Reading

- Read sections 11.1.1 and 11.1.3 of BayesNotes.pdf, and do problems 1, 4 and 8 on page 644;
- read sections 11.2.1, 11.2.2 and 11.2.4 and do problems 1 and 4 on pages 663–664;
- read section 11.4 (pages 691–694) and do problems 1 and 2 on page 699; and
- read section 11.5 (pages 699–704) and do problem 1 on page 709.

Inferences

The paper that won the 2024 IgNobel Prize in Probability summarises an experiment resulting in 350,757 coin flips by 48 people using 211 different coins. It suggests that a coin showing heads when it is flipped into the air will land heads up with probability around 0.51, significantly higher than 0.5, and likewise for tails. A summary of the data in `data-agg.csv` gives for each participant and coin the number of heads when the coin started heads up and tails up. The paper itself uses a variety of prior densities on the probability that the outcome is heads given the coin's initial state. Here we shall attempt to estimate an appropriate prior density using an empirical Bayes approach. We shall assume that the coin tosses are mutually independent, given their initial position (heads up/tails up), though the paper finds some evidence of a learning effect.

(a) Show first that if R conditional on θ is binomial with denominator m and probability θ, and θ has a beta prior density with known $a, b > 0$, then the marginal density of R is $\binom{m}{r} B(a+r, b+m-r)/B(a,b)$, where $B(a,b) = \Gamma(a)\Gamma(b)/\Gamma(a+b)$ is the beta function.

(b) Re-express the prior for θ in terms of its mean $p = a/(a+b)$ and $\nu = a+b$, and hence find its posterior mean and variance in terms of a and b.

(c) If $R_j \mid \theta_j$ ($j = 1, \ldots, n$) are independent binomial variables with denominators m_j and $\theta_1, \ldots, \theta_n$ are a random sample from the beta density with parameters a, b, find the joint marginal density of $R_1, \ldots, R_n$ given a, b.

(d) How would the formulation in (c) need to be completed in a Bayesian hierarchical model? How would this differ from an empirical Bayes approach?

(e) Give a log likelihood from which p and ν can be estimated in an empirical Bayes analysis.

(f) Use the data in `data-agg.csv` to estimate p and ν for the data starting heads up and tails up separately. You can read the data into `R` and make a data frame with the code

```
df <- read.csv("data-agg.csv",header=T)
df$person <- factor(df$person); df$coin <- factor(df$coin)
```

and will then need to maximise the log likelihood from (e); use the `lbeta` function.

(g) Find Bayes factors to compare the following models: M0: $a = b$, i.e., the θ_j are drawn from a distribution symmetric around $p = 0.5$; M1: $a \neq b$, i.e., $p \neq 0.5$; M2: different values of p and ν are appropriate when starting heads up and tails up. Which seems most appropriate?

(h) Compute the posterior means and variances of the θ_j for model you choose in (g), and use them to check the model.

(i) Do you think the conclusions from the study are reasonable?