Biost/Stat 571: Exercise 2

Answer Key

Question 1

TenHave and Uttal (1994) report on the results of a psychology experiment in which children were shown a map giving the location of a hidden toy, and allowed up to three chances to successfully find the object. The goal of the experiment is to compare a group of children presented with a correctly oriented map (control) to a second group that was shown a rotated map. Are the children presented with the rotated map able to compensate and find the hidden object as well as the control children? Each child repeated the test 10 times. During each test, a child will only attempt the maze a second time if they fail the first time. Similarly, a child will only attempt the maze a third time if they fail the second time. If a child is unsuccessful the third time the outcome is an overall failure. Thus the outcome, number of attempts to successfully find the object takes (Y), takes on one of four values with $Y = 4$ denoting all three attempts having failed.

Part a

Figure (1) displays the distribution of the outcome, at each of the ten tests, by treatment group.

It is difficult to discern a "learning effect" over time for the group that was presented the correctly oriented map (control). It does not seem as though the percentage that fail decreases over time, nor that the percentage that pass the first time increases with time. In the treatment group (the group that was presented with the rotated maps), it seems at though there is some evidence of a "learning effect". In particular, the percentage that fail decreases across the 10 tests, while the percentage that pass the first time increases.

Part b

Table 1 shows the number of attempts by treatment group for test 1.

Table 1: Number of attempts by treatment group for test 1.

	Number of Attempts			
	1	2	3	4
Control	18	14	7	7
Rotated	2	6	5	30

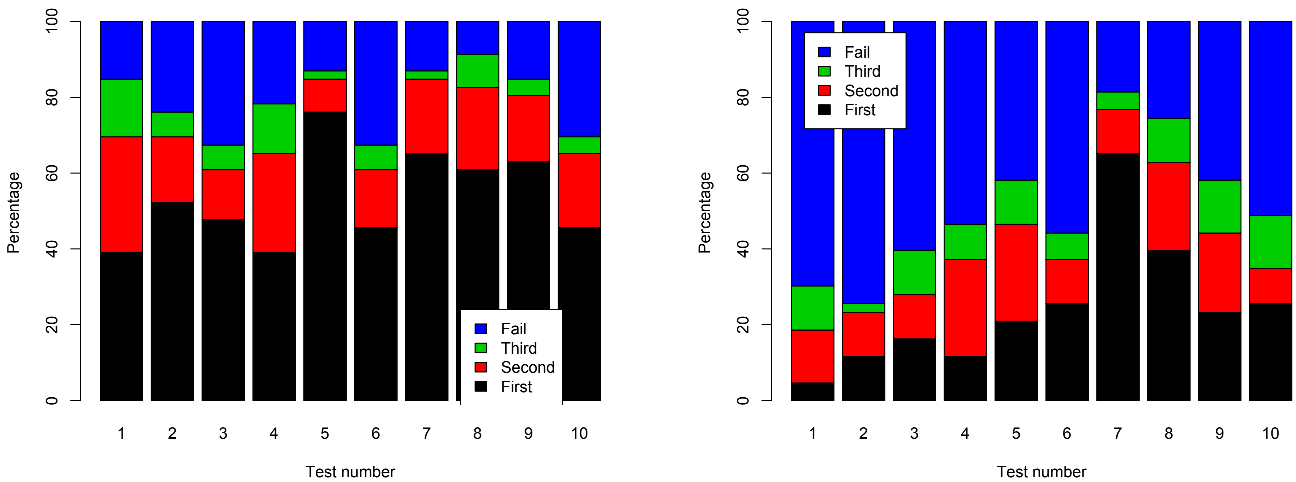

(a) Control Group ($n = 46$) (b) Rotated Group ($n = 43$)

Figure 1: Percentage of the number of attempts per test.

Parts c

Table 2 shows the parameter estimates and 95% confidence intervals from regressing $Y_{i1} = \mathbf{1}[T_{i1} > 1]$ on treatment group. The odds of taking more than one attempt to successfully find the object (or not finding the object at all) is 13.2 (95% CI: (2.83, 61.3)) times higher for the rotated group than the control group. We would obtain the same odds ratio by collapsing columns 2, 3, and 4 in the table presented in (b) and calculating the odds ratio.

Table 2: Parameter estimates and 95% CIs for outcome $Y_{i1} = \mathbf{1}[T_{i1} > 1]$.

	Estimate	95% CI
β_0	0.44	(-0.15, 1.03)
β_1	2.58	(1.04, 4.12)

Part d

Table 3 shows the parameter estimates and 95% confidence intervals from regressing $Y_{i2} = \mathbf{1}[T_{i2} > 2]$ on treatment group. The odds of taking more than two attempts to successfully find the object (or not finding the object at all) is 10.0 (95% CI: (3.71, 27.0)) times higher for the rotated group than the control group. We would obtain the same odds ratio by collapsing columns 1 and 2, and 3 and 4 in the table presented in (b) and calculating the odds ratio.

Table 3: Parameter estimates and 95% CIs for outcome $Y_{i2} = \mathbf{1}[T_{i2} > 2]$.

	Estimate	95% CI
β_0	-0.83	(-0.45, -0.20)
β_1	2.30	(1.31, 3.29)

Part e

Table 4 shows the parameter estimates and 95% confidence intervals from regressing $Y_{i3} = \mathbf{1}[T_{i3} > 3]$ on treatment group. The odds of never successfully finding the object is 12.9 (95% CI: (4.57, 36.2)) times higher for the rotated group than the control group. We would obtain the same odds ratio by collapsing columns 1, 2 and 3 in the table presented in (b) and calculating the odds ratio.

Table 4: Parameter estimates and 95% CIs for outcome $Y_{i3} = \mathbf{1}[T_{i3} > 3]$.

	Estimate	95% CI
β_0	-1.72	(-2.52, -0.91)
β_1	2.55	(1.52, 3.59)

Part f

The proportional odds model that is being fit is given by

$$\text{logit}(P(\mathbf{T}_1 > j)) = \beta_{(0,1)} + \beta_{(0,2)} + \beta_{(0,3)} + \beta_R \text{Rotate}.$$

Hence, β_x is the common log-odds ratios for the logistic regressions performed in (c)-(e), and $\beta_{(0,1)}$, $\beta_{(0,2)}$, and $\beta_{(0,3)}$ are the negatives of their corresponding intercepts in the regressions above. Using the `polr()` function in `R`, we have the following estimates: $(\hat{\beta}_{(0,1)}, \hat{\beta}_{(0,2)}, \hat{\beta}_{(0,3)}, \hat{\beta}_R) = (-0.45, 0.88, 1.64, 2.46)$. As expected, the estimates obtained using the proportional odds model are very similar to those obtained using separate logistic regressions. We have similar interpretations of our model parameters as in parts (c)-(e). $\beta_{(0,j)}$ is the log-odds of not finding the object in the first j attempts in the control group. β_R is the log-odds of "slow" completion comparing the rotated and control groups. Note that the definition of "slow" can be any of $T_1 > 1$, $T_1 > 2$, and $T_1 > 3$. Note also that the odds ratios are independent of how we define "slow" (*ie.* proportional odds).

Part g

Since the odds ratios obtained in parts (c)-(e) comparing the rotated group to the control group are very similar, the proportional odds assumption seems reasonable.

Part h

Table 5 shows the number of attempts by treatment group for test 8.

Table 5: Number of attempts by treatment group for test 1.

	Number of Attempts			
	1	2	3	4
Control	28	10	4	4
Rotated	17	10	5	11

To compare these results to those from test 1, let $\Delta_i = Y_{i8} - Y_{i1}$ denote the difference between the outcomes of the 8^{th} and 1^{st} tests for the i^{th} child. Consequently, a negative change indicates an improvement over time. For example, $\Delta = -3$ corresponds to a score of 1 on the 8^{th} test and a score of 4 on the 1^{st}, indicating improvement. The table below summarizes these differences by treatment group.

Δ	-3	-2	-1	0	1	2	3	Total
Rotated	10	14	5	8	3	2	1	43
Control	5	2	13	18	5	1	2	46

From this we find that 63% (29/43) of the rotated group improved their score from the 1^{st} to the 8^{th} tests, with 10 children passing on the first attempt during the 8^{th} test when they had failed all three attempts during the 1^{st} test. In the control group, only 47% improved their score with a much larger group (42%) maintaining their score between the two tests. It is interesting to note that of the 8 children in the rotated group that maintained their score, 7 maintained a score of 4. However, in the control group, of the 18 that maintained their score 16 maintained scores of 1 or 2. Neither group exhibited much learning between the two tests.

Part i

The proportional odds model that is being fit is given by

$$\text{logit}(P(\mathbf{T}_8 > j)) = \beta_{(0,1)} + \beta_{(0,2)} + \beta_{(0,3)} + \beta_R \text{Rotate}.$$

Using the `polr()` function in `R`, we have the following estimates: $(\hat{\beta}_{(0,1)}, \hat{\beta}_{(0,2)}, \hat{\beta}_{(0,3)}, \hat{\beta}_R) = (0.47, 1.50, 2.12, 0.95)$. The three intercept parameters are much larger than those obtained for test 1, while β_R is much smaller. Since β_R is smaller, this shows that the rotated and control groups are performing more similarly for test 8 than test 1, which is indirect evidence for a learning effect.

Part j

The score equations that yield the maximum likelihood estimate of $\boldsymbol{\beta} = (\beta_{(0,1)}, \beta_{(0,2)}, \beta_{(0,3)}, \beta_1)$ are given by

$$U(\boldsymbol{\beta}) = \sum_i \left(\frac{\partial \pi_i}{\partial \boldsymbol{\beta}}\right)^T (\mathbf{L}\Sigma_i^{-1}\mathbf{L}^T)\mathbf{L}(\mathbf{Y}_i - \pi_i).$$

Pooling $(Y_{i1}(1), Y_{i1}(2), Y_{i1}(3))$, we have

$$U(\boldsymbol{\beta}) = \sum_{i=1}^{n} \left(\frac{\partial \mu_i}{\partial \boldsymbol{\beta}}\right)^T \Sigma_i^{-1}(\mathbf{Y}_i - \mu_i),$$

where $\Sigma_i^{-1} = \text{diag}(\mu_i(1-\mu_i))$. So, we can see that the estimates that satisfy these two scores equations are different. Further, pooling the logistic regressions does not account for the dependence between the observations made on a particular individual (however note that this will only affect the variances, and not the estimates).

Question 2

Part a

To find $\mathbf{B}$ such that $\mathbf{B}(\mathbf{Y} - \boldsymbol{\pi}) = (\mathbf{Y} - \boldsymbol{\mu} \cdot \mathbf{H})$, we solve algebraically (writing $\mathbf{H}$ in terms of the Y_i) to find

$$B_{ij} = \begin{cases} i = j & : & 1 \\ i > j & : & \frac{\pi_i}{1-\gamma_{i-1}} \\ i < j & : & 0 \end{cases},$$

a lower triangular matrix. For example, for 4 cutpoints,

$$B = \begin{pmatrix} 1 & 0 & 0 & 0 \\ \frac{\pi_2}{1-\gamma_1} & 1 & 0 & 0 \\ \frac{\pi_3}{1-\gamma_3} & \frac{\pi_3}{1-\gamma_2} & 1 & 0 \\ \frac{\pi_4}{1-\gamma_3} & \frac{\pi_4}{1-\gamma_3} & \frac{\pi_4}{1-\gamma_3} & 1 \end{pmatrix},$$

where $\pi_i = P(Y = i)$ and $\gamma_i = \sum_{k=1}^{i} \pi_i$.

Part b

$\mathbf{B}$ is constructed so that $\mathbf{B}(\mathbf{Y} - \boldsymbol{\pi}) = (\mathbf{Y} - \boldsymbol{\mu} \cdot \mathbf{H})$. So,

$$\begin{aligned} \text{Cov}(\mathbf{B}(\mathbf{Y} - \boldsymbol{\pi})) &= \text{Cov}(\mathbf{Y} - \boldsymbol{\mu} \cdot \mathbf{H}) \\ &= E[(\mathbf{Y} - \boldsymbol{\mu} \cdot \mathbf{H})(\mathbf{Y} - \boldsymbol{\mu} \cdot \mathbf{H})^T]. \end{aligned}$$

We know that all the off-diagonal elements are zero. As for the diagonal elements, variances of $\mathbf{B}(\mathbf{Y} - \boldsymbol{\pi})$ are the same (for the first element) or less (for the other elements) than the diagonal elements of Σ. This is because $\text{Cov}(\mathbf{B}(\mathbf{Y} - \boldsymbol{\pi})) = \mathbf{B}\Sigma\mathbf{B}^T$. The negative off-diagonals of Σ (a multinomial covariance) are multiplied by the positive off-diagonals of $\mathbf{B}$, and subtracted from $\text{Var}(\mathbf{Y})$, since the diagonal elements of $\mathbf{B}$ are 1.

Part c

We fit a continuation ratio logit model using a single odds ratio for treatment group using code adapted from the WESDR example (shown in the appendix) from the notes (slides 123-1 to 123-4). An analysis of deviance indicates that a common treatment odds ratio is appropriate (p=0.65).

Part d

For test 8, an analysis of deviance indicates that a common treatment odds ratio is appropriate (p=0.98).

Part e

Modeling test 1 responses, the odds of taking longer to complete the task (relative to less time, for any cutpoint) are $7.4(= 1/\exp(-2.0076))$ times higher for the group that sees the rotated maze, relative to the control group. For test 8, the odds ratio is $2.2(= 1/\exp(-0.8072))$. The two groups perform more similarly at time 8 than at time 1. This is indirect evidence for a learning effect. We don't know that either group improved by completing the task in fewer tries - we just know that the gap between the groups is smaller. As in question 1, we can use the EDA to support the idea that the group seeing the rotated map improved its performance.

Question 3

Part a

Based on the plots produced in Question 1, it seems reasonable to allow each treatment group to have its own slope. We can allow each group to have linear slopes, or quadratic slopes. A linear model would include main effects for trial and treatment, and an interaction term between trial and treatment. The corresponding POM model is given by

$$\text{logit}(P(T_{ij} \leq c) = \beta_{(0,c)} + \beta_1 \text{Rotated}_i + \beta_2 j + \beta_3 \text{Rotated}_i \times j.$$

In this model, the primary scientific question would be addressed by the interaction term (which looks at waning). Note that the contrast between the rotated and control groups is given by $\beta_1 + j \times \beta_3$, allowing differences across groups to change over time.

A quadratic model would include main effects for treatment, trial and trial^2, as well as interaction terms between trial and treatment, and trial^2 and treatment. The corresponding POM model is given by

$$\text{logit}(P(T_{ij} \leq c) = \beta_{(0,c)} + \beta_1 \text{Rotated}_i + \beta_2 j + \beta_3 j^2 + \beta_4 \text{Rotated}_i \times j + \beta_4 \text{Rotated}_i \times j^2.$$

Here, the primary scientific question would be addressed by the treatment main effect (as was done in question 2). In this case, the contrast between the rotated and control groups is given by $\beta_1 + j \times \beta_4 + j^2 \times \beta_5$.

Part b

There are several possibles summary methods for T_{ij} that might be used to compare treatment groups regarding evidence for a "learning effect". The first option is to estimate each child's slope (for a regression over time, say), then calculate the average slope in both the rotated and control groups. We can then compare the two groups using a t-test. The second option is to calculate the change in the number of attempts to find the object from test 1 to test 8. We can then perform a t-test comparing the rotated and control groups, or we can use a rank test (e.g. Wilcoxon rank sum test).

Appendix - R Code

```
##########################
## Question 1
##########################

library(MASS)
library(Design)
#
##
####################################################################################
#### Load in the tenhave data, label variables and make all data manipulations ####
####################################################################################
##
#
tenhave <- matrix( scan( "tenhave.data", sep = " " ), ncol = 11, byrow = T )
dimnames( tenhave ) <- list( NULL, c( "rotate", paste( "test", 1:10, sep = "" ) ) )
tenhave <- as.data.frame( tenhave )
#
## Change the coding of the 'rotate' variable so that rotate = 1 is the rotated (intervention) group
#
tenhave$rotate <- 1 - tenhave$rotate
#
## Compute `change' outcome; Test8 - Test1
##     - positive change implies improvement over time
#
tenhave$change <- tenhave$test8 - tenhave$test1
#
##
##################
#### Part (a) ####
##################
##
#
##
#### Look at some individual 'learning curves'
##
#
sample.rotate <- sort( sample( c(1:43), 8 ) )
sample.not <- sort( sample( c(44:89), 8 ) )
par( mfrow = c(4,4) )
for( i in 1:8 ) {
plot( 1:10, tenhave[sample.rotate[i], 2:11], type = "l", ylim = c(0, 4),
xlab = "Test number", ylab = "Outcome", main = paste("ID:", sample.rotate[i], "- rotated") )
}
for( i in 1:8 ) {
plot( 1:10, tenhave[sample.not[i], 2:11], type = "l", ylim = c(0, 4),
```

```
xlab = "Test number", ylab = "Outcome", main = paste("ID:", sample.not[i], "- not") )
}
#
##
#### Calculate percentages within each test; PBT = percent by test
##
#
pbt.rotated <- matrix( 0, ncol = 10, nrow = 4 )                        # n = 43
pbt.not <- matrix( 0, ncol = 10, nrow = 4 )                            # n = 46
for( i in 1:10 )
{
pbt.rotated[,i] <- (table( tenhave[1:43,(i+1)] ) / 43) * 100
pbt.not[,i] <- (table( tenhave[44:89,(i+1)] ) / 46) * 100
}
#
##
#### Produce barplots
##
#
barplot( pbt.not, names = as.character(c(1:10)), xlab = "Test number", ylab = "Percentage",
col = c(1,2,3,4) )
legend( 8.7, 24, c("Fail", "Third", "Second", "First" ), fill = c(4,3,2,1), bg="white" )
barplot( pbt.rotated, names = as.character(c(1:10)), xlab = "Test number", ylab = "Percentage",
col = c(1,2,3,4))
legend( 0.5, 97, c("Fail", "Third", "Second", "First" ), fill = c(4,3,2,1), bg="white" )
#
##
##################
#### Part (b) ####
##################
##
#
with(tenhave, table(rotate, test1))
#
##
##################
#### Part (c) ####
##################
##
#
# create indicator for not succeeding on 1st attempt
tenhave$test1.first<-ifelse(tenhave$test1!=1, 1, 0)
logmod1<-glm(test1.first~rotate, data=tenhave, family=binomial)
summary(logmod1)
exp(logmod1$coef)
exp(confint.default(logmod1))
#
##
```

```
##################
#### Part (d) ####
##################
##
#
# create indicator for not succeeding on 1st or 2nd attempt
tenhave$test1.second<-ifelse(tenhave$test1!=1 & tenhave$test1!=2, 1, 0)
logmod2<-glm(test1.second~rotate, data=tenhave, family=binomial)
summary(logmod2)
exp(logmod2$coef)
exp(confint.default(logmod2))
#
##
##################
#### Part (e) ####
##################
##
#
# create indicator for not succeeding on any attempt
tenhave$test1.fail<-ifelse(tenhave$test1==4, 1, 0)
logmod3<-glm(test1.fail~rotate, data=tenhave, family=binomial)
summary(logmod3)
exp(logmod3$coef)
exp(confint.default(logmod3))
#
##
##################
#### Part (f) ####
##################
##
#
mod.POM1<-polr(factor(test1)~rotate, data=tenhave)
summary(mod.POM1)
#
##
##################
#### Part (h) ####
##################
##
#
with(tenhave, table(rotate, test8))
#
##
##################
#### Part (i) ####
##################
##
#
```

```
mod.POM8<-polr(factor(test8)~rotate, data=tenhave)
summary(mod.POM8)
#
##

##########################
## Question 2
##########################

## construct continuation model data for trials 1 and 8
## (1) construct pairs (Y,H)
nsubj<-nrow(tenhave)
ncuts<-3
y1.crm<-NULL
y8.crm<-NULL
h1.crm<-NULL
h8.crm<-NULL
id<-NULL
for (j in (1:nsubj)) {
yj1<-rep(0,ncuts)
yj8<-rep(0,ncuts)
if (tenhave$test1[j] <= ncuts ) { yj1[tenhave$test1[j]] <- 1 }
if (tenhave$test8[j] <= ncuts ) { yj8[tenhave$test8[j]] <- 1 }
hj1 <- 1-c(0,cumsum(yj1)[1:(ncuts-1)])
hj8 <- 1-c(0,cumsum(yj8)[1:(ncuts-1)])
y1.crm <- c(y1.crm, yj1)
y8.crm <- c(y8.crm, yj8)
h1.crm <- c(h1.crm, hj1)
h8.crm <- c(h8.crm, hj8)
id <- c(id, rep(j, ncuts))
}

## (2) construct intercepts
level <- factor(rep(1:ncuts,nsubj))
int.mat <- NULL
for (j in 1:ncuts) {
intj <- rep(0,ncuts)
intj[j] <- 1
int.mat <- cbind(int.mat, rep(intj, nsubj))
}

#### (3) expand the covariate
rotate.crm <- rep(tenhave$rotate, rep(ncuts, nsubj))

#### (4) drop H=0 and build dataframe
keep1 <- h1.crm==1
keep8 <- h8.crm==1
t1.crm.data <- data.frame(id=id[keep1], y=y1.crm[keep1], level=level[keep1], rotate=rotate.crm[keep1])
```

```
t8.crm.data <- data.frame(id=id[keep8], y=y8.crm[keep8], level=level[keep8], rotate=rotate.crm[keep8])

### model fitting for 2c
fit1 <- glm(y~level+rotate, family=binomial, data=t1.crm.data)
summary(fit1)
fit2 <- glm(y~level*rotate, family=binomial, data=t1.crm.data)
summary(fit2)
anova(fit1,fit2)
1-pchisq(0.85938, 2)

### model fitting for 2d
fit1d <- glm(y~level+rotate, family=binomial, data=t8.crm.data)
summary(fit1d)
fit2d <- glm(y~level*rotate, family=binomial, data=t8.crm.data)
summary(fit2d)
anova(fit1d, fit2d)
1-pchisq(0.052734, 2)
```