QUEEN MARY, UNIVERSITY OF LONDON

MTH6102: Bayesian Statistical Methods

Practical 12

2023-2024

One of the datasets that is included with R, the *sleep* dataset. The data was used as an example (Illustration 1) in the paper *The probable error of a mean* (1908) by Gosset ("Student") in which he derived and explored the one sample t-test. A version of the paper is here.

This dataset doesn't need to be imported, it can accessed from within R.

The data consists of two sets of measurements on 10 people, recording the increase in hours of sleep (compared to a control measurement) after taking each of two types of the drug of hyoscyamine. Typing `sleep` lists the entire dataset. We can put the measurements in vectors as follows:

```
x1 = sleep$extra[sleep$group==1]
x2 = sleep$extra[sleep$group==2]
y = x2-x1
```

The paper uses the difference, measurement 2 minus measurement 1, so we do the same.

Carry out a one-sided t-test on y to see if measurement 2 is greater than measurement 1 (i.e. is the mean of y greater than 0). That is,

$$H_0 : \mu = 0 \quad H_1 : \mu > 0$$

This test assumes that

$$y_1, \ldots, y_n \sim N(\mu, \sigma^2)$$

for unknown σ, and tests in this case if $\mu = 0$. The syntax for a one-sample t-test is

```
t.test(y)
```

By default the test is two-sided. To get a one-sided test against the alternative that $\mu > 0$, the option `alternative="greater"` is needed. The null hypothesis is tested using the p-value. This is the probability of observing a test statistic at least as extreme as the observed one, given H_0 is true. What is the p-value?

Now, to make the comparison with Bayesian methods easier, assume we know the standard deviation $\sigma = 1.2$. Carry out a test of the same one-sided hypothesis $\mu = 0$ vs $\mu > 0$, by assuming that

$$\bar{y} \sim N(0, \sigma^2/n)$$

under the null hypothesis, by using the `pnorm` function.

To check your calculations, there is a package that contains a function `z.test`. The package in "BSDA", which you would need to install, then the syntax is

```
library(BSDA)
z.test(y, sigma.x=s, alternative="greater")
```

where `s` is the standard deviation (assumed to be known).

For a Bayesian version, we don't use p-values but we can use posterior probabilities. We again assume that σ is known and equal to 1.2. The prior distribution is

$$\mu \sim N(\mu_0, \sigma_0^2).$$

Take $\mu_0 = 0$. Work out the posterior distribution parameters for $p(\mu \mid y)$ if $\sigma_0 = 5$. Use these to calculate the posterior probability

$$P(\mu > 0)$$

and

$$P(\mu \leq 0).$$

Repeat this with $\sigma_0 = 100$. How do these posterior probabilities compare to the p-values from the one-sided hypothesis tests?