Cumulative Logit model for mental impairment data:

R code (using package "VGAM"):
(This package uses parameterization $\alpha_j + \beta x$ for the linear predictor.)

a) Fit Model:

```
> mental <-
read.table(file="http://sites.williams.edu/bklingen/files/2012/02/mental.txt",
header=TRUE)
> head(mental)
  impair ses life
1      1   1    1
2      1   1    9
3      1   1    4
4      1   1    3
5      1   0    2
6      1   1    0
> require(VGAM)
> fit <- vglm(impair ~ life + ses, family=cumulative(parallel=TRUE), data=mental)
> summary(fit)

Selected Output:
Coefficients:
             Estimate Std. Error  z value
(Intercept):1 -0.28176    0.62304 -0.45223
(Intercept):2  1.21291    0.65119  1.86260
(Intercept):3  2.20947    0.71719  3.08075
life          -0.31888    0.11944 -2.66973
ses            1.11112    0.61427  1.80884
```

b) Test if life is significant

```
> maxl=logLik(fit)
> maxl
[1] -49.54895
> fit0 <- vglm(impair ~ ses, family=cumulative(parallel=TRUE), data=mental)
> maxl0 <- logLik(fit0)
> maxl0
[1] -53.43718
> LR.stat <- -2*(maxl0 - maxl)
> LR.stat
[1] 7.776457
> 1 - pchisq(LR.stat,df=1)
[1] 0.005293151
```

Likelihood Ratio (LR) test for H0: $\beta_1 = 0$ vs H1: $\beta_1 \neq 0$ (coefficient for life): LR statistic = 7.78 on df = 1: P-value = 0.0053.
The data provide evidence (P-value = 0.0053) of a significant effect of the number of life events on the cumulative log-odds of mental impairment. (Wald Statistic in Chi-square form: (-0.3189/0.1210)^2 = (-2.6697)^2 = 6.95, df = 1: P-value = 0.0084, similar conclusion.)

Interpretation of effect: For both low and high SES adults, the odds of a mental impairment score less than or equal to any category *j* (instead of greater than *j*) decrease by a factor of exp(-0.3189) = 0.73 for every unit increase in the life event score. This is true for all *j* (proportional odds assumption). For instance, when *j* = 1 = "well", the estimated odds of feeling well decrease by a factor of 0.73 for every additional life event. When j = 2, the odds of feeling well or showing mild symptoms of mental impairment (versus moderate symptoms or mental impairment) decrease by a factor of 0.73 for every unit increase in the life event score.

c) Compute CI for the effect if life.

Cannot compute profile likelihood interval directly (package ordinal can compute profile likelihood intervals, see below). From there, we get CI [-0.57184548, -0.09203351] (beware different parameterization): We are 95% confident that the odds of mental impairment below any level j decrease by a factor of at least 0.91 and at most 0.56 for every additional life event.

Wald interval from output above: -0.3188 +/- 1.96*0.119 = [-0.55204 -0.08556] (similar)

d) Fit a more general model that allows an interaction between life and SES. Interpret the interaction.

```
> fit1 <- vglm(impair ~ life + ses + life*ses, family=cumulative(parallel=TRUE),
data=mental)
> summary(fit1)
Coefficients:
              Estimate Std. Error  z value
(Intercept):1  0.098131    0.81107  0.12099
(Intercept):2  1.592521    0.83729  1.90199
(Intercept):3  2.606616    0.90980  2.86504
life          -0.420448    0.19034 -2.20893
ses            0.370876    1.13027  0.32813
life:ses       0.181294    0.23613  0.76777
```

Estimated coefficient for interaction effect: -0.1813

Interpretation: The decrease in the estimated cumulative odds of mental impairment below any level is stronger for those adults with low socioeconomic status (ses=0) than for those with high socioeconomic status (ses=1). (Why? Because the ses coefficient above is positive, and it refers to the high SES group.) In particular, for adults with low SES, the estimated cumulative odds decrease by a factor of exp(-0.4202) = 0.66 for every unit increase in the life events score, compared to a decrease of exp(-0.4204 + 0.1813) = 0.79 for those of high SES.

e) Is the interaction significant?

```
> fit1 <- vglm(impair ~ life + ses + life*ses, family=cumulative(parallel=TRUE),
data=mental)
> LR.stat <- -2*(maxl - logLik(fit1))
> LR.stat
[1] 0.5934586
> 1 - pchisq(LR.stat,df=1)
[1] 0.4410848
```

The coefficient for the interaction term (estimate = 0.1813, SE = 0.238) is not significant. Likelihood ratio test statistic = 0.59, df = 1: P-value = 0.441. The decrease in the estimated odds can be considered the same for both the high and low SES group.

f) Plot the estimated cumulative logits, cumulative probabilities and category probabilities against the life score for each SES category.

```
> ### fitted logits
> life1 <- seq(0,9,1)
> fit.logit.ses0 <- predict(fit,newdata=data.frame(life=life1,ses=0))
> fit.logit.ses1 <- predict(fit,newdata=data.frame(life=life1,ses=1))
> name <- colnames(fit.logit.ses1)
> plot.data <- data.frame(life=rep(life1,3),ses=rep(c("SES = low","SES =
high"),each=3*10), type=rep(name,each=10), logit=c(fit.logit.ses0,fit.logit.ses1))
> head(plot.data)
  life       ses             type      logit
1    0 SES = low logit(P[Y< = 1]) -0.2817575
2    1 SES = low logit(P[Y< = 1]) -0.6006407
3    2 SES = low logit(P[Y< = 1]) -0.9195239
4    3 SES = low logit(P[Y< = 1]) -1.2384071
5    4 SES = low logit(P[Y< = 1]) -1.5572904
6    5 SES = low logit(P[Y< = 1]) -1.8761736
> xyplot(logit~life|ses, group=type, data=plot.data, type="l", auto.key =
list(points=FALSE, lines=TRUE, columns=3))
```

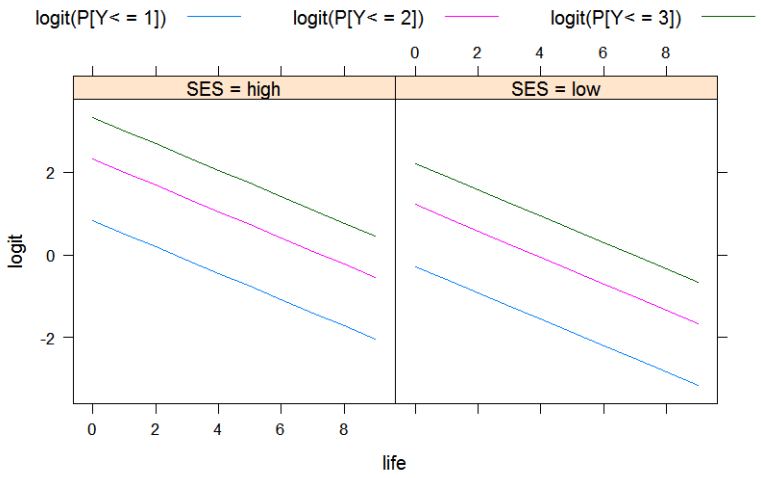


```
> ### fitted cumulative probabilities
> fit.cumprob.ses0 <- predict(fit,newdata=data.frame(life=life1,ses=0),
untransform=TRUE)
> fit.cumprob.ses1 <- predict(fit,newdata=data.frame(life=life1,ses=1),
untransform=TRUE)
> name <- colnames(fit.cumprob.ses1)
> plot.data <- data.frame(life=rep(life1,3),ses=rep(c("SES = low","SES =
high"),each=3*10), type=rep(name,each=10),
cumProb=c(fit.cumprob.ses0,fit.cumprob.ses1))
> head(plot.data)
  life       ses      type   cumProb
1    0 SES = low P[Y< = 1] 0.4300230
2    1 SES = low P[Y< = 1] 0.3541971
3    2 SES = low P[Y< = 1] 0.2850549
4    3 SES = low P[Y< = 1] 0.2247134
5    4 SES = low P[Y< = 1] 0.1740358
6    5 SES = low P[Y< = 1] 0.1328290
> xyplot(cumProb~life|ses, group=type, data=plot.data, type="l", auto.key =
list(points=FALSE, lines=TRUE, columns=3))
```

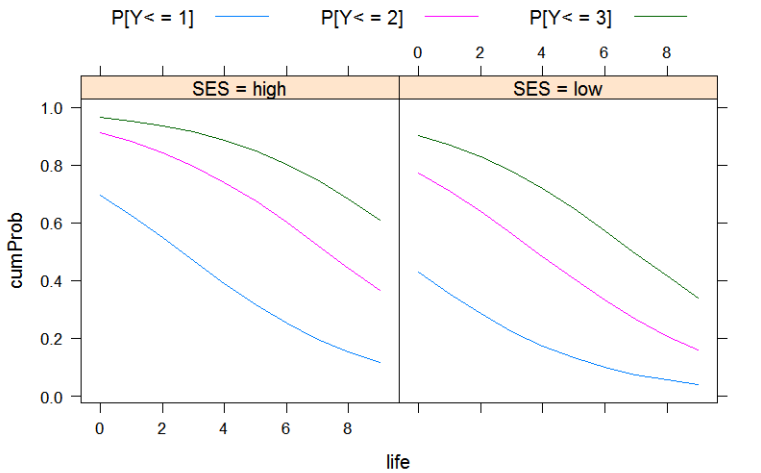


```
> fit.prob.ses0 <- predict(fit,newdata=data.frame(life=life1,ses=0), type="response")
> fit.prob.ses1 <- predict(fit,newdata=data.frame(life=life1,ses=1), type="response")
> name <- colnames(fit.prob.ses1)
> plot.data <- data.frame(life=rep(life1,4),ses=rep(c("SES = low","SES =
high"),each=4*10), type=rep(name,each=10), Prob=c(fit.prob.ses0,fit.prob.ses1))
```

```
> head(plot.data)
  life       ses type      Prob
1    0 SES = low well 0.4300230
2    1 SES = low well 0.3541971
3    2 SES = low well 0.2850549
4    3 SES = low well 0.2247134
5    4 SES = low well 0.1740358
6    5 SES = low well 0.1328290
> xyplot(Prob~life|ses, group=type, data=plot.data, type="l", auto.key =
list(points=FALSE, lines=TRUE, columns=3))
```

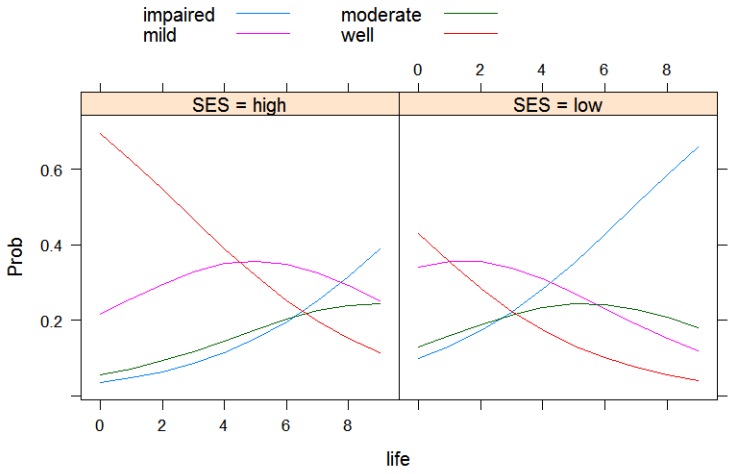

R code (using package "ordinal"):

Attention, package "ordinal" uses the latent variable coding, i.e., linear predictor = $\alpha_j - \beta x$!

a)

```
> mental <- read.table("mental.dat", header=TRUE)
> mental$impair <- factor(mental$impair,
labels=c("well","mild","moderate","impaired"), ordered=TRUE)
> head(mental)
  impair ses life
1   well   1    1
2   well   1    9
3   well   1    4
4   well   1    3
5   well   0    2
6   well   1    0
> require(ordinal)
> fit <- clm(impair ~ life + ses, data=mental)
> summary(fit)
 link  threshold nobs logLik AIC    niter max.grad cond.H
 logit flexible  40   -49.55 109.10 4(0)  3.17e-08 3.6e+02

Coefficients:
     Estimate Std. Error z value Pr(>|z|)
life   0.3189     0.1210   2.635   0.0084 **
ses   -1.1112     0.6109  -1.819   0.0689 .
---
Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
```

```
Threshold coefficients:
                  Estimate Std. Error z value
well|mild          -0.2819     0.6423  -0.439
mild|moderate       1.2128     0.6607   1.836
moderate|impaired   2.2094     0.7210   3.064
```

b)

```
> fit0 <- clm(impair ~ ses, data=mental) # model without life effect
> # or:
> # fit0 <- update(fit, ~ -life)
> anova(fit,fit0)
Likelihood ratio tests of cumulative link models:

     formula:            link: threshold:
fit0 impair ~ ses        logit flexible
fit  impair ~ life + ses logit flexible

     no.par    AIC  logLik LR.stat df Pr(>Chisq)
fit0      4 114.87 -53.437
fit       5 109.10 -49.549  7.7765  1   0.005293 **
---
```

Likelihood ratio statistic: 7.77; P-value for Likelihood ratio test: 0.0053.

c)

```
> confint(fit)
              2.5 %      97.5 %
life   0.09203351 0.57184548
ses   -2.34711898 0.06410755
```

95% Profile Likelihood Confidence for β_1: [0.092;0.572] (Remember, β_1 here is $-\beta_1$ from vglm output)

d)

```
> fit1 <- clm(impair ~ life + ses + ses*life, data=mental)
> summary(fit1)
 link  threshold nobs logLik AIC    niter max.grad cond.H
 logit flexible  40   -49.25 110.50 4(0)  2.30e-08 1.2e+03

Coefficients:
         Estimate Std. Error z value Pr(>|z|)
ses       -0.3709     1.1361  -0.326   0.7441
life       0.4204     0.1864   2.255   0.0241 *
ses:life  -0.1813     0.2383  -0.761   0.4468
```

Estimated coefficient for interaction effect: -0.1813

e)

```
> anova(fit1,fit)
     no.par   AIC  logLik LR.stat df Pr(>Chisq)
fit       5 109.1 -49.549
fit1      6 110.5 -49.252  0.5935  1     0.4411
```

LR-stat = 0.594; P-value = 0.441