# HW3, Theory of Inference 2016/7

Jonathan Rougier
School of Mathematics
University of Bristol UK

In the following questions, I show marks in square brackets, to give you an idea of the approximate tariff the question would carry in an exam.

1. Prove that LP $\rightarrow$ WIP, and that LP $\rightarrow$ WCP. [10 marks]

   You should use this question to practice writing really clear and compelling proofs.

2. A Traffic Inspector sits outside the Bristol Royal Infirmary, on a road known to have a very high level of pollution. In one hour she records 423 passing cars (excluding vans and lorries), noting for each car the number of empty seats.

   (a) Write down a model for this experiment, bearing in mind that the number of cars seen in an hour ($N$) is itself a random quantity. [10 marks]

   (b) State what is meant by an 'ancillary' random quantity, and discuss whether $N$ is ancillary. [10 marks]

   (c) State the Conditionality Principle (CP), and show that it is implied by the Likelihood Principle (LP). [10 marks]

   (d) Treating $N$ as ancillary, how does the Traffic Inspector analyse her results, if she adopts the CP? [10 marks]

3. This question on stopping rules is from David MacKay's *Information Theory, Inference, and Learning Algorithms* (CUP, 2003), sec. 37.2. It is not the kind of question I would set in an exam, except for the last part. You will need to use R to compute the probabilities.

   In an expensive laboratory, Dr Bloggs tosses a coin twelve times and the result is HHHTHHHHTHHT. He is interested to know whether $\Theta$, the probability of tails (T) is not equal to 0.5.

(a) Dr Bloggs consults a Frequentist statistician, who tells him that he needs to compute a $P$-value, namely $\Pr(R \leq 3; n = 12, \theta = 0.5)$, where $R$ is the number of tails. Give the formula for this $P$-value, and show that it is equal to 0.07.

(b) Dr Bloggs is vexed; the $P$-value is not below the threshold of 0.05, which is what he needs to publish in the prestigious Journal of Experimental Coin Psychology. But when he talks to his statistician he discovers that the statistician failed to account for his (Dr Bloggs's) stopping rule. Dr Bloggs had decided to toss the coin until 3 tails had appeared. In that case, says the statistician, $N$, the sample size, is the random quantity, and we need to compute $\Pr(N \geq 12; r = 3, \theta = 0.5)$. Give the formula for $\Pr(N = n; r, \theta)$, and show that the $P$-value is equal to 0.03.

    Hint: To get $r$ tails on the $n$th toss, we need to get exactly $r - 1$ tails in $n - 1$ tosses, and then a tail on the $n$th toss.

(c) Dr Bloggs is delighted, and writes up his result without delay, with the respectably low $P$-value of 0.03. What does the Stopping Rule Principle (SRP) say about this experiment? What do we infer about the Frequentist practice of using $P$-values for inference? What would a Bayesian statistician do in this situation? [15 marks]