

ORIGINAL ARTICLE OPEN ACCESS

# Tractable Ridge Regression for Paired Comparisons

Cristiano Varin<sup>1</sup>  | David Firth<sup>2</sup><sup>1</sup>Department of Environmental Sciences, Informatics and Statistics, Ca' Foscari University, Venice, Italy | <sup>2</sup>Department of Statistics, University of Warwick, Coventry, UK**Correspondence:** Cristiano Varin ([cristiano.varin@unive.it](mailto:cristiano.varin@unive.it))**Received:** 3 October 2024 | **Revised:** 25 October 2024 | **Accepted:** 3 November 2024**Keywords:** Bradley–Terry model | composite likelihood | pairwise empirical Bayes | probit regression | rating | shrinkage | sport tournaments | Thurstone–Mosteller model

## ABSTRACT

Paired-comparison models, such as Bradley–Terry and Thurstone–Mosteller, are commonly used to estimate relative strengths of pairwise compared items in tournament-style data. We discuss estimation of paired-comparison models with a ridge penalty. A new approach is derived which combines empirical Bayes and composite likelihoods without any need to refit the model, as a convenient alternative to cross-validation of the ridge tuning parameter. Simulation studies demonstrate much better predictive accuracy of the new approach relative to ordinary maximum likelihood. A widely used alternative, the application of a standard bias-reducing penalty, is also found to improve appreciably the performance of maximum likelihood; but the ridge penalty, with tuning as developed here, yields greater accuracy still. The methodology is illustrated through application to 28 seasons of English Premier League football.

## 1 | Introduction

The problem of rating a list of items on the basis of a set of paired comparisons arises frequently, in a variety of fields including artificial intelligence (e.g., Rafailov et al. 2023), bibliometrics (e.g., Varin, Cattelan, and Firth 2016), education (e.g., Bartholomew and Jones 2022), forensic science (e.g., Thompson et al. 2018), genetics (e.g., Ma, Wong, and Owen 2012), politics (e.g., Loewen, Rubenson, and Spirling 2012), psychometrics (e.g., Maydeu-Olivares and Böckenholt 2005) and sport (e.g., Glickman and Stern 2017), to name just a few. A classical reference for the statistical analysis of paired comparisons is David (1988). A more recent review is Cattelan (2012), while Aldous (2017) discusses various aspects of paired comparisons from the point of view of applied probability.

In binary paired-comparison models, the probability that item  $i$  beats item  $j$  in a comparison between them is

$$\Pr(i \text{ beats } j) = F(\mu_i - \mu_j), \quad (1)$$

where  $F(\cdot)$  is the cumulative distribution function of a zero-symmetric continuous random variable and  $\mu_i$  is the strength of item  $i$ . The symmetry restriction here is important, to ensure that  $F(\mu_i - \mu_j) + F(\mu_j - \mu_i) = 1$  for all possible values of  $\mu_i$  and  $\mu_j$ . Popular choices for  $F$  are the logistic distribution which gives the Bradley–Terry model (Bradley and Terry 1952) and the normal distribution that corresponds to the Thurstone–Mosteller model (Thurstone 1927; Mosteller 1951). These two models are quite similar in practice, given the well-known correspondence between the logit link assumed in the Bradley–Terry model and the probit link of the Thurstone–Mosteller model (e.g., Agresti 2002, 246–247). The logit and probit functions are nearly proportional except in the tails, and thus, in practice, it is usually difficult to see any appreciable difference in fit between logit and probit models.

Consider  $p$  items. Maximum likelihood estimation of the vector of strength parameters  $\boldsymbol{\mu} = (\mu_1, \dots, \mu_p)^\top$  is well known to be problematic, in part because the estimate of  $\mu_i$  diverges if item  $i$  wins all of its comparisons or loses them all (e.g., Kosmidis and

This is an open access article under the terms of the [Creative Commons Attribution](https://creativecommons.org/licenses/by/4.0/) License, which permits use, distribution and reproduction in any medium, provided the original work is properly cited.

© 2024 The Author(s). *Stat* published by John Wiley & Sons Ltd.

Firth 2021). Bias-reduced maximum likelihood (Firth 1993) overcomes such limitations of maximum likelihood, guaranteeing finiteness and also better frequentist properties of the strength estimates (Kosmidis and Firth 2021). Finite estimates can also be obtained through other forms of penalized maximum likelihood (Mease 2003) or Bayesian methods (Caron and Doucet 2012).

The present paper develops a highly tractable alternative regularization of the likelihood for the Thurstone–Mosteller model, using a ridge penalty that is tuned by using a new ‘pairwise empirical Bayes’ method. The predictive performance of this new approach is studied in simulation experiments and through application to match results from incomplete seasons of English Premier League football.

## 2 | Ridge Regression for Paired Comparisons

We begin by assuming that the outcome of a paired comparison does not depend on the order in which the items are presented, which in sport applications means that there is no home-field effect; and we assume for now also that there are no ties. These assumptions will be relaxed later, in Section 3.2.

Denote the outcome of the paired comparison between items  $i$  and  $j$  as  $Y_{ij} = 1$  if  $i$  wins and  $Y_{ij} = -1$  if  $j$  wins. In this paper, we investigate a natural form of regularization for estimation of the strength parameters, based on the ridge-penalized log likelihood

$$\ell_{\lambda}(\boldsymbol{\mu}) = \ell(\boldsymbol{\mu}) - \frac{\lambda}{2} \sum_{i=1}^p \mu_i^2, \quad (2)$$

where  $\ell(\boldsymbol{\mu})$  is the log likelihood of the paired-comparison model, that is,

$$\ell(\boldsymbol{\mu}) = \sum_{(i,j) \in \mathcal{S}} \log F\{y_{ij}(\mu_i - \mu_j)\}. \quad (3)$$

Here,  $\mathcal{S}$  denotes the observed tournament, that is,  $\mathcal{S}$  is the set of index pairs  $(i, j)$  in which  $y_{ij}$  is observed; and  $\lambda \geq 0$  is a scalar tuning parameter. For a fixed value of  $\lambda$ , computation of the ridge estimates of the strength parameters is a relatively simple optimization, by using, for example, standard quasi-Newton algorithms. The standard approach to determining a good value for  $\lambda$  is cross-validation. An important point is that in paired-comparison data, there is often some balancing structure that should be respected in cross-validated splitting of the dataset; in round-robin tournaments, for example, it will usually make sense to use whole tournament rounds as cross-validation units. When such structure is absent, though, it can often be difficult to identify a suitable set of splits for cross-validation with paired-comparison data.

An alternative to cross-validation is estimation of  $\lambda$  using the empirical Bayes method (Morris 1983). Ridge regression coincides with maximum a posteriori estimation when the strength parameters are uncorrelated zero-mean normal random variables with precision  $\lambda$ ,

$$\boldsymbol{\mu} \sim \mathcal{N}(\mathbf{0}, \lambda^{-1} \mathbf{I}_p), \quad (4)$$

where  $\mathbf{I}_p$  is the identity matrix of size  $p$ . Normality of the strengths is purely a working assumption here, to allow determination of  $\lambda$  via the empirical Bayes method. In the simulation studies reported in Section 4.2, we show that this method of determining  $\lambda$  is robust in the sense that it works well even when the true strengths have appreciably heavier tails than a normal distribution. (There is no need to consider violations of normality due to skewness, because the paired-comparison model is identified by differences  $\mu_i - \mu_j$ , which have a symmetric distribution in any case.)

The empirical Bayes method determines  $\lambda$  by maximizing the marginal likelihood obtained from integrating out the item strengths. This leads to a cumbersome integral of dimension  $p$ ,

$$L(\lambda) = \lambda^{p/2} \int_{\mathbb{R}^p} \exp\{\ell(\boldsymbol{\mu})\} \prod_{i=1}^p \phi(\lambda^{1/2} \mu_i) d\boldsymbol{\mu}, \quad (5)$$

where  $\phi(\cdot)$  is the density of the standard normal random variable.

## 3 | Pairwise Empirical Bayes

### 3.1 | Estimation of the Tuning Parameter

We develop here a simple method that avoids the potentially high-dimensional integration (5) by using composite likelihood methods (Varin, Reid, and Firth 2011). The tuning parameter  $\lambda$  is estimated by maximizing the pairwise log likelihood constructed from all couples of correlated paired comparisons, that is, those couples with one item in common, such as  $(i$  vs  $j)$  and  $(i$  vs  $k)$ . The pairwise log likelihood for  $\lambda$  takes the form

$$\ell_{\text{pair}}(\lambda) = \sum_{r,s \in \mathcal{O}} n_{rs} \log p_{rs}(\lambda). \quad (6)$$

Here,  $\mathcal{O} = \{-1, +1\}$  is the set of possible outcomes of the paired comparisons (which will be later extended in Section 3.2 to include ties), while  $n_{rs} = \sum_{ijk} \mathbf{1}(Y_{ij} = r, Y_{ik} = s)$  is the number of correlated couples of paired comparisons with outcomes  $r$  and  $s$ , and  $p_{rs}(\lambda) = \Pr(Y_{ij} = r, Y_{ik} = s)$  is the corresponding marginal bivariate probability.

In the following, we use the Thurstone–Mosteller model, because it is mathematically more convenient than the Bradley–Terry model and the two models give essentially equivalent fits as already noted in Section 1. The mathematical convenience of the Thurstone–Mosteller model comes directly from the fact that its normal latent-variable representation is conjugate with the normal prior (4) that underpins ridge-penalized maximum likelihood—a fact that greatly simplifies the integrals required by the empirical Bayes method. Under the Thurstone–Mosteller model, the bivariate probabilities  $p_{rs}(\lambda)$  have closed-form expressions, by a well-known result sometimes referred to as the theorem of median dichotomy (Sheppard 1898; Cox and Wermuth 2002). The probability of two wins for the common item  $i$  is

$$\begin{aligned}
 p_{11}(\lambda) &= \Pr(Y_{ij}=1, Y_{ik}=1) \\
 &= \int_0^\infty \int_0^\infty \phi_2\left(u, v; \frac{\lambda^{-1}}{1+2\lambda^{-1}}\right) du dv \\
 &= \frac{(1+\tau)}{4},
 \end{aligned}$$

where  $\phi_2(u, v; \rho)$  is the density function of a bivariate standard normal variable with correlation  $\rho$ , and the Kendall rank correlation

$$\tau = \frac{2}{\pi} \arcsin\left(\frac{\lambda^{-1}}{1+2\lambda^{-1}}\right) \quad (7)$$

takes a value in the interval (0,1/3). The general expression for  $p_{rs}$  depends on whether there is concordance or not in the outcomes of the two paired comparisons:

$$p_{rs}(\lambda) = \begin{cases} \frac{(1+\tau)}{4}, & \text{if } r=s, \\ \frac{(1-\tau)}{4}, & \text{if } r \neq s. \end{cases}$$

A closed-form estimate of  $\lambda$  is obtained via reparameterization of the pairwise likelihood in  $\tau$ ,

$$\ell_{\text{pair}}(\tau) = c \log(1+\tau) + d \log(1-\tau),$$

where  $c$  is the count of concordant correlated couples of paired comparisons ( $Y_{ij}=Y_{ik}$ ) and  $d$  the count of discordant couples ( $Y_{ij} \neq Y_{ik}$ ). The maximum pairwise likelihood estimator of  $\tau$  is then the familiar Kendall  $\tau$  coefficient

$$\hat{\tau} = \frac{c-d}{c+d},$$

and inversion of Equation (7) yields the maximum pairwise likelihood estimate of  $\lambda$ ,

$$\hat{\lambda} = \frac{1-2 \sin(\hat{\tau}\pi/2)}{\sin(\hat{\tau}\pi/2)}. \quad (8)$$

In sparse datasets with only a modest number of comparisons per item, it is helpful to consider making a small-sample adjustment along the lines of the classical ‘rule of succession’ of Laplace (1814). In the present context such an adjustment adds one imaginary concordant couple and one imaginary discordant couple per item, and the corresponding small-sample adjusted estimator of  $\tau$  is

$$\hat{\tau} = \frac{c-d}{c+d+2p}.$$

An alternative interpretation of this adjusted estimator is that it maximizes a penalized version of the pairwise log likelihood, namely,

$$\tilde{\ell}_{\text{pair}}(\tau) = \ell_{\text{pair}}(\tau) + p \log(1-\tau^2). \quad (9)$$

### 3.2 | Order Effect and Ties

The order of presentation of the items is important in many paired-comparison contexts, notable examples being the home-field effect in sport competitions or the white player advantage in chess. In paired-comparison models, such an order effect is typically described through the inclusion of an intercept term (e.g., Agresti, 2002),

$$\Pr(Y_{ij}=1|\boldsymbol{\mu}) = \Phi(\delta + \mu_i - \mu_j), \quad (10)$$

where  $\Phi(\cdot)$  is the distribution function of a standard normal variable.

Another common complication in paired comparisons is the presence of ties, represented by the null outcome  $Y_{ij}=0$  whenever  $i$  and  $j$  are compared but neither beats the other. As discussed in Agresti (1992), ties can be incorporated in paired-comparison models through a cumulative link model. The Thurstone–Mosteller model with an order effect corresponds to the latent continuous process  $Z_{ij} = \delta + \mu_i - \mu_j + \epsilon_{ij}$ , where  $\epsilon_{ij}$  are uncorrelated standard normal variables. The outcome of the comparison depends on a threshold parameter  $\gamma \geq 0$ , so that  $j$  beats  $i$  if  $Z_{ij} < -\gamma$ ,  $i$  ties with  $j$  if  $-\gamma \leq Z_{ij} < \gamma$  and  $i$  beats  $j$  if  $Z_{ij} \geq \gamma$ . The corresponding cumulative probabilities for the three outcomes given the strengths are

$$\Pr(Y_{ij} \leq x|\boldsymbol{\mu}) = \Phi(c_x - \mu_i + \mu_j), \quad x \in \{-1, 0, 1\},$$

for cutpoints  $-\infty = c_{-2} < c_{-1} \leq c_0 < c_1 = +\infty$ , with  $c_{-1} = -\gamma - \delta$  and  $c_0 = \gamma - \delta$ . Ridge estimates of the strength parameters are still computed by maximizing a penalized likelihood as in (2), but now with the log likelihood

$$\ell(\boldsymbol{\mu}) = \sum_{(i,j) \in S} \log\{\Phi(c_{y_{ij}} - \mu_i + \mu_j) - \Phi(c_{y_{ij}-1} - \mu_i + \mu_j)\}.$$

As before, we proceed by empirical Bayes estimation of the model parameters  $\gamma$ ,  $\delta$  and  $\lambda$  from the marginal distribution of the paired outcomes and then compute the ridge estimates of the strength parameters. The marginal probabilities of the three possible outcomes are simply calculated by setting all strength parameters to zero, giving  $\Pr(Y_{ij}=x) = \Phi(c_x) - \Phi(c_{x-1})$ , for  $x \in \{-1, 0, 1\}$ . From these marginal probabilities, consistent estimates of  $\gamma$  and  $\delta$  are readily obtained as

$$\hat{\gamma} = \frac{1}{2} \{\Phi^{-1}(1 - \hat{p}_1) - \Phi^{-1}(\hat{p}_{-1})\}, \quad (11)$$

$$\hat{\delta} = \frac{1}{2} \{\Phi^{-1}(\hat{p}_1) - \Phi^{-1}(\hat{p}_{-1})\}, \quad (12)$$

where  $\hat{p}_{-1} = (n+1)^{-1} \sum_{ij} 1(Y_{ij} = -1)$  and  $\hat{p}_1 = (n+1)^{-1} \sum_{ij} 1(Y_{ij} = 1)$ . In tournaments without ties,  $\hat{\gamma} = 0$  as expected.

Having estimated  $\gamma$  and  $\delta$ , we proceed with estimation of the tuning parameter  $\lambda$  by maximizing the pairwise log likelihood (6), where  $p_{rs}(\lambda)$  now denotes the joint probability of the outcomes of two paired comparisons with the same item appearing in either the first or the second position. For example, the probability that

the first item  $i$  obtains a win and a tie in such a couple of paired comparisons is

$$\begin{aligned} p_{10}(\lambda) &= \Pr(Y_{ij}=1, Y_{ik}=0) \\ &= \int_{c_0}^{c_1} \int_{c_{-1}}^{c_0} \phi_2\left(u, v; \frac{\lambda^{-1}}{1+2\lambda^{-1}}\right) dudv \\ &= \int_{\gamma-\delta}^{\infty} \int_{-\gamma-\delta}^{\gamma-\delta} \phi_2\left(u, v; \frac{\lambda^{-1}}{1+2\lambda^{-1}}\right) dudv. \end{aligned}$$

The other bivariate probabilities involved in the pairwise likelihood are similarly computed. In these probabilities  $\gamma$  and  $\delta$  are replaced with their estimates (11) and (12) so that the pairwise likelihood depends on  $\lambda$  only.

Maximization of the pairwise likelihood for  $\lambda$  is conveniently performed by using the robust algorithm of Brent (2013, ch. 4) on the finite interval domain of  $\tau$  from Equation (7), with subsequent back-transformation as in Equation (8). The same small-sample adjustment as before remains available through the pairwise likelihood penalty of Equation (9).

## 4 | Simulations

### 4.1 | True Strengths Normally Distributed

The performance of the pairwise empirical Bayes method is illustrated first by simulations from the Thurstone–Mosteller model with normally distributed strength parameters. We considered 120 scenarios given by the combination of three values for the number of items  $p \in \{20, 40, 60\}$ , eight values for the tuning parameter  $\lambda \in \{2^{-1}, 2^0, \dots, 2^6\}$  and five sample sizes corresponding to five increasing fractions of the paired comparisons in a double round-robin tournament. More specifically, we simulated the full double round-robin tournament and used the first  $m$  rounds for training and the remaining  $2(p-1) - m$  rounds for testing. The number of training rounds was chosen so as to approximate a specified proportion of the tournament paired comparisons. For example, when we consider 20% of paired comparisons for training and the number of items is  $p = 30$ , the training set is taken to be the first 12 rounds because  $0.2 \times (2 \times 29) = 11.6$ . The five sample fractions considered for training in our simulations are 20%, 30%, 40%, 50% and 80% of the paired comparisons in the double round-robin tournament. For each scenario, 1000 replications were generated and analysed.

The performance of our pairwise empirical Bayes method is compared here with three other methods: (1) maximum likelihood estimation, (2) bias-reduced maximum likelihood as in Firth (1993) and (3) maximum ridge-penalized likelihood with  $\lambda$  determined by cross-validation. Bias-reduced maximum likelihood estimation is computed through the **R** (R Core Team 2024) package **brglm2** (Kosmidis 2023). Because paired-comparison models are specified through pairwise differences of strengths, an arbitrary constraint is needed for identification of the maximum likelihood and bias-reduced maximum likelihood estimates; in the simulations, we used the constraint  $\mu_1 = 0$ . Cross-validated ridge estimates were calculated via the **glmnet** package (Friedman, Hastie, and Tibshirani 2010; Tay,

Narasimhan, and Hastie 2023), using tournament rounds as the assessment unit. Paraphrasing the typical language used in model validation, we could call such a design a ‘leave one round out’ cross-validation.

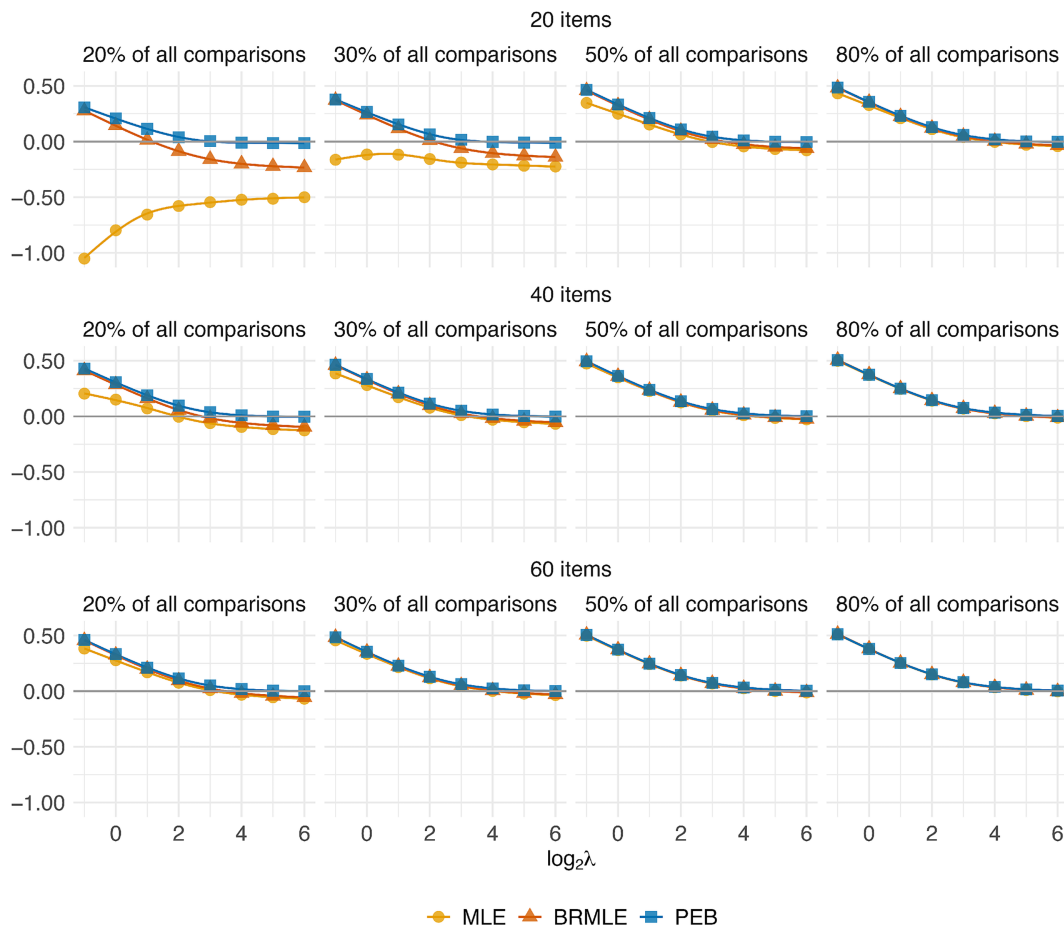
We simulated paired comparisons without ties, in order to focus on the estimation of the strength parameters. Estimation of the tie threshold parameter  $\gamma$  would not be problematic, but bias-reduced maximum likelihood and cross-validated logistic regression are computationally less expensive (for repeated use in a simulation experiment) in paired-comparison models without ties. In our simulations, the order effect parameter was set to  $\delta = 0.2$ , meaning that in approximately 58% of paired comparisons the first item beats the second (because  $\Phi(0.2) \approx 0.58$ ).

The predictive performance of each method is measured by the classical logarithmic score (LS), defined as the negative average log likelihood for the testing paired comparisons computed at the strengths estimated from the training comparisons. As reference, we consider the naive forecast that predicts the result of every paired comparison using the assumed marginal probability of a win for the first item, that is,  $LS_{\text{naive}} = -0.58 \log(0.58) - 0.42 \log(0.42) = 0.68$ . The naive forecast corresponds to predictions for the paired-comparison model with strength parameters all zero,  $\mu_i = 0$  for all  $i = 1, \dots, p$ . With this reference, we define the logarithmic skill score as  $LSS = 1 - LS/LS_{\text{naive}} = 1 - LS/0.68$ . The higher the logarithmic skill score, the better the predictive performance. Better forecasts on average than those obtained by the naive forecast have a logarithmic skill score above zero.

The simulation results are summarized in Figure 1, which displays the average of the simulated logarithmic skill scores. We do not report results based on cross-validation in Figure 1, because they were almost everywhere visually indistinguishable from those based on the adjusted pairwise empirical Bayes method. This finding confirms that pairwise empirical Bayes for the Thurstone–Mosteller model is essentially equivalent to standard cross-validation, but with the substantial computational advantage that no repeated refitting of the model is needed.

We now comment on the results reported in Figure 1. As expected, as the number of items or the fraction of paired comparisons used for training increases, all methods converge to the same predictive performance. Furthermore, as  $\lambda$  increases, estimates of item strengths converge to zero and therefore all methods produce predictions that are on average equivalent to those of the naive forecast. The predictive performance of maximum likelihood estimation is completely unsatisfactory when the training set is small and the number of items is 20, with its average logarithmic skill scores being worse than the naive forecast. The bias-reduced maximum likelihood estimator improves substantially on maximum likelihood in most settings, but with small training sets and small values of  $\lambda$  it still leads to worse logarithmic skill scores on average than the naive forecast.

As shown in Figure 1, pairwise empirical Bayes has uniformly the best predictive performance of all methods considered. Pairwise empirical Bayes was even found to perform relatively well in a still more extreme setting, with  $p = 20$  items and only four paired comparisons per item (approximately 10% of



**FIGURE 1** | Logarithmic skill scores computed for incomplete double round-robin tournaments of different sizes. The columns correspond to various levels of incompleteness (20%, 30%, 50% or 80% of all possible paired comparisons used for training and the rest of paired comparisons for prediction). The rows correspond to different numbers of items (20, 40 and 60). The item strengths are drawn from a zero-mean normal distribution with precision  $\lambda$ . The strengths are estimated by bias-reduced maximum likelihood (BRMLE), maximum likelihood (MLE) and pairwise empirical Bayes (PEB). The reference method is a naive forecast that uses only the assumed marginal probability of a win for the first item and corresponds to a zero logarithmic skill score (grey solid line).

the tournament) used for training; that extreme setting is not displayed in Figure 1 because maximum likelihood performs so poorly that its logarithmic skill score values fall too far below zero.

#### 4.2 | True Strengths Distributed as Student's $t$

The results just described were obtained under the assumption of normally distributed strengths, the assumption used to construct the pairwise empirical Bayes method. Less benign settings for pairwise empirical Bayes are considered next, by drawing the strengths from Student's  $t$  distribution with  $\nu$  degrees of freedom. The simulated strengths are scaled so that their precision continues to be  $\lambda$ ,  $\mu_i \sim t_\nu \sqrt{(\nu - 2)/(\lambda \nu)}$ ,  $\nu > 2$ ,  $i = 1, \dots, p$ , with  $t_\nu$  denoting a Student's  $t$ -distributed variate with  $\nu$  degrees of freedom. We take  $\nu > 2$  since otherwise the variance would be undefined. Drawing the strengths from a  $t$  distribution is of interest because it allows a small number of items to be distinctly stronger or weaker than the rest. Figures S1 and S2 display the average LSs for the same 120 scenarios considered previously, but with strengths simulated from Student's

$t$  distributions with  $\nu = 8$  and  $\nu = 3$  degrees of freedom. As far as predictive performance is concerned, there is no perceptible difference between the  $t_8$  and normal distributions; and even in the more extreme scenario with  $t_3$ -distributed strengths, the differences are rather small. While such simulation evidence is inevitably limited, it does provide some reassurance that the pairwise empirical Bayes method is robust to excess kurtosis in the strengths.

### 5 | Premier League

As an illustration of the ridge method, we consider estimation of the strength of football teams in a round-robin tournament, before the tournament is complete. Since season 1995–1996, the English Premier League has been a double round robin between 20 teams. The current format of the tournament therefore consists of 38 match-weeks. We consider here all the results of the 10,640 matches of the 28 seasons between 1995 and 2023, in order to evaluate the robustness of our method over a long period characterized by substantial changes in (English) football. For each season, we trained the paired-comparison models with



the first 10, 15, 20, 25 and 30 match-weeks and predicted the results of the rest of the tournament.

The problem we consider in this illustration is different from sequential prediction of next week's match results given the team's past results, for which it would be more appropriate to consider dynamic paired-comparison models, such as the state-space approaches of Knorr-Held (2000) and Fahrmeir and Tutz (1994), the exponential smoothing of Cattelan, Varin, and Firth (2013), the time-weighted likelihood of McHale and Morton (2011), the kernel smoothing of Bong et al. (2020) or the recent nonparametric spectral ranker of Tian et al. (2024).

In our illustration, the Thurstone–Mosteller model with ties (which here will be called 'draws', the word used in football) was estimated by maximum likelihood estimation using the **polr** function from the **MASS** package (Venables and Ripley 2002), by bias-reduced maximum likelihood for the cumulative logistic regression model developed in Kosmidis (2014) and implemented in the R function **bpolr** available through the supplementary materials of that paper and by our pairwise empirical Bayes method using R code available at <https://github.com/crisvarin/peb>.

As a reference forecast, we considered the naive predictions calculated from long-term frequencies of home wins, draws and away wins pooled across all 28 seasons of the Premier League. Those seasons saw 46% home wins, 25% draws and 29% away wins out of a total number of 10,640 matches. The LS for the naive forecast is thus  $LS_{\text{naive}} = -0.46 \log(0.46) - 0.25 \log(0.25) - 0.29 \log(0.29) = 1.06$ .

Figure 2 displays boxplots of the collected logarithmic skill scores: Each boxplot is calculated from the 28 logarithmic skill scores for the various seasons, in a given match-week.

The boxplots show that the predictive performance of pairwise empirical Bayes is substantially better than that of maximum likelihood estimation, in particular when the training data are match results from only the first 10 or 15 weeks of the season. At those early stages of the tournament, just as we saw in the simulation studies of Section 4, maximum likelihood very often produces strength-parameter estimates that predict future matches worse than does the naive forecast.


The results here also show that the bias-reduction method of Firth (1993) improves upon maximum likelihood estimation in terms of prediction: In all 28 seasons and for all sizes of the training data, the logarithmic skill scores based on bias-reduced maximum likelihood are better than those of standard maximum likelihood. However, pairwise empirical Bayes typically performs even better than bias-reduced maximum likelihood. For prediction in the early weeks of the season, especially, pairwise empirical Bayes outperforms the other methods by a substantial margin.

## 6 | Final Remarks

The main purposes of this paper were to discuss ridge-penalized estimation of paired-comparison models and to develop the simple and effective pairwise empirical Bayes method for tuning the ridge penalty in the Thurstone–Mosteller model. As expected, the shrinkage provided by a ridge penalty yields predictions with substantially higher precision than standard maximum likelihood estimation. Indeed, when a paired-comparison model is trained with relatively few observations, prediction based on maximum likelihood is actually worse than the naive predictor that uses only the marginal proportions of (home win, away win and draw) outcomes. The poor predictive performance of maximum likelihood when the training set is small is due to overfitting; the shrinkage achieved through a ridge penalty, tuned by the pairwise empirical Bayes method, effectively eliminates that problem.

The simulations of Section 4 show that the amount of shrinkage needed for good predictive performance, especially when the training set and the number of items are both small, is appreciably greater than the shrinkage inherent in the bias-reduced maximum likelihood method of Firth (1993). This is unsurprising: It is well known that unbiasedness typically is not the main consideration for predictive performance.

The authors are not aware of other work focused on estimation of paired-comparison models with a ridge penalty. Some previous work in team sports modelling has considered more traditional linear ridge penalty models to assess the contribution of individual players to a team's performance. A recent example is the adjusted plus/minus regularized model discussed in Matano et al. (2023); see also the references therein.



**FIGURE 2** | Boxplot summary of the distributions, across 28 Premier League seasons, of the logarithmic skill scores computed after 10, 15, 20, 25 and 30 match-weeks. The predictions are computed with strengths estimated using maximum likelihood (MLE), bias-reduced maximum likelihood (BRMLE) and pairwise empirical Bayes (PEB). The reference method is a naive forecast that uses only long-term frequencies of home wins, draws and away wins and corresponds to a zero logarithmic skill score (grey solid line).

## Author Contributions

The two authors worked together on the development of the method, the design of the simulations, the illustrative application and the writing of the article. The first author (CV) implemented the method in the R language and performed all of the reported computations.

## Acknowledgements

The work of both authors was supported by the IRIDE grant 2023 from DAIS, Ca' Foscari University. Open access publishing facilitated by Universita Ca' Foscari, as part of the Wiley - CRUI-CARE agreement.

## Data Availability Statement

The Supporting Information provides the simulation results described in Section 4.2 where strengths are drawn from a Student's  $t$  distribution. The Premier League data and R code to reproduce the illustration of Section 5 are available in the public repository <https://github.com/crisvarin/peb>.

## References

- Agresti, A. 1992. "Analysis of Ordinal Paired Comparison Data." *Journal of the Royal Statistical Society: Series C (Applied Statistics)* 41, no. 2: 287–297.
- Agresti, A. 2002. *Categorical Data Analysis*. London, UK: John Wiley & Sons.
- Aldous, D. 2017. "Elo Ratings and the Sports Model: A Neglected Topic in Applied Probability?." *Statistical Science* 32, no. 4: 616–629.
- Bartholomew, S. R., and M. D. Jones. 2022. "A Systematized Review of Research With Adaptive Comparative Judgment (ACJ) in Higher Education." *International Journal of Technology and Design Education* 32, no. 2: 1159–1190.
- Bong, H., W. Li, S. Shrotriya, and A. Rinaldo. 2020. "Nonparametric Estimation in the Dynamic Bradley–Terry Model." In *Proceedings of the 23rd International Conference on Artificial Intelligence and Statistics*, 3317–3326. Palermo, Italy: PMLR.
- Bradley, R. A., and M. E. Terry. 1952. "Rank Analysis of Incomplete Block Designs: I. The Method of Paired Comparisons." *Biometrika* 39, no. 3/4: 324–345.
- Brent, R. P. 2013. *Algorithms for Minimization Without Derivatives*. Englewood Cliffs N.J. Prentice-Hall.
- Caron, F., and A. Doucet. 2012. "Efficient Bayesian Inference for Generalized Bradley–Terry Models." *Journal of Computational and Graphical Statistics* 21, no. 1: 174–196.
- Cattelan, M. 2012. "Models for Paired Comparison Data: A Review With Emphasis on Dependent Data." *Statistical Science* 27, no. 3: 412–433.
- Cattelan, M., C. Varin, and D. Firth. 2013. "Dynamic Bradley–Terry Modelling of Sports Tournaments." *Journal of the Royal Statistical Society Series C: Applied Statistics* 62, no. 1: 135–150.
- Cox, D. R., and N. Wermuth. 2002. "On Some Models for Multivariate Binary Variables Parallel in Complexity With the Multivariate Gaussian Distribution." *Biometrika* 89, no. 2: 462–469.
- David, H. A. 1988. *The Method of Paired Comparisons* (2nd ed.). London, UK: Griffin.
- Fahrmeir, L., and G. Tutz. 1994. "Dynamic Stochastic Models for Time-Dependent Ordered Paired Comparison Systems." *Journal of the American Statistical Association* 89, no. 428: 1438–1449.
- Firth, D. 1993. "Bias Reduction of Maximum Likelihood Estimates." *Biometrika* 80, no. 1: 27–38.
- Friedman, J., T. Hastie, and R. Tibshirani. 2010. "Regularization Paths for Generalized Linear Models via Coordinate Descent." *Journal of Statistical Software* 33, no. 1: 1–22.
- Glickman, M. E., and H. S. Stern. 2017. "Estimating Team Strength in the NFL." In *Handbook of Statistical Methods and Analyses in Sports*, 129–152. Boca Raton, FL: Chapman and Hall/CRC.
- Knorr-Held, L. 2000. "Dynamic Rating of Sports Teams." *Journal of the Royal Statistical Society: Series D (The Statistician)* 49, no. 2: 261–276.
- Kosmidis, I. 2014. "Improved Estimation in Cumulative Link Models." *Journal of the Royal Statistical Society Series B: Statistical Methodology* 76, no. 1: 169–196.
- Kosmidis, I. 2023. "brglm2: Bias Reduction in Generalized Linear Models." <https://CRAN.R-project.org/package=brglm2>. R package version 0.9.2.
- Kosmidis, I., and D. Firth. 2021. "Jeffreys-Prior Penalty, Finiteness and Shrinkage in Binomial-Response Generalized Linear Models." *Biometrika* 108, no. 1: 71–82.
- Laplace, P. S. 1814. *Essai Philosophique sur les Probabilités*. Paris, France: Courcier.
- Loewen, P. J., D. Rubenson, and A. Spirling. 2012. "Testing the Power of Arguments in Referendums: A Bradley–Terry Approach." *Electoral Studies* 31, no. 1: 212–221.
- Ma, L., W. H. Wong, and A. B. Owen. 2012. "A Sparse Transmission Disequilibrium Test for Haplotypes Based on Bradley–Terry Graphs." *Human Heredity* 73, no. 1: 52–61.
- Matano, F., L. Richardson, T. Pospisil, C. A. Politsch, and J. Qin. 2023. "Augmenting Adjusted Plus-Minus in Soccer With FIFA Ratings." *Journal of Quantitative Analysis in Sports* 19, no. 1: 43–49.
- Maydeu-Olivares, A., and U. Böckenholt. 2005. "Structural Equation Modeling of Paired-Comparison and Ranking Data." *Psychological Methods* 10, no. 3: 285.
- McHale, I., and A. Morton. 2011. "A Bradley–Terry Type Model for Forecasting Tennis Match Results." *International Journal of Forecasting* 27, no. 2: 619–630.
- Mease, D. 2003. "A Penalized Maximum Likelihood Approach for the Ranking of College Football Teams Independent of Victory Margins." *American Statistician* 57, no. 4: 241–248.
- Morris, C. N. 1983. "Parametric Empirical Bayes Inference: Theory and Applications." *Journal of the American Statistical Association* 78, no. 381: 47–55.
- Mosteller, F. 1951. "Remarks on the Method of Paired Comparisons: I. The Least Squares Solution Assuming Equal Standard Deviations and Equal Correlations." *Psychometrika* 16, no. 1: 3–9.
- R Core Team. 2024. *R: A Language and Environment for Statistical Computing*. Vienna, Austria: R Foundation for Statistical Computing. <https://www.R-project.org/>.
- Rafailov, R., A. Sharma, E. Mitchell, S. Ermon, C. D. Manning, and C. Finn. 2023. "Direct Preference Optimization: Your Language Model is Secretly a Reward Model." In *Proceedings of the 37th International Conference on Neural Information Processing Systems*, 53728–53741. New Orleans: Curran Associates, Inc.
- Sheppard, W. F. 1898. "On the Geometric Treatment of the 'Normal Curve' of Statistics With Particular Reference to Correlation and the Theory of Errors." *Proceedings of the Royal Society of London* 62: 170–173.
- Tay, J. K., B. Narasimhan, and T. Hastie. 2023. "Elastic Net Regularization Paths for All Generalized Linear Models." *Journal of Statistical Software* 106, no. 1: 1–31.
- Thompson, W. C., R. H. Grady, E. Lai, and H. S. Stern. 2018. "Perceived Strength of Forensic Scientists—Reporting Statements

About Source Conclusions.” *Law, Probability and Risk* 17, no. 2: 133–155.

Thurstone, L. L. 1927. “A Law of Comparative Judgment.” *Psychological Review* 34, no. 4: 273–286.

Tian, X., J. Shi, X. Shen, and K. Song. 2024. “A Spectral Approach for the Dynamic Bradley–Terry Model.” *Stat* 13, no. 3: e722.

Varin, C., M. Cattelan, and D. Firth. 2016. “Statistical Modelling of Citation Exchange Between Statistics Journals.” *Journal of the Royal Statistical Society Series A: Statistics in Society* 179, no. 1: 1–63.

Varin, C., N. Reid, and D. Firth. 2011. “An Overview of Composite Likelihood Methods.” *Statistica Sinica* 21: 5–42.

Venables, W. N., and B. D. Ripley. 2002. *Modern Applied Statistics with S* (4th ed.). New York: Springer.

### Supporting Information

Additional supporting information can be found online in the Supporting Information section.