STAT 7630: Homework 2
(Due: Thursday, 09/17/2024)

Note: Show all your work for the necessary steps to receive full credit.

Please turn in the HW on paper, hand-written and/or typed. For computational problems, return only the relevant parts of the output with comments/annotations. Questions taken from the textbook are marked with BR for "Bayes Rules!". From the code you are using to answer the problems, turn in the relevant output, and the figures (if requested), preferably printed from the output. No need to turn in your code or long lists of generated samples.

BR Chapter 3:
Do Exercises 3.9, 3.11, and 3.18.

BR Chapter 4:
Do Exercises 4.4, 4.10, 4.16, and 4.19.
Hint for 4.19: The code "`bechdel %>% filter(year==1980)`" will pick out the movies in the data set from the year 1980.

Additional Questions (AQ):
AQ1. 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 = TRUE)
```

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 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 = TRUE)
```

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