

# Bayesian analysis of ranking data with the Extended Plackett-Luce model

Cristina Mollica · Luca Tardella

Received: date / Accepted: date

**Abstract** Multistage ranking models, including the popular Plackett-Luce distribution (PL), rely on the assumption that the ranking process is performed sequentially, by assigning the positions from the top to the bottom one (*forward order*). A recent contribution to the ranking literature relaxed this assumption with the addition of the discrete-valued *reference order* parameter, yielding the novel *Extended Plackett-Luce model* (EPL). Inference on the EPL and its generalization into a finite mixture framework was originally addressed from the frequentist perspective. In this work, we propose the Bayesian estimation of the EPL in order to address more directly and efficiently the inference on the additional discrete-valued parameter and the assessment of its estimation uncertainty, possibly uncovering potential idiosyncratic drivers in the formation of preferences. We overcome initial difficulties in employing a standard Gibbs sampling strategy to approximate the posterior distribution of the EPL by combining the data augmentation procedure and the conjugacy of the Gamma prior distribution with a tuned joint Metropolis-Hastings algorithm within Gibbs. The effectiveness and usefulness of the proposal is illustrated with applications to simulated and real datasets.

**Keywords** Ranking data · Plackett-Luce model · Bayesian inference · Data augmentation · Gibbs sampling · Metropolis-Hastings

## 1 Introduction

Ranking data are common in those experiments aimed at exploring preferences, attitudes or, more generically, choice behavior of a given population towards a set of *items* or *alternatives* [18, 4, 19, 17]. A similar evidence emerges also in the sport context, yielding an ordering of the competitors, for instance players or teams, in terms of their ability or strength, see [5], [15] and [2].

Formally, a *ranking*  $\pi = (\pi(1), \dots, \pi(K))$  of  $K$  items is a sequence where the entry  $\pi(i)$  indicates the rank attributed to the  $i$ -th alternative. Data can be equivalently collected in the ordering format  $\pi^{-1} = (\pi^{-1}(1), \dots, \pi^{-1}(K))$ , such that the generic component  $\pi^{-1}(j)$  denotes the item ranked in the  $j$ -th position. Regardless of the adopted format, ranked observations are multivariate and, more precisely, correspond to permutations of the first  $K$  integers. Indeed, in line with the standards of the ranking literature (some fundamental examples can be found in [3, 10, 1, 8]), the vector notation  $\pi$  (or  $\pi^{-1}$ ) has a dual interpretation. In fact, borrowing from those standards,  $\pi$  is also understood as a bijective mapping  $\pi : I \rightarrow R$  and  $\pi^{-1} : R \rightarrow I$  as its inverse, where  $I = \{1, \dots, K\}$  and  $R = \{1, \dots, K\}$  denote the set of items and the set of ranks both labelled as the first  $K$  integers. This explains why the inverse function notation  $\pi^{-1}$  is used as a symbol representing a generic

---

C. Mollica  
Via del Castro Laurenziano 9, 00161 Roma  
Tel.: +39-338-1823408  
E-mail: cristina.mollica@uniroma1.it

L. Tardella  
Piazzale Aldo Moro 5, 00185 Roma

ordering. Hence, we warn the reader of this twofold meaning, whose correct specific interpretation should be clear from the context.

The statistical literature concerning ranked data modeling and analysis is broadly reviewed in [10] and, more recently, in [1, 8]. Several parametric distributions on the set of permutations  $\mathcal{S}_K$  have been developed and applied to real data. A popular parametric family is the *Plackett-Luce model* (PL), belonging to the class of the so-called *stagewise ranking models*. The PL was introduced by [9] and [13] and has a long history in the ranking literature for its numerous successful applications as well as for still inspiring new research developments. The basic idea is the decomposition of the ranking process into  $K - 1$  stages, concerning the attribution of each position according to the *forward order*, that is, the ordering of the alternatives proceeds sequentially from the most-liked to the least-liked item. The implicit forward order assumption has been released by [11] in the *Extended Plackett-Luce model* (EPL). This relaxation allows for a more flexible dependence structure, hence for a better and possibly more parsimonious fitting of the observed ranking data. The PL extension in [11] relies on the introduction of the *reference order* parameter  $\rho = (\rho(1), \dots, \rho(K))$ , indicating the rank assignment order. Specifically, the generic entry  $\rho(t)$  indicates the rank attributed at the  $t$ -th stage of the ranking process. Thus,  $\rho$  is a discrete parameter given by a permutation of the first  $K$  integers and its estimation was originally considered from the frequentist perspective [11]. However, in that work there was no specific attention to the inferential ability to recover this extra parameter, as well as quantifying the estimation uncertainty. Indeed, the new EPL parameter space is of mixed-type (continuous and discrete) and, as far as the discrete component is concerned, there was little guidance on how to move efficiently towards the global optimal solution and only a multiple starting point approach was conceived. This resulted in a substantial computational burden with increasing  $K$ , which makes it less feasible a bootstrap approach to investigate parameter estimation uncertainty.

In this work, we investigate the estimation of the EPL within the Bayesian domain and detail an original MCMC method to perform approximate posterior inference. In particular, we show how the estimation of the discrete reference order parameter can be effectively addressed by combining the Gibbs sampling (GS) for some components of the parameter space with a Metropolis-Hastings (MH) step based on a joint proposal distribution.

The outline of the article is the following. After a review of the main features of the EPL and the related data augmentation approach with latent variables, the novel Bayesian EPL is introduced in Section 2. The detailed description of the MCMC algorithm to perform approximate posterior inference is presented in Section 3, whereas illustrative applications to both simulated and real ranking data follow in Section 4. Final remarks and hints for future research are discussed in Section 5.

## 2 The Bayesian Extended Plackett-Luce model

### 2.1 Model specification

As described in [11], the probability of a generic ordering under the EPL can be written as

$$\mathbf{P}_{\text{EPL}}(\pi^{-1} | \rho, \underline{p}) = \mathbf{P}_{\text{PL}}(\pi^{-1} \circ \rho | \underline{p}) = \prod_{t=1}^K \frac{p_{\pi^{-1}(\rho(t))}}{\sum_{v=t}^K p_{\pi^{-1}(\rho(v))}} \quad \pi^{-1} \in \mathcal{S}_K, \quad (1)$$

where the symbol  $\circ$  indicates the composition of permutations. Hereinafter, we will shortly refer to the probability distribution in (1) as  $\text{EPL}(\rho, \underline{p})$ . The positive quantities  $\underline{p} = (p_1, \dots, p_K)$ , known as *support parameters*, are proportional to the probabilities for each item to be selected at the first stage of the ranking process and, hence, to be ranked in the position indicated by the first entry of  $\rho$ . Expression (1) reveals that, under the EPL assumption, the probability of a generic ordering is equal to the product of the conditional probability masses of the items selected at each stage, which are obtained by normalizing the corresponding support parameters with respect to the alternatives still available in the choice set. This explains the analogy of the EPL ranking construction with the sampling without replacement of the alternatives in order of preference. The standard PL is a special instance of the EPL with  $\rho = \rho_{\text{F}} = (1, 2, \dots, K)$ , i.e., the *identity permutation* also named

*forward order*; the backward PL is the special case with  $\rho = \rho_B = (K, K-1, \dots, 1) = (K+1) - \rho_F$ , i.e., the *backward order*. A simple example of EPL distribution is described in the Appendix.

As in [12], the data augmentation with the latent quantitative variables  $\underline{y} = (y_{st})$  for  $s = 1, \dots, N$  and  $t = 1, \dots, K$  crucially contributes to make the Bayesian inference for the EPL analytically more tractable. Let  $\underline{\pi}^{-1} = \{\pi_s^{-1}\}_{s=1}^N$  be the observed sample of  $N$  orderings. The complete-data model can be specified as follows

$$\begin{aligned} \pi_s^{-1} | \rho, \underline{p} &\stackrel{\text{iid}}{\sim} \text{EPL}(\rho, \underline{p}) & s = 1, \dots, N, \\ y_{st} | \pi_s^{-1}, \rho, \underline{p} &\stackrel{i}{\sim} \text{Exp} \left( \sum_{\nu=t}^K p_{\pi_s^{-1}(\rho(\nu))} \right) & t = 1, \dots, K, \end{aligned}$$

where the auxiliary variables  $y_{st}$ 's are assumed to be conditionally independent on each other and exponentially distributed with rate parameter equal to the normalization term of the EPL. The complete-data likelihood turns out to be

$$L_c(\rho, \underline{p}, \underline{y}) = \prod_{i=1}^K p_i^N e^{-p_i \sum_{s=1}^N \sum_{t=1}^K y_{st} \delta_{sti}}, \quad (2)$$

where

$$\delta_{sti} = \begin{cases} 1 & \text{if } i \in \{\pi_s^{-1}(\rho(t)), \dots, \pi_s^{-1}(\rho(K))\}, \\ 0 & \text{otherwise,} \end{cases}$$

with  $\delta_{s1i} = 1$  for all  $s = 1, \dots, N$  and  $i = 1, \dots, K$ .

In order to setup all the ingredients needed for a fully Bayesian analysis, we must specify a prior distribution on the unknown parameters  $\rho$  and  $\underline{p}$ . In the absence of a genuinely subjective prior information, we conveniently adopt prior independence and the following marginal distributions for  $\rho$  and  $\underline{p}$

$$\begin{aligned} \rho &\sim \text{Unif} \{ \mathcal{S}_K \} \\ p_i &\stackrel{\text{iid}}{\sim} \text{Ga}(c, d) \quad i = 1, \dots, K. \end{aligned}$$

The adoption of independent Gamma densities for the support parameters is motivated by the conjugacy with the model, as apparent by checking the form of the likelihood (2). In our analysis, we considered the hyperparameter setting  $c = d = 1$ .

### 3 Bayesian estimation of the EPL via MCMC

#### 3.1 Attempts towards posterior sampling

In this section, we describe an original MCMC algorithm to carry out Bayesian inference for the EPL. We propose a tuned joint Metropolis-within-Gibbs sampling (TJM-within-GS) as simulation-based method to approximate the posterior distribution. Its distinguishing feature is the use of a suitably tuned MH algorithm relying on a joint proposal on the mixed-type parameter components  $(\rho, \underline{p})$ , combined with two other kernels which act more specifically on the discrete component  $\rho$  and then on the continuous  $(\underline{p}, \underline{y})$ . In fact, although partitioning the parameter vector of the augmented space by using the components of  $(\rho, \underline{p}, \underline{y})$  makes it possible to derive standard full-conditionals for a GS, this ends up being a difficult-to-implement and unsuccessful strategy in practice. On one hand, the discrete full-conditional for  $\rho$  involves a support with a rapidly-increasing cardinality with respect to  $K$ . On the other hand, the meaning of the support parameters is strictly related to the reference order (mainly to its first component) and this means that, when one keeps  $\underline{p}$  fixed, the full-conditional of  $\rho$  is very unlikely to achieve high posterior density regions.

Indeed, if one aims at implementing an MCMC algorithm for a distribution with a mixed-type nature, one can try to split the task by separating the update of the continuous  $(\underline{p}, \underline{y})$ , for which the full-conditionals can be easily derived, from the update of the discrete  $\rho$ . For the latter, in fact,

a full conditional could be derived only in examples where  $K$  is small and perhaps a Metropolis update could be used instead.

In order to setup a proposal distribution for the reference order parameter, we could rely on a sequential urn model not necessarily with the same urn composition. In that case, let us denote with  $\underline{\lambda} = (\lambda_{tj})$  the conditional multinomial probabilities governing the sequential rank attribution path such that, for each stage  $t = 1, \dots, K$ , one has

$$\lambda_{tj} = \mathbf{P}(\rho(t) = j | \rho(1), \dots, \rho(t-1)) \quad j \in R_t,$$

where  $R_t = R \setminus \{\rho(1), \dots, \rho(t-1)\}$  is the set of positions which are still available and can be assigned at stage  $t$ . Of course, at stage  $t = 1$ , one has  $R_1 = R$ ; at stage  $t = K$ , one has  $\lambda_{Kj} = 1$  for the last available rank  $j \in R_K$ . A possible proposal distribution to be employed in the MH step for sampling the reference order  $\rho$  could have the following form

$$\mathbf{P}(\rho) = \prod_{t=1}^K \mathbf{P}(\rho(t) | \rho(1), \dots, \rho(t-1)) = \prod_{t=1}^K \mathbf{P}(\underline{c}_t | \underline{c}_1, \dots, \underline{c}_{t-1}) = \prod_{t=1}^K \prod_{j \in R_t} \lambda_{tj}^{c_{tj}},$$

where  $\underline{c}_t = (c_{tj})_{j \in R_t}$  indicates the vector with binary components defined as follows

$$c_{tj} = I_{[\rho(t)=j]} = \begin{cases} 1 & \text{at stage } t \text{ rank } j \text{ is assigned,} \\ 0 & \text{otherwise.} \end{cases}$$

Nevertheless, preliminary implementations of a MH step of this kind on synthetic data revealed some critical issues in achieving high posterior density regions. Some attempts are reported in Section 1 of the supplementary material. This suggested, for the reasons described above, that a joint proposal distribution of the reference order and the support parameters allows for an adequate mixing of the resulting Markov chain.

This argument inspired us a working MCMC algorithm which combines three different Markov kernels in the spirit of hybrid strategies described in Section 2.4 of [16]. Of course, each Markov kernel has our target posterior as invariant distribution. The most original of these kernels relies on the MH algorithm with a joint proposal strategy for proposing the update of  $(\rho, \underline{p})$ . The other two kernels are nothing but simple GS kernels limited to specific blocks of the parameter components.

### 3.2 Tuned Joint Metropolis-Hastings step

Since the marginal distribution of the observed items ranked in the first stage basically embodies a considerable portion of the information on the support parameters, we decided to decompose the joint proposal distribution starting from the first component of the reference order, immediately followed by all the support parameter components conditionally in  $\rho(1)$ . In fact, in this way, the chances of proposing acceptable candidates are indeed increased. In order to sample candidate values for  $\rho$  and  $\underline{p}$  simultaneously, we devised a Metropolis kernel  $\mathcal{K}_{\text{TJM}}$  based on a joint proposal distribution  $g(\rho, \underline{p})$  with a specific decomposition of the dependence structure, given by

$$g(\rho, \underline{p}) = g(\rho(1)) \times g(\underline{p} | \rho(1)) \times \prod_{t=2}^K g(\rho(t) | \underline{p}, \rho(1), \dots, \rho(t-1)). \quad (3)$$

The dependence structure in (3) shows that, after drawing the first component of  $\rho$ , the proposal can exploit the sample evidence on the support parameters to guide the simulation of the remaining candidate entries of the reference order. In fact, it is expected that, if  $\rho(1)$  is the rank assigned at the first stage, the observed marginal frequencies of the items ranked in position  $\rho(1)$  can be regarded as an estimate of the support parameters  $\underline{p}$ . In this way, the two simulated parameter vectors are linked to each other and the joint proposal should better mimic the target posterior density and, thus, get a better mixing chain. Candidate values  $(\tilde{\rho}, \tilde{\underline{p}})$  are jointly generated according to the following scheme:

1. *sample the first component of  $\tilde{\rho}$  (stage  $t = 1$ ):*

$$\tilde{c}_1 \sim \text{Multinomial}(1, \tilde{\lambda}_1 = (\tilde{\lambda}_{1j})_{j \in R}) \Rightarrow \tilde{\rho}(1) = \prod_{j \in R} j^{\tilde{c}_{1j}}.$$

In our application, we set  $\tilde{\lambda}_{1j} = \mathbf{P}(\tilde{\rho}(1) = j) = \mathbf{P}(\tilde{c}_{1j} = 1) = 1/K$  for all  $j \in R$ ;

2. *sample the support parameters:*

$$\tilde{p} | \tilde{\rho}(1) \sim \text{Dirich}(\alpha_0 \times r_{\tilde{\rho}(1)}),$$

where Dirich denotes the Dirichlet distribution,  $\alpha_0$  is a scalar tuning parameter and  $r_{\tilde{\rho}(1)}$  is the marginal distribution of the items with rank  $\tilde{\rho}(1)$ , that is, the vector collecting the relative frequencies of the times that each item  $i = 1, \dots, K$  has been ranked in position  $\tilde{\rho}(1)$ . Specifically, the  $i$ -th entry of  $r_{\tilde{\rho}(1)}$  is

$$r_{\tilde{\rho}(1)i} = \frac{1}{N} \sum_{s=1}^N I_{[\pi_s^{-1}(\tilde{\rho}(1))=i]}$$

and, in our analysis, we set  $\alpha_0 = 500$ ;

3. *sample iteratively the remaining entries of  $\tilde{\rho}$  (from stage  $t = 2$  to stage  $t = K - 1$ ):*

in this step, we deem it convenient to split the sampling of the remaining entries of  $\rho$  sequentially. Unfortunately, it is computationally demanding to compute the conditional distribution of the reference order component at stage  $t$  given the reference order components selected up to stage  $t - 1$ , the support parameters and the data. Hence, in order to surrogate it, we exploit the marginal bivariate distributions of the items in the two consecutive stages  $(t - 1, t)$ , resulting from the observed  $K$ -variate contingency table. Indeed, once selected the reference order component at stage  $t - 1$ , we compute the observed contingency table  $\tilde{\tau}_{tj}$  having as first margin the item placed at the selected reference order component  $\tilde{\rho}(t - 1)$  and as second margin the item placed at the reference order component  $\tilde{\rho}(t)$  which can be possibly selected at the next stage among the remaining positions in  $R_t$ . The generic entry of the contingency table is

$$\tilde{\tau}_{ii'tj} = \sum_{s=1}^N I_{[\pi_s^{-1}(\tilde{\rho}(t-1))=i, \pi_s^{-1}(j)=i']},$$

corresponding to the actually observed joint frequency counting how many times each item  $i$  in the previous stage is followed by any other item  $i'$  ranked  $j$ -th at the next stage.

We then compare these frequencies with the corresponding expected frequencies  $\tilde{E}_{ii't}$  under the EPL by using a Monte Carlo approximation

$$\tilde{\eta}_s^{-1}(1), \dots, \tilde{\eta}_s^{-1}(t) | \tilde{p} \stackrel{i}{\sim} \text{PL}(\tilde{p}) \quad s = 1, \dots, N,$$

$$\tilde{E}_{ii't} = \sum_{s=1}^N I_{[\tilde{\eta}_s^{-1}(t-1)=i, \tilde{\eta}_s^{-1}(t)=i']}$$

and compute the following distances

$$\tilde{d}_{tj} = \sum_{i=1}^K \sum_{i'=1}^K (\tilde{\tau}_{ii'tj} - \tilde{E}_{ii't})^2.$$

In fact, these different distances  $(\tilde{d}_{tj})_{j \in R_t}$  are intuitively ranked according to how compatible, hence likely, is the attribution of the  $j$ -th position at stage  $t$  in the light of the observed data. The above distances are then suitably scaled as  $\tilde{z}_{tj} = \frac{\tilde{d}_{tj}}{\sum_{j' \in R_t} \tilde{d}_{tj'}}$  to ensure unit sum. Note that the position corresponding to the smallest distance should be the one to be proposed with the highest probability. Moreover, we should maintain the disparities among the alternative distances so that, when the distances are very similar, the probabilities should be also very similar. To achieve these goals, we simply maintain the same values  $\tilde{z}_{tj}$  reallocating them to the positions  $j \in R_t$  according to the reverse order, so that if  $j_1$  is the position corresponding to the

smallest rescaled distance  $\tilde{z}_{tj_1}$ , after the rearrangement it is assigned a value  $\tilde{q}_{tj_1} = \max_{j \in R_t} \tilde{z}_{tj}$ . Similarly, if  $j_2$  is the position corresponding to the second smallest rescaled distance  $\tilde{z}_{tj_2}$ , it is assigned a value  $\tilde{q}_{tj_2} = \max_{j \in R_t \setminus j_1} \tilde{z}_{tj}$  and so on. The rearranged relative distances  $\{\tilde{q}_{tj}, j \in R_t\}$  are used as surrogate approximation of the conditional probabilities corresponding to the target distribution. Indeed, we define the proposal probability of assigning position  $j$  at stage  $t$  as

$$\tilde{\lambda}_{tj} = \frac{\tilde{q}_{tj} + h}{\sum_{j' \in R_t} (\tilde{q}_{tj'} + h)},$$

where  $h \in (0, 1/K)$  is a tuning parameter introduced to avoid a too sparse multinomial distribution. As default value we set  $h = 0.1$ . Finally, for  $t = 2, \dots, K-1$ , we sample

$$\tilde{c}_t | \tilde{c}_1, \dots, \tilde{c}_{t-1} \sim \text{Multinomial}(1, \tilde{\lambda}_t = (\tilde{\lambda}_{tj})_{j \in R_t}) \Rightarrow \tilde{\rho}(t) | \tilde{\rho}(1), \dots, \tilde{\rho}(t-1) = \prod_{j \in R_t} j^{\tilde{c}_{tj}}.$$

The resulting joint proposal probability of the candidate values is

$$g(\tilde{\rho}, \tilde{p}) = \text{Dirich}(\tilde{p} | \alpha_0 \times r_{\tilde{\rho}(1)}) \prod_{t=1}^K \prod_{j \in R_t} \tilde{\lambda}_{tj}^{\tilde{c}_{tj}}.$$

Hence, if we denote the observed-data likelihood with  $L(\rho, \underline{p})$ , the acceptance probability is equal to

$$\alpha' = \min \left\{ \frac{g(\rho^{(l)}, \underline{p}^{(l)})}{g(\tilde{\rho}, \tilde{p})} \frac{L(\tilde{\rho}, \tilde{p}) \prod_{i=1}^K \text{Ga}(\tilde{p}_i | c, d)}{L(\rho^{(l)}, \underline{p}^{(l)}) \prod_{i=1}^K \text{Ga}(p_i^{(l)} | c, d)}, 1 \right\}$$

and the MH step ends with the classical acceptance/rejection of the candidate pair

$$(\rho', \underline{p}') = \begin{cases} (\tilde{\rho}, \tilde{p}) & \text{if } \log(u') < \log(\alpha'), \\ (\rho^{(l)}, \underline{p}^{(l)}) & \text{otherwise,} \end{cases}$$

where  $u' \sim \text{Unif}(0, 1)$  and  $(\rho^{(l)}, \underline{p}^{(l)})$  is the current pair.

In order to promote a better mixing, we combine the just illustrated Metropolis kernel  $\mathcal{K}_{\text{TJM}}$  by composing it with two additional kernels, labelled as  $\mathcal{K}_{\text{SM}}$  (Swap Move) and  $\mathcal{K}_{\text{GS}}$  (Gibbs Sampling). Hence, the MCMC simulation is based on the composition kernel  $\mathcal{K} = \mathcal{K}_{\text{TJM}} \circ \mathcal{K}_{\text{SM}} \circ \mathcal{K}_{\text{GS}}$ . Indeed, the  $\mathcal{K}_{\text{SM}}$  kernel focuses on local moves of the discrete component  $\rho$ , whereas  $\mathcal{K}_{\text{GS}}$  aims at improving the mixing of the continuous component. The kernel  $\mathcal{K}_{\text{SM}}$  is illustrated in the next section, while  $\mathcal{K}_{\text{GS}}$  is just a full GS cycle involving the  $(\underline{p}, \underline{y})$  components and is detailed in Section 3.4.

### 3.3 Swap move

We remind that the values  $(\rho', \underline{p}')$  have to be regarded as temporary parameter drawings. We decided to accelerate the exploration of the parameter space by including an intermediate kernel  $\mathcal{K}_{\text{SM}}$  of the  $\rho$  component only, that attempts a possible local move with respect to the current value. We label the possibly successful update of this kernel as *Swap Move* (SM). In fact, the idea relies on a random swap of two adjacent components of  $\rho'$ . Specifically, one first simulate

$$t^* \sim \text{Unif}\{1, \dots, K-1\}$$

and then define the new candidate as

$$\rho'' = (\rho'(1), \dots, \rho'(t^* + 1), \rho'(t^*), \dots, \rho'(K)).$$

Finally, by computing the acceptance probability as

$$\alpha'' = \min \left\{ \frac{L(\rho'', \underline{p}')}{L(\rho', \underline{p}')}, 1 \right\},$$

the sampled value of the reference order at the  $(l+1)$ -th iteration turns out to be

$$\rho^{(l+1)} = \begin{cases} \rho'' & \text{if } \log(u'') < \log(\alpha''), \\ \rho' & \text{otherwise,} \end{cases}$$

where  $u'' \sim \text{Unif}(0, 1)$ .

### 3.4 Tuned Joint Metropolis-within-Gibbs-sampling

At the generic iteration  $(l + 1)$ , the TJM-within-GS iteratively alternates the following simulation steps

$$\begin{aligned} \rho^{(l+1)}, \underline{p}' | \underline{\pi}^{-1}, \rho^{(l)}, \underline{p}^{(l)} &\sim \mathcal{K}_{\text{TJM}} \circ \mathcal{K}_{\text{SM}}, \\ y_{st}^{(l+1)} | \pi_s^{-1}, \rho^{(l+1)}, \underline{p}' &\sim \text{Exp} \left( \sum_{i=1}^K \delta_{sti}^{(l+1)} p'_i \right), \\ p_i^{(l+1)} | \underline{\pi}^{-1}, \underline{y}^{(l+1)}, \rho^{(l+1)} &\sim \text{Ga} \left( c + N, d + \sum_{s=1}^N \sum_{t=1}^K \delta_{sti}^{(l+1)} y_{st}^{(l+1)} \right). \end{aligned}$$

In the last two lines, the full-conditional of the unobserved continuous variables  $y$ 's is derived from the construction of the complete-data model specified in Section 2.1, whereas the full-conditional of the support parameters is deduced by the partial conjugate structure, requiring a straightforward update of the corresponding Gamma priors.

## 4 Illustrative applications

### 4.1 Simulated data

In order to verify the efficacy of our MCMC strategy, as well as the ensuing inferential ability of the proposed Bayesian framework, we have setup the following simulation plan: we considered a grid of simulation settings combining different number of items  $K \in \{5, 10, 20\}$  and sample sizes  $N \in \{50, 200, 1000, 10000\}$ . For each distinct pair  $(K, N)$ , we replicated 100 times the simulation of datasets from the EPL model by varying the parameter configuration: for each replication  $R = 1, \dots, 100$ , after fixing a true reference order  $\dot{\rho}^{(R)}$  and a true support parameter vector  $\underline{\dot{p}}^{(R)}$  by drawing  $\dot{\rho}^{(R)}$  uniformly in  $\mathcal{S}_K$  and all the components  $\dot{p}_i^{(R)}$  i.i.d. from a uniform distribution, we simulated  $N$  i.i.d. orderings of  $K$  items from the EPL( $\dot{\rho}^{(R)}, \underline{\dot{p}}^{(R)}$ ), forming the sample matrix  $\underline{\pi}_{(R)}^{-1}$ .

For each replication, we ran the TJM-within-GS described in Section 3 for a total of 10000 iterations and 2000 were discarded as burn-in phase. As a starting point for the MCMC simulation, we randomly selected  $\rho$  and all the components  $p_i$ . The resulting 8000 simulations were used to approximate the posterior distribution. Several measures were computed for assessing the inferential performance of the proposed MCMC sampler to recover the actual parameter values  $(\dot{\rho}^{(R)}, \underline{\dot{p}}^{(R)})$  and the true EPL modal ranking  $\dot{\sigma}^{(R)}$  of  $K$  the items. The ability of the approximated marginal posterior distributions to infer the known EPL parameters and modal sequences was quantified with five criteria and by averaging the results over the 100 replicated datasets  $\underline{\pi}_{(R)}^{-1}$ :

1. the percentage of matches between the true reference order  $\dot{\rho}^{(R)}$  and the most frequently simulated reference order  $\hat{\rho}^{(R)}$

$$\%_{\text{rec}\hat{\rho}} = \sum_{R=1}^{100} I_{[\dot{\rho}^{(R)} = \hat{\rho}^{(R)}]};$$

2. average posterior probability on the matching estimate  $\hat{\rho}^{(R)} = \dot{\rho}^{(R)}$

$$\bar{\mathbf{P}}(\rho = \hat{\rho} | \underline{\pi}^{-1}) = \frac{1}{\%_{\text{rec}\hat{\rho}}} \sum_{R=1}^{100} I_{[\dot{\rho}^{(R)} = \hat{\rho}^{(R)}]} \mathbf{P}(\rho = \dot{\rho}^{(R)} | \underline{\pi}_{(R)}^{-1});$$

3. average relative Kendall distance between  $\dot{\rho}^{(R)}$  and  $\hat{\rho}^{(R)}$

$$\bar{d}_{\text{K}}(\dot{\rho}, \hat{\rho}) = \frac{1}{100} \sum_{R=1}^{100} d_{\text{K}}(\dot{\rho}^{(R)}, \hat{\rho}^{(R)});$$

4. average scaled total variation distance between  $\underline{\hat{p}}^{(R)}$  and  $\underline{\hat{p}}^{(R)} = \mathbb{E}[\underline{p} | \rho = \hat{\rho}^{(R)}, \underline{\pi}_{(R)}^{-1}]$ , given by

$$\bar{d}_{\text{TV}}(\underline{\hat{p}}, \underline{\hat{p}}) = \frac{1}{\% \text{rec}_{\hat{\rho}}} \sum_{R=1}^{100} I_{[\hat{\rho}^{(R)} = \hat{\rho}^{(R)}]} d_{\text{TV}}(\underline{\hat{p}}^{(R)}, \underline{\hat{p}}^{(R)})$$

where  $d_{\text{TV}}(\underline{\hat{p}}^{(R)}, \underline{\hat{p}}^{(R)}) = \frac{1}{2} \sum_{i=1}^K |\hat{p}_i^{(R)} - \hat{p}_i^{(R)}|$ .

5. average relative Kendall distance between  $\hat{\sigma}^{(R)}$  and  $\hat{\sigma}^{(R)}$

$$\bar{d}_{\text{K}}(\hat{\sigma}, \hat{\sigma}) = \frac{1}{100} \sum_{R=1}^{100} d_{\text{K}}(\hat{\sigma}^{(R)}, \hat{\sigma}^{(R)}),$$

where  $\hat{\sigma}^{(R)}$  is the true EPL modal ranking and  $\hat{\sigma}^{(R)}$  is the *maximum a posteriori* (MAP) estimate.

The first three criteria concern the estimation of  $\rho$ , whereas the last two measures are used to assess the inferential performance, respectively, for  $\underline{p}$  and  $\sigma$ . Note that, since the range of the Kendall distance depends on  $K$ , we rescaled it by  $K(K-1)/2$  in order to have a relative index ranging over  $[0, 1]$  regardless of  $K$ .

Additionally, we have verified whether the choice of the starting point of the MCMC simulation, especially the starting reference order, could make a difference in terms of convergence speed and its possible impact on the inferential ability to recover the true  $\hat{\rho}$ . So, for all the same simulated datasets  $\underline{\pi}_{(R)}^{-1}$ , we re-ran our MCMC sampling starting from the *known* true  $\hat{\rho}^{(R)}$  instead of a random starting reference order. Indeed, on average, we have verified no noteworthy change of inferential ability, as testified by the values of the performance criteria reported in Table 2 of the supplementary material. This can be interpreted as a sign of an MCMC algorithm with a suitable mixing behaviour not strongly dependent on the starting point.

Regarding the estimation of  $\rho$ , for each replication  $R$ , we focussed on the approximated marginal posterior distribution of the reference order, denoted as  $\mathbf{P}(\rho | \underline{\pi}_{(R)}^{-1})$ , and used the corresponding posterior mode  $\hat{\rho}^{(R)}$  as the point MAP estimate of  $\hat{\rho}^{(R)}$ . In the first column of Table 1, we can appreciate how frequently the posterior mode  $\hat{\rho}^{(R)}$  matches the true reference order  $\hat{\rho}^{(R)}$  in terms of the percentage of true recoveries ( $\% \text{rec}_{\hat{\rho}}$ ). We note that this percentage consistently grows with  $N$  and, on average, a larger portion of posterior mass  $\mathbf{P}(\rho = \hat{\rho} | \underline{\pi}^{-1})$  is assigned to the matching mode (second column of Table 1). Moreover, by considering all the replications, even those in which there is no match, the posterior mode ensures a consistently decreasing (with  $N$ ) average relative Kendall distance between the true reference order and the estimated one (third column of Table 1). This means that the whole posterior distribution is consistently concentrating around  $\hat{\rho}^{(R)}$ .

Also for the parameters  $\underline{p}$  and  $\sigma$ , the inferential results improve with increasing  $N$ , although with a slower rate for larger values of  $K$  (forth and fifth columns of Table 1). In fact, if we look at a fixed  $N$ , we observe that the overall inferential properties worsen for increasing values of  $K$ , as expected. Indeed, for larger values of  $K$ , the cardinality of the finite space  $\mathcal{S}_K$  of the discrete parameter  $\rho$  grows factorially with  $K$  and this increases the chance that a slightly more diffuse posterior distribution might miss to select the true  $\hat{\rho}$  as the MAP estimate. Moreover, the diffuseness of the posterior is also related to the more likely closeness of some components of the true support parameters  $\hat{p}^{(R)}$ , which weakens the identifiability of the reference order. In fact, in the extreme case when all the support parameter components are the same, we know that the reference order is not identifiable.

Finally, we carried out a sensitivity analysis to investigate the possible effect of the hyperparameter and tuning value setting on the performance of our MCMC algorithm. The reader can find a detailed description of this study in the supplementary material.

## 4.2 Application to the sport data

We applied our Bayesian EPL to the **sport** dataset included in R package **Rankcluster** [6], where  $N=130$  students at the University of Illinois were asked to rank  $K = 7$  sports in order of preference: 1=Baseball, 2=Football, 3=Basketball, 4=Tennis, 5=Cycling, 6=Swimming and 7=Jogging. Prior



Table 1: Inferential performance of the TJM-within-GS algorithm on simulated data with different settings of  $K$  and  $N$  and random initialization.

$(K, N)$	$\% \text{rec}_{\hat{\rho}}$	$\bar{\mathbf{P}}(\rho = \hat{\rho}   \underline{\pi}^{-1})$	$\bar{d}_K(\hat{\rho}, \hat{\rho})$	$\bar{d}_{\text{TV}}(\hat{p}, \hat{p})$	$\bar{d}_K(\hat{\sigma}, \hat{\sigma})$
(5,50)	48	0.521	0.212	0.060	0.288
(5,200)	79	0.798	0.065	0.033	0.138
(5,1000)	97	0.975	0.008	0.014	0.087
(5,10000)	100	1.000	0.000	0.006	0.048
(10,50)	1	0.123	0.329	0.161	0.411
(10,200)	20	0.426	0.156	0.077	0.350
(10,1000)	64	0.738	0.062	0.044	0.259
(10,10000)	100	0.998	0.000	0.032	0.201
(20,50)	0	–	0.459	–	0.496
(20,200)	0	–	0.348	–	0.472
(20,1000)	3	0.070	0.195	0.087	0.449
(20,10000)	76	0.826	0.031	0.075	0.412

to the Bayesian EPL analysis, sample heterogeneity has been preliminary investigated by using the EPL mixture methodology presented by [11], that suggested the presence of two preference groups labelled as  $G_1$  and  $G_2$  with sizes  $N_1 = 62$  and  $N_2 = 68$  (best fitting 2-component EPL mixture with  $\text{BIC}=2131.25$ ). We then decided to estimate the Bayesian EPL separately on the two clusters. Furthermore, our modeling proposal has been compared with two special cases thereof: the forward (standard) PL and the backward PL, respectively labelled as  $\text{EPL-}\rho_{\text{F}}$  and  $\text{EPL-}\rho_{\text{B}}$ . To inspect the fitting gain obtained with the EPL parametric class, the popular *Deviance Information Criterion* (DIC) developed by [14] was adopted.

To fit the Bayesian EPL, we ran the TJM-within-GS for a total of 20000 iterations and 10000 samples were discarded as burn-in phase. The MCMC algorithm was launched with four randomly dispersed starting points to explore the mixing performance of the sampler over the parameter space. Indeed, the four chains were found to be consistent with respect to the initial values. In this regard, Figure 1 shows the traceplots of the log-posterior values for the four MCMC chains launched with alternative starting values for the two subsamples, where it is apparent that the MCMC algorithm exhibits a good mixing behavior and a satisfactory convergence regardless of the initialization. The graphical inspection was confirmed by Geweke’s convergence diagnostics (equal to 1.29, 1.29, 0.33 and 1.68 for  $G_1$  and to 0.71, 1.54, 0.16 and 1.56 for  $G_2$ ) and Gelman and Rubin’s convergence diagnostics (equal to 1.002 for  $G_1$  and to 1.004 for  $G_2$ ). Convergence diagnostics for the traces of the support parameters are shown in the supplementary material.

The DIC for the inference on the first subsample  $G_1$  highlights evidence in favor of the conventional PL, corresponding to the minimum value highlighted in bold in the first column of Table 2. Indeed, this result is recovered through the estimation of the more general EPL, whose DIC is very

Table 2: DIC values for the Bayesian EPL models fitted to the two subsamples of the `sport` dataset. Optimal (minimum) values of the criterion are highlighted in bold.

	$G_1$	$G_2$
$\text{EPL-}\rho_{\text{F}}$	<b>898.25</b>	1125.42
$\text{EPL-}\rho_{\text{B}}$	934.56	1083.95
EPL	898.60	<b>1021.48</b>

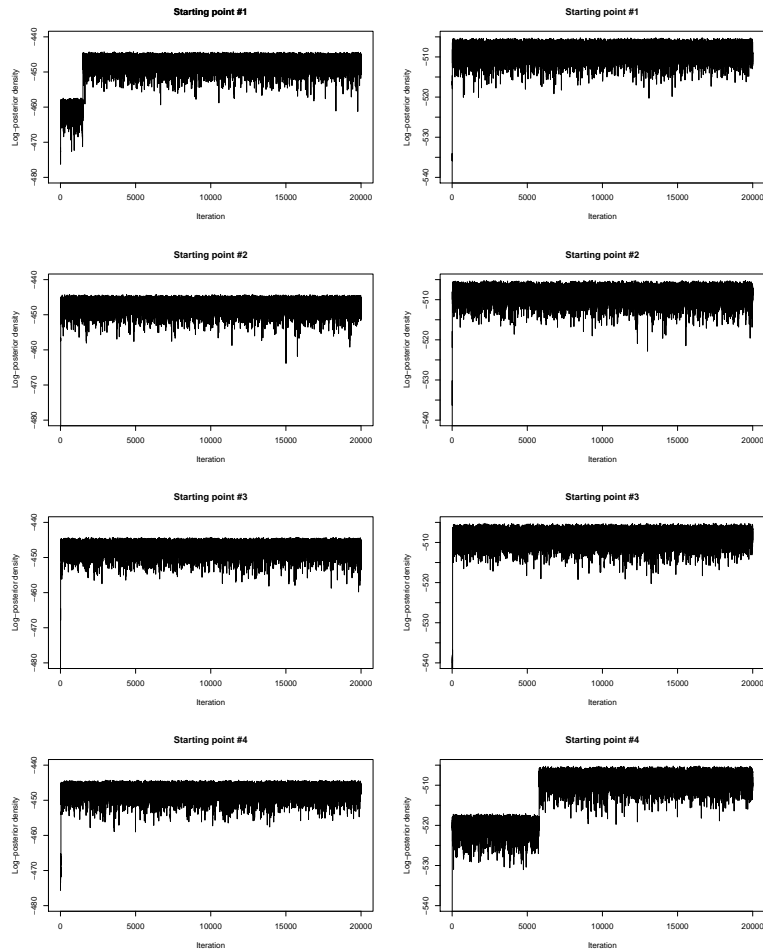


Fig. 1: Traceplots of the log-posterior values for the MCMC chains launched with four randomly dispersed starting values for the two subsamples of the `sport` dataset: log-posterior values for the subsample  $G_1$  with  $N_1 = 62$  units are plotted in the left panel, whereas the log-posterior values for the subsample  $G_2$  with  $N_2 = 68$  units are plotted in the right panel.

close to that of the PL. For the second subsample  $G_2$ , instead, the DIC indicates a significant improvement of the fit associated to the EPL family with respect to the two considered restrictions thereof (second column of Table 2).

Figure 2 shows the approximations of the posterior distribution on the reference order for the two subsamples of the `sport` dataset. The corresponding posterior modal reference orders turned out to be  $\hat{\rho}_1 = (2, 1, 3, 5, 4, 6, 7)$  and  $\hat{\rho}_2 = (5, 6, 4, 7, 3, 1, 2)$ , i.e., a reference order very similar to the forward one and an alternative reference sequence suggesting a preference elicitation starting from the middle-bottom positions. The posterior means and standard deviations of the group-specific support parameters conditionally on the MAP estimates  $\hat{\rho}_1$  and  $\hat{\rho}_2$  are detailed in Table 3. Those obtained from the EPL- $\rho_F$  and the EPL- $\rho_B$  are reported in the supplementary material. The MAP estimates of the group-specific modal orderings under the EPL scenario turn out to be (2=Football, 1=Baseball, 3=Basketball, 5=Cycling, 4=Tennis, 6=Swimming, 7=Jogging) and (6=Swimming, 7=Jogging, 4=Tennis, 3=Basketball, 2=Football, 1=Baseball, 5=Cycling) indicating opposite preferences in the two subsamples towards team and individual sports. Actually, we verified that these modal sequences are very similar to those found with the clustering approach described by [10].

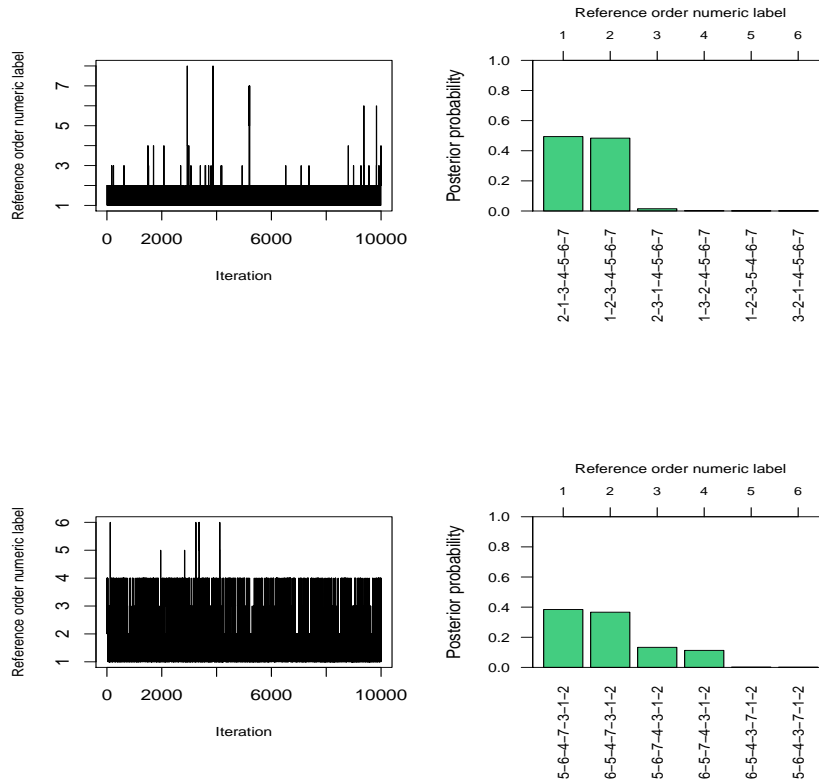


Fig. 2: Traceplots (left) and top posterior probabilities (right) for the reference order parameter estimated from the two subsamples of the `sport` dataset. In the traceplots, sampled reference orders are renamed with numeric labels according to the decreasing order of the posterior probability masses, so label 1 is associated to the MAP reference order. The correspondence between sampled reference orders and numeric labels is shown on the horizontal axes of the barplots.

#### 4.3 Application to the sushi data

For the second real data application, we considered a subsample of the popular `sushi` data concerning preferences on  $K = 10$  types of sushi, namely 1=Shrimp, 2=Sea eel, 3=Tuna, 4=Squid, 5=Sea urchin, 6=Salmon roe, 7=Egg, 8=Fatty tuna, 9=Tuna roll, 10=Cucumber roll. See [7] for further details on the sushi dataset. Specifically, we focused on the  $N=134$  sequences regarding rankers who declared to have most longly lived in the 39th prefecture until 15 years old.

Besides the forward and backward PL, also the Bayesian distance-based model with the Kendall distance (DB-Kend) was considered as a possible competitor of the EPL to learn preferences on the sushi types. The latter was fitted to the data by employing the functions of the R package `BayesMallows`, recently released on CRAN [18] and including the entire sushi dataset of 5000 complete rankings. As in the previous application, four chains of the TJM-within-GS procedure were launched with randomly dispersed starting points and the DIC was used to select the optimal model.

We checked the effective achievement of convergence of the MCMC output with the analysis of the traceplots and the computation of relevant diagnostics. For the log-posterior values we obtained Geweke's z-scores equal to 1.73, 1.61, 0.87 and 0.64 and Gelman and Rubin's potential scale reduction factor equal to 1.001. Convergence diagnostics for the traces of the support parameters are shown in the supplementary material. We then focussed on model comparison and parameter interpretation. For the sushi data analysis, we noticed a greater variability associated

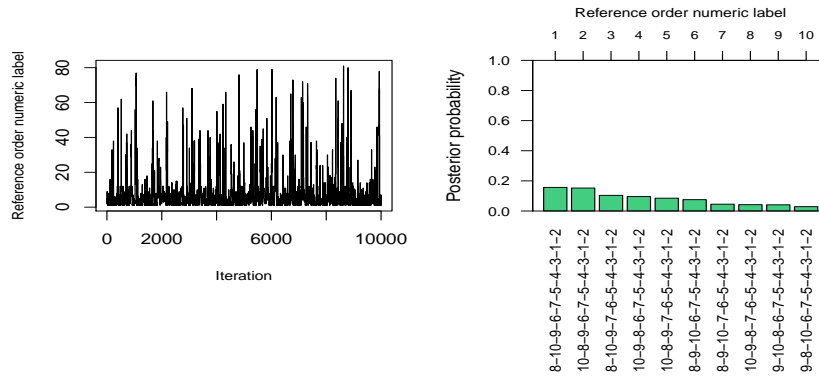


Fig. 3: Traceplots (left) and top posterior probabilities (right) for the reference order parameter estimated from the subsample of the `sushi` dataset. In the traceplots, sampled reference orders are renamed with numeric labels according to the decreasing order of the posterior probability masses, so label 1 is associated to the MAP reference order. The correspondence between sampled reference orders and numeric labels is shown on the horizontal axes of the barplots.

to the posterior distribution of the reference order, where probability masses turn out to be diffuse on a neighbourhood close to the backward path (Figure 3). This variability represents the uncertainty on the reference order estimation which is overlooked by the restricted  $EPL-\rho_B$ . The DIC indicates a remarkable superiority of  $EPL-\rho_B$  over  $EPL-\rho_F$  (DIC=3819.65 vs DIC=3885.94), but the more flexible EPL class, which properly accounts for the uncertainty on  $\rho$ , produces a further improvement of the final fit (DIC=3815.07). The new Bayesian EPL outperforms also the DB-Kend (DIC=3838.95). The posterior means and standard deviations of the EPL support parameters are reported in Table 4. Those obtained from the  $EPL-\rho_F$  and the  $EPL-\rho_B$  are reported in the supplementary material. The MAP estimate of the modal ordering under the EPL scenario is equal to (5=Sea urchin, 8=Fatty tuna, 1=Shrimp, 6=Salmon roe, 3=Tuna, 4=Squid, 2=Sea eel, 10=Cucumber roll, 9=Tuna roll, 7=Egg).

## 5 Conclusions

We have addressed some relevant issues in modelling and inferring choice behavior and preferences. The standard PL for complete rankings relies on the hypothesis that the probability of a particular ordering does not depend on the subset of items from which one can choose (Luce's Axiom).

Table 3: Posterior means and standard deviations (s.d.) of the EPL support parameters for the two subsamples of the `sport` dataset.

Sport	$G_1$		$G_2$	
	$\hat{p}_i$	s.d.	$\hat{p}_i$	s.d.
Baseball	0.3743	(0.040)	0.2642	(0.031)
Football	0.2102	(0.029)	0.3194	(0.035)
Basketball	0.1888	(0.027)	0.1584	(0.022)
Tennis	0.0697	(0.012)	0.0795	(0.012)
Cycling	0.0856	(0.014)	0.0910	(0.014)
Swimming	0.0457	(0.008)	0.0501	(0.009)
Jogging	0.0256	(0.006)	0.0374	(0.007)

Although this can be considered a very strong assumption, the widespread use of the PL suggested to exploit it as a building block to gain more flexibility. In particular, [11] explored the possibility of adding flexibility with the help of two main ideas: i) the use of an additional discrete parameter, the reference order, specifying the order of the ranks sequentially assigned by the individual and which should be inferred from the data; ii) the finite mixture of PL distributions enriched with the reference order parameter (EPL mixture).

In this paper, we have focussed on i) and developed a methodology for a fully Bayesian estimation of the EPL distribution. We have devised a hybrid strategy by combining appropriately tuned MH kernels and the GS. This allows for a successful exploration of the whole mixed-type parameter space. The resulting Bayesian inference turned out to be effective in inferring the true reference order and provided a suitable MCMC simulation which can be used to quantify the uncertainty on the EPL parameters. We stress that the previous attempts to infer on the underlying reference order were limited to the maximum likelihood point estimation via EM algorithm, which required a brute-force multiple starting point strategy in order to attempt to reach the global maximum. We note that, in order to assess the uncertainty on the reference order parameter, one should employ a bootstrap approach which, of course, implies a greater computational effort. Moreover, as  $K$  increases, a fixed number of multiple starting points becomes a factorially decreasing (hence negligible) fraction of all the possible initializations and, obviously, this hampers the possibility of achieving eventually the global optimum in applications with larger  $K$ . On the other hand, the hybrid MCMC strategy provided valid results which are less sensitive to the starting reference order. Hence, our Bayesian methodology provides both an appropriate way to assess parameter uncertainty and a computational improvement for inferring the EPL.

Simulation studies explored the sensitivity of TJM-within-GS to the prior hyperparameters as well as tuning constants. They also evaluated its efficacy in recovering the true EPL parameters generating the data. Finally, the novel parametric approach was successfully applied to two real datasets concerning the analysis of preference patterns.

Moreover, our model setup gains more insights on the sequential mechanism of formation of preferences and whether it privileges a more or less naturally ordered assignment of the most extreme ranks. In other words, we show how it is possible to assess with a suitable statistical approach the formation of ranking preferences and answer through a statistical model the following questions: “What do I start ranking first? And what do I do then?”.

For possible future developments, several directions can be contemplated to further extend the Bayesian EPL. First, the methodology can be generalized for the accommodation of partial orderings and the introduction of item-specific and individual covariates that can improve the characterization of the preference behavior. Moreover, a Bayesian EPL mixture could fruitfully support more efficiently the identification of a parsimonious cluster structure in the sample.

Table 4: Posterior means and standard deviations (s.d.) of the EPL support parameters for the subsample of the `sushi` dataset.

Sushi type	$\hat{p}_i$	s.d.
Shrimp	0.0538	(0.006)
Sea eel	0.0799	(0.008)
Tuna	0.0760	(0.007)
Squid	0.0974	(0.009)
Sea urchin	0.0453	(0.005)
Salmon roe	0.0729	(0.007)
Egg	0.1591	(0.014)
Fatty tuna	0.0380	(0.004)
Tuna roll	0.1271	(0.012)
Cucumber roll	0.2504	(0.020)

**Acknowledgements** We are deeply grateful to both anonymous referees, whose comments and suggestions allowed us to improve the article. This work has been supported by Sapienza Università di Roma, grant RP11816436B15B6B.

## Appendix

### An example of EPL model

Here is a simple example to clarify the EPL formulation introduced in (1). Without loss of generality, let us suppose that an EPL( $\hat{\rho}, \hat{p}$ ) with parameter values  $\hat{\rho} = (4, 1, 3, 2)$  and  $\hat{p} = (0.4, 0.3, 0.2, 0.1)$  represents the true data generating mechanism. Under this EPL scenario, the positions are assigned according to an alternating preference scheme: the item selected at the first stage corresponds to the least-liked alternative ( $\hat{\rho}(1) = 4$ ); at the second stage, the most-liked item is specified ( $\hat{\rho}(2) = 1$ ); the item ranked at the third stage is the one receiving the third position ( $\hat{\rho}(3) = 3$ ) and, finally, the remaining alternative of the fourth stage is placed second in order of preference ( $\hat{\rho}(4) = 2$ ). Regarding the support parameters, they reflect a decreasing first-stage choice probability such that  $\hat{p}_i \propto (K + 1) - i$ . Hence, the chance of being ranked last reduces when proceeding from alternative 1 up to alternative 4: item 1 is more likely to be chosen at the first step and, thus, to be ranked last, followed in the order by item 2, 3 and 4.

Since the rank assignment order  $\rho$  is not restricted to the identity permutation  $\rho_{\mathbb{F}}$  as in the PL, a generic ordering  $\pi^{-1}$  does not coincide with the sequence  $\eta^{-1} = \pi^{-1} \circ \rho$  listing the items selected at each stage of the ranking process. For example, the considered EPL( $\hat{\rho}, \hat{p}$ ) implies that observing the ordering  $\pi^{-1} = (3, 1, 4, 2)$  corresponds to the sequential item selections indicated by the composition below

$$\begin{aligned} \eta^{-1} &= \pi^{-1} \circ \hat{\rho} = (\pi^{-1}(\hat{\rho}(1)), \pi^{-1}(\hat{\rho}(2)), \pi^{-1}(\hat{\rho}(3)), \pi^{-1}(\hat{\rho}(4))) \\ &= (\pi^{-1}(4), \pi^{-1}(1), \pi^{-1}(3), \pi^{-1}(2)) = (2, 3, 4, 1), \end{aligned}$$

that is, one chooses item 2 at the first stage, item 3 at the second stage, item 4 at the third stage and item 1 at the last stage.

Equation (1) states that the probability mass associated to  $\pi^{-1} = (3, 1, 4, 2)$  under the specified EPL( $\hat{\rho}, \hat{p}$ ) can be computed from the PL distribution after rearranging the components of  $\pi^{-1}$  according to the reference order  $\hat{\rho}$ :

$$\mathbf{P}_{\text{EPL}}(\pi^{-1} = (3, 1, 4, 2) | \hat{\rho}, \hat{p}) = \mathbf{P}_{\text{PL}}(\eta^{-1} = (2, 3, 4, 1) | \hat{p}) = \frac{0.3}{1} \cdot \frac{0.2}{0.4 + 0.2 + 0.1} \cdot \frac{0.1}{0.4 + 0.1} \cdot 1 \approx 0.017.$$

## References

1. Alvo M, Yu PL (2014) Statistical methods for ranking data. Springer, New York
2. Caron F, Doucet A (2012) Efficient Bayesian inference for Generalized Bradley-Terry models. *J Comput Graph Statist* 21 (1), 174–196
3. Critchlow DE, Fligner MA, Verducci JS (1991) Probability models on rankings. *J Math Psychol* 35 (3), 294–318
4. Gormley IC, Murphy TB (2006) Analysis of Irish third-level college applications data. *J Royal Stat Soc A* 169 (2), 361–379
5. Henery RJ (1981) Permutation probabilities as models for horse races. *J Royal Stat Soc B* 43 (1), 86–91
6. Jacques J, Grimprez Q, Biernacki C (2014) Rankcluster: An R package for clustering multivariate partial rankings. *R J* 6 (1), 101–110
7. Kamishima T (2003) Nantonac collaborative filtering: recommendation based on order responses. *Proceedings of the ninth ACM SIGKDD international conference on Knowledge discovery and data mining*, ACM, 583–588
8. Liu Q, Crispino M, Scheel I, Vitelli V, Frigessi A (2019) Model-Based Learning from Preference Data. *Annu Rev Stat Appl* 6, 329–354
9. Luce RD (1959) Individual choice behavior: A theoretical analysis. John Wiley & Sons Inc, New York
10. Marden JI (1995) Analyzing and modeling rank data. Vol. 64 of *Monographs on Statistics and Applied Probability*. Chapman & Hall, New York
11. Mollica C, Tardella L (2014) Epitope profiling via mixture modeling of ranked data. *Stat Med* 33 (21), 3738–3758
12. Mollica C, Tardella L (2017) Bayesian mixture of Plackett-Luce models for partially ranked data. *Psychometrika* 82 (2), 442–458
13. Plackett RL (1975) The analysis of permutations. *J Royal Stat Soc C* 24 (2), 193–202
14. Spiegelhalter DJ, Best NG, Carlin BP, Van Der Linde A (2002) Bayesian measures of model complexity and fit. *J Royal Stat Soc B* 64 (4), 583–639
15. Stern H (1990) Models for distributions on permutations. *J Am Stat Assoc* 85 (410), 558–564
16. Tierney L (1994) Markov chains for exploring posterior distributions. *Ann Stat* 22 (4), 1701–1762
17. Vigneau E, Courcoux P, Semenou M (1999) Analysis of ranked preference data using latent class models. *Food Qual Prefer* 10 (3), 201–207
18. Vitelli V, Sørensen Ø, Crispino M, Frigessi A, Arjas E (2018) Probabilistic preference learning with the Mallows rank model. *J Mach Learn Res* 18 (158), 1–49
19. Yu PLH, Lam KF, Lo SM (2005) Factor analysis for ranked data with application to a job selection attitude survey. *J Royal Stat Soc A* 168 (3), 583–597