Animal Fight Club
An Analysis of Violent Interactions between Caribou

Philipp Burckhardt

August 12, 2014

1 Introduction

Most living caribou reside in North America, where they have a total population of 3.5 million (Britannica, 2006). Both male and female caribou grow antlers, and frequently compete in sparring activities in order to establish a dominance hierarchy among their flock. The fighting patterns of Caribou have attracted some scientific interest, with studies conducted on how and by whom fights are initiated (Barrette and Vandal, 1990).

The present analysis fills a gap insofar as it uses statistical modelling techniques to develop a ranking scheme which orders the fighting abilities of individual caribou and allows to assess the influence of personal traits on fighting skills. We use a data set which contains information about 823 aggressive interactions among 20 caribou, where for each caribou its *age* and *gender* have been recorded as explanatory variables.

The paper is structured as follows: At the beginning, data exploration is undertaken and the symmetry of the matrix of encounters is assessed. In Section 3, a simple Bradley-Terry model is presented as a means to model the contests between the caribou. Section 4 presents the results. Next, the model is adapted in Section 5 to incorporate caribou-specific information. Finally, Section 6 comments on the findings and provides an outlook for further research.

2 Data Exploration

The recorded fighting encounters of caribou are depicted in Figure 1. The number of fights between caribou range from 0 to 22, with most caribou fighting at least once with each other: Of all possible one-on-one duels ($\binom{20}{2} = 190$), there are only 12 combinations which did not take place. An initial look at Figure 1 might suggest otherwise, though: There appear to be many blank spots without circles. However, this only means that many caribou did not win against a specific opponent, not that no fight between the two of them took place. Having this in mind and observing that there are some caribou with a huge number of won fights on their record while others have almost none (like caribous 2, 7 and 19), there seems to be considerable variation in fighting skills among the considered group of caribou.

Recall that we have two goals in mind: First, we want to order the caribou with respect to their fighting abilities. Albeit it is quite easy to order the caribou depending on the percentage of won fights, this ordering is potentially deceptive: Although a lot of wins are generally associated with a higher ability, one has to take the quality of the respective opponents into account. This analysis is best undertaken using a statistical modelling approach, and is hence deferred to Section 3. Second, we want to relate fighting ability to the two explanatory variables *age* and *gender*.

In Figure 2, the number of won battles is displayed conditional on age. One not surprising observation is that the younger caribou either do not engage in many fights or often lose. In the following analysis, it will be shown that the latter is true, with winning percentages among the younger ones often below 20%.

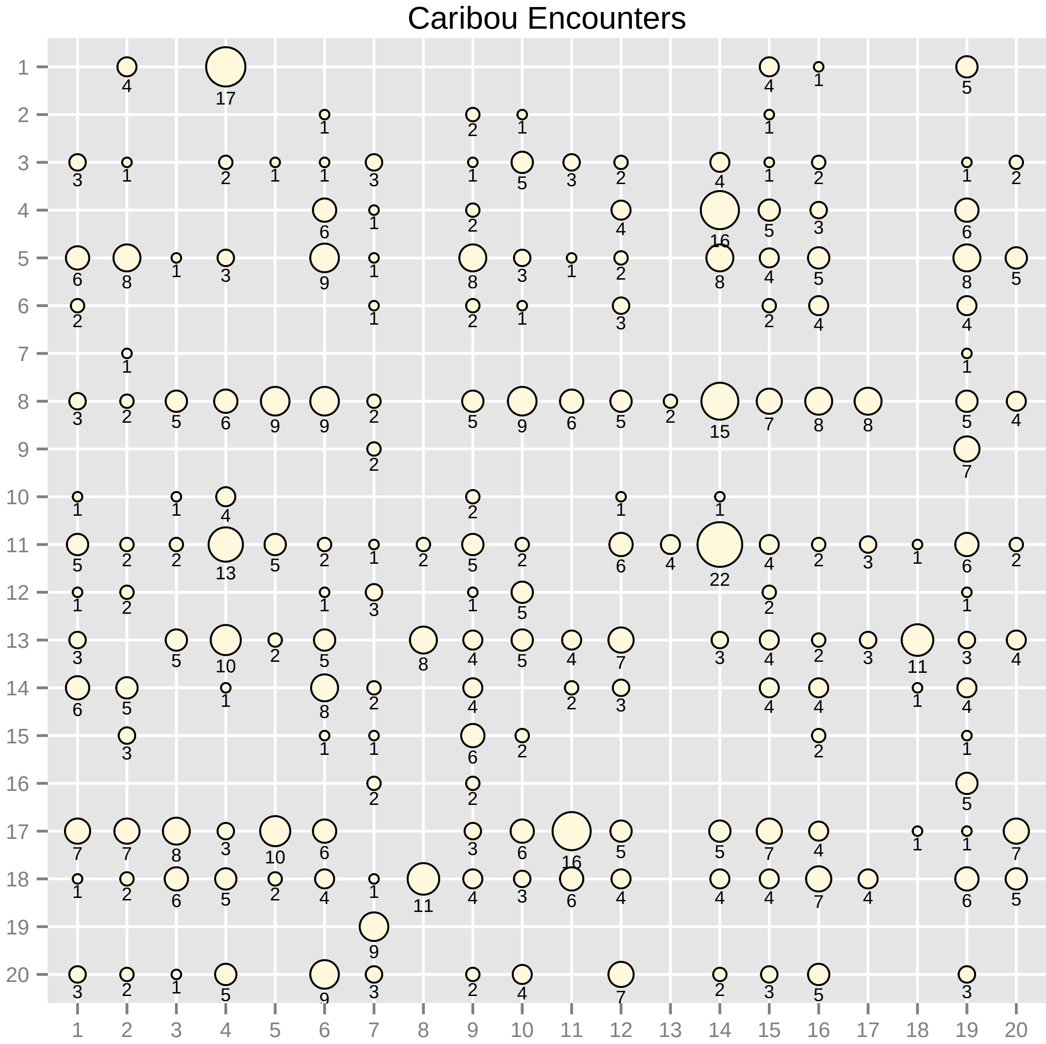

Figure 1: Balloon plot of the encounters of the 20 monitored caribou. For the i-th caribou, the values in the i-th row show the number of won fights against opponent j, where j denotes the column index. As an example, caribou 11 won 22 times against caribou 14, which is the highest number of lost battles any caribou had to endure against a single opponent.

The color of the bars in Figure 2 denotes the gender of the respective caribou, which again reveals a few interesting patterns. For example, it is interesting to observe that the data set does not contain any male caribou with an age of four or older: There is not a single blue box for that group. This reflects the life expectancy of caribou, which is about four years for male and around 10 years for female caribou. Given this background information, it is no mystery anymore that female caribou end up winning many fights when aged four or more: They just hit the phase in life where they are at the zenith of their abilities.

To investigate these relationships further and to sort the caribou depending on fighting skills, we now turn to statistical modelling. However, we first check whether the matrix of encounters in Figure 1 satisfies the symmetry property, which would lead to a very simple generalized linear model.

Following the definition Agresti (1990) for square contingency tables, a sufficient condition for symmetry is that

$$\pi_{ij} = \pi_{ji} \qquad \forall i \neq j,$$

where $\{\pi_{ij}\}$ denotes the $I \times I$ joint distribution of outcomes π_{ij}, which in our case denotes the probability of the i-th caribou to defeat the j-th one when engaged in a violent interaction. What would be the consequence if symmetry was satisfied in the present case? The answer is simple: it would mean that between any two caribou i and j, their respective chances of winning were equal. This does not seem to be a reasonable assumption for the present data set, as the graph in Figure 1 clearly shows that some caribou are much stronger than others.

However, to formally test for symmetry requires some care, and some of the usually applied methods don't work due to the sparseness of the matrix of encounters. For example, Bowker (1948) proposes the following test statistic:

$$\sum_{i>j}\sum \frac{(n_{ij} - n_{ji})^2}{n_{ij} + n_{ji}}, \tag{1}$$

in which n_{ij} marks the observed cell counts. For tables with dimension $I = 2$, this simplifies to McNemar's test (McNemar, 1947). Bowker (1948) showed that under the null hypothesis of symmetry, this test statistic follows a chi-squared distribution with $I(I-1)/2$ degrees of freedom. A significant result implies that the marginal frequencies are not homogeneous and therefrom that the table does not possess the symmetry property.

However, the approximation to the chi-squared distribution is doubtful for small frequencies and furthermore Equation (1) is undefined for the matrix of caribou encounters since $n_{ij} + n_{ji} = 0$ in many instances. There are some crude approximations to deal with these issues, including adding small values to each cell count and correcting for degrees of freedom, see Agresti (1990).

Looking at the problem from a different perspective, it can be seen that symmetry can be expressed as a log-linear model provided that $\pi_{ij} > 0$ for all i, j (Agresti, 1990). For the expected frequencies $\mu_{ij} = n\pi_{ij}$, this log-linear model has the form

$$\log \mu_{ij} = \lambda + \lambda_i + \lambda_j + \lambda_{ij},$$

where all $\lambda_{ij} = \lambda_{ji}$. To ensure identifiability, constraints need to be placed on the parameter values. Agresti (1990) shows that a simpler expression is then given by

$$\log \mu_{ij} = \lambda_{ij} \text{ with } \lambda_{ij} = \lambda_{ji}. \tag{2}$$

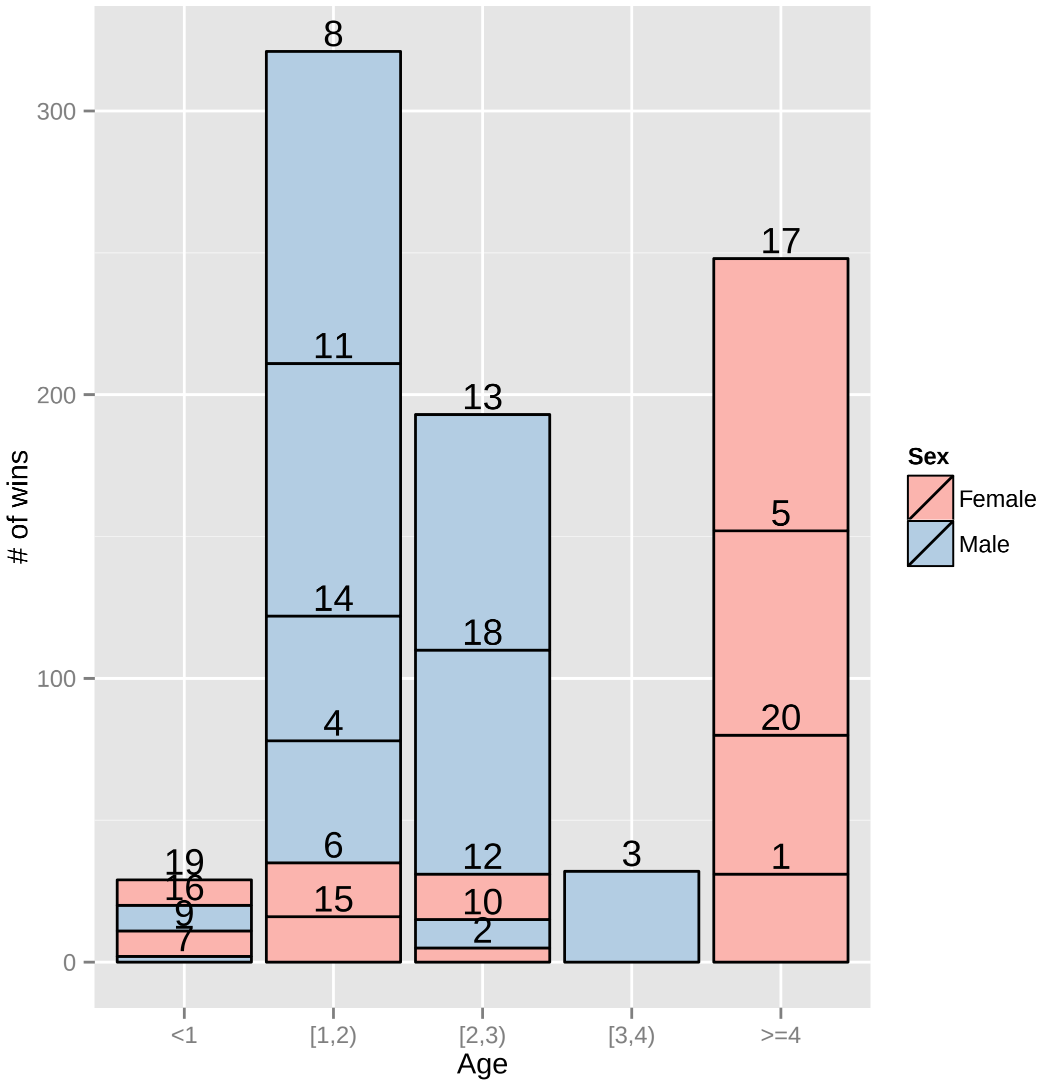

Figure 2: Bar plots of the number of won fights depending on age and gender. Individual bars for every caribou are stacked on each other, with the labels on top of every bar disclosing their id.

This observation enables us to the test the assumption of symmetry as follows: We fit the model in Equation (2) on the observed cell counts. The residual deviance of this model compares it to the saturated model, in which the constraint $\lambda_{ij} = \lambda_{ji}$ is lifted and exactly one parameter is fitted for each observation. It can be shown that the residual deviance satisfies

$$D(P) = 2\left[l(\text{saturated}) - l(P)\right] \sim \chi_{I\times(I-1)-p}$$

under H_0, where p is the number of different regression coefficients λ. Carrying out the model fitting process using maximum likelihood estimation, we obtain a residual deviance of 1008 under 190 degrees of freedom. This signals a very bad fit, and the associated p-value of 0 provides overwhelming evidence against the null hypothesis that the fit of our model is satisfactory. But recall that the fitted model directly followed from the symmetry property, which therefore must likewise be given up.

This is not a huge sacrifice, though. The use cases of the symmetry model are fairly limited (Agresti, 1990). Instead, a more commonly used assumption is that of quasi-symmetry, which proclaims that the following relationship holds (Caussinus, 1966):

$$\pi_{ij}\pi_{jk}\pi_{ki} = \pi_{ji}\pi_{kj}\pi_{ik} \quad \forall i, j, k$$

This assumption, which does not lend itself to any immediate intuition, provides the ground for the Bradley-Terry model. As Agresti (1990) notes, Fienberg and Larntz (1976) show that the Bradley-Terry model is a logit formulation for the quasi-symmetry model. In the following, we employ the Bradley-Terry model to investigate the fighting abilities of the caribou.

3 Bradley-Terry Model

The model which was developed by Bradley and Terry (1952) postulates the following form for the log-odds of winning for player i:

$$\text{logit}\left(\Pr\left[i \text{ beats } j\right]\right) = \lambda_i - \lambda_j, \tag{3}$$

where we interpret the coefficients λ as representing the ability of the respective players. This interpretation is sensible due to the following reasoning: First notice that Equation (3) is equivalent to

$$\frac{\Pr\left[i \text{ beats } j\right]}{\Pr\left[j \text{ beats i}\right]} = \frac{\exp\left(\lambda_i\right)}{\exp\left(\lambda_j\right)}. \tag{4}$$

Since the exponential is a monotonic function, $\lambda_i > \lambda_j$ always implies that player i has higher chances of winning compared to player j and vice verse. In light of this, it makes sense to interpret λ or $exp(\lambda)$ as measures of the ability of the players.

To fit the model depicted in Equation (3), we use statistical programming environment R, in particular the *BradleyTerry2* package (Firth, 2005).

4 Results

Parameter estimates for λ are displayed in Table 1. Estimation is carried out by maximum likelihood. Notice that the model in Equation 3 is overidentified and can only be fitted after certain restrictions have been placed on the coefficients. In this case, we have chosen caribou 7, the player with the lowest winning rate, to serve as the baseline category, constraining its coefficient to zero.

In the table, the caribou are ordered with respect to their estimated λs. To facilitate comparison with the empirical winning percentages, these are also displayed. As one can see, there is in general a large agreement between the two: A caribou with a high percentage of won fights is also very likely to have a large estimated ability coefficient. However, there are some exceptions: Caribou 3 has a very low winning rate but at the same time exhibits the 7-th highest λ coefficient. How can that be? A look at Figure 1 proves once again to be informative: Looking at the 8-th column, we can see that caribou 3 lost most battles against caribou with very high abilities: caribou 17, 18 and 8 all have very high ability estimates. The Bradley-Terry model behind these estimates fits the data adequately, as a residual deviance of 203.81 on 159 degrees of freedom signals.

ID	Win %	λ_i	S.D.
13	0.93	16.50	1.89
18	0.85	15.63	1.87
8	0.84	15.01	1.85
17	0.84	14.58	1.84
11	0.70	13.50	1.83
5	0.71	12.36	1.82
3	0.52	12.27	1.85
20	0.63	11.31	1.81
1	0.43	9.38	1.75
14	0.35	9.11	1.73
4	0.38	9.04	1.73
10	0.18	8.13	1.78
12	0.25	8.08	1.76
6	0.23	7.69	1.73
15	0.24	7.05	1.71
2	0.11	6.15	1.79
16	0.16	4.91	1.60
9	0.15	3.55	1.29
19	0.12	1.43	0.79
7	0.06	0.00	0.00

Table 1: Table of winning percentages for each caribou and results of fitting the Bradley-Terry model without covariates

From the results of Table 1, one can now predict outcomes of fights between the combatants. To give an example, let us calculate the probability that caribou 13 defeats caribou 18 to get a sense about how much the first and second placed caribou differ in their abilities.

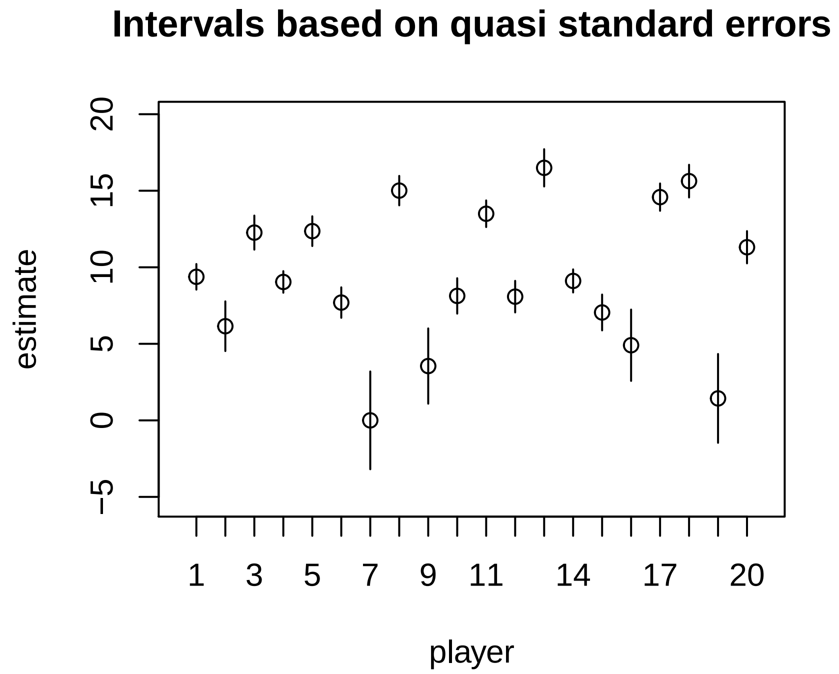

Figure 3: Estimated relative fighting skills of caribou

The result is

$$\Pr(13 \text{ beats } 18) = \frac{\exp(16.50)}{\exp(16.50) + \exp(15.63)} \approx 0.7,$$

meaning that we can make a fairly strong statement that caribou 13 will win provided that our model holds.

To test whether the abilities of any two caribou are equal, one can calculate the difference $\lambda_i - \lambda_j$ and test whether it is equal to zero. Doing this inference step requires full access to the covariance matrix of the coefficients since we have

$$Var(\lambda_i - \lambda_j) = Var(\lambda_i) + Var(\lambda_j) - 2 \times Cov(\lambda_i, \lambda_j)$$

from basic probability theory. Since this can get cumbersome, Firth (2004) developed the concept of quasi-variances, which allows for easier calculation as well as direct comparisons between any two players without the need to resort to a baseline category. Estimated comparison intervals are displayed in Figure 3.

Based on quasi standard errors, it is now easy to infer whether say the third caribou and the fifth differ from each other with respect to ability, i.e. we want to test whether $\lambda_5 - \lambda_3 = 0$. Their quasi-standard errors are 0.49 and 0.56, respectively. By using the familiar Pythagorean theorem, one can then calculate the standard error of $\lambda_5 - \lambda_3 = 0$ as $(0.49^2 + 0.56^2)^{1/2} = 0.74$. The test statistic can be calculated as $(12.36 - 12.27)/0.74 = 0.12$, which corresponds to a p-value of 0.9 for the two-tailed test assuming a normal distribution of the test statistics. This provides strong evidence in favour of the null hypothesis that the two caribou have the same abilities, which we cannot reject at a significance level of $\alpha = 0.05$.

Table 2: Output for the regression model of λ_i on the explanatory variables *age* and *gender*.

	lambda
Age[1,2)	7.136^{***}
	(1.934)
Age[2,3)	8.050^{***}
	(1.997)
Age[3,4)	7.918^{**}
	(3.411)
Age>=4	11.310^{***}
	(2.242)
SexMale	3.750^{**}
	(1.579)
Constant	0.597
	(1.681)
Observations	20
R^2	0.702
Adjusted R^2	0.595
Residual Std. Error	2.968 (df = 14)
F Statistic	6.588^{***} (df = 5; 14)
Notes:	***Significant at the 1 percent level.
	**Significant at the 5 percent level.
	*Significant at the 10 percent level.

Returning to the earlier posed question of how ability is related to *age* and *gender*, we propose to run a regression of the estimated λ_i with these as covariates. Specifically, we assume the usual fixed-effects model in standard linear regression:

$$\lambda = X\beta + \varepsilon, \quad \varepsilon \sim N\left(0, \sigma^2 I_n\right),$$

where the error terms follow an 20-dimensional multivariate normal with mean zero and uncorrelated error terms with constant variance. The results from that regression are depicted in Table 2. As can be seen, the model is clearly significant (F statistic equals 6.588). The high R^2 signals that the considered variables account for 70% of the variance in λ. All regression coefficients are significant at the 5% level. Since *age* has different levels, (joint) significance testing was conducted with ANOVA (p-value almost zero).

To interpret the regression coefficients, notice for example that being a male caribou is associated with a multiplicative effect of $exp(3.750) = 42.52$ on the odds of winning compared to being female, as can be seen from Equation (4).

5 Extension

As an alternative, one can also directly incorporate the covariates into the Bradley-Terry model. This leads to a mixed-effects model, in which both *age* and *gender* enter as fixed effects and random effects are included for each caribou. That is, one assumes that for each λ, the following relationship holds:

$$\lambda_i = x_i^\intercal \beta + U_i.$$

Inserting this into Equation (3) yields

$$\text{logit}\left(\Pr\left[i \text{ beats } j\right]\right) = x_i^\intercal \beta + U_i - x_j^\intercal \beta - U_j, \tag{5}$$

where $U_{i,j} \sim N\left(0, \sigma^2\right)$. Fitting this models gives the parameter estimates in Table 3, which are very similar to those earlier obtained. Using random-effects for the caribou has the advantage that inference for caribou outside of the sample population is eased, whereas the previously employed technique of including indicator variables for each caribou necessarily limits the scope of the analysis to the sample at hand.

Table 3: Fixed-Effects of the model in Equation (5).

	age[1,2)	age[2,3)	age[3,4)	age>=4	sexMale
Fixed Effects	6.699	7.558	7.342	10.550	3.456

6 Conclusion

Using exploratory analysis and statistical modelling via the Bradley-Terry model, we have analyzed the fighting hierarchies among a group of 20 traced caribou. We found ample evidence of different fighting abilities among the caribou, where being male and older was found to increase the odds of winning a fight.

Further research should be conducted about the fighting patterns among caribou, as the gained information might provide wider insights into the social structure and behaviour of wildlife caribou. More information should be collected about the individual caribou, as this could significantly improve the fit of the considered methods. Also, spatial statistics could be used to model the roaming behaviour of the caribou and their encounters among each other.

References

AGRESTI, A. (1990): *Categorical Data Analysis*, vol. 49 of *Wiley Series in Probability and Mathematical Statistics*, Wiley.

BARRETTE, C. AND D. VANDAL (1990): "Sparring, relative antler size, and assessment in male caribou," *Behav. Ecol. Sociobiol.*, 26, 383–387.

BOWKER, A. (1948): "A test for symmetry in contingency tables," *J. Am. Stat. Assoc.*, 43, 572–574.

BRADLEY, R. AND M. TERRY (1952): "Rank analysis of incomplete block designs: I. The method of paired comparisons," *Biometrika*, 39, 324–345.

BRITANNICA, E. (2006): *Britannica Concise Encyclopedia*, vol. 440, Encyclopaedia Britannica.

CAUSSINUS, H. (1966): "Contribution a l'analyse statistique des tableaux de correlation," 29, 77–183.

FIENBERG, S. AND K. LARNTZ (1976): "Log linear representation for paired and multiple comparisons models," *Biometrika*, 245–254.

FIRTH, D. (2004): "Quasi-variances," *Biometrika*, 91, 65–80.

——— (2005): "Bradley-Terry models in R," *J. Stat. Softw.*, 48.

MCNEMAR, Q. (1947): "Note on the sampling error of the difference between correlated proportions or percentages." *Psychometrika*, 12, 153–157.

Appendix

R Code

```
set.seed(1215)  # for reproducibility

# use of knitr (Sweave successor) to create this
# report
opts_chunk$set(fig.align = "center", cache = FALSE,
    message = FALSE, echo = FALSE, eval = TRUE)
options(replace.assign = TRUE, width = 85)

Encounters <- read.table(file = "Table2.txt")

Caribou <- data.frame(id = 1:20, Sex = factor(x = c(0,
    0, 1, 1, 0, 0, 1, 1, 0, 1, 1, 0, 1, 1, 0, 1, 0,
    1, 0, 0), labels = c("Female", "Male")), Age = c(5,
    3, 4, 2, 5, 2, 1, 2, 1, 3, 2, 3, 3, 2, 2, 1, 5,
    3, 1, 5))

Caribou$Age <- factor(Caribou$Age, labels = c("<1",
    "[1,2)", "[2,3)", "[3,4)", ">=4"))

# symmetry check

x <- numeric()

bowker.test <- function(x) {
    cs <- 0
    for (j in 1:20) {
        for (i in (j + 1):21) {
            if (i == 21)
                return(cs) else {
                if (x[i, j] + x[j, i] == 0)
                  cs <- cs + 0 else cs <- cs + (x[i, j] - x[j, i])^2/(x[i,
                  j] + x[j, i])
            }
        }
    }
}
# data preparation
pairwise <- combn(20, 2)

player1 <- pairwise[1, ]
player2 <- pairwise[2, ]

win1 <- vector()
win2 <- vector()

for (i in 1:length(player1)) {
    win1 <- c(win1, Encounters[player1[i], player2[i]])
    win2 <- c(win2, Encounters[player2[i], player1[i]])
}

player1 <- factor(player1, levels = 1:20)
player2 <- factor(player2, levels = 1:20)

caribou.bt <- data.frame(player1 = player1, player2 = player2,
    win1 = win1, win2 = win2)

# no of encounters which did not happen
no.without.fights <- sum(rowSums(cbind(win1, win2)) ==
    0)


# due to sparseness, chi-squared approximation in
# bowker test is not valid -> apply custom
# procedure:
```

```
caribou.symmetry <- melt(caribou.bt, id.vars = c("player1",
    "player2"))
caribou.symmetry$lambda <- factor(c(1:190, 1:190))
caribou.symmetry$lambda2 <- factor(1:380)

pois.fit1 <- glm(formula = value ~ lambda, data = caribou.symmetry,
    family = poisson)

# large residual deviance confirms that model is
# not adequate, with p-value of zero

pvalue <- 1 - pchisq(q = pois.fit1$deviance, df = pois.fit1$df.residual)

### Bradley-Terry model

BT.fit <- BTm(outcome = cbind(win1, win2), formula = ~player,
    id = "player", player1 = player1, player2 = player2,
    data = caribou.bt, refcat = "7")



player1.df <- data.frame(player = player1, age = Caribou$Age[player1],
    sex = Caribou$Sex[player1])

player2.df <- data.frame(player = player2, age = Caribou$Age[player2],
    sex = Caribou$Sex[player2])

caribou.bt.cov <- data.frame(win1 = win1, win2 = win2)
caribou.bt.cov$player1 <- player1.df
caribou.bt.cov$player2 <- player2.df

# fit extended bradley-terry model from equation
# (5)

BT.fit.cov <- BTm(outcome = cbind(win1, win2), formula = ~age +
    sex + (1 | player), id = "player", player1 = player1.df,
    player2 = player2.df, data = caribou.bt.cov, x = TRUE)


# plot of Caribou Encounters (Figure 1)
rownames(Encounters) <- paste0(1:20)
colnames(Encounters) <- paste0(1:20)

encs <- as.matrix(Encounters)
names(dimnames(encs)) <- c("row", "column")
df1 <- melt(encs, value.name = "count")
df1$row <- factor(df1$row, levels = paste0(20:1))
df1$column <- factor(df1$column, levels = paste0(1:20))

p <- ggplot(aes(x = column, y = row), data = df1) +
    geom_point(aes(size = count), shape = 21, colour = "black",
        fill = "cornsilk") + scale_size_area(max_size = 8,
    guide = FALSE)
p <- p + geom_text(data = subset(df1, df1$count > 0),
    aes(y = as.numeric(row) - sqrt(count)/8, label = count),
    vjust = 1, color = "black", size = 2.5)
p <- p + labs(x = "", y = "", title = "Caribou Encounters") +
    theme(text = element_text(size = 10))
p
# produces Figure 2
won.encs <- apply(Encounters, 1, sum)
Caribou$Wins <- won.encs
df.barplot <- arrange(Caribou, Age, Wins)
df.barplot <- ddply(df.barplot, "Age", transform, label_y = cumsum(Wins))
df.barplot$id <- as.character(df.barplot$id)
ggplot(aes(x = Age, y = Wins), data = df.barplot) +
    geom_bar(stat = "identity", aes(fill = Sex), color = "black") +
    labs(y = "# of wins") + geom_text(aes(x = Age,
    y = label_y, label = id), vjust = -0.2) + scale_fill_brewer(palette = "Pastel1")
```

```
# calculate values for Table 1
coefs <- rbind(summary(BT.fit)$coefficients[1:6, 1:2],
    c(0, 0), summary(BT.fit)$coefficients[7:19, 1:2])

WinPercentage <- lapply(1:20, FUN = function(x) {
    sum(Encounters[x, ])/(sum(Encounters[x, ]) + sum(Encounters[,
        x]))
})

coefs.df <- data.frame(id = 1:20, WinPercentage = unlist(WinPercentage))
coefs.df$beta <- coefs[, 1]
coefs.df$sd <- coefs[, 2]
coefs.df <- coefs.df[order(coefs.df$beta, decreasing = TRUE),
    ]
colnames(coefs.df) <- c("ID", "Winning Percentage",
    "$9+3$", "S.D.")

my.xtable <- xtable(coefs.df, digits = 2, caption = paste0("Table of winning percentages for each caribou "
    "and results of fitting the Bradley-Terry model without covariates"),
    label = "tab:brad1", align = "c|c|l|ll|")

print(my.xtable, include.rownames = FALSE, include.colnames = FALSE,
    add.to.row = list(pos = list(0), command = "ID & Win \\% & $\\lambda_i $ & S.D.  \\\\ "))
# calculate quasi-standard errors and plot them
caribou.qv <- qvcalc(BTabilities(BT.fit))
plot(caribou.qv)
# produce regression output in Table 2
Caribou$lambda <- coefs[, 1]
stargazer(lm(lambda ~ Age + Sex, data = Caribou), style = "aer",
    title = "Output for the regression model of $\\lambda_i$ on the explanatory variables \\emph{age} and \
    label = "tab:stargazer")
# fixed effects of model from Equation (5)
fixed.effs <- BT.fit.cov$coefficients[1:5]
fixed.effs <- as.data.frame(t(fixed.effs))
rownames(fixed.effs) <- "Fixed Effects"
stargazer(fixed.effs, summary = FALSE, title = "Fixed-Effects of the model in Equation (\\ref{eq:mixed}).",
    style = "aer", label = "tab:mixed")
```