StAN Exercise Sheet 4

Richard D. Gill
Mathematical Institute, University of Leiden, Netherlands
http://www.math.leidenuniv.nl/~gill

25 October, 2012

1 Confidence intervals, bootstrap

1.1 Binomial distribution

Cf. discussion around formula (3.27) of Feigelson and Babu, and cf. English Wikipedia article on Binomial distribution, section on Confidence Intervals.
According to a recent article in "NRC Handelsblad" the Belgian Euro coins are biased. To test this, I had my students spin Belgian 1 Euro coins $n = 250$ times. The "1 Euro" side (let me call that "tails") finished up on top on $X = 140$ of the spins, the outcome was "heads" the other 90 times. I'm interested in finding a 95% confidence interval for the probability p of tails on a single spin. You may assume that $X \sim \mathrm{Bin}(n, p)$.

(1) Compute an approximate 95% confidence interval for p using the approximate normal distribution of the maximum likelihood estimator of p and the observed Fisher information in the data.

(2) Using propagation of error, show that the variance of $\arcsin \sqrt{X/n} \approx 1/4n$. This suggests use of the approximate 95% confidence interval $\sin^2(\arcsin(\sqrt{\widehat{p}}) \pm 1/\sqrt{n})$ for p. Compute this interval for our coin problem, and compare with the result of (1).

(3) Suppose $X \sim \mathrm{Bin}(n, p)$ and we observe $X = x$. Clearly, the chance is less than 2.5% that we observe an outcome below the lower 0.025-quantile of the distribution of X, and it is also less than 2.5% that we observe an outcome above the upper 0.025-quantile of the distribution of X. Thus excepting a 2.5% chance of bad luck, p cannot be so large, that the observed value x of X is such that $\mathrm{P}_p(X \leq x) \leq 0.025$. Excepting a 2.5% chance of bad luck, p can't be so small that $\mathrm{P}_p(X \geq x) \leq 0.025$. These observations lead to the following proposal for a guaranteed at least 95% confidence interval for p: the interval of values from the smallest p such that $\mathrm{P}_p(X \geq x) \geq 0.025$ up to the largest p such that $\mathrm{P}_p(X \leq x) \geq 0.025$. Determine these two values for our coin problem, and compare with

the results of (1) and (2).

(4) Perform a small simulation experiment with `R` to investigate the coverage probability of the three methods for deriving a confidence interval for p, when n is small, e.g., 20, and for a range of values of p.

(5) Can you figure out parametric or non-parametric bootstrap methods for obtaining a confidence interval for p? If you manage to come up with any proposals, please compare them to the earlier proposals, perhaps using a simulation experiment.

1.2 Normal distribution

Cf. discussion around formula (3.25) of Feigelson and Babu.
Consider the problem of finding a confidence interval for the unknown mean of a normal distribution given an i.i.d. sample of size n from such a distribution, mean and variance unknown.

Denote by $\overline{X}$ the average of n observations, and by S the sample standard deviation, defined, as is conventionally done, by $S^2 = \sum(X_i - \overline{X})^2/(n-1)$.

(1) Argue that the probability distribution of $T = \sqrt{n}(\overline{X} - \mu)/S$ depends only on n, not on μ or σ^2, the parameters of the normal distribution in question.
Hint: suppose I subtract a constant from every observation, and divide every observation by another constant; suppose I apply simultaneously the same operations to μ. I now recompute T based on the transformed observations and transformed value of μ. What is the relationship between the old and the new T?

(2) Show that confidence intervals for μ derived from the parametric bootstrap of $\sqrt{n}(\overline{X} - \mu)/S$ are identical to the confidence intervals (3.25), case 3, of Feigelson and Babu.

(3) Confidence intervals for μ could also be derived from the non-parametric bootstrap distribution of $\sqrt{n}(\overline{X} - \mu)$. Supposing that n is large, one would expect that this bootstrap distribution would be approximately Gaussian. What are the exact mean and exact variance of the non-parametric bootstrap distribution of $\sqrt{n}(\overline{X} - \mu)$? Show that using the obvious normal approximation to the non-parametric bootstrap distribution of $\sqrt{n}(\overline{X} - \mu)$ results in the confidence interval (3.25), case 2, of Feigelson and Babu, except that the S needs to be redefined with denominator n instead of $n - 1$.