Exercise B: Priors and Updating

Benedikt Ehinger

3. November 2016

Videos C+D+E: Bayesian Parameter Estimation
Recommended Reading Material:

- Doing Bayesian Data Analysis Chapter 5 + 6

For this exercise we will make use of R. If you don't know R, don't worry! I recommend to use RStudio, so google it, download, install and run it. Then you simply have to run the file (ßource") it (you can also read DBDA Chapter 3 for a short R-Tutorial). R is very similar to python & matlab in syntax and functioning. After sourcing the function once (to put them into your environment) you can interact in the console with the program, as you can in python and matlab.

1 Priors for Bernoulli Distributions

1.1 Exercise 1

[1] Open the program *BernGridExample.R*. You will in addition need *BernGrid.R* and *DBDA2E-utilities_exB.R* from StudIP.

You will notice there are several examples of using the function *BernGrid*. Run the script. For each example explain what the example illustrates. Try to find real situation of when the specific prior could be appropriate.

Hints: Look at Figures 5.2, 5.3 and 6.5 DBDA. Two of the examples involve a single flip, with the only difference between the examples being whether the prior is uniform or contains only two extreme options. The point of those two examples is to show that a single datum implies little when the prior is vague, but a single datum can have strong implications when the prior allows only two very different possibilities.

[1] Source: Doing Bayesian Data Analysis 5.4

2 Iterative Updating

2.1 Exercise 2

For this exercise, use the R function *BernBeta.R*. (Don't forget to source the function before calling it.) Notice that the function returns the posterior beta values each time it is called, so you can use the returned values as the prior values for the next function call.

1. Start with a prior distribution that expresses some uncertainty that a coin is fair: $beta(\theta|4,4)$. Flip the coin once; suppose we get a head. What is the posterior distribution?

2. Use the posterior from the previous flip as the prior for the next flip. Suppose we flip again and get a head. Now what is the new posterior? Hint: If you type

   ```
   post = BernBeta( c(4,4) , c(1) )
   ```

   for the first part, then you can type

   ```
   post =BernBeta( post , c(1) )
   ```

   for the next part.)

3. Using that posterior as the prior for the next flip, flip a third time and get a tail. Now what is the new posterior? Hint:

   ```
   post = BernBeta( post , c(0) )
   ```

4. Do the same three updates but in the order T, H, H instead of H, H, T. Is the final posterior distribution the same for both orderings of the flip results?

2.2 Exercise 3

Suppose an election is approaching, and you are interested in knowing whether the general population prefers candidate A or candidate B. There is a just published poll in the newspaper, which states that of 100 randomly sampled people, 58 preferred candidate A and the remainder preferred candidate B.

1. Suppose that before the newspaper poll, your prior belief was a uniform distribution. What is the 95% HDI on your beliefs after learning of the newspaper poll results?

2. You want to conduct a follow-up poll to narrow down your estimate of the population's preference. In your follow-up poll, you randomly sample 100 other people and find that 57 prefer candidate A and the remainder prefer candidate B. Assuming that peoples' opinions have not changed between polls, what is the 95% HDI on the posterior?