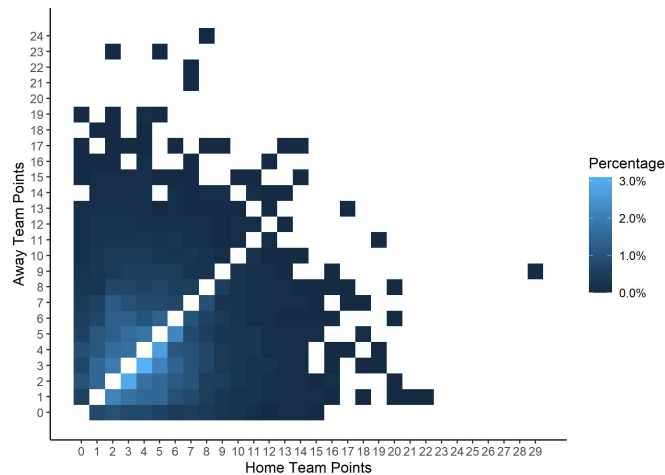


Fig. 10: Analysis of MLB data: distribution of scores in the MLB data

In Figure 11 we find the Index of dispersion (ID) for the MLB data. In this case, we observe a natural tendency for data to be over-dispersed, especially in the number of points scored by the home team. Furthermore, the index of dispersion varies for specific teams. For instance, when playing at home, the ID of the number of runs of the Toronto Blue Jays (TOR) is greater than 1 regardless of the team they are playing against (TOR column, left heatmap). However, the number of runs of some opponents (ARI, SEA, SFN, WAS) could be slightly equi- or under-dispersed (TOR column, right heatmap). These analyses support the proposed model, as we observe significant variation in the dispersion of goals or points among different teams, and suggest that the effect differs when playing at home or away.

Baseball, in comparison to other sports, is known for having a smaller HA (Jones, 2015). In the dataset analyzed, we found that the average number of runs scored by the home team before the pandemic, during the pandemic, and after the pandemic were 4.82, 4.74, and 4.61, respectively. Conversely, the average number of runs scored by the away team before the pandemic, during the pandemic, and after the pandemic were 4.84, 4.55, and 4.46, respectively. Figure 12 provides a visual representation of the average runs per

² Data extracted from: <https://www.retrosheet.org/>

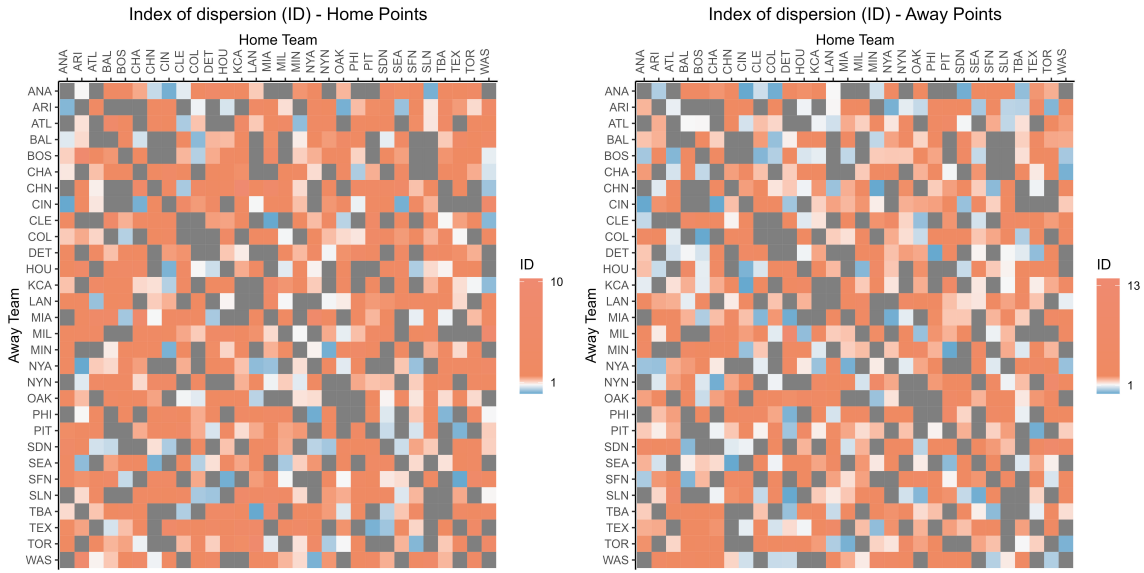


Fig. 11: Analysis of MLB data: Indices of dispersion (ID)

game for all teams during different stages of the COVID-19 pandemic. Comparing the pandemic period to the 2019 season, we observe that 17 out of the 30 teams decreased the number of runs scored when playing at home, while 15 teams decreased the number of runs scored when playing away. However, the decrease was more pronounced in the average number of runs scored by away teams, suggesting a potential increase in the home advantage.

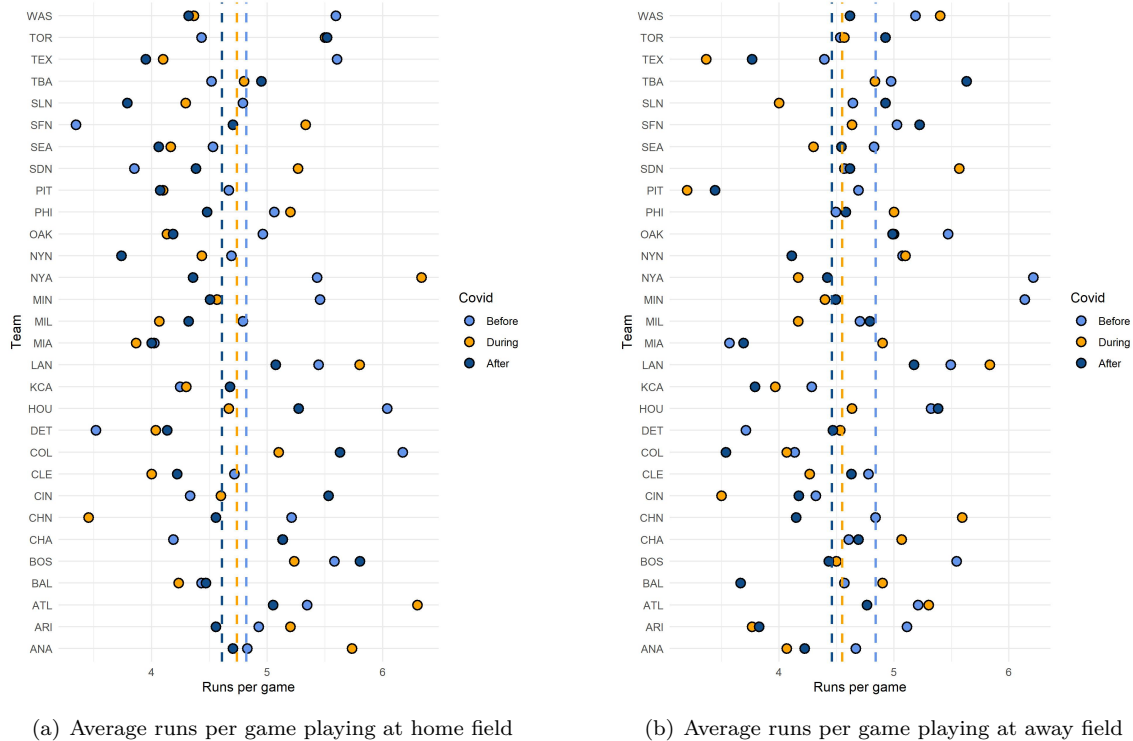


Fig. 12: Analysis of MLB data: average team runs per game by COVID-19 stage

Additionally, in Figure 13, we displayed the distribution of the home difference of runs across seasons. A positive difference indicates a win for the home time, while a negative difference indicates a win for the away team. With this, we observed a different trend in the 2020 Season, with the proportion of home team wins increasing to 55.12% compared to 52.94% in the 2019 season, before decreasing again to 53.8% in 2021. Similar to the soccer case, an ANOVA test was performed trying to detect any significant difference in the mean of home points differences across seasons. The resulting p-value of 0.357 indicates that there are no significant differences detected by the test.




Fig. 13: Analysis of MLB data: distribution of the differences between home team runs and away team runs

5.2 Model fitting

Although the data reveals no apparent correlation between the scores of the home team and the away team, it does exhibit significant over-dispersion that can be effectively captured by the proposed CMP model. To fit the model, we assumed again weak prior information, setting $\beta_0 = \gamma_0 = 0$, $B_0 = G_0 = 0.1I$, $\nu_0 = 50$, and $R_0 = I$. We used 180K MCMC iterations with a burn-in of 50K iterations, which resulted in an effective sample size greater than 400 for all the parameters. The resulting posterior predictive distribution is displayed in Figure 14, illustrating a well-fitting distribution that closely aligns with the observed data.

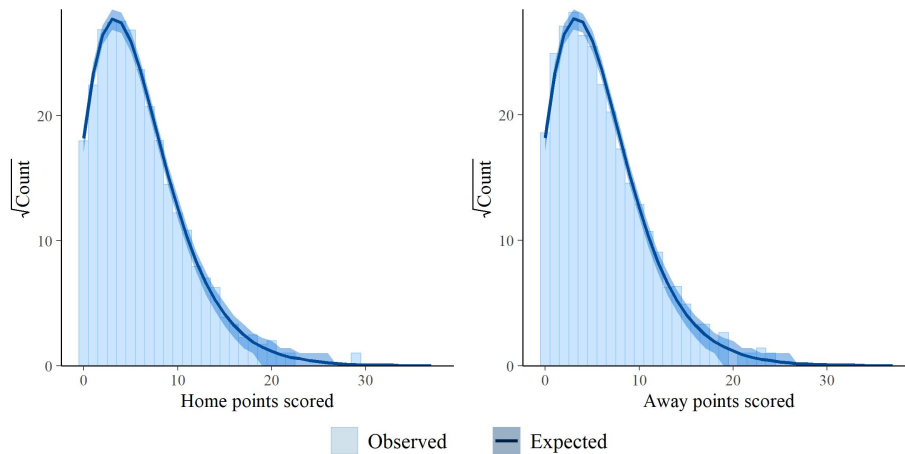


Fig. 14: Analysis of MLB data: empirical distribution and the posterior predictive distribution of the home goals and away points in MLB data

In Figure 15, we can observe the strength of defense and attack for each team when playing on their home field and when playing away. Notably, the Los Angeles Dodgers (LAN), Atlanta Braves (ATL), and

Houston Astros (HOU), winners of the last three World Series Championships, are among the teams with strong performance in both home and away games. On the other hand, the Colorado Rockies (COL) exhibit a contrasting pattern. While they have one of the weakest attacks when playing away, they possess the strongest attack when playing at home. This can be attributed to the unique circumstances of their home stadium, Coors Field, which is situated at an elevation of 5,183 feet above sea level. The Colorado Rockies may benefit from their familiarity and adaptation to playing under these conditions, which could account for their notable advantage when competing at home. An opposite behavior is observed with the Seattle Mariners (SEA), as their strength in attacking is better on the road than at home. During 2022, they also scored more runs away than at home, emerging as one of MLB’s top road teams with 41 wins. Manager Scott Servais suggested that some road ballparks are more conducive to home runs than their home ballpark (T-Mobile Park) (Kramer, 2022). Estimation of the shape parameters are reported in the Appendix.

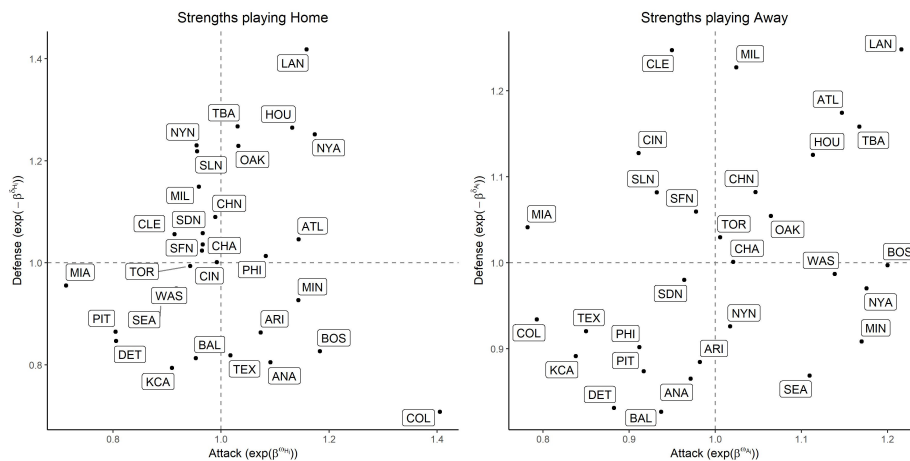


Fig. 15: Analysis of MLB data: attack and defense strength of each team

As depicted in Figure 16, the HA effect demonstrated an increase during the 2020 season compared to both the 2019 and 2021 seasons. Notably, studies such as Higgs and Stavness (2021) and Chiu and Chang (2022) suggest that, despite a slight increase in the HA effect in baseball, it was not significantly affected during the COVID-19 pandemic. It is important to acknowledge a limitation of this analysis, namely that we only considered one season before the pandemic; thus, the observed difference during the COVID-19 season may not be statistically significant when considering previous years. However, these findings lend support to the hypothesis that the HA effect in baseball behaves differently compared to other sports, where it was negatively impacted by the pandemic (Higgs and Stavness, 2021; McCarrick et al., 2021).

Our observations align with the prevailing notion that the HA in baseball tends to be smaller compared to soccer. In soccer, the HA effect is often more pronounced and can have a significant impact on game outcomes. This difference can be attributed to various factors, including cultural differences and the behavior of fans in stadiums. In baseball, the HA effect is thought to be more associated with factors such as familiarity with the home field or travel fatigue, as suggested by Chiu and Chang (2022). Players may have a better understanding of the unique characteristics of their home field, providing them with a slight advantage. Additionally, the physical demands of travel in baseball, with teams frequently playing on the road, may contribute to a smaller HA effect. Another perspective suggests that the dynamics of baseball as a sport contribute to a smaller HA effect compared to team sports like soccer. Baseball is often considered more individualistic, with the actions of individual players, such as pitchers and batters, playing a significant role in the outcome of the game. This individual nature of the sport may reduce the impact of the HA effect, which is more pronounced in team sports where the collective dynamics and energy from the crowd can influence player performance (Jones, 2015; Losak and Sabel, 2021). These theories highlight the complex




Fig. 16: Analysis of MLB data: distribution of the HA for pre, during, and post-COVID-19 season in the MLB

interplay of factors contributing to the HA effect in baseball, which may differ from those observed in other sports.

In a similar manner to the analysis of soccer data, we conducted a visual comparison of the fitting performance of our proposed CMP model with that of the Negative Binomial and Poisson models using rootograms, as shown in Figure 17. The rootograms provide a visual representation of the deviations between the observed data and the expected values from each model. Upon inspection, it becomes apparent that the Poisson model is not suitable for accurately modeling the baseball data due to the evident over-dispersion present in the dataset. Conversely, the CMP model demonstrates a strong performance and is a compelling alternative to the widely used Negative Binomial model in this context. The rootograms exhibit fewer and smaller deviations for the CMP model, indicating a better fit to the considered datasets. To complement this analysis, we calculated the Deviance Information Criterion (DIC) for all the models. Results are presented in Table 4. This underscores the robustness of the proposed CMP model in accommodating various types of data dispersion and its ability to achieve a favorable fitting outcome.

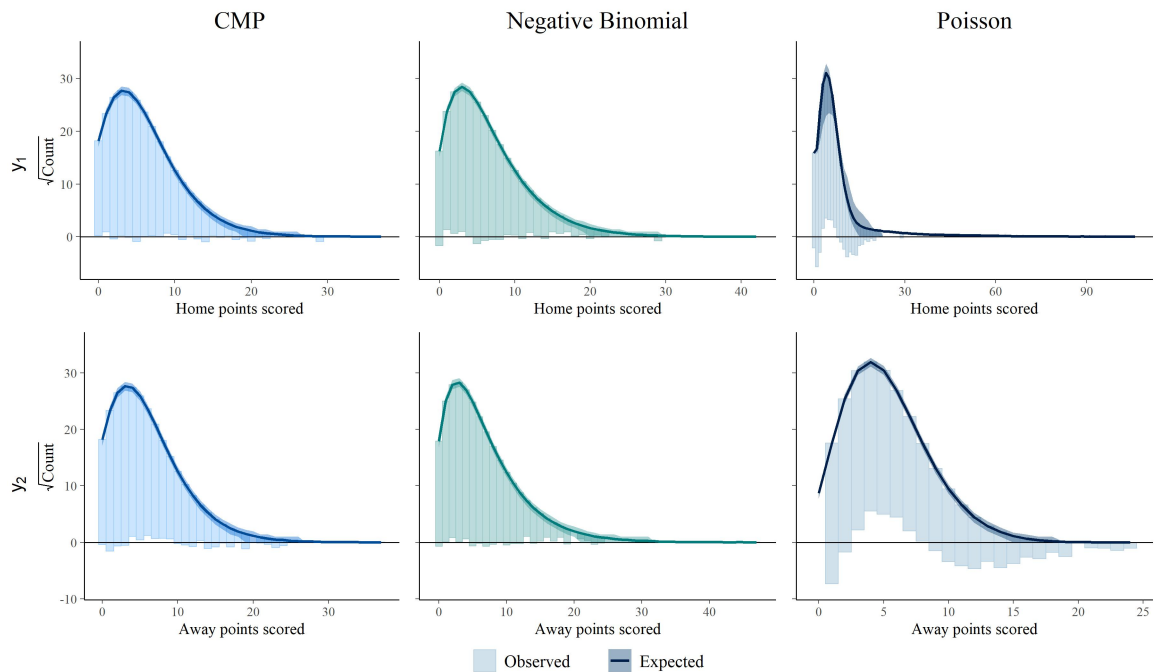


Fig. 17: Analysis of MLB data: posterior predictive rootogram of home team points y_1 and away team points y_2 for the CMP, Negative Binomial and Poisson models

Tab. 4: Analysis of MLB data: Deviance Information Criterion (DIC) for the CMP, Negative Binomial (NB) and Poisson models. Bolded numbers represent the lowest DIC for each response

Model	y_1	y_2
CMP	27943.03	28135.32
NB	28516.47	28564.52
Poisson	65994.64	30906.11

6 Discussion and conclusion

In this paper, we have demonstrated the practical application of a bivariate CMP model for analyzing count data in sports, with a specific focus on investigating the HA effect in soccer and baseball, while also considering its variations during the COVID-19 pandemic. The flexibility of the proposed model allows to effectively handle count data exhibiting different levels of dispersion, including under-dispersion, over-dispersion, and equi-dispersion. To account for the potential correlation between home and away scores, we have incorporated random effects into the model. We have demonstrated the utilization of the rejection sampler proposed by Benson and Friel (2021) and the Exchange algorithm presented by Murray et al. (2012) to perform inference. These methods offer computational efficiency in the estimation of the model parameters.

Results of our analyses have indicated that the proposed CMP model is robust and capable of effectively modeling count data with any level of dispersion. Indeed, data applications and simulations have shown that our model is not only capable of adapting to different types of data, consistently matching or outperforming the fitting of standard Negative Binomial and Poisson models, but also reliable in the estimation of the true parameters regardless of the dispersion of the data. It is worth noting, however, that the proposed model increases the computational time per iteration with respect to fitting simpler models. For example, under the one season scenario, 1000 iterations of our model require approximately 30 seconds, while other models take about 15 seconds. Moreover, due to the implementation of the Exchange Algorithm, a larger number of iterations is required to achieve a given effective sample size. We have demonstrated that when modeling under or over-dispersed data, estimations from the standard models can be biased due to the dispersion of the data in contrast to the proposed CMP model. This highlights the reliability and versatility of our model when modeling data from various leagues, seasons, or sports, where the level of dispersion may vary.

In our analyses of the soccer and baseball data, we found that the HA in soccer was negatively affected during the COVID-19 pandemic but subsequently recovered as restrictions were lifted. In contrast, we found that the HA in baseball, in general, is smaller than in soccer, supporting the various theories that have been put forth to explain this phenomenon, including cultural differences and the behavior of fans in stadiums. It is worth noting that in baseball, if the home team is winning after eight and a half innings, then the home team is excluded from batting in the bottom half of the ninth inning. This particular rule could slightly deflate the estimated HA in baseball. It is also important to note that the model's performance can be sensitive to the choice of random effects, particularly the variance components. Therefore, careful consideration should be given to selecting appropriate priors that reflect prior knowledge or assumptions about the variability of the random effects.

In summary, our model offers a flexible framework that can be applied to various settings in sport and beyond sports, enabling the modeling of different types of count data and the potential inclusion of more than two response variables. The code for implementing the model is available (GitHub repository: <https://github.com/mauroflorez/CMP> and as the R package *MultRegCMP* at CRAN).

References

- Backlund, J., & Johdet, N. (2018). A Bayesian approach to predict the number of soccer goals: Modeling with Bayesian Negative Binomial regression. *Linköping University, The Division of Statistics and Machine Learning. (Dissertation)*. Available at: <https://urn.kb.se/resolve?urn=urn:nbn:se:liu:diva-149028>
- Baio, G., & Blangiardo, M. (2010). Bayesian hierarchical model for the prediction of football results. *Journal of Applied Statistics*, 37(2), 253-264.
- Benson, A., & Friel, N. (2021). Bayesian inference, model selection and likelihood estimation using fast rejection sampling: the Conway-Maxwell-Poisson distribution. *Bayesian Analysis*, 16(3), 905-931.
- Benz, L. S., & Lopez, M. J. (2021). Estimating the change in soccer's home advantage during the COVID-19 pandemic using bivariate Poisson regression. *AStA Advances in Statistical Analysis*, 1-28.
- Boshnakov, G., Kharrat, T., & McHale, I. G. (2017). A bivariate Weibull count model for forecasting association football scores. *International Journal of Forecasting*, 33(2), 458-466.
- Brooks, S. P., & Gelman, A. (1998). General methods for monitoring convergence of iterative simulations. *Journal of computational and graphical statistics*, 7(4), 434-455.
- Chanialidis, C., Evers, L., Neocleous, T., & Nobile, A. (2018). Efficient Bayesian inference for COM-Poisson regression models. *Statistics and computing*, 28(3), 595-608.
- Chiu, Y., & Chang, C. (2022). Major League Baseball during the COVID-19 pandemic: does a lack of spectators affect home advantage? *Humanities and Social Sciences Communications*, 9(1), 1-6.
- Conway, R. W., & Maxwell, W. L. (1962). A queuing model with state dependent service rates. *Journal of Industrial Engineering*, 12(2), 132-136.
- Dixon, M. J., & Coles, S. G. (1997). Modelling Association Football Scores and Inefficiencies in the Football Betting Market. *Journal of the Royal Statistical Society: Series C (Applied Statistics)*, 46(2), 265-280.
- Fedrizzi, G., Canal, L., & Micciolo, R. (2022). UEFA EURO 2020: a pure game of chance? *arXiv preprint arXiv:2203.07531*.
- Guikema, S. D., & Goffelt, J. P. (2008). A flexible count data regression model for risk analysis. *Risk Analysis: An International Journal*, 28(1), 213-223.
- Higgs, N., & Stavness, I. (2021). Bayesian analysis of home advantage in North American professional sports before and during COVID-19. *Scientific Reports*, 11(1), 1-11.
- Jones, M. B. (2015). The Home Advantage in Major League Baseball. *Perceptual and motor skills*, 121(3), 791-804.
- Karlis, D., & Ntzoufras, I. (2000). On modelling soccer data. *Student*, 3(4), 229-244.
- Karlis, D., & Ntzoufras, I. (2003). Analysis of sports data by using bivariate Poisson models. *Journal of the Royal Statistical Society: Series D (The Statistician)*, 52(3), 381-393.
- Karlis, D., & Ntzoufras, I. (2009). Bayesian modelling of football outcomes: using the Skellam's distribution for the goal difference. *IMA Journal of Management Mathematics*, 20(2), 133-145.
- Kleiber, C., & Zeileis, A. (2016). Visualizing count data regressions using rootograms. *The American Statistician*, 70(3), 296-303.
- Kramer, D. (2022). 3 reasons for Seattle's recent surge. *MLB*. Retrieved from <https://www.mlb.com/news/mariners-playoff-odds-surgin>. Accessed: 2024-4-18.
- Lee, A. J. (1997). Modeling scores in the Premier League: is Manchester United really the best? *Chance*, 10(1), 15-19.
- Losak, J. M., & Sabel, J. (2021). Baseball Home Field Advantage Without Fans in the Stands. *International Journal of Sport Finance*, 16(3).
- Lopez, M. J. (2016). Persuaded under pressure: Evidence from the National Football League. *Economic Inquiry*, 54(4), 1763-1773.
- Maher, M. J. (1982). Modelling association football scores. *Statistica Neerlandica*, 36(3), 109-118.
- McCarrick, D., Bilalic, M., Neave, N., & Wolfson, S. (2021). Home Advantage during the COVID-19 Pandemic in European football. *Psychology Of Sport And Exercise*. 56 pp. 102013.
- McHale, I., & Scarf, P. (2011). Modelling the dependence of goals scored by opposing teams in international soccer matches. *Statistical Modelling*, 11(3), 219-236.
- Murray, I., Ghahramani, Z., & MacKay, D. (2012). MCMC for doubly-intractable distributions. In *Proceedings of the Twenty-Second Conference on Uncertainty in Artificial Intelligence* (pp. 359-366).
- Payne, E. H., Gebregziabher, M., Hardin, J. W., Ramakrishnan, V., & Egede, L. E. (2018). An empirical approach to determine a threshold for assessing overdispersion in Poisson and negative binomial models for count data. *Communications in statistics: Simulation and computation*, 47(6), 1722-1738.
- Pettersson-Lidbom, P., & Priks, M. (2010). Behavior under social pressure: empty Italian stadiums and referee bias. *Economics Letters*, 108(2), 212-214.
- Piancastelli, L. S., Friel, N., Barreto-Souza, W., & Ombao, H. (2023). Multivariate Conway-Maxwell-Poisson Distribution: Sarmanov Method and Doubly-Intractable Bayesian Inference. *Journal of Computational and Graphical Statistics*, 32(2), 483-500.

- Price, K., Cai, H., Shen, W., & Hu, G. (2022). How Much Does Home Field Advantage Matter in Soccer Games? A Causal Inference Approach for English Premier League Analysis. *arXiv preprint arXiv:2205.07193*.
- Reade, J., Schreyer, D., & Singleton, C. (2022). Eliminating supportive crowds reduces referee bias. *Economic Inquiry*, 60(3), 1416–1436.
- Reep, C., Pollard, R., & Benjamin, B. (1971). Skill and chance in ball games. *Journal of the Royal Statistical Society: Series A (General)*, 134(4), 623-629.
- Shmueli, G., Minka, T. P., Kadane, J. B., Borle, S., & Boatwright, P. (2005). A useful distribution for fitting discrete data: revival of the Conway-Maxwell-Poisson distribution. *Journal of the Royal Statistical Society: Series C (Applied Statistics)*, 54(1), 127-142.
- Spiegelhalter, D. J., Best, N. G., Carlin, B. P., & Van Der Linde, A. (2002). Bayesian Measures of Model Complexity and Fit. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 64(4), 583-639.
- Tilp, M., & Thaller, S. (2020). COVID-19 Has Turned Home Advantage Into Home Disadvantage in the German Soccer Bundesliga. *Frontiers in sports and active living*, 2, 593499.
- Thomas, R. (2019). West Ham are better when they don't have the ball, which is why they're thriving away from home. *The Athletic*. Available at: <https://theathletic.com/1467224/2019/12/19/west-ham-are-better-when-they-dont-have-the-ball-which-is-why-theyre-thriving-away-from-home/>. Accessed: 2024-4-18.
- Vihola, M. (2012). Robust adaptive Metropolis algorithm with coerced acceptance rate. *Statistics and Computing*, 22, 997-1008.

Appendix

A. Proof of Equation (3)

Proof. Let $\delta_i = (\delta_{i1}, \delta_{i2})$, where $\delta_{ij} = \exp(b_{ij})$, then $\delta_i = \exp(b_i) \sim LN_2(\mu, \Sigma)$ with $\mu = \exp(0.5 \cdot \text{diag}(D))$, and $\Sigma = \text{diag}(\mu)(\exp(D) - \mathbf{1}\mathbf{1}^T)\text{diag}(\mu)$. Also, assume that $\lambda_{ij} = \exp(x_{ij}^T \beta_j)$, then $y_{ij} | \lambda_{ij}, \delta_{ij}, \nu_{ij} \sim \text{CMP}(\lambda_{ij} \delta_{ij}, \nu_{ij})$. Finally, let $\tilde{\lambda}_{ij} = \lambda_{ij} \mu$, $\tilde{\lambda}_i = (\tilde{\lambda}_{i1}, \tilde{\lambda}_{i2})$, and $\tilde{\Lambda}_i = \text{diag}(\tilde{\lambda}_i)$. By law of iterated expectations and using the asymptotic approximations of Shmueli et al. (2005), we have that $E(y_i | \beta, \gamma, D) \approx \tilde{\lambda}_i + \frac{1}{2\nu_i} - \frac{1}{2}$, and $\text{var}(y_i | \beta, D) \approx V_i^{-1} \tilde{\Lambda}_i (\exp(D) - \mathbf{1}\mathbf{1}^T) \tilde{\Lambda}_i$, $V_i = \text{diag}(\nu_{i1}, \nu_{i2})$. Thus

$$\begin{aligned} \text{cov}(y_{i1}, y_{i2}) &\approx \tilde{\lambda}_{i1} (\exp(d_{12}) - 1) \tilde{\delta}_{i2} \\ &\approx \lambda_{i1} \exp(0.5d_{11}) (\exp(d_{12}) - 1) \lambda_{i2} \exp(0.5d_{22}). \end{aligned}$$

Clearly, (3) can be positive or negative depending on the sign of d_{12} , i.e, the non-diagonal element of D . \square

B. Prior Sensitivity analysis

We assess the performance of the bivariate Conway-Maxwell-Poisson (CMP) model under various prior specifications. Specifically, we focus on the Over-dispersed scenario of 1 Season discussed in Section 3 and examine different combinations of values for B_0 , G_0 , ν_0 , and R_0 . These combinations are summarized in Table 5. To monitor convergence, we calculate the multivariate potential scale reduction factor (\hat{R}) Brooks and Gelman (1998) using 3 chains and also check the sampler's efficiency by calculating the effective sample size (ESS).

Tab. 5: Prior sensitivity: values considered in each scenario to assess the sensitivity of the model's performance to the specification of prior hyperparameters

Scenarios	A	B	C	D
B_0	0.1I	I	3I	10I
G_0	0.1I	I	3I	10I
ν_0	30	10	10	5
R_0	I	I	0.1I	0.1I

Moving on to Table 6, we present the mean squared error (MSE) values for μ_j across all scenarios. Consistent with expectations, scenario A demonstrates the best recovery of the observed values. Again, this outcome can be attributed to the small variance of the priors employed in this scenario. Conversely, scenarios C and D exhibit the poorest performance in terms of recovering the observed values.

Tab. 6: Mean Squared Error (MSE) of μ - Prior sensitivity analysis

Scenarios	A	B	C	D
μ_1	0.71	1.01	1.19	1.2
μ_2	0.25	0.25	0.22	0.545

Scenarios A and B exhibit more stable estimations, with all \hat{R} values below 1.2 (see Table 7). Indeed, some of the trace plots for chains in scenarios C and D demonstrated some mild convergence issues (plots not shown). Additionally, the effective sample size (ESS) analysis in Table 8 indicates higher algorithm efficiency in scenario A, while scenarios C and D show lower efficiency. These findings are likely connected to the

Tab. 7: Prior sensitivity analysis: multivariate potential scale factor (\hat{R})

Scenarios	A		B		C		D	
\hat{R}	y_1	y_2	y_1	y_2	y_1	y_2	y_1	y_2
β	1.07	1.06	1.15	1.16	1.2	1.21	4.22	1.82
γ	1.05	1.05	1.19	1.17	1.22	1.24	4.74	5.04
b	1.02	1.02	1.03	1.03	1.22	1.24	4.22	3.91

Tab. 8: Prior sensitivity analysis: Effective Sample Size (ESS)

Scenarios	A		B		C		D	
\hat{R}	y_1	y_2	y_1	y_2	y_1	y_2	y_1	y_2
β	350.44	408.54	123.25	131.13	98.48	97.13	7.26	10.71
γ	506.3	511.11	195.37	192.4	106.45	111.09	47.51	43.48
b	9366	9555.62	9175.4	9286.9	2324.947	2351.463	2306.58	2329.84

specified values of ν_0 and R_0 , as well as the larger covariance matrices B_0 and G_0 . Notably, the chains for random effects display larger \hat{R} values and smaller ESS in scenarios C and D compared to scenarios A and B. This discrepancy may be attributed to the expected value of the inverse of the random effects' covariance matrix D , which is determined by $\nu_0 R_0$. Consequently, larger values are permitted for the random effects in scenarios C and D, despite the expectation of low heterogeneity in this simulation scenario.

C. Shape parameters γ

In Figure 18, we observe the shape parameters estimation for the Premier League data analysis. In the left sub-figure we have the estimation for parameters associated with the Home Goals. Strong teams, such as Chelsea, Liverpool, and Manchester City, are located below 0, indicating that when other teams play against them, the dispersion parameter will be lower, implying a larger variance in their scored goals (over-dispersion). We observe a similar effect analyzing the Away Goals. Conversely, weaker teams such as Norwich or Brighton have their shape parameter above 0, indicating that when teams play against them, the goals will be more under-dispersed. On the other hand, in the x-axis, Manchester United and Manchester City are below 0, which means that the number of goals they score tends to be more over-dispersed compared to the other teams. We see the opposite trend with Liverpool, where the goals scored tend to be more under-dispersed. We can compare this with respect to over or under-dispersion because the intercept effect was 0. It is noticeable that the effects are different in the Home Goals and Away scores, supporting the modeling approach we assumed.

Similarly, we plot the estimations for the MLB analysis in Figure 19. One of the noticeable cases is Miami Marlins (MIA), with a negative effect on the x-axis in the Home Points and a positive effect in the Away Points. This means that at Home, MIA tends to have scores more dispersed, while Away, their scores will be less dispersed in comparison to the others. This highlights the ability of our model to adapt to data of different types and to accommodate different phenomena and mechanics observed in the different teams and sports.

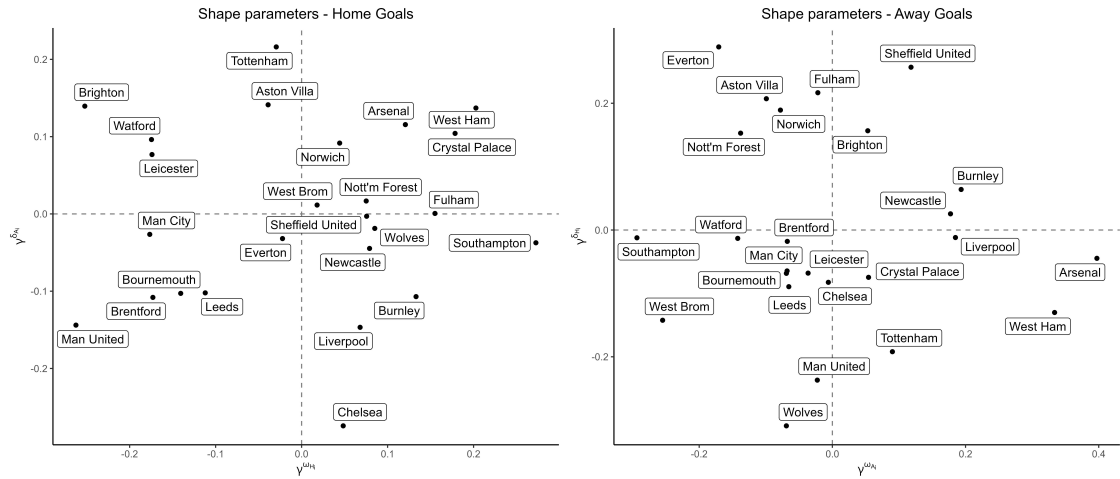


Fig. 18: Analysis of Premier League data: Shape parameters




Fig. 19: Analysis of MLB data: Shape parameters