

1. In the 1996 General Social Survey, for males age 30 and over, the following was true about respondents:
 - 11% of those in the lowest income quantile were college graduates.
 - 19% of those in the second income quantile were college graduates.
 - 31% of those in the third income quantile were college graduates.
 - 53% of those in the highest income quantile were college graduates.

In class we found $P(Q_1|G)$, the probability that a randomly selected college graduate falls in the lowest income quartile. Also find $P(Q_2|G)$, $P(Q_3|G)$, and $P(Q_4|G)$. Discuss how this distribution compares to the unconditional distribution $\{P(Q_1), P(Q_2), P(Q_3), P(Q_4)\}$.

2. Do Exercise 2.2 in the Chapter 2 Exercises of the textbook at bayesrulesbook.com.
3. Do Exercise 2.8 in the Chapter 2 Exercises of the textbook at bayesrulesbook.com.
4. Do Exercise 2.9 in the Chapter 2 Exercises of the textbook at bayesrulesbook.com.
5. **Required for Graduate Students; Extra Credit for Undergrads:** Do Exercise 2.11 in the Chapter 2 Exercises of the textbook at bayesrulesbook.com.
6. **Required for Graduate Students; Extra Credit for Undergrads:** Do Exercise 2.13 in the Chapter 2 Exercises of the textbook at bayesrulesbook.com.
7. Frank and Jerry each suspect that a particular coin is biased and that the probability θ of “heads” is actually greater than 0.5. They propose two different experiments to test this belief statistically. Frank will toss the coin 10 times and record as his test statistic the number of heads obtained in the 10 trials. Jerry will toss the coin repeatedly until 2 tails have been obtained in total. At this point Jerry will use as his test statistic the number of heads he has obtained in his series of tosses.
 - (a) It turns out that of Frank’s 10 tosses, 8 were heads. Either derive on paper the P-value for Frank’s test, or argue carefully that the R code
`1 - pbinom(7, 10, 0.5, lower.tail=T)`
produces the correct P-value for Frank’s test. You can type `help(pbinom)` in R for details about this R function.
 - (b) It turns out that at the point that Jerry gets his second tail, he has accumulated 8 heads. Either derive on paper the P-value for Jerry’s test, or argue carefully that the R code
`1 - pnbinom(7, 2, 0.5, lower.tail=T)`
produces the correct P-value for Jerry’s test. You can type `help(pnbinom)` in R for details about this R function.

(c) Would the conclusions of Frank's and Jerry's significance tests agree (assuming a 0.05 significance level in each test)? Explain.

(d) Bayesian Betty decides to assume a (noninformative) Uniform(0, 1) prior for θ of

$$p(\theta) = 1, \quad 0 \leq \theta \leq 1.$$

Show that Bayesian inference based on the posterior distribution for θ will be exactly the same whether, as her observed data, she uses the results of Frank's experiment or uses the results of Jerry's experiment. Discuss how this relates to the Likelihood Principle.

8. Do Exercise 3.9 in the Chapter 3 Exercises of the textbook at bayesrulesbook.com.
9. Do Exercise 3.11 in the Chapter 3 Exercises of the textbook at bayesrulesbook.com.
10. Do Exercise 3.18 in the Chapter 3 Exercises of the textbook at bayesrulesbook.com.