i Cover Page

Department of Mathematical Sciences

Examination paper for ST2304 Statistical modelling for biologists and biotechologists

Examination date: 1st June 2022

Examination time (from-to): 09:00 – 13:00

Permitted examination support material: All support material is allowed

Academic contact during examination: Bob O'Hara
Phone: 915 54 416

Technical support during examination: Orakel support services
Phone: 73 59 16 00

OTHER INFORMATION

Do not open Inspera in multiple tabs, or log in on multiple devices, simultaneously. This may lead to errors in saving/submitting your answer.

Get an overview of the question set before you start answering the questions.

Read the questions carefully, make your own assumptions and specify them in your answer. Only contact academic contact if you think there are errors or insufficiencies in the question set.

Cheating/Plagiarism: The exam is an individual, independent work. Examination aids are permitted, but make sure you follow any instructions regarding citations. During the exam it is not permitted to communicate with others about the exam questions or distribute drafts for solutions. Such communication is regarded as cheating. All submitted answers will be subject to plagiarism control. *Read more about cheating and plagiarism here.*

Weighting: Weighting of the questions is given for each question.

ABOUT SUBMISSION

Answering in Inspera: If the question set contains questions that are not upload assignment, you must answer them directly in Inspera. In Inspera, your answers are saved automatically every 15 seconds.

NB! We advise against pasting content from other programs (other than as plain text), as this may cause loss of formatting and/or entire elements (e.g. images, tables).

Automatic submission: Your answer will be submitted automatically when the examination time expires and the test closes, as long as you have answered at least one question. This will happen even if you do not click “Submit and return to dashboard” on the last page of the question set. You can reopen and edit your answer as long as the test is open. If no questions are answered by the time the examination time expires, your answer will not be submitted. This is considered as “did not attend the exam”.

Withdrawing from the exam: If you become ill during the exam or wish to submit a blank answer/withdraw from the exam for another reason, go to the menu in the top right-hand corner and click “Submit blank”. This cannot be undone, even if the test is still open.

Accessing your answer post-submission: You will find your answer in Archive when the examination time has expired.

Good luck!

i Hardness Description

Foresters can be interested in all sorts of things. One is how hard their wood is (you don't want to use soft wood as your floor, because it will be dented the first time you sit down on a chair). They can test this by pushing a steel ball into the wood, but would prefer something easier and less destructive. So one thing that has been looked at it wood density: if there is a strong correlated, it could be used as a proxy to predict hardness.

Here we can look at some data on hardness and density, to see what the relationship is, and how best to predict hardness from measures of density. This can be done with regression. Although the positive relationship is obvious, if we want to get good predictions we need to check that the model fits the data well.

This is the data:

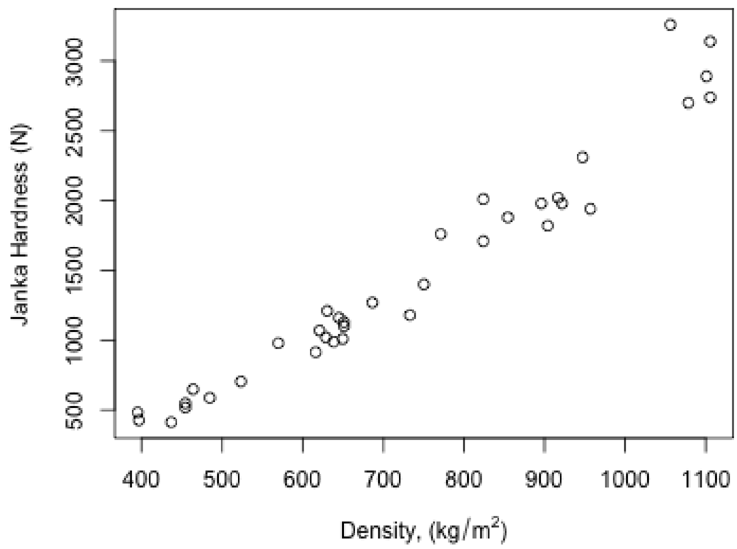

The first model was a simple linear regression. It gave the following summary:

```
##
## Call:
## lm(formula = hardness ~ density, data = janka)
##
## Residuals:
##     Min      1Q Median      3Q     Max
##    -338     -97    -16      93     625
##
## Coefficients:
##             Estimate Std. Error t value Pr(>|t|)
## (Intercept) -1160.50     108.58     -11    2e-12 ***
## density         3.59       0.14      25   <2e-16 ***
## ---
## Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
##
## Residual standard error: 180 on 34 degrees of freedom
## Multiple R-squared:  0.95,	Adjusted R-squared:  0.95
## F-statistic: 6.4e+02 on 1 and 34 DF,  p-value: <2e-16
```

1 Simple linear model Slope

What percentage of the variance is explained by the model (to the nearest percent)? **95**

What is the slope of the model, to 1 decimal place? **3.6** .

Maximum marks: 2

2 Simple Linear Model Prediction

Based on this model, what is the predicted hardness of a piece of wood with a density of 1000 kg/m^2? Answer to the nearest integer.

2430

Maximum marks: 1

3 SLM violations

The simple model was checked to see if the assumptions of the model were reasonable

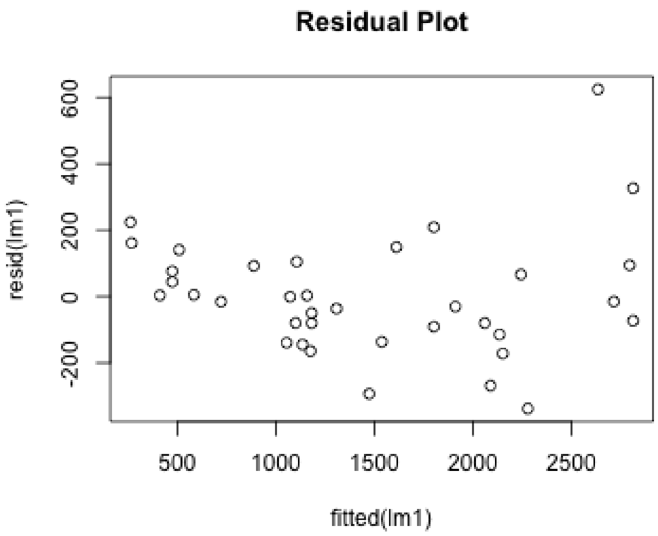

From studying this residual plot, which assumptions do you think are violated?
Select one or more alternatives:

- [X] Linearity
- [ ] Normality
- [ ] None, it looks OK
- [?] Outliers
- [X] Homoscedasticity (i.e. equal variance)

Maximum marks: 2

4 SLM: What to do

Based on this plot, what would you do next?
(this may not be what was actually done, and what comes next may not be the best thing to do)

Fill in your answer here

Format **B** *I* U x_2 x^2 I_x Ω Σ

I think the model is is not that bad, but it looks a little curved and heteroscedastic. However, based on the residual plot, it looks like the variance get bigger with higher mean. To fix this we could try a Box-Cox transformation, with a power that was less than 1. So either a log transformation or a square root transformation would be good.

(notes: (1) suggesting adding a quadratic term to cure the non-linearity also gets some marks, (2) there might be an outlier, but it isn't obvious that it's an outlier rather than just part of the heteroscedasticity.)

Words: 0

Maximum marks: 4

It was decided to try a square root transformation of hardness (the response). Fitting the model gave the following summary

```
##
## Call:
## lm(formula = sqrt(hardness) ~ density, data = janka)
##
## Residuals:
##    Min     1Q Median     3Q    Max
##   -3.5   -1.1   -0.3    1.0    4.8
##
## Coefficients:
##             Estimate Std. Error t value Pr(>|t|)
## (Intercept)    2.259      1.094       2     0.05 *
## density        0.047      0.001      33   <2e-16 ***
## ---
## Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
##
## Residual standard error: 2 on 34 degrees of freedom
## Multiple R-squared:   1, Adjusted R-squared:   1
## F-statistic: 1e+03 on 1 and 34 DF,  p-value: <2e-16
```

5 Sqrt Prediction

Based on this model, what is the predicted hardness of a piece of wood with a density of 1000 kg/m^2?

Answer to the nearest integer

2426 .

Maximum marks: 1

6 Sqrt violations

The model was checked to see if the assumptions of the model were reasonable

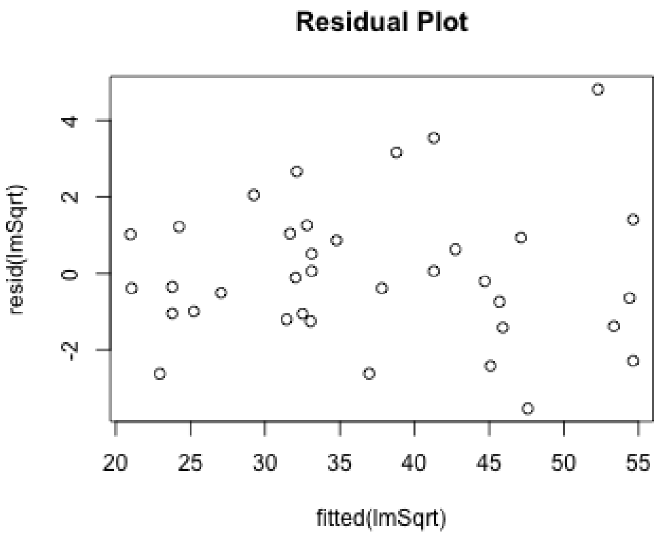

From studying this residual plot, which assumptions do you think are violated?
Select one or more alternatives:

- ☐ It looks OK
- ☐ Linearity
- ☐ Outliers
- ☐ Normality
- ☒ Homoscedasticity (i.e. equal variance)

Maximum marks: 2

Next, a log transformation of hardness (the response) was tried. This gave the following summary

```
##
## Call:
## lm(formula = log(hardness) ~ density, data = janka)
##
## Residuals:
##    Min     1Q Median     3Q    Max
## -0.336 -0.087  0.002  0.085  0.233
##
## Coefficients:
##             Estimate Std. Error t value Pr(>|t|)
## (Intercept)  5.20858    0.08050      65   <2e-16 ***
## density      0.00263    0.00011      25   <2e-16 ***
## ---
## Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
##
## Residual standard error: 0.14 on 34 degrees of freedom
## Multiple R-squared:  0.95,   Adjusted R-squared:  0.95
## F-statistic: 6.2e+02 on 1 and 34 DF,  p-value: <2e-16
```

7 Log transformation Prediction

Based on this transformed model, what is the predicted hardness of a piece of wood with a density of 1000 kg/m^2? Answer to the nearest integer.

2537

(note that a natural log transformation was used)

Maximum marks: 1

8 log violations

The model was checked to see if the assumptions of the model were reasonable

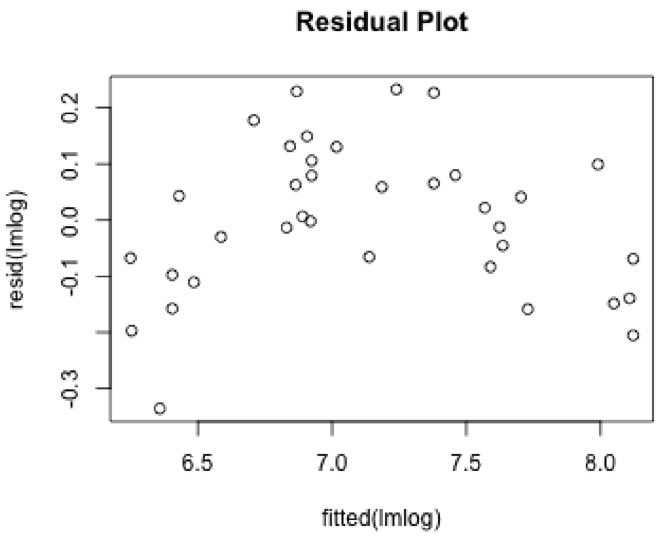

From studying this residual plot, which assumptions do you think are violated?
Select one or more alternatives:

- [ ] Homoscedasticity (i.e. equal variance)
- [ ] Normality
- [ ] It looks OK
- [ ] Outliers
- [x] Linearity

Maximum marks: 2

At the suggestion of another statistician, a generalised linear model was used. This assumed a Gamma distribution, and had a square root link function.

We do not need the details of the gamma distribution: it has to be positive, and the variance increases with the mean. The inverse of the square root link function is the square.

Fitting this model gave the following summary:

```
lmGamma <- glm(hardness ~ density, family=Gamma("sqrt"), data=janka)
print(summary(lmGamma), digits=1)
```

```
##
## Call:
## glm(formula = hardness ~ density, family = Gamma("sqrt"), data = janka)
##
## Deviance Residuals:
##    Min      1Q  Median      3Q     Max
##  -0.22   -0.07   -0.02    0.05    0.17
##
## Coefficients:
##             Estimate Std. Error t value Pr(>|t|)
## (Intercept)    1.867      0.904       2     0.05 *
## density        0.048      0.001      34   <2e-16 ***
## ---
## Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
##
## (Dispersion parameter for Gamma family taken to be 0.009716208)
##
##     Null deviance: 11.26788  on 35  degrees of freedom
## Residual deviance:  0.32876  on 34  degrees of freedom
## AIC: 453
##
## Number of Fisher Scoring iterations: 4
```

9 Janka Gamma Model Prediction

Based on this model, what is the predicted hardness of a piece of wood with a density of 1000 kg/m^2? Answer to the nearest integer

2487

Maximum marks: 1

10 Janka Gamma violations

The simple model was checked to see if the assumptions of the model were reasonable

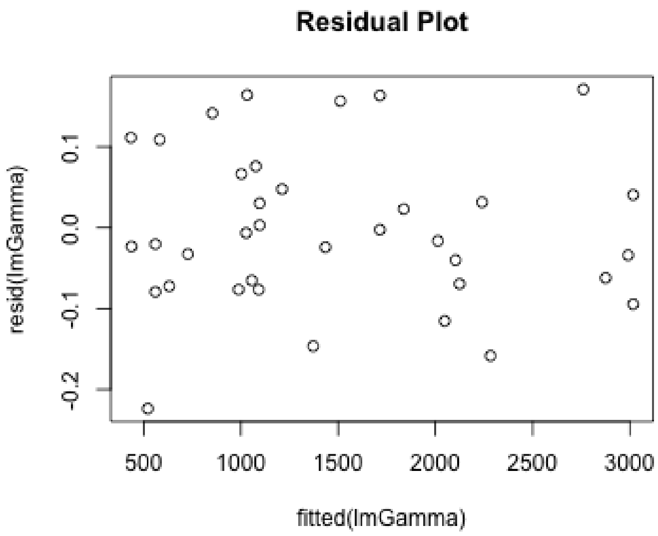

From studying this residual plot, which assumptions do you think are violated?
Select one or more alternatives:

- [ ] Outliers
- [ ] Heteroscedasticity
- [x] It looks OK
- [ ] Normality
- [ ] Linearity

Maximum marks: 2

Now we have tried several models, we want to compare them, and decide how well they describe the data.

These are the estimated coefficients:

Model	Intercept	Slope
Simple	-1160.5	3.6
Square Root transformation	2.26	0.047
Log transformation	5.21	0.0026
Gamma	1.87	0.048

11 Why not R2?

Why can we not compare the R^2 for this model to the value for the untransformed model?

Fill in your answer here

Format B I U x_2 x^2 I_x Ω Σ

("this model" should have been "these models")

The different models all have different transformations. Comparing the Gamma model to the untransformed model is actually possible, but (a) you need to dig in and calculate it yourself, (b) the Gamma model assumes that the variance is heteroscedastic, which makes the interpretation a pain, and (c) the untransformed model optimises in the fitting the R^2, so it will be larger by definition. (I would not expect anyone to give point (c): it relies on quite a more technical understanding of the statistical methods).

Comparing to the other models, we can't directly compare the R^2 from the untransformed model to the R^2 for (say) the log transformed model because they are for different transformations of the data: y vs log(y). We could, of course, transform back to one scale, but that would lead us back to the previous paragraph.

(note: this was intended to be a hard question)

Words: 0

Maximum marks: 4

12 Compare Predictions

How different are the model predictions for a sample of hardness 1000 kg/m^3? (i.e. compare the predictions you have made).

Fill in your answer here

Format B I U x_2 x^2 I_x Ω Σ

The predictions all look fairly similar, ranging from 2426N to 2487N (i.e. difference of 61N). The main difference is that the gamma model is about 50N bigger than the others, which are all very similar.

But how important is this? We can get an idea by looking at the plot of the data. For the hardest wood (i.e. in the top right of the plot of the data) the variation is over about 500N. So the maximum difference is about a tenth of this, which suggests that the variation in predictions is probably not too important, and if there is an important difference it is between the Gamma model and the others.

Words: 0

Maximum marks: 4

13 Compare Intercepts

How similar/different are the intercepts? (note: you might have to transform them). How relevant are these for comparing the models, and how well they describe the data?

Fill in your answer here

Format **B** *I* U x_2 x^2 I_x Ω Σ

The intercept on the scale of the data are

Simple	**-1160.5**
Square root	**5.1**
log	**183.1**
gamma	**3.5**

So the square root and gamma intercepts are similar (probably because the gamma also use a square root transformation). The simple regression has an intercept that is very negative.

These are not very relevant: the data vary between about 400 and 1100, so are a long way from zero, where the intercept is. Wood of density zero is a vacuum, which would be useless as a floor.

Of course, we still need the intercept in the model, but the underlying idea is that this is a local approximation: as long as the model works over the range of the data, we should be happy.

Words: 0

Maximum marks: 6

14 Assessment of the Models

On the basis of these analyses explain which model you would prefer, and why. In particular:

- which model looks best?
- how much difference do the different models make?
- are there any other analyses or calculations you would like to have seen?

Fill in your answer here

Format | **B** *I* U x_2 x^2 | I_x | Ω | Σ

Based on the residual plots, the Gamma model looks the best: the residuals look nice. It is a more difficult model to understand, though.

Based on the predictions for high density, the different models do not seem that different. The intercepts look more different, but that might not be important.

Thus, one thing we would like to see is predictions for other densities (but would you want to be asked to calculate them as well in an exam?). It may be that the models are similar at high density, but diverge at low densities.

The standard errors and confidence intervals of the predictions would also be useful to see. A lot of the issues surround heteroscedasticity, so we should see this in the predictions. For example, the simple analysis probably over-estimates the uncertainty at low densities, snd under-estimates it at high density.

It would also be useful to look at probability plots, e.g. to have another check to see if the data point that looks like it might be an outlier actually is one.

Words: 0

Maximum marks: 10

You have been hired as a consultant by Project Ekorn to assess their performance. Project Ekorn have been trying to stop a mysterious global group of squirrels called Cyber Squirrel Operations from taking out electric power systems. You have been asked to look at how successful these counter-operations have been. Obviously there are other reasons for power systems to go down (birds, storms, octopus etc.), so you want to look to see if the proportion of attacks by squirrels has changed, and also if this is effective in reducing the numbers of people affected.

You have collected data from around the world on the following variables:
- whether the outage was caused by a squirrel
- the year of the event
- the region (some countries have been combined into a larger region)
- the number of people affected, and
- the duration of the event

We expect the country to have an effect (as some countries do not have many squirrels), but are interested in the other factors.

First we will look at the binary response of whether an attack was caused by a squirrel or not.

(note: this is based on real data. The explanation has been exaggerated a lot)

15 Choose Squirrel Models

If we want to test whether the proportion of squirrel attacks has changed over time (=Year), what models would we compare?

Null model

Select one alternative

- ○ Year + Country
- X Country
- ○ Year*Country
- ○ Year

Alternative Model

Select one alternative

- ○ Country
- ○ Year*Country
- X Year + Country
- ○ Year

Maximum marks: 2

16 Squirrel Model Comparison type

What approach is being taken to the model selection in this problem (i.e. asking if the proportion of squirrel attacks has changed over time)?

Select one alternative:

- Exploratory, to get a model that predicts the data well
- X Comfirmatory, to test a hypothesis
- Exploratory, to get a model that explains the data well

Maximum marks: 1

17 Test Results

When the relevant test was done, we got the following result:

```
## Model 1: REDACTED
## Model 2: REDACTED
##   Resid. Df Resid. Dev Df Deviance  Pr(>Chi)
## 1      2524     3134.3
## 2      2523     3104.7  1   29.601 5.307e-08 ***
```

Does this suggest an effect of year?

Select one alternative

- [ ] Can't tell
- [x] Yes
- [ ] No

What statistics tell you this?

Select one or more alternatives

- [ ] Resid. Df
- [ ] Resid. Dev
- [ ] Df
- [ ] Deviance
- [x] Pr(>Chi)

Why is this not enough information to assess the effect of year, and what more would you want to be told?

This tells us that we are unlikely to see this data (or something more extreme) if the null hypothesis is true, but it does not tell us what direction the effect does in, or how strong it is. For that we would need to see the parameter estimates, as well as an estimate of the uncertainty (the standard error or a confidence interval).

Maximum marks: 2

18 Need More Information

When the relevant test was done, we got the following result:

```
## Model 1: REDACTED
## Model 2: REDACTED
##   Resid. Df Resid. Dev Df Deviance  Pr(>Chi)
## 1      2524     3134.3
## 2      2523     3104.7  1   29.601 5.307e-08 ***
```

Why is this not enough information to assess the effect of year on the probability that an attack was by a squirrel, and what more would you want to be told?

Fill in your answer here

Format B I U x₂ x² Ix

Σ

This tells us that it is very unlikely to get the data, or something more extreme, if the null hypothesis was true. So we could conclude that there is an effect. But this does not tell us what direction the effects goes in (does the probability increase of decrease over time?) or how big the effect is. We have a lot of data (>2500 df), so even a very small effect could be statistically significant.

Thus, we also need to see the model coefficients (& their standard errors/confidence intervals).

Words: 0

Maximum marks: 4

In order to try to get a good model, the following model was fitted:

```
sq1 <- glm(Squirrel ~ Region + YearS + lnAffected + lnDuration,
           data=Squirrels, family = binomial())
```

where YearS is Year - 2000 (so the intercept is at 2000 AD), and lnAffected and lnDuration are the natural logs of the number of people affected and the length of time the power was out (in hours). The reference level is the USA, so the other country effects are contrasts to that. The response, Squirrel, is a binary response: it is 1 if the attack was by a squirrel, 0 if it was not.

This gave the following summary:

```
##
## Call:
## glm(formula = Squirrel ~ Region + YearS + lnAffected + lnDuration,
##     family = binomial(), data = Squirrels)
##
## Deviance Residuals:
##      Min        1Q    Median        3Q       Max
## -2.29739  -1.19724  -0.00052   1.09060   2.60735
##
## Coefficients:
##                      Estimate Std. Error z value Pr(>|z|)
## (Intercept)          1.501505   0.366771   4.094 4.24e-05 ***
## RegionAsia          -4.017684   1.015025  -3.958 7.55e-05 ***
## RegionCanada        -0.941575   0.344321  -2.735 0.006246 **
## RegionEurope        -1.760425   0.511984  -3.438 0.000585 ***
## RegionSouth America -15.785693 589.340327  -0.027 0.978631
## RegionUK            -1.439707   0.700209  -2.056 0.039772 *
## YearS               -0.086268   0.024656  -3.499 0.000467 ***
## lnAffected          -0.220496   0.077955  -2.829 0.004677 **
## lnDuration           0.009776   0.078283   0.125 0.900613
## ---
## Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
##
## (Dispersion parameter for binomial family taken to be 1)
##
##     Null deviance: 1089.62  on 785  degrees of freedom
## Residual deviance:  973.11  on 777  degrees of freedom
##   (1791 observations deleted due to missingness)
## AIC: 991.11
##
## Number of Fisher Scoring iterations: 14
```

19 link function

Which link function was used here?

Select one alternative:

- identity
- ☒ logit
- cloglog
- square root
- log
- probit

Maximum marks: 1

20 Why year - 2000?

Why was (Year - 2000) used in the model, rather than Year? How does it affect the other parameter estimates?

Fill in your answer here

Format B I U x_2 x^2 I_x Ω Σ

Year - 2000 was used because the analysis was done by someone with a small brain, and this makes the parameters easier to interpret. The intercept becomes the value at 2000 rather than 1 BC (as 0 AD does not even exist).

The only parameter estimate that is affected is the intercept. Because the model is

eta_i = a + b(y_i-2000) + c x_i

(where a is the intercept, b is the slope of the year effect, y is year, and x_i is another covariate, with effect c), the model can be written as

eta_i = (a - 2000b) + b y_i + c x_i

So the intercept becomes (a - 2000b). When we look at the estimates, we see that this is 174, rather than 1.49. Whilst the model is the same, a value of 1.49 looks more sane.

Words: 0

Maximum marks: 6

21 Squirrel Effects Interpretation

Based on this summary, has there been a change in the proportion of squirrel attacks, and if so how much of an effect? Does this mean that actions against squirrels have been effective?

Are there any other interesting effects? (ignore South America for this question: it is doing something weird and obscure)

Fill in your answer here

Format B I U x_2 x^2 I_x Ω Σ

Yes, there has been a change in squirrel attacks, decreasing by about 8% every year. This suggests that anti-squirrel actions have been effective (although it might also means that the squirrels have just got bored), but one way of another things have been getting better. Of course it might be that the octopodes are increasing their attacks.

The summary also shows that if more people are affected, it is less likely to be due to a squirrel: an increase of people affected by 10% decreases the probability that it was a squirrel by about 2%.

The regions all have a smaller proportion of squirrel attacks than USA (the intercept), if we ignore South America. For example, Canada has about the same effect as waiting 12 years in the USA

(Note: South America's estimate is weird because the best estimate is -infinity, i.e. no squirrel attacks, so R makes the number as small as possible. The standard error calculation is then also utterly wrong, due to something called the Hauck-Donner effect. If you ever want to be left alone at a party, the Hauck-Donner effect is a great conversation topic)

Words: 0

Maximum marks: 10

In order to try to get a good model, the following model was fitted:

```
sq1 <- glm(Squirrel ~ Region + YearS + lnAffected + lnDuration,
           data=Squirrels, family = binomial())
```

where YearS is Year - 2000 (so the intercept is at 2000 AD), and lnAffected and lnDuration are the natural logs of the number of people affected and the length of time the power was out (in hours). The reference level is the USA, so the other country effects are contrasts to that.

This gave the following summary:

```
##
## Call:
## glm(formula = Squirrel ~ Region + YearS + lnAffected + lnDuration,
##     family = binomial(), data = Squirrels)
##
## Deviance Residuals:
##    Min      1Q  Median      3Q     Max
##     -2      -1       0       1       3
##
## Coefficients:
##                     Estimate Std. Error z value Pr(>|z|)
## (Intercept)             1.50       0.37     4.1    4e-05 ***
## RegionAsia             -4.02       1.01    -4.0    8e-05 ***
## RegionCanada           -0.94       0.34    -2.7    0.006 **
## RegionEurope           -1.76       0.51    -3.4    6e-04 ***
## RegionSouth America   -15.79     589.34     0.0    0.979
## RegionUK               -1.44       0.70    -2.1    0.040 *
## YearS                  -0.09       0.03    -3.5    5e-04 ***
## lnAffected             -0.22       0.08    -2.8    0.005 **
## lnDuration              0.01       0.08     0.1    0.901
## ---
## Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
##
## (Dispersion parameter for binomial family taken to be 1)
##
##     Null deviance: 1089.62  on 785  degrees of freedom
## Residual deviance:  973.11  on 777  degrees of freedom
##   (1791 observations deleted due to missingness)
## AIC: 991.1
##
## Number of Fisher Scoring iterations: 14
```

22 Europe Prediction

What is the probability that a power outage was caused by a squirrel under the following conditions?

- **Region**: Europe
- **Year**: 2020 (note that we use 2000 as the intercept, i.e. YearS = year - 2000)
- **lnAffected**: 1
- **lnDuration**: 0

The log odds (i.e. the value on the linear scale):

-2.28

The probability: **0.09**

Maximum marks: 2

As well as trying to find out if the proportion of squirrel attacks have changed, you have also been asked to look at the effects of attacks by different groups, in particular on the numbers that are affected (if squirrels only affect a few people, but birds affect many more, then perhaps we should be more worried bird effects).

So, we looked at several factors to try to understand what influenced the numbers of people affected by each attack. Several models were tried, and their fits to the data summarised in the following table:

Model	AIC	BIC	R^2 (%)
lnDuration	113.0	118.9	1.0
YearS + lnDuration	92.9	104.6	1.9
Countries + lnDuration	73.2	114.2	3.0
Countries + YearS + lnDuration	49.4	96.3	4.0
Operative + lnDuration	61.0	166.4	4.3
Operative + YearS + lnDuration	35.4	146.7	5.3
Operative + Countries + lnDuration	15.0	155.5	6.4
Operative + Countries + YearS + lnDuration	-11.2	135.2	7.4

23 Which Statistic

Which statistic would you use to decide on the best model to explain and understand what is influencing the number of people affected?

Select one alternative:

- Something else
- ☒ BIC
- R^2
- AIC

Maximum marks: 2

24 Which model is best?

Which model do you think is best, according to the criterion you chose?

Select one alternative:

- X Countries + YearS + lnDuration
- lnDuration
- Operative + lnDuration
- Countries + lnDuration
- YearS + lnDuration
- Operative + Countries + YearS + lnDuration
- Operative + YearS + lnDuration
- Operative + Countries + lnDuration

Maximum marks: 2

This is the summary of the full model (which may not be the best one). The reference levels are squirrels (for Operative) and USA (for Region).

```
##
## Call:
## lm(formula = lnAffected ~ Operative + Region + YearS + lnDuration,
##     data = Squirrels)
##
## Residuals:
##    Min     1Q Median     3Q    Max
##   -4.2   -0.3    0.2    0.6    2.3
##
## Coefficients:
##                           Estimate Std. Error t value Pr(>|t|)
## (Intercept)                   0.28       0.14     1.9    0.057 .
## OperativeAnimal              -0.34       0.28    -1.2    0.219
## OperativeBat                 -0.26       0.74    -0.4    0.724
## OperativeBeaver               0.43       0.36     1.2    0.230
## OperativeBird                 0.14       0.09     1.6    0.116
## OperativeEagle                0.64       0.59     1.1    0.283
## OperativeMarten               1.49       0.52     2.9    0.004 **
## OperativeMonkey               0.92       1.03     0.9    0.371
## OperativeOther Mammal         0.10       0.31     0.3    0.735
## OperativeOther Vertebrate     0.48       0.17     2.8    0.006 **
## OperativePossum              -0.34       0.40    -0.8    0.402
## OperativeRaccoon              0.41       0.16     2.5    0.013 *
## OperativeRat                 -0.11       0.37    -0.3    0.767
## OperativeUnknown              0.24       0.16     1.5    0.131
## RegionAsia                    0.42       0.18     2.3    0.020 *
## RegionCanada                  0.26       0.16     1.6    0.109
## RegionEurope                 -0.23       0.24    -1.0    0.322
## RegionSouth America           1.02       0.44     2.3    0.019 *
## RegionUK                     -0.58       0.31    -1.8    0.067 .
## YearS                        -0.03       0.01    -2.9    0.004 **
## lnDuration                    0.10       0.04     2.6    0.009 **
## ---
## Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
##
## Residual standard error: 1 on 765 degrees of freedom
##   (1791 observations deleted due to missingness)
## Multiple R-squared:  0.07,   Adjusted R-squared:  0.05
## F-statistic:  3 on 20 and 765 DF,  p-value: 8e-06
```

25 Affected Assessment

Based on these results, what are the important effects on the variation in the number of people being affected? Consider both the model you chose, the parameter estimates and other summaries of the model.

Fill in your answer here

Format | B I U x_2 x^2 | I_x | | Ω |

First up, the R^2 is 7%, so this model does not explain a lot, and perhaps mostly demonstrates the desperation of the exam writer to find some data.

Beyond that, there are effects of year and (log)duration, and there is variation in the operatives and regions. The intercept levels are squirrels and the USA, so (for example) martens affect e^1.49 > 4 times as many people as squirrels do: but the confidence interval is wide from about 1.6 to 12 times more (exp(1.49-1.96*0.52) to exp(1.49+1.96*0.52)).

To get some feel for the effects of year and duration, we can note that the smallest region effect is -0.23, i.e. a reduction of about 20%. So it would take about 0.23/0.03 = 7.7 years to have the same effect. For duration, it would take an increase in duration of exp(0.23/0.10) = 10 times to have the same effect. So these seem of the same scale (roughly!).

Words: 0

Maximum marks: 10