Lecture 22: Bayesian hypothesis testing

- I defined prior odds, posterior odds and Bayes factors.
- I noted:

$$\text{Posterior Odds} = \text{Prior Odds} \times \text{Bayes Factor}$$

- The Bayes factor in favour of a null hypothesis H_o is

$$\frac{f(X|H_0)}{f(X|H_a)}$$

- I worked out the details for a $N(\mu,1)$ example where π_0 is the prior probability of the null $H_o : \mu = 0$ and the conditional prior on μ given $\mu \neq 0$ is $N(0,\tau^2)$.
- I observed that for $X = 4$ for instance a frequentist would reject while if τ were big enough a Bayesian would prefer the null to the alternative.

Informative priors

- I discussed the paper by Bélisle et al (*Can J Statist*, 2002) which has huge Bayes factors.
- In one case comparing a model in which a large number of variances are equal to a model in which they can vary the BF in favour of all equal is enormous.
- Surprising because all equal is a submodel.
- Explained by the fact that the prior on the many variances model is *informative*.
- It specifies iid variances. Large number so mean of these variances is a priori nearly a known constant.
- But fitted values have an average different from that constant.
- So pick submodel.

Coverage in text and notes

- Chapters 11 and 12.
- Course slides on Bayes estimation.