

Partial rank data with the `hyper2` package: likelihood functions for generalized Bradley-Terry models

Robin K. S. Hankin

Auckland University of Technology

Abstract

Here I present the `hyper2` package for generalized Bradley-Terry models and analyze two situations: single scull rowing, and the competitive cooking game show MasterChef Australia. A number of natural statistical hypotheses may be tested straightforwardly using the software. This vignette is based on (Hankin 2017).

Keywords: Dirichlet distribution, hyperdirichlet distribution, combinatorics, R, multinomial distribution, constrained optimization, Bradley-Terry.

1. Introduction

2. Introduction: the Bradley-Terry model

The Bradley-Terry model for datasets involving paired comparisons has wide uptake in the R community. However, existing functionality¹ is restricted to paired comparisons. The canonical problem is to consider n players who compete against one another; the basic inference problem is to estimate numbers $\mathbf{p} = (p_1, \dots, p_n)$, $p_i \geq 0$, $\sum p_i = 1$ which correspond to player “strengths”. Information about the p_i may be obtained from the results of paired comparisons between the players.

Applications are legion. The technique is widely used in a competitive sport context (Turner and Firth 2012), in which matches are held between two opposing individuals or teams. It can also be applied to consumer choice experiments in which subjects are asked to choose a favourite from among two choices (Hatzinger and Ditttrich 2012), in which case the p_i are known as “worth parameters”.

If player i competes against player j , and wins with probability P_{ij} then the likelihood function for p_1, \dots, p_n corresponding to a win for i is $\frac{p_i}{p_i + p_j}$. As Turner and Firth (2012) point out, this may be expressed as

$$\text{logit}(P_{ij}) = \log p_i - \log p_j$$

¹In theory, the `hyperdirichlet` package (Hankin 2010) provides similar functionality but it is slow and inefficient. It is limited to a small number of players and cannot cope with the examples considered here.

and this fact may be used to estimate \mathbf{p} using generalized linear models. However, consider the case where three competitors, i , j , and k compete. The probability that i wins is then $\frac{p_i}{p_i+p_j+p_k}$ (Luce 1959); but there is no simple way to translate this likelihood function into a GLM. However, working directly with the likelihood function for \mathbf{p} has several advantages which are illustrated below. The resulting likelihood functions may readily be generalized to accommodate more general order statistics, as in a race. In addition, likelihood functions may be specified for partial order statistics; also, observations in which a *loser* is identified may be given a likelihood function using natural R idiom.

2.1. Further generalizations

Observing the winner w from a preselected set of competitors \mathcal{C} has a likelihood function of $p_w / \sum_{i \in \mathcal{C}} p_i$. But consider a more general situation in which two disjoint teams \mathcal{A} and \mathcal{B} compete; this would have likelihood $\sum_{i \in \mathcal{A}} p_i / \sum_{i \in \mathcal{A} \cup \mathcal{B}} p_i$. Such datasets motivate consideration of likelihood functions $\mathcal{L}(\cdot)$ with

$$\mathcal{L}(\mathbf{p}) = \prod_{s \in \mathcal{O}} \left(\sum_{i \in s} p_i \right)^{n_s} \quad (1)$$

where \mathcal{O} is a set of observations and s a subset of $[n] = \{1, 2, \dots, n\}$; numbers n_s are integers which may be positive or negative. The approach adopted by the **hyperdirichlet** package is to store each of the 2^n possible subsets of $[n]$ together with an exponent:

$$\prod_{s \in 2^{[n]}} \left(\sum_{i \in s} p_i \right)^{n_s} . \quad (2)$$

but this was noted as being needlessly memory intensive and slow; it is limited, in practice, to $n \leq 9$. Use of the more efficient equation 1 necessitates some mechanism of keeping track of which subsets of $[n]$ have nonzero powers.

3. The hyper2 package

One such mechanism is furnished by the C++ Standard Template Library’s “map” class (Musser, Derge, and Saini 2009) to store and retrieve elements. A *map* is an associative container that stores values indexed by a key, which is used to sort and uniquely identify the values. In the package, the key is a (STL) **set** of strictly positive integers $\leq n$. The relevant **typedef** statements are:

```
typedef set<unsigned int> bracket;
typedef map<bracket, double> hyper2;
```

In the STL, a map object stores keys and associated values in whatever order the software considers to be most propitious. This allows faster access and modification times but the order in which the maps, and indeed the elements of a set, are stored is not defined. In the case of likelihood functions such as Equation 1, this is not an issue because both multiplication and addition are associative and commutative operations. One side-effect of using this system is that the order of the bracket-power key-value pairs is not preserved.

Topalov	Anand	Karpov	total
22	13	-	35
-	23	12	35
8	-	10	18
30	36	22	88

Table 1: Results of 88 chess matches (dataset `chess` in the `aylmer` package) between three Grandmasters; entries show number of games won up to 2001 (draws are discarded). Topalov beats Anand 22-13; Anand beats Karpov 23-12; and Karpov beats Topalov 10-8

4. The package in use

Consider the `Chess` dataset of the `hyperdirchlet` package, in which matches between three chess players are tabulated (table 1). The Bradley-Terry model (Bradley and Terry 1952) is appropriate here (Caron and Doucet 2012), and the `hyper2` package provides a likelihood function for the strengths of the players, p_1, p_2, p_3 with $p_1 + p_2 + p_3 = 1$. A likelihood function might be

$$\frac{p_1^{30} p_2^{36} p_3^{22}}{(p_1 + p_2)^{35} (p_2 + p_3)^{35} (p_1 + p_3)^{18}}$$

Using the `hyper2` package, the R idiom to create this likelihood function is as follows:

```
> chess <- hyper2()
> chess["Topalov"] <- 30
> chess["Anand" ] <- 36
> chess["Karpov" ] <- 22
> chess[c("Topalov", "Anand" )] <- 35
> chess[c("Anand", "Karpov" )] <- 35
> chess[c("Karpov", "Topalov")] <- 18
> chess
```

```
log(Anand^36 * (Anand + Karpov)^35 * (Anand + Topalov)^35 * Karpov^22 *
(Karpov + Topalov)^18 * Topalov^30)
```

The internal representation of `hyper2` objects is that of a `map`, an associative array implemented as part of the standard template library (STL) of C++. It is thus an unordered collection of (*key*, *value*) pairs. In this case the *value* is the power and the *key* is the content of the bracket, which is a subset of `{Topalov, Anand, Karpov}`. The `stl` provides a `set` class for efficient manipulation of such objects, and this is used in the package.

One side-effect of using this system is that the order of the bracket-power key-value pairs is not preserved; above, note how the terms appear in an essentially random order. Also note that the use of the `set` class means that the R idiom is insensitive to the order of the terms within a bracket:

```
> chess[c("Karpov", "Topalov")]
```

```
log((Karpov + Topalov)^18)
> chess[c("Topalov", "Karpov")]
log((Karpov + Topalov)^18)
>
```

The package can calculate log-likelihoods:

```
> loglik(c(1/3, 1/3), chess)
[1] -132.4
```

[the first argument of function `loglik()` is a vector of length 2, third element of \mathbf{p} being the “fillup” value (Aitchison 1986)]; the gradient of the log-likelihood is given by function `gradient()`:

```
> gradient(chess, c(1/3, 1/3))
[1] 43.5 -24.0
```

Such functionality allows the rapid location of the maximum likelihood estimate for \mathbf{p} :

```
> maxp(chess)
  Anand  Karpov Topalov
0.4532  0.2380  0.3088
```

5. Men’s single sculling in the 2016 Summer Olympic Games

In this section, I will take results from the 2016 Summer Olympic Games and create a likelihood function for the finishing order in Men’s single sculling. In Olympic sculling, up to six individual competitors race a small boat called a scull over a course of length 2 km; the object is to cross the finishing line first. Note that actual timings are irrelevant, given the model, as the sufficient statistic is the order in which competitors cross the finishing line. The 2016 Summer Olympics is a case in point: the gold and silver medallists finished less than 5 milliseconds apart, corresponding to a lead of ~ 2.5 cm.

We may use a generalization of the Bradley-Terry model due to Luce (1959), in which the probability of competitor i winning in a field of $j = 1, \dots, n$ is

$$\frac{p_i}{p_1 + \dots + p_n}.$$

However, there is information in the whole of the finishing order, not just the first across the line. Once the winner has been identified, then the runner-up may plausibly be considered to

be the winner among the remaining competitors; and so on down the finishing order. Without loss of generality, if the order of finishing were 1, 2, 3, 4, 5, 6, then a suitable likelihood function would be, following Plackett (1975):

$$\frac{p_1}{p_1 + p_2 + p_3 + p_4 + p_5 + p_6} \cdot \frac{p_2}{p_2 + p_3 + p_4 + p_5 + p_6} \cdot \frac{p_3}{p_3 + p_4 + p_5 + p_6} \cdot \frac{p_4}{p_4 + p_5 + p_6} \cdot \frac{p_5}{p_5 + p_6} \cdot \frac{p_6}{p_6} \quad (3)$$

We may represent the result of Heat 1 as

fournier > cabrera > bhokanal > saensuk > kelmelis > teilemb

(a full list of the finishing order for all 25 events is given in the package as `rowing.txt`). Incorporating the information from Heat 1 into a likelihood function corresponding to equation 3 is straightforward using the `rankvec_likelihood()` function:

```
> heat1 <- c("fournier", "cabrera", "bhokanal", "saensuk", "kelmelis", "teilemb")
> H <- rankvec_likelihood(heat1)
> H
```

```
log( bhokanal * (bhokanal + cabrera + fournier + kelmelis + saensuk +
teilemb)^-1 * (bhokanal + cabrera + kelmelis + saensuk + teilemb)^-1 *
(bhokanal + kelmelis + saensuk + teilemb)^-1 * cabrera * fournier *
kelmelis * (kelmelis + saensuk + teilemb)^-1 * (kelmelis + teilemb)^-1
* saensuk)
```

(variable `heat1` shows the finishing order for Heat 1: Fournier came first, then Cabrera, etc). We can add the information from heat 2 easily:

```
> heat2 <- c("drysdale", "molnar", "esquivel", "garcia", "khafaji", "monasterio")
> H <- H + rankvec_likelihood(heat2)
> head(H)
```

```
log( bhokanal * (bhokanal + cabrera + fournier + kelmelis + saensuk +
teilemb)^-1 * (bhokanal + cabrera + kelmelis + saensuk + teilemb)^-1 *
(bhokanal + kelmelis + saensuk + teilemb)^-1 * cabrera * drysdale)
```

Again observe that object `H` includes its terms in no apparent order. Although it would be possible to incorporate information from subsequent heats in a similar manner, the package includes a ready-made dataset, `rowing`:

```
> head(rowing) # see rowing.Rd
```

```
log(banna^4 * (banna + boudina + cabrera + molnar + obreno +
rivarola)^-1 * (banna + boudina + cabrera + molnar + rivarola)^-1 *
(banna + boudina + molnar + rivarola)^-1 * (banna + cabrera + campbell
+ grant + hoff)^-1 * (banna + cabrera + campbell + grant + hoff +
szymczyk)^-1)
```

```
> dotchart(rowing_maxp)
```




Figure 1: Maximum likelihood estimate for the strengths of the 32 competitors in the Men's singles sculls in the 2016 Summer Olympics

Finding the maximum likelihood estimate for the parameter $p_{\text{banna}}, \dots, p_{\text{zambrano}}$ is straightforward using the `maxp()` function, provided with the package (Figure 1). The optimization routine has access to derivatives which means that the estimate is found very quickly.

Figure 1 shows very directly that the competitor with the highest strength is Drysdale, the gold medallist for this event. The bronze and silver medallists were Synek and Martin respectively, whose estimated strengths were second and third highest in the field.

6. MasterChef Australia

MasterChef Australia is a game show in which amateur cooks compete for a title (Wikipedia 2017a). From a statistical perspective the contest is interesting because the typical show format is to identify the *weakest* player, who is then eliminated from the competition. Here, results from MasterChef Australia Series 6 (Wikipedia 2017b) will be analysed; an extended discussion of the data used is given in the package at `masterchef.Rd`.

We wish to make inferences about the contestants' generalized Bradley-Terry strengths p_1, \dots, p_n , $\sum p_i = 1$. One informative event was a team challenge in which the contestants were split randomly into two teams, red and blue:

```
> team_red <- c("Jamie", "Tracy", "Ben", "Amy", "Rena", "Georgia")
> team_blue <- c("Brent", "Laura", "Emelia", "Colin", "Kira", "Tash")
```

We may represent the fact that the red team won as

$$\{\text{Jamie} + \text{Tracy} + \text{Ben} + \text{Amy} + \text{Rena} + \text{Georgia}\} \succ \{\text{Brent} + \text{Laura} + \text{Emelia} + \text{Colin} + \text{Kira} + \text{Tash}\} \quad (4)$$

The likelihood function for the observation given in equation 4 would be

$$\frac{p_{\text{Jamie}} + p_{\text{Tracy}} + p_{\text{Ben}} + p_{\text{Amy}} + p_{\text{Rena}} + p_{\text{Georgia}}}{p_{\text{Jamie}} + p_{\text{Tracy}} + p_{\text{Ben}} + p_{\text{Amy}} + p_{\text{Rena}} + p_{\text{Georgia}} + p_{\text{Brent}} + p_{\text{Laura}} + p_{\text{Emelia}} + p_{\text{Colin}} + p_{\text{Kira}} + p_{\text{Tash}}} \quad (5)$$

To generate a likelihood function in R, we need to set up a `hyper2` object with appropriate contestants:

```
> H <- hyper2(pnames = c(
+   "Amy", "Ben", "Brent", "Colin", "Emelia",
+   "Georgia", "Jamie", "Kira", "Laura", "Rena",
+   "Sarah", "Tash", "Tracy"))
> H
```

```
log()
```

Object `H` is a uniform likelihood function. The package R idiom for incorporating likelihood from Equation 5 is straightforward and natural:

```
> H[team_red] <- +1
> H[c(team_red,team_blue)] <- -1
> H
```

```
log((Amy + Ben + Brent + Colin + Emelia + Georgia + Jamie + Kira +
Laura + Renae + Tash + Tracy)^-1 * (Amy + Ben + Georgia + Jamie + Renae
+ Tracy))
```

(Sarah did not take part). The above idiom makes it possible to define likelihood functions for observations that have a peculiar probability structure, and I give two examples below.

One event involved eight competitors who were split randomly into four teams of two. The show format was specified in advance as follows: The teams were to be judged, and placed in order. The two top teams were to be declared safe, and the other two teams sent to an elimination trial from which an individual winner and loser were identified, the loser being obliged to leave the competition. The format for this event is also typical in MasterChef.

The observation was that Laura and Jamie's team won, followed by Emelia and Amy, then Brent and Tracy. Ben and Renae's team came last:

$$\{Laura + Jamie\} \succ \{Emelia + Amy\} \succ \{Brent + Tracy\} \succ \{Ben + Renae\} \quad (6)$$

A plausible likelihood function can be generated using the standard assumption (Hankin 2010) that the competitive strength of a team is the sum of the strengths of its members. We can then combine this assumption with the order statistic technique of Plackett (1975):

$$\frac{\frac{p_{Laura} + p_{Jamie}}{p_{Laura} + p_{Jamie} + p_{Emelia} + p_{Amy} + p_{Brent} + p_{Tracy} + p_{Ben} + p_{Renae}}{p_{Emelia} + p_{Amy}}}{p_{Emelia} + p_{Amy} + p_{Brent} + p_{Tracy} + p_{Ben} + p_{Renae}} \cdot \frac{p_{Brent} + p_{Tracy}}{p_{Brent} + p_{Tracy} + p_{Ben} + p_{Renae}} \quad (7)$$

We would like to incorporate information from this observation into object `H`, which is a likelihood function for the two-team challenge discussed above. The corresponding R idiom is natural:

```
> blue <- c("Laura","Jamie") # first
> yellow <- c("Emelia","Amy") # second
> green <- c("Brent","Tracy") # third
> red <- c("Ben","Renae") # fourth
> H[blue] <- 1
> H[c(blue,yellow,green,red)] <- -1
> H[yellow] <- 1
> H[c(yellow,green,red)] <- -1
> H[green] <- 1
> H[c(green,red)] <- -1
> H
```


$\log((\text{Amy} + \text{Ben} + \text{Brent} + \text{Colin} + \text{Emelia} + \text{Georgia} + \text{Jamie} + \text{Kira} + \text{Laura} + \text{Renaë} + \text{Tash} + \text{Tracy})^{-1} * (\text{Amy} + \text{Ben} + \text{Brent} + \text{Emelia} + \text{Jamie} + \text{Laura} + \text{Renaë} + \text{Tracy})^{-1} * (\text{Amy} + \text{Ben} + \text{Brent} + \text{Emelia} + \text{Renaë} + \text{Tracy})^{-1} * (\text{Amy} + \text{Ben} + \text{Georgia} + \text{Jamie} + \text{Renaë} + \text{Tracy}) * (\text{Amy} + \text{Emelia}) * (\text{Ben} + \text{Brent} + \text{Renaë} + \text{Tracy})^{-1} * (\text{Brent} + \text{Tracy}) * (\text{Jamie} + \text{Laura}))$

We may incorporate subsequent observations relating to the elimination trial among the four competitors comprising the losing two teams. The observation was that Laura won, and Renaë came last, being eliminated. We might write

$$\{\text{Laura}\} \succ \{\text{Brent, Tracey, Ben}\} \succ \{\text{Renaë}\}, \quad (8)$$

which indicates that Laura came first, then Brent/Tracey/Ben in some order, then Renaë came last. For this observation a likelihood function, following [Plackett \(1975\)](#), might be

$$\mathcal{L}(p_1, p_2, p_3, p_4, p_5) = P(p_1 \succ p_2 \succ p_3 \succ p_4 \succ p_5 \cup p_1 \succ p_2 \succ p_4 \succ p_3 \succ p_5 \cup \dots) \quad (9)$$

$$= P\left(\bigcup_{[abc]} p_1 \succ p_a \succ p_b \succ p_c \succ p_5\right) \quad (10)$$

$$= \frac{p_1}{p_1 + p_2 + p_3 + p_4 + p_5} \cdot \frac{p_2}{p_2 + p_3 + p_4 + p_5} \cdot \frac{p_3}{p_3 + p_4 + p_5} \cdot \frac{p_4}{p_4 + p_5}$$

$$+ \frac{p_1}{p_1 + p_2 + p_4 + p_3 + p_5} \cdot \frac{p_2}{p_2 + p_4 + p_3 + p_5} \cdot \frac{p_4}{p_4 + p_3 + p_5} \cdot \frac{p_3}{p_3 + p_5}$$

$$+ \frac{p_1}{p_1 + p_3 + p_2 + p_4 + p_5} \cdot \frac{p_3}{p_3 + p_2 + p_4 + p_5} \cdot \frac{p_2}{p_2 + p_4 + p_5} \cdot \frac{p_4}{p_4 + p_5}$$

$$+ \dots$$

where Laura's strength is shown as p_1 etc for brevity. The R idiom is as follows:

```
> L <- ggr1(H,
+         winner      = "Laura",
+         btm4        = c("Brent", "Tracy", "Ben"),
+         eliminated  = "Renaë")
```

Arguments to `ggr1()` are disjoint subsets of the players, the subsets themselves being passed in competition order from best to worst. Object `L` includes information from the team challenge (via first argument `H`) and the elimination results. It is a list of length $3! = 6$ objects of class `hyper2`, each of which gives a [Luce](#) likelihood function for a consistent total ordering of the competitors.

A final example (taken from *MasterChef* series 8, week 10) is given as a generalization of the [Luce](#) likelihood. The format was as follows. Eight contestants were split randomly into four teams of two, the top two teams being declared safe. Note that the likelihood interpretation differs from the previous team challenge, in which the observation was an unambiguous team ranking: here, there is only a partial ranking of the teams and one might

expect this observation to be less informative. Without loss of generality, the result may be represented as

$$\{p_1 + p_2, p_3 + p_4\} \succ \{p_5 + p_6, p_7 + p_8\} \quad (11)$$

and a likelihood function on p_1, \dots, p_8 for this observation might be

$$\begin{aligned} \mathcal{L}(p_1, \dots, p_8) &= \mathbb{P} \left(\{p_1 + p_2\} \succ \{p_3 + p_4\} \succ \{p_5 + p_6\} \succ \{p_7 + p_8\} \cup \right. \\ &\quad \{p_1 + p_2\} \succ \{p_3 + p_4\} \succ \{p_7 + p_8\} \succ \{p_5 + p_6\} \cup \\ &\quad \{p_3 + p_4\} \succ \{p_1 + p_2\} \succ \{p_5 + p_6\} \succ \{p_7 + p_8\} \cup \\ &\quad \left. \{p_3 + p_4\} \succ \{p_1 + p_2\} \succ \{p_5 + p_6\} \succ \{p_7 + p_8\} \right) \quad (12) \\ &= \frac{p_1 + p_2}{p_1 + p_2 + p_3 + p_4 + p_5 + p_6 + p_7 + p_8} \cdot \frac{p_3 + p_4}{p_3 + p_4 + p_5 + p_6 + p_7 + p_8} \cdot \frac{p_5 + p_6}{p_5 + p_6 + p_7 + p_8} \\ &\quad + \dots + \\ &\quad \frac{p_3 + p_4}{p_3 + p_4 + p_1 + p_2 + p_7 + p_8 + p_5 + p_6} \cdot \frac{p_1 + p_2}{p_3 + p_4 + p_7 + p_8 + p_5 + p_6} \cdot \frac{p_7 + p_8}{p_7 + p_8 + p_5 + p_6} \end{aligned}$$

6.1. Maximum likelihood estimation

The package provides an overall likelihood function for all informative judgements in the series on the final 13 players in object `masterchef_series6`, documented at `masterchef.Rd`. We may assess a number of related hypotheses using the package. The first step is to calculate the likelihood for the hypothesis that all players are of equal strength:

```
> data("masterchef")
> n <- 13 # 13 players
> equal_strengths <- rep(1/n, n-1)
> like_series(equal_strengths, masterchef) # see masterchef.Rd
```

```
[1] -79.21
```

The strengths of the 13 players may be estimated using standard maximum likelihood techniques. This requires constrained optimization in order to prevent the search from passing through inadmissible points in p-space:

```
> UI <- rbind(diag(n-1), -1) # p_i >= 0
> CI <- c(rep(0, n-1), -1) # p_1 + ... + p_{n-1} <= 1
> constrOptim( # maxp() does not work for masterchef_series6
+   theta = equal_strengths, # starting point for optimization
+   f = function(p){-like_series(p, masterchef)}, # see masterchef.Rd
+   ui=UI, ci=CI,
+   grad=NULL)
```

```
> masterchef_maxp
```

Amy	Ben	Brent	Colin	Emelia	Georgia	Jamie	Kira
1.086e-01	7.458e-02	1.344e-01	2.820e-02	1.170e-01	6.850e-09	1.065e-01	2.056e-02
Laura	Renae	Sarah	Tash	Tracy			
2.751e-01	4.071e-02	2.804e-02	1.142e-09	6.637e-02			

```
> dotchart(masterchef_maxp)
```




Figure 2: Maximum likelihood estimate for the strengths of the top 13 competitors in Series 6 of *MasterChef Australia*

The resulting maximum likelihood estimate, `masterchef_maxp` in the package, is shown pictorially in Figure 2. The support at the precalculated evaluate is

```
> like_series(indep(masterchef_maxp), masterchef)

[1] -66.26
```

and this allows us to test the hypothesis of equal player strengths: by Wilks's theorem, $-2\log\Lambda$ has an asymptotic null distribution of χ_{12}^2 . This corresponds to a p -value of

```
> pchisq(2*(78.7-66.2),df=12,lower.tail=FALSE)

[1] 0.01482
```

showing that the observations do constitute evidence for differential player strengths. Figure 2 suggests that Laura, the runner-up, is actually a stronger competitor than the winner, Brent. We may assess this statistically by finding the maximum likelihood for \mathbf{p} , subject to the constraint that $p_{\text{Laura}} \leq p_{\text{Brent}}$:

```
> UI <- rbind(UI,c(0,0,1,0,0,0,0,0,-1,0,0,0,0)) # Brent >= Laura
> CI <- c(CI,0)
> ans2 <-
+ constrOptim( # maxp() does not work for masterchef_series6
+   theta = equal_strengths,
+   f = function(p){-like_series(p,masterchef_series6)}, # see masterchef.Rd
+   grad=NULL,
+   ui = UI, ci=CI)
```

Object `maxp_masterchef6_constrained` in the package is the result of the above optimization, at which point the likelihood is

```
> like_series(indep(masterchef_constrained_maxp), masterchef)

[1] -67.42
```

The two estimates differ by about 1.18, less than the two-units-of-support criterion of [Edwards \(1992\)](#); alternatively, one may observe that the likelihood ratio is not in the tail region of its asymptotic distribution (χ_1^2) as the p -value is about 0.12. This shows that there is no strong evidence for Laura's competitive strength being higher than that of Brent. Similar techniques can be used to give a profile likelihood function; the resulting support interval for Laura's strength is $[0.145, 0.465]$, which does not include $\frac{1}{13} \simeq 0.077$, the mean player strength.

However, further work would be needed to make statistically robust inferences from these findings. Even given equal competitive ability, one would expect the player with the highest estimated Bradley-Terry strength to have an elevated support interval; and it is not clear to what extent Wilks, being asymptotic, is relevant to the data at hand.

7. Conclusions

Several generalization of Bradley-Terry strengths are appropriate to describe competitive situations in which order statistics are sufficient.

The `hyper2` package is introduced, providing a suite of functionality for generalizations of the partial rank analysis of Critchlow (1985). The software admits natural R idiom for translating commonly occurring observations into a likelihood function.

The package is used to calculate maximum likelihood estimates for generalized Bradley-Terry strengths in two competitive situations: Olympic rowing, and *MasterChef Australia*. The estimates for the competitors' strengths are plausible; and several meaningful statistical hypotheses are assessed quantitatively.

References

- Aitchison J (1986). *The Statistical Analysis of Compositional Data*. The Blackburn Press.
- Bradley RA, Terry ME (1952). “The Rank Analysis of Incomplete Block Designs I. The Method of Paired Comparisons.” *Biometrika*, **39**, 324–345.
- Caron F, Doucet A (2012). “Efficient Bayesian inference for generalized Bradley-Terry models.” *Journal of Computational and Graphical Statistics*, **21**(1), 174–196.
- Critchlow DE (1985). *Metric methods for analyzing partially ranked data*. Springer-Verlag, New York.
- Edwards AWF (1992). *Likelihood (Expanded Edition)*. John Hopkins.
- Hankin RKS (2010). “A Generalization of the Dirichlet Distribution.” *Journal of Statistical Software*, **33**(11), 1–18. doi:10.18637/jss.v033/i11.
- Hankin RKS (2017). “Partial rank data with the `hyper2` package: likelihood functions for generalized Bradley-Terry models.” *The R Journal*, **9**(2), 429–439.
- Hatzinger R, Dittrich R (2012). “prefmod: An R package for modeling preferences based on paired comparisons, rankings, or ratings.” *Journal of Statistical Software*, **48**, 1–31.
- Luce R (1959). *Individual choice behaviour: A theoretical analysis*. Wiley, New York.
- Musser DR, Derge GJ, Saini A (2009). *STL Tutorial and Reference Guide: C++ Programming with the Standard Template Library*. 3rd edition. Addison-Wesley Professional. ISBN 0321702123, 9780321702128.
- Plackett RL (1975). “The analysis of permutations.” *Applied Statistics*, **24**, 193–202.
- Turner H, Firth D (2012). “Bradley-Terry Models in R: The BradleyTerry2 Package.” *Journal of Statistical Software*, **48**(9), 1–21. doi:10.18637/jss/v048/i09.
- Wikipedia (2017a). “MasterChef Australia — Wikipedia, The Free Encyclopedia.” [Online; accessed 1-February-2017], URL https://en.wikipedia.org/w/index.php?title=MasterChef_Australia&oldid=761024807.

Wikipedia (2017b). “MasterChef Australia (series 6) — Wikipedia, The Free Encyclopedia.” [Online; accessed 1-February-2017], URL [https://en.wikipedia.org/w/index.php?title=MasterChef_Australia_\(series_6\)&oldid=762395535](https://en.wikipedia.org/w/index.php?title=MasterChef_Australia_(series_6)&oldid=762395535).

Affiliation:

Robin K. S. Hankin
Auckland University of Technology
E-mail: hankin.rob@gmail.com

