

Course EPIB-683 - Intermediate Bayesian Analysis for the Health Sciences

Assignment 2

1. In the first two questions we will consider a generic clinical trial setup, and calculate sample sizes using first standard frequentist and then Bayesian methods. In the third question, we will compare the information received from each.

The trial setup is as follows: Two systolic blood pressure lowering drugs, A and B, are being compared. Thus, the main outcome is the amount by which blood pressure is lowered by each drug on average, which is assumed normally distributed. The prior information is such that Drug A is thought to lower blood pressure by 10 mm Hg on average, with a prior sample size equivalent of 100 subjects, while Drug B is thought to lower blood pressure by an average of 12 mm Hg, with a prior sample size equivalent of 50 subjects. The standard deviation of this blood pressure lowering is thought to be the same in both groups, with prior 95% credible interval from 10 mm Hg to 20 mm Hg (so that the range of the variance runs from 100 to 400, and the precisions then run from 1/400 to 1/100, in other words, the precision interval is (0.0025, 0.01). It is desired to estimate the true difference between the amount of blood pressure lowering in A compared to B to an accuracy of ± 1 mm Hg (total interval width 2) with a 95% interval.

1. Use standard frequentist methods to calculate the sample size required for each group. Assume equal sized groups, and use the midpoints of the prior intervals given above as point estimates. You can plug your numbers into the sample size formula below

$$N_{per\ group} = \frac{4z_{1-\alpha/2}^2(\sigma_1^2 + \sigma_2^2)}{w^2}$$

where σ_1 and σ_2 are the estimates of the standard deviations in groups 1 and 2, respectively, $z_{1-\alpha/2}$ is the usual normal percentile (1.96 for a 95% interval), and w is the total length of the desired interval.

You can check your answer using the frequentist sample size calculator found [here](#) (very easy to download and install, and trivially easy to use):

<http://www.medicine.mcgill.ca/epidemiology/Joseph/software/SampleSizeCalculator/samplesize.zip>

2. We will now redo the sample size calculations using a variety of Bayesian sample size criteria.

(a) The first step is to convert the prior information on σ_1 and σ_2 given above into gamma prior distributions. While one can use a normal approximation to the gamma density and fit a gamma density by matching gamma parameters to the desired mean and standard deviation, there is a program on the course web page that takes the desired 95% interval as input, and provides the exact required gamma parameters. Note that it is the reciprocal of the variance, or the precision, that follows a gamma, so some conversion is needed for each endpoint. Use this program to calculate the required gamma prior inputs for both σ_1 and σ_2 (after converting the SD interval endpoints to precision endpoints as required).

(b) Install and load the R package called `SampleSizeMeans`, and use it to fill in the following table of sample sizes (assume equal sized groups, check for “equalvar” in the function name), first row is done for you as an example, see code below):

Bayesian sample size Criterion	Sample Size
Average Coverage Criterion (ACC)	1509
Average Length Criterion (ALC)	
MBL Average Coverage Criterion (MBLACC)	
MBL Average Length Criterion (MBLALC)	
Modified Worst Outcome Criterion (MODWOC(prob = 0.5))	
Modified Worst Outcome Criterion (MODWOC(prob = 0.9))	
MBL Modified Worst Outcome Criterion (MBLMODWOC(prob = 0.5))	
MBL Modified Worst Outcome Criterion (MBLMODWOC(prob = 0.9))	

Hints:

1. To fill in the first line of the table, you need to type something like

```
> mudiff.acc.equalvar(2, 8.475, 1506.17, 100, 50, level = 0.95, equal = TRUE)
[1] 1509 1509
```

so that 1509 subjects are needed in each group here (your result may vary a bit, as a Monte Carlo algorithm is used within the program, but should be very close to the size given above). Note that the program may take a few seconds or even a few minutes to run, depending on the speed of your computer.

2. Refer to the article on the course web page on Bayesian sample size for details about how to interpret each criterion.
3. We will now compare the information available from each of these sample size calculations. How do the sample sizes from the frequentist approach calculated in question 1 compare to the Bayesian approaches calculated in question 2? Take some examples of Bayesian sample sizes that are larger/smaller than the frequentist approach, and explain why they are larger or smaller.
4. In this question we will calculate a Bayesian sample size for a Poisson outcome via simulation. Suppose we would like to estimate the number of arrivals at a busy hospital emergency room in a given day to within an accuracy of ± 1 (total width of 2) arrivals with a 95% credible interval. The 95% prior interval for the Poisson parameter is $\text{gamma}(\text{shape}=144, \text{scale} = 2.083)$, which has mean close to 300 visits per day, with a standard deviation close to 25 visits per day.

We will slowly zero in on the required size by simulating the average length of the 95% interval across results for a succession of sample sizes, each time moving the guess at the sample size higher or lower, depending on the current average length.

- (a) Assuming the prior given above will be used for posterior analyses, what would the posterior density be if Poisson outcomes of (x_1, x_2, \dots, x_n) visits are observed over n independent days? (Hint: Recall that a gamma prior density is conjugate to a Poisson likelihood.)
- (b) Continuing from (a), once the posterior density is known, you can obtain the upper and lower limits of the posterior 95% credible interval using quantiles of the gamma posterior density. Once you have these limits, subtracting the lower limit from the upper limit provides the length of the 95% interval. Create R code that takes as input n Poisson observations and prior gamma density parameters, and outputs the length of the 95% posterior credible interval for the resulting posterior density.

(c) Create an R function that takes as input

1. Gamma prior parameters
2. A sample size, say n (representing the number of days the emergency room visits are counted)

and provides the average length of the posterior credible interval as output. To do this, your function needs to simulate n Poisson random variables, taking each from a new Poisson parameter generated from the gamma prior, finding the length of the gamma posterior credible interval for each data set, and taking the average length of these intervals. For each sample size, calculate the average length over 1000 simulations.

(d) Using the function from (c), find the required sample size by using it successively, slowly zeroing in on the required sample size. In a table, report the path along which you found the optimal sample size, each line of the table giving the sample size attempted, and the average credible interval length for the sample size.

5. Repeat question 4 above, but using a mixed Bayes/likelihood approach. Therefore, the prior information used to predict future data remains the same, but the prior used for purposes of calculating the posterior density for each data set becomes non-informative. For this purpose, you can use a `gamma(shape=0.001, scale = 1000)` density (note that the `rate = 1/scale` parameter is what is used in WinBUGS, so this is equivalent to the familiar `gamma(0.001, 0.001)` prior density often used in WinBUGS).