

Quant I - Problem Set IV

Hypothesis Testing and Goodness of Fit

Kenneth Benoit

Assigned: Thursday, November 5th, 2009

Due: Thursday, November 19th, 2009

1. Consider the case of a well-known athlete who is tested positive for using a certain forbidden drug. In 6 out of a 100 cases, the test fails to detect the use of the drug. In 8 out of a 100 cases the test indicates the use of the drug although the athlete is “clean”. Insider knowledge suggests that 5% of athletes use that particular prohibited drug. Assuming that the athlete tested positive was chosen at random, what is the conditional probability that she/he indeed uses the drug? (Please do not just provide a number, but show how you arrive at the solution.)

2. Verzani p. 236 Q 8.27

(Note: Against which alternative hypothesis should you test the null?)

Additional question: Briefly state what it means that the differences are (not) significant at the .01 level. Include a precise interpretation of the p -value.

3. Verzani p. 247 Q 8.30

Check the help file referring to the dataset and set numeric values that don't represent actual age to NA.

Important: Please do not use the `t.test` command in this task, but calculate this “manually”, step-by-step, with the help of the respective formulae.

Hints: You can assume that the variances are equal. You will probably need to specify the `na.rm=TRUE` option to make R ignore the missing values in some instances. You have to take into account the fact that there are cases with missing values when calculating the number of observations in the groups.

Extra question: Briefly interpret the result and include a precise interpretation of the p -value.

4. The dataset `ines2002selection.RData` contains a selection of variables from the Irish National Election Study 2002 (in form of a dataframe called `d`). In this question, we will compare thermometer scores across gender groups.

We are interested in the popularity of Gerry Adams, the Sinn Fein leader. The variable `adamsth` is a so-called thermometer score, where respondents indicated how they feel about Gerry Adams. The scale ranges from 0 to 100, where increasing values reflect higher ratings.

- (a) Inspect the distribution of the thermometer scores by gender. Use summary statistics and an appropriate plot.
 - (b) Why is it better not use a t-test in this case?
 - (c) Use a Wilcoxon rank-sum-test to test whether the center of the distribution of the variable in the male group is the same as the center in the female group. Briefly interpret, include a precise interpretation of the p -value.
5. Chi-square test. Use the same INES 2002 dataset for this question. The variable `nounit` represents attitudes towards the aim of Irish unification. The scale ranges from 0 (Insist on a United Ireland now) to 10 (Abandon the aim of a United Ireland altogether).
 - (a) The variable `gender` is a factor. Remove the level `Dontknow` (see Verzani p. 137 on how to do this).
 - (b) Cross-tabulate `gender` and the `United Ireland`-variable. Perform a standard chi-square test, step-by-step in R according to the formula. One way to do this is to make a table of the observed frequencies as well as create an empty table that is then “filled” with the expected values.
 - (c) Verify your result with `chisq.test` and briefly interpret the result. Include a precise interpretation of the p -value.
 - (d) Try simulating the p -value to test whether the parametric version appears to be accurate. What might account for the differences between the simulated and the parametrically computed p -values?