

R Exercises

1. **Regression Analysis:** The data frame `States` in the `car` package contains education data pertaining to the 50 states of the United States plus Washington D.C. These data are from Moore, *The Basic Practice of Statistics, Second Edition* (Freeman, 2000). The following variables are included:

<code>region</code>	U.S. Census regions: ENC, etc. (See the help file for the data frame.)
<code>pop</code>	Population, in 1000s.
<code>SATV</code>	Average score of graduating high-school students in the state on the <i>math</i> component of the Scholastic Aptitude Test (a standard university admission exam).
<code>SATM</code>	Average score in the state on the <i>verbal</i> component of the SAT.
<code>percent</code>	Percentage of graduating high-school students in the state who took the SAT exam.
<code>dollars</code>	State expenditures on public education, in \$1000s per student.
<code>pay</code>	Average teacher's salary in the state, in \$1000s.

- a. Regress `SATM` *separately* on each of `percent` and `dollars`.
 - b. Regress `SATM` on *both* `percent` and `dollars`. Do the regressions in parts (a) and (b) differ importantly? If so, why?
2. **Dummy-Variable Regression:** The data given in the data frame `Burt` in the `car` package, on the IQs of 27 pairs of identical twins reared apart, were reported by Sir Cyril Burt (1966). (These “data” are wholly fraudulent.) One twin in each pair was raised by his or her biological parents; the other twin was raised in a foster home. In each case, Burt recorded (i.e., made up) the “social class” to which the twins’ biological parents belonged.
- a. Explore the data graphically by plotting `IQbio` (as the response variable) against `IQfoster`, using a different symbol and plotting a separate linear regression line for each social class.
 - b. Then regress the IQ of the twins reared by their biological parents (`IQbio`) on the IQ of the twins reared by foster parents (`IQfoster`), dummy variables to represent the three social classes (`class`), and regressors for the interaction between foster-twin IQ and social class. *Suggestion:* You may want to re-order the categories of the factor `class` so that they are in their natural order rather than in the (default) alphabetic order:

```
class <- factor(class, levels=c("low", "medium", "high"))
```
 - c. Test the interaction between foster-twin IQ and social class. If the interaction proves to be non-significant, test the partial effects of foster-twin IQ and social class on biological-twin IQ. Compute the appropriate incremental F-tests using the `Anova` function in the `car` package. This test should make you suspicious about the provenance of the data. Why?

3. **Logistic Regression:** Employing a sample of 1643 men between the ages of 20 and 24 from the U.S. National Longitudinal Survey of Youth, Powers and Xie (2000) investigate the relationship between high-school graduation and parents' education, race, family income, number of siblings, family structure, and a test of academic ability. The data set, in the file `Powers.txt`, contains the following variables:

<code>hsgrad</code>	Whether the respondent was graduated from high school by 1985 (Yes or No)
<code>nonwhite</code>	Whether the respondent is black or Hispanic (Yes or No)
<code>mhs</code>	Whether the respondent's mother is a high-school graduate (Yes or No)
<code>fhs</code>	Whether the respondent's father is a high-school graduate (Yes or No)
<code>income</code>	Family income in 1979 (in \$1000s) adjusted for family size
<code>asvab</code>	Standardized score on the Armed Services Vocational Aptitude Battery ¹
<code>nsibs</code>	Number of siblings
<code>intact</code>	Whether the respondent lived with both biological parents at age 14 (Yes or No)

Following Powers and Xie perform a logistic regression of `hsgrad` on the other variables in the data set. The logistic regression in the previous problem assumes that the partial relationship between the log-odds of high-school graduation and number of siblings is linear. Test for nonlinearity by fitting a model that treats `nsibs` as a factor (just use `factor(nsibs)` in place of `nsibs` in the model), performing an appropriate likelihood-ratio test. In the course of working this problem, you should discover two errors in the data. Deal with the errors in a reasonable manner. Does the result of the test change?

4. **Poisson Regression:** Long (1990, 1997) investigates factors affecting the research productivity of doctoral students in biochemistry. The response variable in this investigation, `art`, is the number of articles published by the student during the last three years of his or her PhD programme. The explanatory variables are as follows:

<code>fem</code>	dummy variable: 1 if female, 0 if male
<code>mar</code>	dummy variable: 1 if married, 0 if not
<code>kid5</code>	number of children five years old or younger
<code>phd</code>	prestige of PhD department
<code>ment</code>	number of articles published by mentor during last three years

Long's data (on 915 biochemists) are in the file `Long.txt`. The variable names listed above are those employed by Long, and appear in the first row of the data file.

- Examine the distribution of the response variable, `art`. Based on this distribution, does it appear promising to model these data by linear least-squares regression, perhaps after transforming the response?
- Following Long, perform a Poisson regression of `art` on the explanatory variables.
- Refit Long's model allowing for overdispersion (i.e., using the `quasipoisson` family in R). Does this make a difference to the results?

¹ Apparently scores on this test were standardized to a mean of 0 and standard deviation of 1 for the NLSY sample as a whole; as you can verify, the mean and standard deviation in this subsample differ somewhat from 0 and 1, respectively