



Estimating the change in soccer's home advantage during the Covid-19 pandemic using bivariate Poisson regression

Luke S. Benz¹ · Michael J. Lopez²

Received: 29 December 2020 / Accepted: 15 July 2021
© Springer-Verlag GmbH Germany, part of Springer Nature 2021

Abstract

In wake of the Covid-19 pandemic, 2019–2020 soccer seasons across the world were postponed and eventually made up during the summer months of 2020. Researchers from a variety of disciplines jumped at the opportunity to compare the rescheduled games, played in front of empty stadia, to previous games, played in front of fans. To date, most of this post-Covid soccer research has used linear regression models, or versions thereof, to estimate potential changes to the home advantage. However, we argue that leveraging the Poisson distribution would be more appropriate and use simulations to show that bivariate Poisson regression (Karlis and Ntzoufras in *J R Stat Soc Ser D Stat* 52(3):381–393, 2003) reduces absolute bias when estimating the home advantage benefit in a single season of soccer games, relative to linear regression, by almost 85%. Next, with data from 17 professional soccer leagues, we extend bivariate Poisson models estimate the change in home advantage due to games being played without fans. In contrast to current research that suggests a drop in the home advantage, our findings are mixed; in some leagues, evidence points to a decrease, while in others, the home advantage may have risen. Altogether, this suggests a more complex causal mechanism for the impact of fans on sporting events.

Keywords Bivariate Poisson · Soccer · Home advantage · Covid-19

✉ Luke S. Benz
lukesbenz@gmail.com

Michael J. Lopez
Michael.Lopez@nfl.com

¹ Medidata Solutions, Inc., New York, USA

² National Football League, Skidmore College, Saratoga Springs, USA

1 Introduction

Why do home teams in sport win more often than visiting teams? Researchers from, among other disciplines, psychology, economics, and statistics, have long been trying to figure that out.¹

One popular suggestion for the home advantage (HA) is that fans impact officiating (Moskowitz and Wertheim 2012). Whether it is crowd noise (Unkelbach and Memmert 2010), duress (Buraimo et al. 2010; Lopez 2016), or the implicit pressure to appease (Garicano et al. 2005), referee decision-making appears cued by factors outside of the run of play. If those cues tend to encourage officials to make calls in favor of the home team, it could account for some (or all) of the benefit that teams receive during home games.

A unique empirical approach for understanding HA contrasts games played in empty stadia to those played with fans, with the goal of teasing out the impact that fans have on HA. If fans impact referee decision making, it stands that an absence of fans would decrease HA. As evidence, both Pettersson-Lidbom and Priks (2010) (using Italian soccer in 2007) and Reade et al. (2020) (two decades of European soccer) found that games without fans resulted in a lower HA.

The coronavirus (Covid-19) pandemic resulted in many changes across sport, including the delay of most 2019–2020 soccer seasons. Beginning in March of 2020, games were put on pause, eventually made up in the summer months of 2020. Roughly, the delayed games account for about a third of regular season play. Make-up games were played as “ghost games”—that is, in empty stadia—as the only personnel allowed at these games were league, club, and media officials. These games still required visiting teams to travel and stay away from home as they normally would, but without fans, they represent a natural experiment with which to test the impact of fans on game outcomes.

Within just a few months of these 2020 “ghost games”, more than 10 papers have attempted to understand the impact that eliminating fans had on game outcomes, including scoring, fouls, and differences in team performances. The majority of this work used linear regression to infer causal claims about changes to HA. By and large, research overwhelmingly suggests that the home advantage decreased by a significant amount, in some estimates by an order of magnitude of one-half (McCarrick et al. 2020). In addition, most results imply that the impact of no fans on game outcomes is homogeneous with respect to league.

The goal of our paper is to expand the bivariate Poisson model (Karlis and Ntzoufras 2003) in order to tease out any impact of the lack of fans on HA. The benefits of our approach are plentiful. First, bivariate Poisson models consider home and visitor

¹ Initial research into the home advantage included, among other sources, Schwartz and Barsky (1977), Courneya and Carron (1992) and Nevill and Holder (1999). Works in psychology (Agnew and Carron 1994; Unkelbach and Memmert 2010), economics (Forrest et al. 2005; Dohmen and Sauermann 2016), and statistics (Buraimo et al. 2010; Lopez et al. 2018) are also recommended.

outcomes simultaneously. This helps account for correlations in outcomes (i.e., if the home team has more yellow cards, there is a tendency for the away team to also have more cards), and more accurately accounts for the offensive and defensive skill of clubs (Thompson 2018). We simulate soccer games at the season level and compare regression models (including bivariate Poisson) with respect to home advantage estimates. We find that the mean absolute bias in estimating a home advantage when using linear regression models is about six times larger when compared to bivariate Poisson. Second, we separate out each league when fit on real data, in order to pick up on both (i) inherent differences in each league's HA and (ii) how those differences are impacted by "ghost games." Third, we use a Bayesian version of the bivariate Poisson model, which allows for probabilistic interpretations regarding the likelihood that HA decreased within each league. Fourth, in modeling offensive and defensive team strength directly in each season, we can better account for scheduling differences pre- and post-Covid with respect to which teams played which opponents. Altogether, findings are inconclusive regarding a drop in HA post-Covid. While in several leagues a drop appears more than plausible, in other leagues, HA actually increases.

The remainder of this paper is outlined as follows. Section 2 reviews post-Covid findings, and 3 describes our implementation of the bivariate Poisson model. Section 4 uses simulation to motivate the use of bivariate Poisson for soccer outcomes, Sects. 5 and 6 explore the data and results, and Sect. 7 concludes.

2 Related literature

To date, we count 11 efforts that have attempted to estimate post-Covid changes to soccer's HA. The estimation of changes to HA has varied in scope (the number of leagues analyzed ranges from 1 to 41), method, and finding. Table 1 summarizes these papers, highlighting the number of leagues compared, whether leagues were treated together or separately, methodology (split into linear regression or correlation based approaches), and overview of finding. For clarity, we add a row highlighting the contributions of this manuscript.

Broadly, methods consider outcome variables Y as a function of T and T' , the home advantages pre- and post-Covid, respectively, as well as W , where W possibly includes game and team characteristics. Though it is infeasible to detail choice of W and Y across each of the papers, a few patterns emerge.

Several papers consider team strength, or proxies thereof, as part of W . This could include fixed effects for each team (Ferraresi et al. 2020; Cueva 2020; Bryson et al. 2020), other proxies for team strength (McCarrick et al. 2020; Fischer and Haucap 2020b; Krawczyk et al. 2020), and pre-match betting odds (Endrich and Gesche 2020). The Cueva (2020) research is expansive and includes 41 leagues across 30 countries, and likewise finds significant impacts on home and away team fouls, as

Table 1 Comparison of post-Covid research on home advantage in football

Paper	Leagues	Method	Finding
Sors et al. (2020)	8 (Together)	Correlation	Drop in HA
Leitner and Richlan (2020b)	8 (Together)	Correlations	Drop in HA
Endrich and Gesche (2020)	2 (Together)	Linear Regression	Drop in HA
Fischer and Haucap (2020b)	3 (Separate)	Linear Regression	Mixed
Dilger and Vischer (2020)	1 (NA)	Linear Regression, Correlations	Drop in HA
Krawczyk et al. (2020)	4 (Separate)	Linear Regression	Mixed
Ferraresi et al. (2020)	5 (Together)	Linear Regression	Drop in HA
Reade et al. (2020)	7 (Together)	Linear Regression	Drop in HA
Jiménez Sánchez and Lavín (2020)	8 (Separate)	Linear Regression, Correlations	Mixed
Scoppa (2020)	10 (Together)	Linear Regression	Drop in HA
Cueva (2020)	41 (Together)	Linear Regression	Drop in HA
McCarrick et al. (2020)	15 (Together)	Linear, Poisson Regression	Drop in HA
Bryson et al. (2020)	17 (Together)	Linear, Poisson Regression	Mixed
Benz and Lopez (this manuscript)	17 (Separate)	Bivariate Poisson Regression	Mixed

HA home advantage. Correlation-based approaches include Chi-square and Mann–Whitney tests. Linear regression includes univariate OLS-based frameworks and *t*-tests. Poisson regression assumes univariate Poisson. Papers are sorted by method and number of leagues. Note that this manuscript is the first paper among those listed which employs a Bayesian framework for model fitting

well as foul differential. Other pre-match characteristics in *W* include if the game is a rivalry and team travel (Krawczyk et al. 2020), as well as match referee and attendance (Bryson et al. 2020).

Choices of *Y* include metrics such as goals, goal differential, points (3/1/0), yellow cards, yellow card differential, whether or not each team won, and other in-game actions such as corner kicks and fouls. Several authors separately develop models for multiple response variables. Linear regression and versions of these models including *t*-tests stand out most common approaches for modeling *Y*. This includes models for won/loss outcomes (Cueva 2020), goal differential (Bryson et al. 2020; Krawczyk et al. 2020), and fouls (Scoppa 2020). Two authors model goals with Poisson regression (McCarrick et al. 2020; Bryson et al. 2020). McCarrick et al. (2020) used univariate Poisson regression models of goals, points and fouls, finding that across the entirety of 15 leagues, the home advantage dropped from 0.29 to 0.15 goals per game, while Bryson et al. (2020) found a significant drop in yellow cards for the away team using univariate Poisson regression.

In addition to choice of *Y*, *W*, and method, researchers have likewise varied with the decision to treat each league separate. As shown in Table 1, all but three papers have taken each of the available leagues and used them in a single statistical model. Such an approach boasts the benefit of deriving an estimated change in HA that can be broadly applied across soccer, but requires assumptions that (i) HA is

homogeneous between leagues and (ii) differences in HA post-Covid are likewise equivalent.

Our approach will make two advances that none of the papers in Table 1 can. First, we model game outcomes using an expanded version of the bivariate Poisson regression model, one originally designed for soccer outcomes (Karlis and Ntzoufras 2003). This model allows us to control for team strength, account for and estimate the correlation in game outcomes, and better model ties. Second, we will show that the assumption of a constant HA between leagues is unjustified. In doing so, we highlight that the frequent choice of combining leagues into one uniform model has far-reaching implications with respect to findings.

3 Methods

Poisson regression models assume the response variable has a Poisson distribution, and models the logarithm of response as a linear combination of explanatory variables.

Let Y_{Hi} and Y_{Ai} be outcomes observed in game i for the home (H_i) and away teams (A_i), respectively. For now we assume Y_{Hi} and Y_{Ai} are goals scored, but will likewise apply a similar framework to yellow cards. The response (Y_{Hi}, Y_{Ai}) is bivariate Poisson with parameters $\lambda_{1i}, \lambda_{2i}, \lambda_{3i}$ if

$$(Y_{Hi}, Y_{Ai}) = BP(\lambda_{1i}, \lambda_{2i}, \lambda_{3i}), \tag{1}$$

where $\lambda_{1i} + \lambda_{3i}$ and $\lambda_{2i} + \lambda_{3i}$ are the goal expectations of Y_{Hi} and Y_{Ai} , respectively, and λ_{3i} is the covariance between Y_{Hi} and Y_{Ai} . As one specification, let

$$\begin{aligned} \log(\lambda_{1i}) &= \mu_{ks} + T_k + \alpha_{H,ks} + \delta_{A,ks}, \\ \log(\lambda_{2i}) &= \mu_{ks} + \alpha_{A,ks} + \delta_{H,ks}, \\ \log(\lambda_{3i}) &= \gamma_k. \end{aligned} \tag{2}$$

In Model (2), μ_{ks} is an intercept term for expected goals in season s (which we assume to be constant), T_k is a home advantage parameter, and γ_k is a constant covariance, all of which correspond to league k . The explanatory variables used to model λ_{1i} and λ_{2i} correspond to factors likely to impact the home and away team's goals scored, respectively. Above, λ_{1i} is a function of the home team's attacking strength ($\alpha_{H,ks}$) and away team's defending strength ($\delta_{A,ks}$), while λ_{2i} is a function of the away team's attacking strength ($\alpha_{A,ks}$) and home team's defending strength ($\delta_{H,ks}$), all corresponding to league k during season s . For generality, we refer to α_{ks} and δ_{ks} as general notation for attacking and defending team strengths, respectively. In using μ_{ks} , α_{ks} and δ_{ks} are seasonal effects, centered at 0, such that $\alpha_{ks} \sim N(0, \sigma_{att,k}^2)$ and $\delta_{ks} \sim N(0, \sigma_{def,k}^2)$.

If $\lambda_3 = 0$ in Eq. 2, then $Y_H \perp\!\!\!\perp Y_A$, and the bivariate Poisson reduces to the product of two independent Poisson distributions. Using observed outcomes in soccer from

1991, Karlis and Ntzoufras (2003) found that assuming independence of the Poisson distributions was less suitable for modeling ties when compared to using bivariate Poisson. More recently, however, Groll et al. (2018) suggest using $\lambda_3 = 0$, as there are now fewer ties when compared to 1991. Structural changes to professional soccer—leagues now reward three points for a win and one point for a tie, instead of two points for a win and one point for a tie—are likely the cause, and thus using $\lambda_3 = 0$ in models of goal outcomes is more plausible.

There are a few extensions of bivariate Poisson to note. Karlis and Ntzoufras (2003) propose diagonally inflated versions of Model (1) and also included team indicators for both home and away teams in λ_3 , in order to test if the home or away teams controlled the amount of covariance in game outcomes. However, models fit on soccer goals did not warrant either of these additional parameterizations. Baio and Blangiardo (2010) use a Bayesian version bivariate Poisson that explicitly incorporates shrinkage to team strength estimates. Additionally, Koopman and Lit (2015) allows for team strength specifications to vary stochastically within a season, as in a state-space model (Glickman and Stern 1998). Though Model's (2) and (3) cannot pick up team strengths that vary within a season, estimating these trends across 17 leagues could be difficult to scale; Koopman and Lit (2015), for example, looked only at the English Premier League. Inclusion of time-varying team strengths, in addition to an assessment of team strengths post-Covid versus pre-Covid, is an opportunity for future work.

3.1 Extending bivariate Poisson to changes in the home advantage

3.1.1 Goal outcomes

We extend Model (2) to consider post-Covid changes in the HA for goals using Model (3).

$$\begin{aligned} (Y_{Hi}, Y_{Ai}) &= BP(\lambda_{1i}, \lambda_{2i}, \lambda_{3i}), \\ \log(\lambda_{1i}) &= \mu_{ks} + T_k \times (I_{\text{pre-Covid}}) + T'_k \times (I_{\text{post-Covid}}) + \alpha_{H,ks} + \delta_{A,ks}, \\ \log(\lambda_{2i}) &= \mu_{ks} + \alpha_{A,ks} + \delta_{H,ks}, \\ \log(\lambda_{3i}) &= \gamma_k, \end{aligned} \quad (3)$$

where T'_k is the post-Covid home advantage in league k , and $I_{\text{pre-Covid}}$ and $I_{\text{post-Covid}}$ are indicator variables for whether or not the match took place before or after the restart date shown in Table 3. Of particular interest will be the comparison of estimates of T_k and T'_k .

3.1.2 Yellow card outcomes

A similar version, Model (4), is used for yellow cards. Let Z_{Hi} and Z_{Ai} be the yellow cards given to the home and away teams in game i . We assume Z_{Hi} and Z_{Ai} are bivariate Poisson such that

$$\begin{aligned}
 (Z_{Hi}, Z_{Ai}) &= BP(\lambda_{1i}, \lambda_{2i}, \lambda_{3i}), \\
 \log(\lambda_{1i}) &= \mu_{ks} + T_k \times (I_{\text{pre-Covid}}) + T'_k \times (I_{\text{post-Covid}}) + \tau_{H,ks}, \\
 \log(\lambda_{2i}) &= \mu_{ks} + \tau_{A,ks}, \\
 \log(\lambda_{3i}) &= \gamma_k,
 \end{aligned}
 \tag{4}$$

where $\tau_{ks} \sim N(0, \sigma_{\text{team},k}^2)$. Implicit in Model (4), relative to Models (2) and (3), is that teams control their own yellow card counts, and not their opponents', and that tendencies for team counts to correlate are absorbed in λ_{3i} .

3.1.3 Model fits in Stan

We use Stan, an open-source statistical software designed for Bayesian inference with MCMC sampling, for each league k , and with models for both goals and yellow cards. We choose Bayesian MCMC approaches over the EM algorithm (Karlis and Ntzoufras 2003; Karlis et al. 2005) to obtain both (i) posterior distributions of the change in home advantage and (ii) posterior probabilities that home advantage declined in each league. No paper referenced in Table 1 assessed HA change probabilistically.

We fit two versions of Models (3) and (4), one with $\lambda_3 = 0$, and a second with $\lambda_3 > 0$. For models where $\lambda_3 = 0$, prior distributions for the parameters in Models (3) and (4) are as follows. These prior distributions are non-informative and do not impose any outside knowledge on parameter estimation.

$$\begin{aligned}
 \mu_{ks} &\sim N(0, 25), \\
 \alpha_{ks} &\sim N(0, \sigma_{\text{att},k}^2), \\
 \delta_{ks} &\sim N(0, \sigma_{\text{def},k}^2), \\
 \tau_{ks} &\sim N(0, \sigma_{\text{team},k}^2), \\
 \sigma_{\text{att},k} &\sim \text{Inverse-Gamma}(1, 1), \\
 \sigma_{\text{def},k} &\sim \text{Inverse-Gamma}(1, 1), \\
 \sigma_{\text{team},k} &\sim \text{Inverse-Gamma}(1, 1), \\
 T_k &\sim N(0, 25), \\
 T'_k &\sim N(0, 25)
 \end{aligned}$$

For models w/ $\lambda_3 > 0$, empirical Bayes priors were used for T_k, T'_k in order to aid in convergence. Namely, let \hat{T}_k and \hat{T}'_k be the posterior mean estimate of pre-Covid and post-Covid HA for from league k , respectively, from the corresponding model with $\lambda_3 = 0$. We let

$$\begin{aligned}
 \bar{T} &= \text{mean}(\{\hat{T}_1, \dots, \hat{T}_{17}\}) \\
 \bar{T}' &= \text{mean}(\{\hat{T}'_1, \dots, \hat{T}'_{17}\}) \\
 s &= 3 \times \text{SD}(\{\hat{T}_1, \dots, \hat{T}_{17}\}) \\
 s' &= 3 \times \text{SD}(\{\hat{T}'_1, \dots, \hat{T}'_{17}\})
 \end{aligned}$$

Priors T_K , T'_k and γ_k for the variants of Model (3) and (4) w/ $\lambda_3 > 0$ are as follows:

$$\begin{aligned} T_k &\sim N(\bar{T}, s^2) \\ T'_k &\sim N(\bar{T}', s'^2) \\ \gamma_k &\sim N\left(0, \frac{1}{2}\right) \text{ (Goals)} \\ \gamma_k &\sim N(0, 2) \text{ (Yellow Cards)} \end{aligned}$$

The priors on T_K and T'_k are weakly informative; the variance in the priors is 9 times as large as the variance in the observed variance in $\{\hat{T}_1, \dots, \hat{T}_{17}\}$ estimated in the corresponding $\lambda_3 = 0$ model variation. As γ_k represents the correlation term for goals/yellow cards, and exists on the log-scale, the priors are not particularly informative, and they allow for values of λ_3 that far exceed typical number of goals and yellow cards per game. Overall, our use of priors is not motivated by a desire to incorporate domain expertise, and instead the use of Bayesian modeling is to incorporate posterior probabilities as a tool to assess changes in HA.

For models with $\lambda_3 = 0$, Models (3) and (4) were fit using 3 parallel chains, each made up of 7000 iterations, and a burn in of 2000 draws. When $\lambda_3 > 0$ was assumed, Models (3) and (4) were fit using 3 parallel chains, each with 20,000 iterations, and a burn-in of 10,000 draws. Parallel chains were used to improve the computation time needed to draw a suitable number of posterior samples for inference. Posterior samples were drawn using the default Stan MCMC algorithm, Hamiltonian Monte Carlo (HMC) with No U-Turn Sampling (NUTS) (Stan Development Team 2019).

To check for model convergence, we examine the \hat{R} statistic (Gelman and Rubin 1992; Brooks and Gelman 1998) for each parameter. If \hat{R} statistics are near 1, that indicates convergence (Gelman et al. 2013). To check for the informativeness of a parameter's posterior distribution, we use effective sample size (ESS, Gelman et al. 2013), which uses the relative independence of draws to equate the posterior distribution to the level of precision achieved in a simple random sample.

For goals, we present results for Model (3) with $\lambda_3 = 0$, and for yellow cards, we present results with Model (4) and $\lambda_3 > 0$. Henceforth, any reference to those models assumes such specifications, unless explicitly stated otherwise. All data and code for running and replicating our analysis are found at https://github.com/lbenz730/soccer_ha_covid.

4 Simulation

4.1 Simulation overview

Most approaches for evaluating bivariate Poisson regression have focused on model fit (Karlis et al. 2005) or prediction. For example, Ley et al. (2019) found bivariate Poisson matched or exceeded predictions of paired comparison models, as judged

by rank probability score, on unknown game outcomes. Tsokos et al. (2019) also compared paired comparison models to bivariate Poisson, with a particular focus on methods for parameter estimation, and found the predictive performances to be similar. As will be our suggestion, Tsokos et al. (2019) treated each league separately to account for underlying differences in the distributions of game outcomes. Bivariate Poisson models have also compared favorably with betting markets (Koopman and Lit 2015).

We use simulations to better understand accuracy and operation characteristics of bivariate Poisson and other models in terms of estimating soccer's home advantage. There are three steps to our simulations; (i) deriving team strength estimates, (ii) simulating game outcomes under assumed home advantages, and (iii) modeling the simulated game outcomes to estimate that home advantage. Exact details of each of these three steps are shown in "Appendix"; we summarize here.

We derive team strength estimates to reflect both the range and correlations of attacking and defending estimates found in the 17 professional soccer leagues in our data. As in Thompson (2018), team strength estimates are simulated across single seasons of soccer using the bivariate normal distribution. To assess if the correlation of team strengths (abbreviated as ρ^*) effects home advantage estimates, we use $\rho^* \in \{-0.8, -0.4, 0\}$ (teams that typically score more goals also allow fewer goals).

Two data generating processes are used to simulate home and away goal outcomes. The first reflects Model (2), where goals are simulated under a bivariate Poisson distribution. The second reflects a bivariate normal distribution. Although bivariate Poisson is more plausible for soccer outcomes (Karlis and Ntzoufras 2003), using bivariate normal allows us to better understand how a bivariate Poisson model can estimate HA under alternative generating processes. For both data generating processes, we fix a simulated home advantage T^* , for $T^* \in \{0, 0.25, 0.5\}$, to roughly reflect ranges of goal differential benefits for being the home team, as found in Figure 1 of Bryson et al. (2020).

Three candidate models are fit. First, we use linear regression models of goal differential as a function of home and away team fixed effects and a term for the home advantage, versions of which were used by Bryson et al. (2020), Scoppa (2020), Krawczyk et al. (2020) and Endrich and Gesche (2020). Second, we use Bayesian paired comparison models, akin to Tsokos et al. (2019) and Ley et al. (2019), where goal differential is modeled as a function of differences in team strength, as well as the home advantage. Finally, we fit Model (2) with $\lambda_3 = 0$. Recall that when $\lambda_3 = 0$, the bivariate Poisson in Eq. 2 reduces to the product of two independent Poisson distributions. The $\lambda_3 = 0$ bivariate Poisson model variant was chosen for use in simulations given that such a choice has proven suitable for modeling goals outcomes in recent year (Groll et al. 2018), and furthermore the $\lambda_3 = 0$ variant of Model (3) will be presented in Sect. 6.1.2.

A total of 100 seasons were simulated for each combination of ρ^* and T^* using each of the two data generating process, for a total of 1800 simulated seasons worth of data.

Table 2 Mean absolute bias (MAB) and mean bias (MB) in 1800 estimates of the home advantage in a single season of soccer games between 20 teams, 100 for each combination of data generating process, team strength correlation (ρ^*) and home advantage (T^*)

Model	$\rho^* = -0.8$		$\rho^* = -0.4$		$\rho^* = 0$	
	MAB	MB	MAB	MB	MAB	MB
$T^* = 0$						
Data generating process: bivariate Poisson						
Bivariate Poisson	0.058	-0.005	0.051	-0.005	0.053	-0.003
Paired comparisons	0.065	-0.005	0.058	-0.005	0.059	-0.003
Linear regression	0.399	0.020	0.403	-0.090	0.382	-0.029
Data generating process: bivariate normal						
Bivariate Poisson	0.058	0.006	0.060	-0.010	0.061	0.007
Paired comparisons	0.059	0.006	0.061	-0.010	0.062	0.008
Linear regression	0.460	0.036	0.480	-0.070	0.446	0.032
$T^* = 0.25$						
Data generating process: bivariate Poisson						
Bivariate Poisson	0.061	0.001	0.061	0.000	0.064	0.015
Paired comparisons	0.075	0.034	0.075	0.034	0.082	0.049
Linear regression	0.424	0.100	0.474	0.036	0.425	-0.054
Data generating process: bivariate normal						
Bivariate Poisson	0.073	-0.019	0.068	-0.017	0.084	-0.015
Paired comparisons	0.074	-0.015	0.068	-0.013	0.085	-0.010
Linear regression	0.485	-0.070	0.454	0.070	0.427	-0.006
$T^* = 0.5$						
Data generating process: bivariate Poisson						
Bivariate Poisson	0.065	0.001	0.072	-0.012	0.071	-0.004
Paired comparisons	0.094	0.069	0.091	0.047	0.089	0.056
Linear regression	0.453	0.138	0.485	0.036	0.529	0.083
Data generating process: bivariate normal						
Bivariate Poisson	0.070	-0.021	0.067	0.007	0.063	-0.004
Paired comparisons	0.070	-0.015	0.069	0.013	0.063	0.002
Linear regression	0.549	0.060	0.450	-0.021	0.549	-0.042

Estimates produced using linear regression, paired comparison, and bivariate Poisson regression models. The mean absolute bias for bivariate Poisson regression compares favorably; when the data generating process of goal outcomes is bivariate Poisson, bivariate Poisson models most accurately estimate the home advantage. Furthermore, when the data generating process of goal outcomes is bivariate normal, bivariate Poisson and paired comparison models perform similarly, with the bivariate Poisson model slightly more accurate

4.2 Simulation results

Table 2 shows mean absolute bias (MAB) and mean bias (MB) of home advantage estimates from each of the three candidate models (linear regression, paired

comparison, and bivariate Poisson) under the two data generating processes (bivariate Poisson and bivariate normal). Each bias is shown on the goal difference scale.

When goal outcomes are simulated using the bivariate Poisson distribution, bivariate Poisson model estimates of home advantage average an absolute bias of roughly 0.06–0.07, and range from 11 to 31% lower than estimated home advantages from paired comparison models. Furthermore, for large advantages of home advantage, the paired comparison is directionally biased and tends to overestimate home advantage.

Both bivariate Poisson and paired comparison models compare favorably to linear regression. The absolute biases from linear regression models vary between 0.40 and 0.53 and tend to increase with larger home advantages. More generally, when using these models across a full season's worth of soccer games, one could expect the estimate of the home advantage from a linear regression (with home and away team fixed effects) to be off by nearly half a goal (in unknown direction), which is about six times the amount of bias shown when estimating using bivariate Poisson.

When goal outcomes are simulated using the bivariate normal distribution, bivariate Poisson and paired comparison models capture the known home advantage with equivalent accuracy (mean absolute bias' within $\pm 3\%$, with bivariate Poisson slightly better). Linear regression performs poorly under these goal outcome models, with an average absolute bias range from 0.427 to 0.549).

Overall, there do not seem to be any noticeable patterns across ρ^* , the range of correlation between team strengths.

5 Data

The data used for this analysis are comprised of games from 17 soccer leagues in 13 European countries spanning 5 seasons between 2015 and 2020. The leagues selected for use in this analysis were among the first leagues to return to play following a suspension of the season to the Covid-19 pandemic. Typically, European countries have hierarchies of leagues (also referred to as divisions, tiers, or flights), with teams competing to be promoted to a better league and/or to avoid being relegated to the league below. For each of the 13 countries used in this analysis, the top league in that country was selected. Additionally, 2nd tier leagues were included for England, Spain, Italy and Germany, the countries among the “Big 5” European soccer to resume domestic competition (the final of the “Big 5” countries, France, cancelled the conclusion of its leagues' 2019–2020 seasons). Only games from intra-league competition were used in this analysis, and games from domestic inter-league cup competitions (such as England's FA Cup), and inter-country competitions (such as the UEFA Champions League), were dropped. A full summary of the leagues and games used in this paper is presented in Table 3.

Table 3 Breakdown of leagues used in analysis

League	Country	Tier	Restart date	Pre-Covid games	Post-Covid games	# of Team-seasons
German Bundesliga	Germany	1	2020-05-16	1448	82	90
German 2. Bundesliga	Germany	2	2020-05-16	1449	81	90
Danish Superliga	Denmark	1	2020-05-31	1108	74	68
Austrian Bundesliga	Austria	1	2020-06-02	867	63	54
Portuguese Liga	Portugal	1	2020-06-03	1440	90	90
Greek Super League	Greece	1	2020-06-06	1168	58	78
Spanish La Liga 2	Spain	2	2020-06-10	2233	129	110
Spanish La Liga	Spain	1	2020-06-11	1790	110	100
Turkish Super Lig	Turkey	1	2020-06-13	1460	70	90
Swedish Allsvenskan	Sweden	1	2020-06-14	960	198	80
Norwegian Eliteserien	Norway	1	2020-06-16	960	175	80
English Premier League	England	1	2020-06-17	1808	92	100
Italy Serie B	Italy	2	2020-06-17	2046	111	105
Swiss Super League	Switzerland	1	2020-06-19	836	65	50
Russian Premier Liga	Russia	1	2020-06-19	1136	64	80
English League Championship	England	2	2020-06-20	2673	113	120
Italy Serie A	Italy	1	2020-06-20	1776	124	100

Data consist of 5 most recent seasons between 2015 and 2020. # of games corresponds to sample sizes for goals model. Due to different levels of missingness between goals and yellow cards in the data, 5 leagues have a smaller # of games in their respective pre-Covid yellow card sample, while 1 league has a smaller # of games in its post-Covid yellow card sample. Restart date refers the date that the league resumed play after an interrupted 2019–2020 season or delayed start of 2020 season (Norway/Sweden)

Data were scraped from Football Reference (Sports Reference 2020) on 2020-10-28. For each league, the five most recent seasons worth of data were pulled, not including the ongoing 2020–2021 season. For 15 of the 17 leagues, this included the 2015–2016 season through the 2019–2020 season. Unlike the majority of European Soccer leagues, which run from August to May, the top flights in Sweden and Norway run from March to November. These leagues never paused an ongoing season due to the Covid-19 pandemic, but rather delayed the start of their respective 2020 seasons until June. As a result, the data used in this analysis are five full seasons worth of data for all the leagues outside of Sweden and Norway, while those two countries have four full seasons of data, plus all games in the 2020 season through 2020-10-28.

Throughout this analysis, we refer to pre-Covid and post-Covid samples. For each league, the pre-Covid sample constitutes all games prior to the league's restart date, listed in Table 3, while the post-Covid sample includes all games which comprised of all games on or after the league's restart date. In nearly all cases, the league's restart date represents a clean divide between games that had fans in attendance and games that did not have any fans in attendance due to Covid restrictions. One

exception is a German Bundesliga game between Borussia Monchengladbach and Cologne on 2020-03-11 that was played in an empty stadium just before the German Bundesliga paused its season. Additionally, seven games in Italy Serie A were played under the same circumstances. While leagues returned from their respective hiatuses without fans in attendance, some, such as the Danish Superliga, Russian Premier League, and Norwegian Eliteserien, began allowing very reduced attendance by the end of the sample.

Unfortunately, attendance numbers attained from Football Reference were not always available and/or accurate, and as such, we cannot systematically identify the exact number games in the sample that had no fans in attendance prior to the league suspending games, or the exact number of games in the post-Covid sample that had fans in attendance. Related, there are several justifications for using the pre-Covid/post-Covid sample split based on league restart date:

1. Any number of games in the pre-Covid sample without fans in attendance is minute compared to the overall size of any league's pre-Covid sample.
2. Several month layoffs with limited training are unprecedented, and possibly impact team strengths and player skill, which in turn may impact game results in the post-Covid sample beyond any possible change in home advantage.
3. Any games in a league's post-Covid sample played before fans have attendances significantly reduced compared to the average attendance of a game in the pre-Covid sample.
4. The majority of leagues do not have a single game in the post-Covid sample with any fans in attendance, while all leagues have games in the post-Covid sample without fans.

Recently started games in the 2020–2021 season are not considered, as leagues have diverged from one another in terms of off-season structure and policies allowing fans to return to the stands.

Games where home/away goals were unavailable were removed for the goals model, and games where home/away yellow cards were unavailable were removed for the yellow cards model. The number of games displayed in Table 3 reflects the sample sizes used in the goals model. The number of games where goal counts were available always matched or exceeded the number of games where yellow card counts were available. Across 5 leagues, 92 games from the pre-Covid samples used in Model (3) were missing yellow card counts, and had to be dropped when fitting Model (4) (2 in Italy Serie B, 2 in the English League Championship, 12 in the Danish Superliga, 34 in the Turkish Super Lig, and 42 in Spanish La Liga 2). 4 games had to be dropped from the Russian Premier's Leagues post-Covid sample for the same reason.

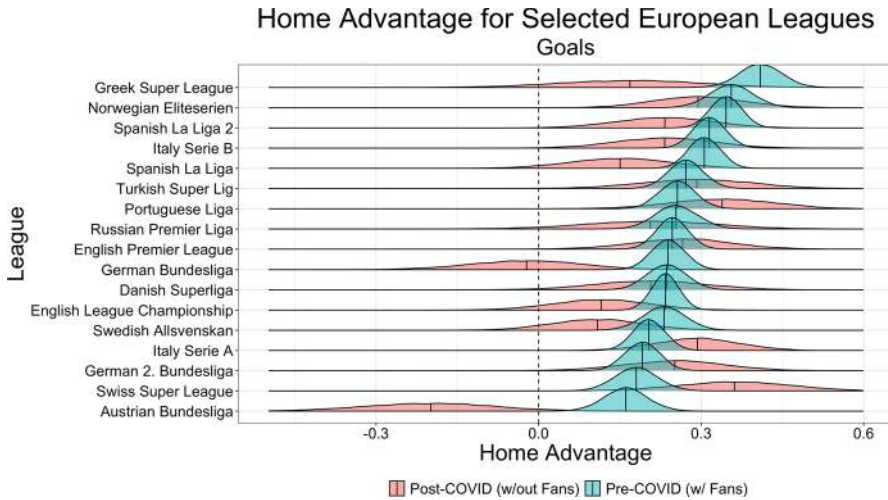


Fig. 1 Posterior distributions of T_k and T'_k , the pre-Covid and post-Covid HAs for goals. Larger values of T_k and T'_k indicate larger home advantages. Prior to the Covid-19 pandemic, the Greek Super League and Norwegian Eliteserien had the largest home advantages for goals, while the Austrian Bundesliga and Swiss Super League had the smallest home advantages for goals. Across the 17 leagues in the sample, a range of differences exist between posterior distributions of T_k and T'_k

6 Results

6.1 Goals

6.1.1 Model fit

Results from goals Model (3), using $\lambda_3 = 0$ for all leagues, are shown below. We choose Model (3) with $\lambda_3 = 0$ because, across our 17 leagues of data, the correlation in home and away goals per game varied between -0.16 and 0.07 .

Using this model, all \hat{R} statistics ranged from 0.9998 to 1.003 , providing strong evidence that the model properly converged. Additionally, the effective sample sizes are provided in Table 6. ESS are sufficiently large, especially HA parameters of interest T_k and T'_k , suggesting enough draws were taken to conduct inference.

Figure 3 (in “Appendix”) shows an example of posterior means of attacking (α_{ks}) and defensive (δ_{ks}) team strengths for one season of the German Bundesliga. In Fig. 3, the top team (Bayern Munich) stands out with top offensive and defensive team strength metrics. However, the correlation between offensive and defensive team strength estimates is weak, reflecting the need for models to incorporate both aspects of team quality.

Table 4 Comparison of posterior means for pre-Covid and post-Covid goals HA parameters from Model (3), \hat{T}_k and \hat{T}'_k , respectively

League	\hat{T}_k	\hat{T}'_k	$\hat{T}'_k - \hat{T}_k$	% Change	$P(T'_k < T_k)$
Austrian Bundesliga	0.161	- 0.202	- 0.363	- 225.7	0.999
German Bundesliga	0.239	- 0.024	- 0.263	- 110.2	0.995
Greek Super League	0.409	0.167	- 0.243	- 59.3	0.972
Spanish La Liga	0.306	0.149	- 0.157	- 51.3	0.959
English League Championship	0.234	0.114	- 0.119	- 51.1	0.912
Swedish Allsvenskan	0.231	0.108	- 0.123	- 53.3	0.907
Spanish La Liga 2	0.346	0.232	- 0.114	- 32.9	0.903
Italy Serie B	0.315	0.232	- 0.083	- 26.4	0.825
Norwegian Eliteserien	0.356	0.295	- 0.061	- 17.1	0.745
Russian Premier Liga	0.254	0.204	- 0.050	- 19.6	0.655
Danish Superliga	0.236	0.206	- 0.030	- 12.9	0.610
Turkish Super Lig	0.271	0.290	0.019	7.0	0.419
English Premier League	0.246	0.264	0.018	7.2	0.416
German 2. Bundesliga	0.191	0.249	0.058	30.5	0.266
Portuguese Liga	0.256	0.338	0.082	32.2	0.194
Italy Serie A	0.204	0.292	0.088	43.4	0.125
Swiss Super League	0.180	0.362	0.182	101.1	0.043

Larger values of T_k and T'_k indicate larger home advantages. Relative and absolute differences between \hat{T}'_k and \hat{T}_k are also shown. Probabilities of decline in HA without fans, $P(T'_k < T_k)$, are estimated from posterior draws. We estimate the probability of a decline in HA without fans to exceed 0.9 in 7 of 17 leagues, and to exceed 0.5 in 11 of 17 leagues

6.1.2 Home advantage

The primary parameters of interest in Model (3) are T_k and T'_k , the pre- and post-Covid home advantages for each league k , respectively. These HA terms are shown on a log-scale, and represent the additional increase in the home team's log goal expectation, relative to a league average (μ_{ks}), and after accounting for team and opponent (α_{ks} and δ_{ks}) effects.²

Posterior distributions for T_k and T'_k are presented in Fig. 1. Clear differences exist between several of the 17 leagues' posterior distributions of T_k . For example, the posterior mean of T_k in the Greek Super League is 0.409, or about 2.5 times that of the posterior mean in the Austrian Bundesliga (0.161). The non-overlapping density curves between these leagues add further support for our decision to estimate T_k separately for each league, as opposed to one T across all of Europe.

Table 4 compares posterior means of T_k (denoted \hat{T}_k) with those of T'_k (denoted \hat{T}'_k) for each of the 17 leagues. Posterior means for HA without fans are smaller than

² In our simulations in Sect. 4, we transformed HA estimates to the goal difference scale, in order to compare to estimates from linear regression.

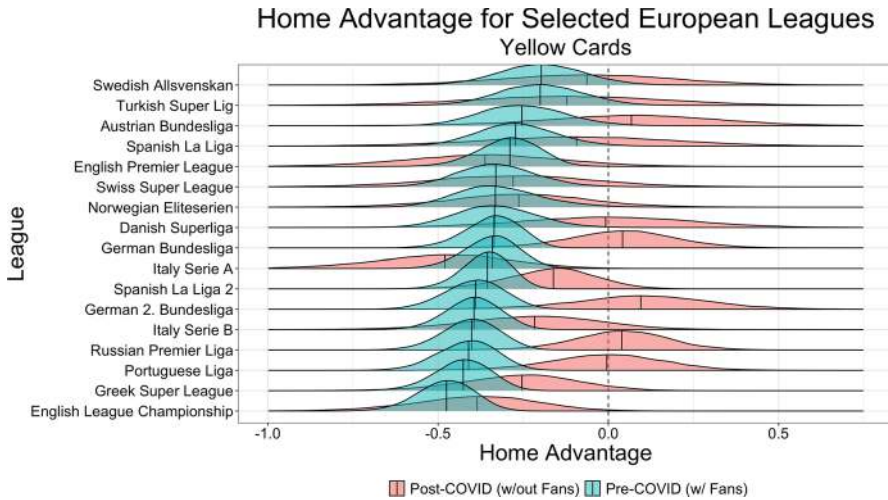


Fig. 2 Posterior distributions of T_k and T'_k , the pre-Covid and post-Covid HAs for yellow cards. Smaller (i.e., more negative) values of T_k and T'_k indicate larger home advantages. Prior to the Covid-19 pandemic, the English League Championship and Greek Super League had the largest home advantages for yellow cards, while the Swedish Allsvenskan and Turkish Super Lig had the smallest home advantages for yellow cards. Across the 17 leagues in the sample, a range of differences exist between posterior distributions of T_k and T'_k

the corresponding posterior mean of HA w/ fans ($\hat{T}'_k < \hat{T}_k$) in 11 of the 17 leagues. In the remaining 6 leagues, our estimate of post-Covid HA is larger than pre-Covid HA ($\hat{T}'_k > \hat{T}_k$).

Our Bayesian framework also allows for probabilistic interpretations regarding the likelihood that HA decreased within each league. Posterior probabilities of HA decline, $P(T'_k < T_k)$, are also presented in Table 4. The 3 leagues with the largest declines in HA, both in absolute and relative terms, were the Austrian Bundesliga ($\hat{T}_k = 0.161, \hat{T}'_k = -0.202$), the German Bundesliga ($\hat{T}_k = 0.229, \hat{T}'_k = -0.024$), and the Greek Super League ($\hat{T}_k = 0.409, \hat{T}'_k = 0.167$). The Austrian Bundesliga and German Bundesliga were the only 2 leagues to have post-Covid posterior HA estimates below 0, perhaps suggesting that HA disappeared in these leagues altogether in the absence of fans. We find it interesting to note that among the leagues with the 3 largest declines in HA are the leagues with the highest (Greek Super League) and lowest (Austrian Bundesliga) pre-Covid HA.

We estimate the probability the HA declined with the absence of fans, $P(T'_k < T_k)$, to be 0.999, 0.995, and 0.972 in the top flights in Austria, Germany, and Greece respectively. These 3 leagues, along with the English League Championship (0.912), Swedish Allsvenskan (0.907), and both tiers in Spain (0.959 for Spanish La Liga, 0.903 for Spanish La Liga 2) comprise seven leagues where we estimate a decline in HA with probability at least 0.9.

Two top leagues—the English Premier League ($\hat{T}_k = 0.246, \hat{T}'_k = 0.264$) and Italy Serie A ($\hat{T}_k = 0.204, \hat{T}'_k = 0.292$)—were among the six leagues with estimated post-Covid HA greater than pre-Covid HA. The three leagues

Table 5 Comparison of posterior means for pre-Covid and post-Covid yellow cards HA parameters from Model (4), \hat{T}_k and \hat{T}'_k , respectively

League	\hat{T}_k	\hat{T}'_k	$\hat{T}'_k - \hat{T}_k$	% Change	$P(T'_k > T_k)$
Russian Premier Liga	-0.404	0.037	0.441	109.1	0.997
German Bundesliga	-0.340	0.039	0.379	111.4	0.986
Portuguese Liga	-0.415	-0.008	0.406	98.0	0.984
German 2. Bundesliga	-0.392	0.090	0.482	123.0	0.982
Spanish La Liga 2	-0.359	-0.169	0.190	52.9	0.917
Danish Superliga	-0.331	-0.010	0.321	96.9	0.878
Austrian Bundesliga	-0.251	0.063	0.314	125.1	0.863
Greek Super League	-0.429	-0.261	0.168	39.2	0.829
Italy Serie B	-0.397	-0.223	0.174	43.8	0.799
Spanish La Liga	-0.269	-0.094	0.176	65.3	0.719
Swedish Allsvenskan	-0.196	-0.063	0.132	67.6	0.682
English League Championship	-0.478	-0.393	0.085	17.7	0.675
Norwegian Eliteserien	-0.323	-0.266	0.057	17.8	0.615
Turkish Super Lig	-0.199	-0.122	0.077	38.8	0.599
Swiss Super League	-0.327	-0.282	0.045	13.7	0.581
English Premier League	-0.293	-0.366	-0.073	-24.9	0.376
Italy Serie A	-0.344	-0.489	-0.145	-42.1	0.240

In the context of yellow cards, smaller (i.e., more negative) values of T_k and T'_k indicate larger home advantages. Relative and absolute differences between \hat{T}'_k and \hat{T}_k are also shown. Probabilities of decline in HA without fans, $P(T'_k > T_k)$, are estimated from posterior draws. We estimate the probability of a decline in HA without fans to exceed 0.9 in 5 of 17 leagues, and to exceed 0.5 in 15 of 17 leagues

with largest increase in HA without fans were the Swiss Super League ($\hat{T}_k = 0.180, \hat{T}'_k = 0.362$), Italy Serie A ($\hat{T}_k = 0.204, \hat{T}'_k = 0.292$), and the Portuguese Liga ($\hat{T}_k = 0.256, \hat{T}'_k = 0.338$).

Figure 4 (provided in “Appendix”) shows the posterior distributions of $T_k - T'_k$, the change in goals home advantage, in each league. Though this information is also partially observed in Table 4 and Fig. 2, the non-overlapping density curves for the change in HA provide additional evidence that post-Covid changes were not uniform between leagues.

Fitting Model (3) with $\lambda_3 > 0$ did not noticeably change inference with respect to the home advantage. For example, the probability that HA declined when assuming $\lambda_3 > 0$ was within 0.10 of the estimates shown in Table 4 in 14 of 17 leagues. In only one of the leagues did the estimated probability of HA decline exceed 0.9 with $\lambda_3 = 0$ and fail to exceed 0.9 with $\lambda_3 > 0$ (Swedish Allsvenskan: $P(T'_k < T_k) = 0.907$ w/ $\lambda_3 = 0$ and 0.897 w/ $\lambda_3 > 0$).

6.2 Yellow cards

6.2.1 Model fit

The yellow cards model presented in this paper is Model (4), using $\lambda_3 > 0$ for all leagues. Unlike with goals, where there was inconsistent evidence of a correlation in game-level outcomes, the correlation in home and away yellow cards per game varied between 0.10 and 0.22 among the 17 leagues.

\hat{R} statistics for Model (3) ranged from 0.9999 to 1.013, providing strong evidence that the model properly converged. Effective sample sizes (ESS) for each parameter in Model (4) are provided in Table 7. ESS are sufficiently large, especially HA parameters of interest T_k and T'_k , suggesting enough draws were taken to conduct inference.

6.2.2 Home advantage

As with Model (3) in Sect. 6.1.2, the primary parameters of interest in Model (4) are T_k and T'_k , the pre- and post-Covid home advantages for each league k , respectively. Unlike with goals, where values of T_k are positive, teams tend to want to avoid yellow cards, and thus estimates of T_k are < 0 . Related, a post-Covid decrease in yellow card HA is reflected by $T_k < T'_k$.

As in Sect. 6.1.2, T_k and T'_k correspond to a log-scale and represent the additional increase on the home team's log yellow card expectation, relative to a league average (μ_{ks}) after accounting for team and opponent (τ_{ks}) tendencies. Additionally, note that the same value of T_k represents a larger home advantage in a league where fewer cards are shown (i.e., smaller μ_{ks}).

Posterior distributions for T_k and T'_k are presented in Fig. 2. Posterior means of T_k range from -0.196 (Swedish Allsvenskan) to -0.478 (English League Championship).

Table 5 compares posterior means of T_k (denoted \hat{T}_k) with those of T'_k (denoted \hat{T}'_k) for each of the 17 leagues for the yellow cards model. Posterior means for T_k are smaller than that the corresponding posterior mean of T'_k , ($\hat{T}_k < \hat{T}'_k$) in 15 of the 17 leagues, suggesting that yellow card HA declined in nearly every league examined in the absence of fans.

The two leagues with the largest declines in HA, both in absolute and relative terms, were the German 2. Bundesliga ($\hat{T}_k = -0.392$, $\hat{T}'_k = -0.090$) and the Austrian Bundesliga ($\hat{T}_k = -0.251$, $\hat{T}'_k = 0.063$). In addition the top Austrian division and the 2nd German division, $\hat{T}'_k > 0$ in the German Bundesliga ($\hat{T}_k = -0.340$, $\hat{T}'_k = 0.039$) and Russian Premier League ($\hat{T}_k = -0.404$, $\hat{T}'_k = 0.037$).

Posterior probabilities of HA decline, $P(T'_k > T_k)$, are also presented in Table 5. This probability exceeds 0.9 in 5 of 17 leagues: Russian Premier Liga (.997), German Bundesliga (0.986), Portuguese Liga (0.984), German 2. Bundesliga (0.982), and Spanish La Liga 2 (0.917).

Alternatively, $\hat{T}_k > \hat{T}'_k$ in 2 leagues, the English Premier League ($\hat{T}_k = -0.293, \hat{T}'_k = -0.366$) and Italy Serie A ($\hat{T}_k = -0.344, \hat{T}'_k = -0.489$). However, given the overlap in the pre-Covid and post-Covid density curves, this does not appear to be a significant change.

Figure 5 (provided in “Appendix”) shows the posterior distributions of $T_k - T'_k$, the change in yellow card home advantage, in each league. In Fig. 5, there is little, if any, overlap between estimates of the change in Serie A's yellow card home advantage, and, for example, the change in German 2. Bundesliga and the Russian Premier League, adding to evidence that the post-Covid changes in HA are not uniform across leagues.

Fitting Model (4) with $\lambda_3 = 0$ changed inference with respect to the home advantage slightly more than was the case between the two variants of Model (3). For example, the probability that HA declined when assuming $\lambda_3 = 0$ was within 0.10 of the estimates shown in Table 5 in only 9 of 17 leagues. With $\lambda_3 = 0$, we estimated the probability HA declined to be 0.979 in the Austrian Bundesliga and 0.944 in the Danish Super Liga, compared to 0.863 and 0.874, respectively, with $\lambda_3 > 0$. Other notable differences include the English Premier League and Italy Serie A, whose estimated probability of HA decline rose from 0.075 and 0.073 to 0.376 and 0.240, respectively. Such differences are to be expected given the much larger observed correlation in yellow cards as compared to goals, and suggest that failure to account for correlation in yellow cards between home and away teams might lead to faulty inference and incorrect conclusions about significant decreases (or increases) in home advantage.

6.3 Examining goals and yellow cards simultaneously

To help characterize the relationship between changes in our two outcomes of interest, Fig. 6 (shown in the “Appendix”) shows the pre-Covid and post-Covid HA posterior means of each of goals and yellow cards in the 17 leagues. The origin of the arrows in Fig. 6 is the posterior mean of HA for pre-Covid yellow cards and goals, and the tip of the arrow is the posterior mean of post-Covid HA for yellow cards and goals.

Of the 17 leagues examined in this paper, 11 fall into the case where yellow cards and goals both experienced a decline in HA. In four leagues, the German Bundesliga, Spanish La Liga 2, Greek Super League, and Austrian Bundesliga, the probability that HA declined was greater than 0.8 for both outcomes of interest.

Despite the posterior mean HA for goals being higher post-Covid when compared to pre-Covid, the Turkish Super Lig, German 2. Bundesliga, Portuguese Liga, and Swiss Super League show a possible decrease in yellow card HA. For example, we estimate the probably that HA for yellow cards declined to be 0.984 for the Portuguese Liga and 0.982 for the German 2. Bundesliga.

Both the English Premier League and Italy Serie A show posterior mean HAs that are higher for both outcomes. Of the four countries where multiple leagues were

examined, only Spain's pair of leagues showed similar results (decline in HA for both outcomes). No leagues had showed posterior means with a lower HA for goals but not for yellow cards.

7 Discussion

Our paper utilizes bivariate Poisson regression to estimate changes to the home advantage (HA) in soccer games played during the summer months of 2020, after the outbreak of Covid-19, and relative to games played pre-Covid. Evidence from the 17 leagues we looked is mixed. In some leagues, evidence is overwhelming that HA declined for both yellow cards and goals. Alternatively, other leagues suggest the opposite, with some evidence that HA increased. Additionally, we use simulation to highlight the appropriateness of bivariate Poisson for home advantage estimation in soccer, particularly relative to the oft-used linear regression.

The diversity in league-level findings highlights the challenges in reaching a single conclusion about the impact of playing without fans, and implies that alternative causal mechanisms are also at play. For example, two of the five major European leagues are the German Bundesliga and Italy's Serie A. In the German Bundesliga, evidence strongly points to decreased HA (> 99% with goals), which is likely why Fischer and Haucap (2020a) found that broadly backing away Bundesliga teams represented a profitable betting strategy. But in Serie A, we only find a 10% probability that HA decreased with goal outcomes. Comparing these two results does not mesh into one common theme. Likewise, Figs. 1, 2, 4 and 5 imply that both (i) HA and (ii) changes in HA are not uniform by league.

Related, there are other changes post Covid-19 outbreak, some of which differ by league. These include, but are not limited to:

1. Leagues adopted rules allowing for five substitutions, instead of three substitutions per team per game. This rule change likely favors teams with more depth (potentially the more successful teams) and suggests that using constant estimates of team strength pre-Covid and post-Covid could be inappropriate.³
2. Certain leagues restarted play in mid-May, while others waited until the later parts of June. An extra month away from training and club facilities could have impacted team preparedness.
3. Covid-19 policies placed restrictions on travel and personal life. When players returned to their clubs, they did so in settings that potentially impacted their training, game-plans, and rest. Additionally, all of these changes varied by country, adding credence to our suggestion that leagues be analyzed separately.

Taken wholly, estimates looking at the impact of HA post-Covid are less of a statement about the cause and effect from a lack of fans (McCarrick et al. 2020; Bryson et al. 2020), and as much about changes due to both a lack of fans *and* changes

³ As shown in Table 3, however, we are limited by the number of post-Covid games in each league.

to training due to Covid-19. Differences in the latter could more plausibly be responsible for the heterogeneous changes we observe in HA post-Covid.

Given league-level differences in both HA and change in HA, we do not recommend looking at the impact of “ghost games” using single number estimate alone. However, a comparison to McCarrick et al. (2020), who suggest an overall decline in per-game goals HA from 0.29 to 0.15 (48%), is helpful for context. As shown in Table 4, our median league-level decline in goals HA, on the log scale, is 0.07. Extrapolating from Model 3, assuming attacking and defending team strengths of 0, and using the average posterior mean for μ_k , averaged across the 17 leagues, this equates to a decline in the per-game goals HA from 0.317 to 0.243 (23%). This suggests the possibility that, when using bivariate Poisson regression, the overall change in HA is attenuated when compared to current literature.

We are also the first to offer suggestions on the simultaneous impact of HA for yellow cards and goals. While traditional soccer research has used yellow cards as a proxy for referee decisions relating to benefits for the home team, we find that it is not always the case that changes in yellow card HA are linked to changes to goal HA. In two leagues, German 2. Bundesliga and Portuguese Liga, there are overwhelming decreases in yellow card HA (probabilities of a decrease of at least 98% in each), but small increases in the net HA given to home team goals. Among other explanations, this suggests that yellow cards are not directly tied to game outcomes. It could be the case that, for example, visiting teams in certain leagues fouled less often on plays that did not impact chances of scoring or conceding goals. Under this hypothesis, yellow cards are not a direct proxy for a referee-driven home advantage, and instead imply changes to player behavior without fans, as suggested by Leitner and Richlan (2020a). Alternatively, having no fan support could cause home players to incite away players less frequently. Said FC Barcelona star Lionel Messi (Reuters 2020), “It’s horrible to play without fans. It’s not a nice feeling. Not seeing anyone in the stadium makes it like training, and it takes a lot to get into the game at the beginning.”

Finally, we use simulations to highlight limitations of using linear regression with goal outcomes in soccer. The mean absolute bias in HA estimates is roughly six times higher when using linear regression, relative to bivariate Poisson. Absolute bias when estimating HA using bivariate Poisson also compares favorably to paired comparison models. Admittedly, our simulations are naive, and one of our two data generating processes for simulated game outcomes aligns with the same Poisson framework as the one we use to model game results. This, however, is supported by a wide body of literature, including Reep and Benjamin (1968), Reep et al. (1971), Dixon and Coles (1997), and Karlis and Ntzoufras (2000). Despite this history, linear regression remains a common tool for soccer research (as shown in Table 1); as an alternative, we hope these findings encourage researchers to consider the Poisson distribution.

Appendix

Simulation details

Team strengths

Attacking (α_{t^*}) and defensive (δ_{t^*}) team strength estimates stem from a bivariate normal distribution, as in Thompson (2018), such that $(\alpha_{t^*}, \delta_{t^*}) \sim$ bivariate normal (μ, Σ) where $\mu = (0, 0)$ and $\Sigma = \begin{bmatrix} 0.35^2 & (\rho^*)0.35^2 \\ (\rho^*)0.35^2 & 0.35^2 \end{bmatrix}$ for $t^* = 1 \dots 20$, where 20 is the number of simulated teams. Estimates in Σ correspond to relative gaps in observed soccer team strength (see Fig. 3). In our simulations, we use $\rho^* \in \{-0.8, -0.4, 0\}$, reflecting the range of correlations in scoring and defending strength (negative correlations infer that teams that score more goals also give up fewer goals). As is custom in professional soccer, we assume each team played each opponent twice, once at home and once away, yielding 380 total games per season (Figs. 3, 4, 5, 6; Tables 6, 7).

Simulating goals

We use two data generating processes for goals, bivariate Poisson (BVP) and bivariate normal (BVN).

Under BVP, we use Model (2) to generate λ_{1i^*} and λ_{2i^*} for each i^* , where $i^* = 1 \dots 380$. We assume $\lambda_3 = 0$, $\mu = 0$, and $T = T^*$, where $T^* \in \{0, 0.25, 0.5\}$ is a simulated home advantage. Using the `rpois()` function in R we simulated goals for both the home ($Y_{Hi^*}^* \sim \text{Pois}(\lambda_{1i^*})$) and away ($Y_{Ai^*}^* \sim \text{Pois}(\lambda_{2i^*})$) in each of the 380 games.

Simulating under bivariate normal (BVN) requires a few steps to ensure goal outcomes roughly correspond to soccer games. First, we use rounded, truncated normal distributions for simulations via the `round()` and `truncnorm()` functions in R, respectively. The mean home (Y_H^*) and away (Y_A^*) goals come from univariate truncated normal distributions with $\mu_{Hi^*} = 0.2 + \alpha_{Hi^*} + \delta_{Ai^*}$ and $\mu_{Ai^*} = 0.2 + \alpha_{Ai^*} + \delta_{Hi^*}$, respectively, and variances of $\sigma^{*2} = 1.75^2$. The lower bounds on both truncated normal distributions are -0.49 . Here, expectations and variances are designed to roughly reflect observed goal outcomes and the lower bounds ensure goal outcomes are positive. For simulations with no home advantage, home and away expectations are identical, as above. For simulations with home advantages of 0.25 and 0.5, a goal is randomly added to the home team's total with probability 0.25 and 0.50, respectively. Although this approach is admittedly unconventional, it yields goal and home advantage outcomes that roughly reflect observed data.

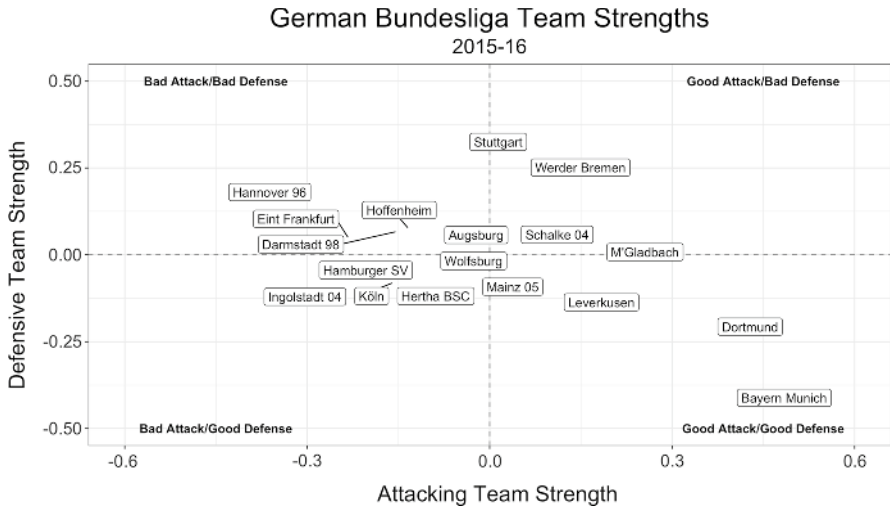


Fig. 3 Posterior means of attacking (α_{ks}) and defensive strengths (δ_{ks}) for teams in the German Bundesliga (k) during 2015–2016 season (s). The casual soccer fan will note familiar powers such as Bayern Munich and Borussia Dortmund as having the best estimates of overall team strength. However, examining the posterior means of teams’ attacking and defensive strengths makes apparent that, in general, teams may be strong in one facet of the game but not the other. Stuttgart, for example, finished in the top third of the German Bundesliga in terms of goals scored, yet were relegated, conceding the most goals in the league. That same season, Ingolstadt 04, on the other hand, scored the 2nd fewest goals in the league, but had a top four defense on the basis of goals conceded. The fact that correlation between α_{ks} and δ_{ks} can be weak demonstrates the need consider both terms in Model (3)

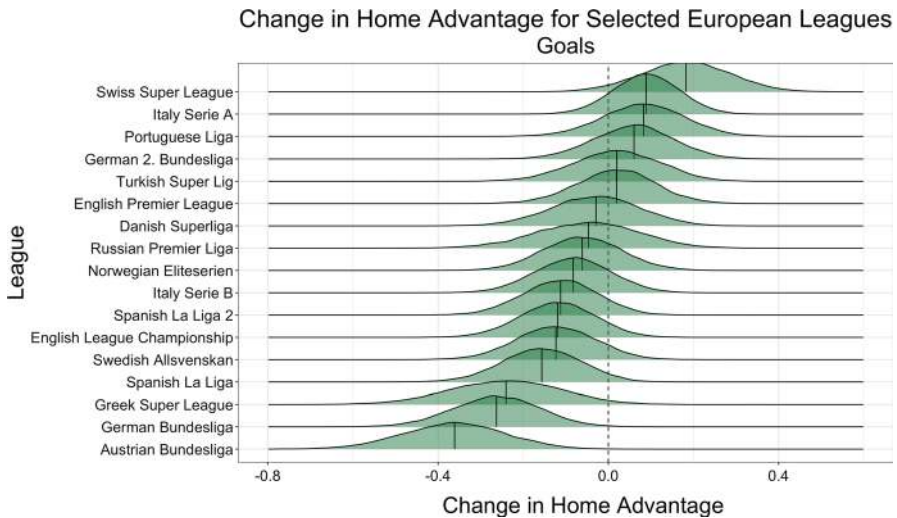


Fig. 4 Posterior distributions of $T_k - T'_k$, the change in goals home advantage. Negative values of $T_k - T'_k$ reflect a decrease in home advantage, while positive values reflect an increase in home advantage. Across the 17 leagues in the sample, a range of differences exist between posterior distributions of $T_k - T'_k$

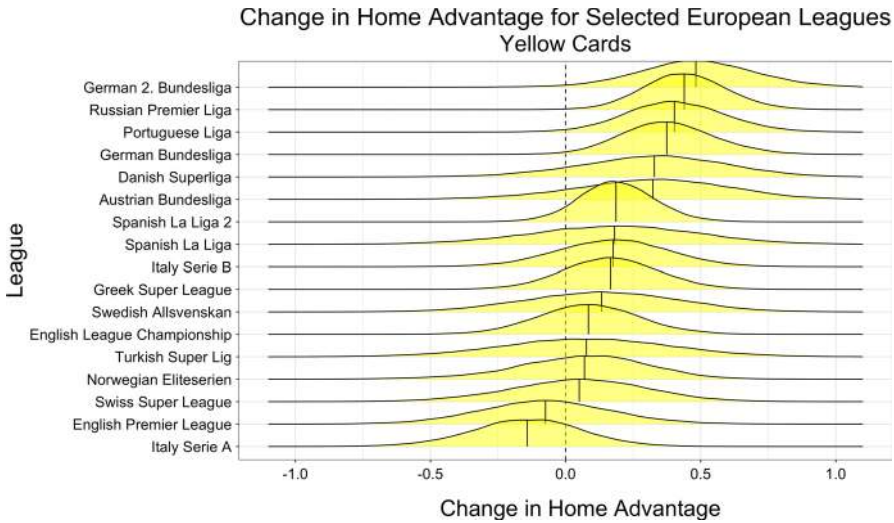


Fig. 5 Posterior distributions of $T_k - T'_k$, the change in yellow card home advantage. Positive values of $T_k - T'_k$ reflect a decrease in home advantage, while negative values reflect an increase in home advantage. Across the 17 leagues in the sample, a range of differences exist between posterior distributions of $T_k - T'_k$

Model candidates

Three candidate models are fit on each of the BVP and BVN data generating processes. First, we fit the bivariate Poisson model shown in Model (2), assuming no covariance, and using 2 parallel chains, 5000 iterations, and a burn in of 2000 draws.

Second, we use ordinary least squares to fit a linear regression model of goal differential, using team-level fixed effects for the home and away teams, as well as a home advantage term. Letting $D_{i*Hi*Ai*} = Y_{Hi*}^* - Y_{Ai*}^*$ be the goal difference in simulated game i^* , we fit Model (5) below,

$$D_{i*Hi*Ai*} = \alpha + home_{i*Hi*} \times I(home = Hi^*) + away_{i*Ai*} \times I(away = Ai^*) + \epsilon_{i*Hi*Ai*}. \tag{5}$$

In Model (5), α is the home advantage, and $home_{i*Hi*}$ and $away_{i*Ai*}$ are fixed effects for the home and away teams, respectively.

Third, we fit a Bayesian paired comparison model, such that

$$D_{i*Hi*Ai*} = \alpha + \theta_{Hi*} - \theta_{Ai*} + \epsilon_{i*Hi*Ai*}, \tag{6}$$

using prior distributions $\theta \sim N(0, \sigma_{team}^2)$, $\alpha \sim N(0, 100)$, and $\sigma_{team} \sim \text{Inverse-Gamma}(1, 1)$, using 2 parallel chains, 5000 iterations, and a burn in of 2000 draws.

A total of 900 seasons were simulated using each of BVN and BVP data generating processes (100 season for each combination of ρ^* and T^*).

Change in Home Advantages: Goals and Yellow Cards

Arrows reflect Pre-Covid to Post-Covid posterior means in HA

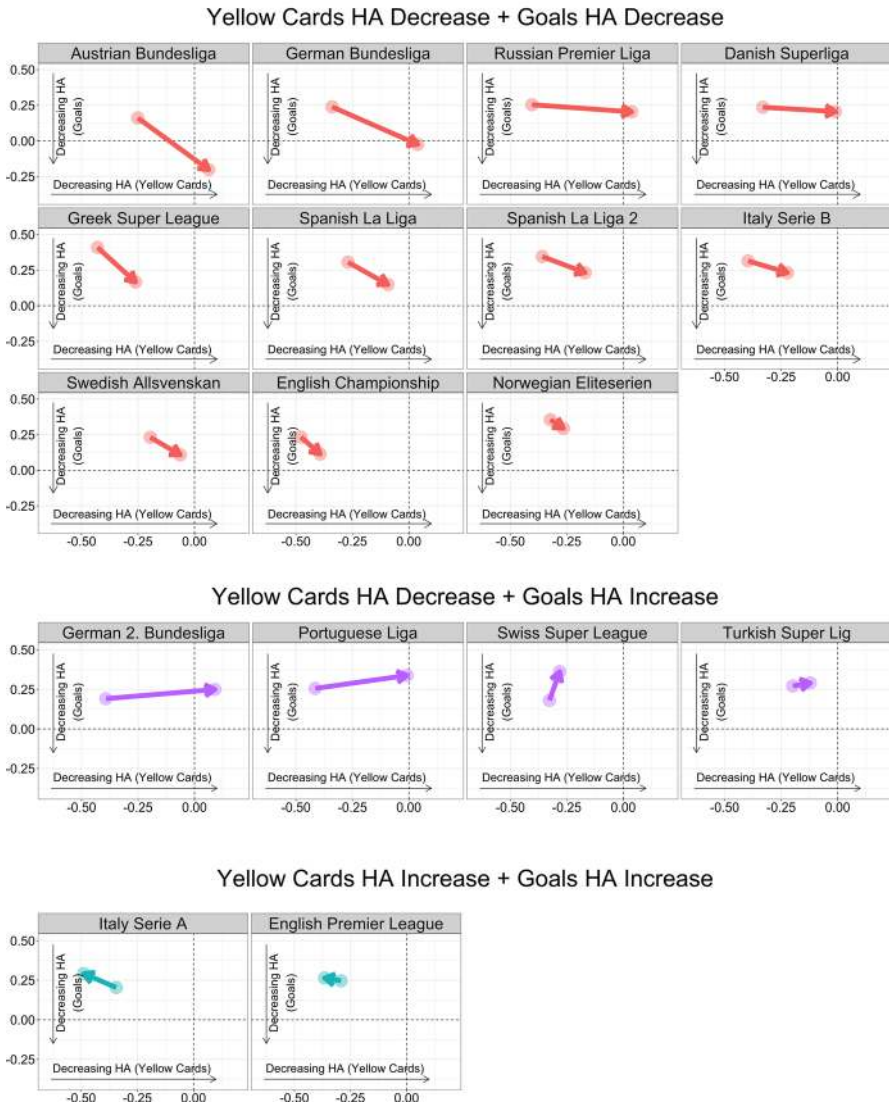


Fig. 6 For each league, an arrow is displayed, originating at the point (pre-Covid YC HA posterior mean, pre-Covid goals HA posterior mean) and terminating at the point (post-Covid YC HA posterior mean, post-Covid goals HA posterior mean). These arrows allow for visual comparison of both the direction and magnitude of any change in HA in the two outcomes of interest. Arrows pointing down denote a decline in goals HA, and arrows pointing right reflect a decline in yellow cards HA. Of 17 leagues, 11 experienced a decline in HA in both outcomes, 4 leagues experienced an increase in goals HA and a decline in yellow cards HA, 2 experienced an increase in both outcomes, and no leagues experienced a decrease in goals HA and an increase in yellow cards HA

Table 6 Effective sample sizes (ESS) of posterior draws for parameters from Model (3), with $\lambda_3 = 0$

League	T_k	T'_k	μ_{ks}	α_{ks}	δ_{ks}	$\sigma_{att,k}$	$\sigma_{def,k}$
Austrian Bundesliga	33,485	38,599	8268	16,076	23,617	15,612	10,214
Danish Superliga	24,990	23,781	9965	18,872	21,732	10,363	7517
English League Championship	33,130	33,600	13,079	30,217	28,961	9000	9856
English Premier League	33,301	34,014	5838	15,962	23,314	16,252	10,443
German 2. Bundesliga	15,460	16,159	9187	16,584	19,958	6153	2461
German Bundesliga	32,620	37,899	8370	20,156	29,073	16,389	8085
Greek Super League	30,760	29,910	6437	15,768	20,427	14,997	12,580
Italy Serie A	31,899	33,764	6988	18,508	23,720	16,826	11,311
Italy Serie B	31,332	34,283	16,172	31,542	32,692	8383	4687
Norwegian Eliteserien	32,307	29,220	13,841	28,978	30,993	10,642	5949
Portuguese Liga	30,775	32,751	5898	16,217	21,875	15,729	10,634
Russian Premier Liga	28,916	32,465	10,177	22,623	25,440	12,714	11,205
Spanish La Liga	33,778	35,011	6506	18,778	22,536	16,283	12,295
Spanish La Liga 2	20,489	24,182	12,721	23,981	24,046	5081	4813
Swedish Allsvenskan	33,677	31,641	9290	26,587	23,395	11,153	12,115
Swiss Super League	28,506	29,041	9054	16,032	23,589	12,510	7184
Turkish Super Lig	29,032	30,581	10,470	23,238	26,978	10,772	8600

For parameters that vary by season (μ_{ks}) or team and season (α_{ks}, δ_{ks}), mean ESS values are presented

Table 7 Effective sample sizes (ESS) of posterior draws for parameters from Model (4), with $\lambda_3 > 0$

League	T_k	T'_k	μ_{ks}	τ_{ks}	$\sigma_{team,k}$	γ_k
Austrian Bundesliga	5157	25,191	1567	31,518	993	431
Danish Superliga	10,080	29,180	2263	32,471	961	1293
English League Championship	2956	36,611	2166	34,905	1838	1748
English Premier League	7354	21,696	976	27,820	659	727
German 2. Bundesliga	5202	31,952	528	34,198	501	979
German Bundesliga	3262	43,122	1006	31,308	911	1096
Greek Super League	3337	39,304	2026	35,878	2061	1539
Italy Serie A	7205	13,276	1619	21,431	1158	963
Italy Serie B	6460	35,809	2145	37,667	1657	1333
Norwegian Eliteserien	1304	22,198	454	29,584	371	226
Portuguese Liga	6360	44,860	1542	36,093	1402	1755
Russian Premier Liga	6933	44,941	1247	36,397	1365	1490
Spanish La Liga	5951	19,471	789	31,242	428	243
Spanish La Liga 2	3197	38,104	2034	39,014	1849	1367
Swedish Allsvenskan	22,669	28,121	4988	37,235	968	1760
Swiss Super League	10,103	31,241	2426	38,500	1572	1109
Turkish Super Lig	20,777	24,891	5737	39,653	723	2545

For parameters that vary by season (μ_{ks}) or team and season (τ_{ks}) mean ESS values are presented

Funding Not Applicable.

Data availability All data used in this project are open source, and come from Football Reference (Sports Reference 2020). We make our cleaned, analysis-ready dataset available at https://github.com/lbenz730/soccer_ha_covid/tree/master/fbref_data.

Code availability All code for scraping data, fitting models, and conducting analyses has been made available for public use at https://github.com/lbenz730/soccer_ha_covid.

Declaration

Conflict of interest The authors declare that they have no conflict of interest. The authors would like to note that this work is not endorsed, nor associated with Medidata Solutions, Inc.

References

- Agnew, G.A., Carron, A.V.: Crowd effects and the home advantage. *Int. J. Sport Psychol.* **66**, 6 (1994)
- Baio, G., Blangiardo, M.: Bayesian hierarchical model for the prediction of football results. *J. Appl. Stat.* **37**(2), 253–264 (2010)
- Brooks, S.P., Gelman, A.: General methods for monitoring convergence of iterative simulations. *J. Comput. Graph. Stat.* **7**(4), 434–455 (1998)
- Bryson, A., Dolton, P., Reade, J.J., Schreyer, D., Singleton, C.: Causal effects of an absent crowd on performances and refereeing decisions during covid-19 (2020)
- Buraimo, B., Forrest, D., Simmons, R.: The 12th man? Refereeing bias in English and German soccer. *J. R. Stat. Soc. Ser. A Stat. Soc.* **173**(2), 431–449 (2010)
- Courneya, K.S., Carron, A.V.: The home advantage in sport competitions: a literature review. *J. Sport Exerc. Psychol.* **14**(1), 66 (1992)
- Cueva, C.: Animal Spirits in the Beautiful Game. Testing Social Pressure in Professional Football During the Covid-19 Lockdown (2020)
- Dilger, A., Vischer, L.: No Home Bias in Ghost Games (2020)
- Dixon, M.J., Coles, S.G.: Modelling association football scores and inefficiencies in the football betting market. *J. R. Stat. Soc. Ser. C Appl. Stat.* **46**(2), 265–280 (1997)
- Dohmen, T., Sauerermann, J.: Referee bias. *J. Econ. Surv.* **30**(4), 679–695 (2016)
- Endrich, M., Gesche, T.: Home-bias in referee decisions: evidence from “ghost matches” during the covid19-pandemic. *Econ. Lett.* **197**, 109621 (2020)
- Ferraresi, M., Gucciardi, G., et al.: Team performance and audience: experimental evidence from the football sector. *Tech. rep* (2020)
- Fischer, K., Haucap, J.: Betting Market Efficiency in the Presence of Unfamiliar Shocks: The Case of Ghost Games During the Covid-19 Pandemic (2020a)
- Fischer, K., Haucap, J.: Does Crowd Support Drive the Home Advantage in Professional Soccer? Evidence from German Ghost Games During the Covid-19 Pandemic (2020b)
- Forrest, D., Beaumont, J., Goddard, J., Simmons, R.: Home advantage and the debate about competitive balance in professional sports leagues. *J. Sports Sci.* **23**(4), 439–445 (2005)
- Garicano, L., Palacios-Huerta, I., Prendergast, C.: Favoritism under social pressure. *Rev. Econ. Stat.* **87**(2), 208–216 (2005)
- Gelman, A., Rubin, D.B.: Inference from iterative simulation using multiple sequences. *Stat. Sci.* **7**(4), 457–472 (1992)
- Gelman, A., Carlin, J.B., Stern, H.S., Dunson, D.B., Vehtari, A., Rubin, D.B.: *Bayesian Data Analysis*, 3rd edn. CRC Press, Boca Raton (2013)
- Glickman, M.E., Stern, H.S.: A state-space model for national football league scores. *J. Am. Stat. Assoc.* **93**(441), 25–35 (1998)
- Groll, A., Kneib, T., Mayr, A., Schaubberger, G.: On the dependency of soccer scores—a sparse bivariate Poisson model for the UEFA European football championship 2016. *J. Quant. Anal. Sports* **14**(2), 65–79 (2018)

- Jiménez Sánchez, Á., Lavín, J.M.: Home advantage in European soccer without crowd. *Soccer Soc.* **66**, 1–14 (2020)
- Karlis, D., Ntzoufras, I.: On modelling soccer data. *Student* **3**(4), 229–244 (2000)
- Karlis, D., Ntzoufras, I.: Analysis of sports data by using bivariate Poisson models. *J. R. Stat. Soc. Ser. D Stat* **52**(3), 381–393 (2003)
- Karlis, D., Ntzoufras, I., et al.: Bivariate Poisson and diagonal inflated bivariate Poisson regression models in r. *J. Stat. Softw.* **14**(10), 1–36 (2005)
- Koopman, S.J., Lit, R.: A dynamic bivariate Poisson model for analysing and forecasting match results in the English premier league. *J. R. Stat. Soc. Ser. A Stat. Soc.* **66**, 167–186 (2015)
- Krawczyk, M., Strawiński, P., et al.: Home advantage revisited. did covid level the playing fields? *Tech. rep* (2020)
- Leitner, M.C., Richlan, F.: Analysis System for Emotional Behavior in Football (aseb-f): Professional Football Players' Emotional Behavior in Ghost Games in the Austrian Bundesliga. Draft version 1 05-08-2020 (2020a)
- Leitner, M.C., Richlan, F.: No Fans-No Home Advantage. Sport Psychological Effects of Missing Supporters on Football Teams in European Top Leagues (2020b)
- Ley, C., Wiele, TVd., Eetvelde, H.V.: Ranking soccer teams on the basis of their current strength: a comparison of maximum likelihood approaches. *Stat. Model.* **19**(1), 55–73 (2019)
- Lopez, M.J.: Persuaded under pressure: evidence from the national football league. *Econ. Inquiry* **54**(4), 1763–1773 (2016)
- Lopez, M.J., Matthews, G.J., Baumer, B.S., et al.: How often does the best team win? A unified approach to understanding randomness in north American sport. *Ann. Appl. Stat.* **12**(4), 2483–2516 (2018)
- McCarrick, D., Bilalic, M., Neave, N., Wolfson, S.: Home Advantage During the Covid-19 Pandemic in European Football (2020)
- Moskowitz, T., Wertheim, L.J.: *Scorecasting: The Hidden Influences Behind How Sports are Played and Games are Won*. Three Rivers Press, CA (2012)
- Nevill, A.M., Holder, R.L.: Home advantage in sport. *Sports Med.* **28**(4), 221–236 (1999)
- Pettersson-Lidbom, P., Priks, M.: Behavior under social pressure: empty Italian stadiums and referee bias. *Econ. Lett.* **108**(2), 212–214 (2010)
- Reade, J.J., Schreyer, D., Singleton, C.: Echoes: What happens when football is played behind closed doors? Available at SSRN 3630130 (2020)
- Reep, C., Benjamin, B.: Skill and chance in association football. *J. R. Stat. Soc. Ser. A Gener.* **131**(4), 581–585 (1968)
- Reep, C., Pollard, R., Benjamin, B.: Skill and chance in ball games. *J. R. Stat. Soc. Ser. A Gener.* **134**(4), 623–629 (1971)
- Reuters: Lionel Messi Says Playing Without Fans is 'Horrible and Ugly' (2020). https://www.eurosport.com/football/liga/2020-2021/lionel-messi-says-playing-without-fans-is-horrible-and-ugly-after-barcelona-star-collects-pichichi_sto8042397/story.shtml
- Schwartz, B., Barsky, S.F.: The home advantage. *Soc. Forces* **55**(3), 641–661 (1977)
- Scoppa, V.: Social Pressure in the Stadiums: Do Agents Change Behavior Without Crowd Support? (2020)
- Sors, F., Grassi, M., Agostini, T., Murgia, M.: The sound of silence in association football: Home advantage and referee bias decrease in matches played without spectators. *Eur. J. Sport Sci.* **6**, 1–21 (2020)
- Sports Reference: Football reference (2020). <https://fbref.com/en/>
- Stan Development Team: Stan reference manual (2019). https://mc-stan.org/docs/2_26/reference-manual/index.html
- Thompson, J.: Soccer predictions using bayesian mixed effects models. *Tech. rep.* (2018). <https://wjakethompson.github.io/soccer/intro.html>. Accessed Dec 2020
- Tsokos, A., Narayanan, S., Kosmidis, I., Baio, G., Cucuringu, M., Whitaker, G., Király, F.: Modeling outcomes of soccer matches. *Mach. Learn.* **108**(1), 77–95 (2019)
- Unkelbach, C., Memmert, D.: Crowd noise as a cue in referee decisions contributes to the home advantage. *J. Sport Exerc. Psychol.* **32**(4), 483–498 (2010)

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.