Statistics Examples Sheet 2

This examples sheet covers material of lectures 6–10 and is appropriate for your second supervision. A copy of this sheet can be found at: `http://www.statslab.cam.ac.uk/~rrw1/stats/`

1. (Lecture 6, Hypothesis testing) A positive random variable X has density function

$$f(x \mid \theta) = \frac{\theta}{(\theta + x)^2}, \quad x > 0$$

where $\theta > 0$ is an unknown parameter. Find the best test of size 0.05 of $H_0 : \theta = 1$ vs $H_1 : \theta = 2$, and show that the probability of Type II error is 19/21.

2. (Lecture 6, Hypothesis testing) Let $X_1, X_2, \dots, X_n$ be IID random variables, each with the Poisson distribution of parameter θ (and therefore of mean θ and variance θ). Find the best size α test of $H_0 : \theta = 1$ against $H_1 : \theta = 1.21$. By using the Central Limit Theorem to approximate the distribution of $\sum_i X_i$ show that the smallest value of n required to make $\alpha = 0.05$ and $\beta \le 0.1$ (where α and β are the Type I and Type II error probabilities) is somewhere near 213.

3. (Lecture 7, Further hypothesis testing) Find the best size α test of $H_0 : \theta = \theta_0$ vs $H_1 : \theta = \theta_1$ ($> \theta_0$) and write down an expression for the power function when $X_1, \dots, X_n$ are IID exponential random variables, with parameter θ.

Use χ^2 tables to find the least sample size which will allow us to test $H_0 : \theta = 1$ against $H_1 : \theta = 3$ with $\alpha = 5\%$ and $\beta \le 10\%$. Describe the appropriate critical region numerically. [*Hint: Recall the equivalence of the gamma*$(n/2, 1/2)$ *and* χ^2_n *distributions.*]

4. (Lecture 7, Further hypothesis testing) Let $X_1, X_2, \dots, X_n$ be a sample of size n from the distribution with probability density function

$$\frac{1}{2}\lambda e^{-\lambda|x|}, \quad -\infty < x < \infty, \quad \lambda > 0,$$

where λ is unknown. It is desired to test the hypothesis $H_0 : \lambda = \lambda_0$ against $H_1 : \lambda = \lambda_1$ where $\lambda_1 > \lambda_0$. Find the test that minimizes the sum of the probabilities of the two types of error. Denote the size of this test by α. Is this the most powerful test of size α for testing $H_0 : \lambda = \lambda_0$ against $H_1 : \lambda = \lambda_2$, for each $\lambda_2 > \lambda_0$?

5. (Lecture 8, Generalized likelihood ratio tests) Let $X_1, \dots, X_n$ be an IID random sample from an exponential distribution $\mathcal{E}(\theta_1)$, with density

$$f(x \mid \theta_1) = \theta_1 e^{-\theta_1 x}, \quad x > 0$$

and let $Y_1, \dots, Y_n$ be an independent IID random sample from $\mathcal{E}(\theta_2)$.

Find the form of the likelihood ratio test for testing $H_0 : \theta_1 = \theta_2$ against $H_1 : \theta_1 \ne \theta_2$. Show that the test can be expressed in terms of the statistic

$$T = \textstyle\sum_i X_i / (\sum_i X_i + \sum_i Y_i).$$

By showing that when H_0 is true the distribution of T does not depend on $\theta = \theta_1 = \theta_2$, construct a test of exact size α of H_0 against H_1 based on T.

6. (Lecture 8, Generalized likelihood ratio tests) The data $x_1, \dots, x_n$ has been observed and it is known that x_i is a sample from a Poisson distribution with an unknown mean λ_i. It is desired to test $H_0 : \lambda_1 = \cdots = \lambda_n$ against a general alternative hypothesis that the λ_i are arbitrary. Show that, on the basis of the generalized likelihood ratio test, H_0 should be rejected for large values of the test statistic

$$2 \sum_{i=1}^{n} x_i \log(x_i / \bar{x}),$$

where $\bar{x} = (1/n) \sum_i x_i$. Show that this statistic is approximately

$$\frac{1}{\bar{x}} \sum_{i=1}^{n} (x_i - \bar{x})^2 .$$

What would you conclude for data $(3, 5, 1, 6, 5)$?

7. (Lecture 9, Chi-squared tests of categorical data)

(i) It is known that an observation made on a certain system will yield a result falling into one of k categories, $C_1, C_2, \dots, C_k$. To each value of the real parameter θ in a given interval corresponds a probability distribution $p_i(\theta)$ $(i = 1, 2, \dots, k)$. The null hypothesis is made that, for some unknown value of θ, the probability of a result falling into category C_i is $p_i(\theta)$. A total of n observations is made, of which n_i fall into category C_i $(i = 1, 2, \dots, k)$. Show carefully how to use a χ^2 distribution in testing the above hypothesis, and describe how you would carry out the test, using a statistic of the form

$$\sum_{i=1}^{k} (n_i - e_i)^2 / e_i,$$

where e_i is the expected number of observations in category C_i under (an appropriate case of) the null hypothesis.

(ii) A scientist gets observations in three categories, across which he suspects a linear trend of probability, in which case

$$p_1(\theta) = \frac{1}{3} - \theta, \quad p_2(\theta) = \frac{1}{3}, \quad p_3(\theta) = \frac{1}{3} + \theta$$

for some value of θ such that $-\frac{1}{3} < \theta < \frac{1}{3}$. The numbers observed are $n_1 = 12$, $n_2 = 24$, $n_3 = 24$. Test the hypothesis of linearity.

8. (Lecture 9, Chi-squared tests of categorical data) A machine produces plastic articles in bunches of three articles at a time. The process is rather unreliable, and quite a few defective articles are observed. In an experimental run of the machine, 512 bunches were produced. Of these, the numbers of bunches with $i = 0, 1, 2, 3$ defective articles were 213 ($i = 0$), 228 ($i = 1$), 57 ($i = 2$), and 14 ($i = 3$). Test the hypothesis that each article has a constant (but unknown) probability θ of being defective, independently of all other articles.

9. (Lecture 9, Chi-squared tests of categorical data) From each of six batches of seed a random sample of 100 seeds was selected for sowing. The numbers of seeds that failed to germinate in the six samples of 100 seeds were

$$12, 20, 9, 17, 24, 16.$$

Test the hypothesis that the proportion of non-germinating seeds was the same for all batches.

10. (Lecture 9, Chi-squared tests of categorical data) The following data is for clinical trials of old and new treatments for a disease. Exactly 1,100 patients were chosen to receive each treatment.

	Survive	Die	Total
Old	505	595	1,100
New	195	905	1,100
Total	700	1,500	2,200

Test the hypotheis that the old and new treatments are equally successful. On this basis, which treatment would you prefer?

In fact, the trials took place at two hospitals, for which the data is given below. Doctors at Hospital A, a famous research hospital, designed the trial. Their patients tend to be more seriously ill and they also used the new treatment more often. Do you wish to revise your conclusion about the relative effectiveness of the two treatments?

	Hospital A		
	Survive	Die	Total
Old	5	95	100
New	100	900	1,000
Total	105	995	1,100

	Hospital B		
	Survive	Die	Total
Old	500	500	1,000
New	95	5	100
Total	595	505	1,100

What lesson do you learn from this example?

11. (Lecture 9, Chi-squared tests of categorical data) A random sample of 59 people from the planet Krypton yielded the following results:

		Eye-colour	
		1 (Blue)	2 (Brown)
Sex	1 (Male)	19	10
	2 (Female)	9	21

Professor A had believed that sex and eye-colour are independent factors on Krypton. After doing a χ^2 test of his hypothesis against the alternative

$$H_1 : \sum_i \sum_j p_{ij} = 1,$$

he finds that he has to reject it at the 5% level and even at the 1% level.

Professor B has always believed the much stronger hypothesis that $p_{ij} = \frac{1}{4}$ for all i and j. After doing a χ^2 test of his hypothesis against H_1, he finds that he does not need to reject it at the 5% level.

You are asked for expert statistical advice: firstly, to check the calculations, and secondly, to comment on whether or not Statistics is an absurd subject.

12. (Lecture 9, Chi-squared tests of categorical data) In Exercises 10 and 11 you have performed Pearson χ^2 tests. Explain carefully how the form of the null and alternative hypotheses in these exercises differ.

13. (Lecture 10, Distributions of the sample mean and variance) Statisticians A and B obtain independent IID samples, $X_1, \dots, X_{10}$ and $Y_1, \dots, Y_{17}$ respectively, from a $N(\mu, \sigma^2)$ distribution, for which μ and σ^2 are both unknown. They estimate (μ, σ^2) by $(\bar{X}, S_{XX}/9)$ and $(\bar{Y}, S_{YY}/16)$, respectively. Given that $\bar{X} = 5.5$ and $\bar{Y} = 5.8$, which statistician's estimate of σ^2 is more probable to have exceeded the true value of σ^2 by more than 50%? Find this probability (approximately) in each case.

Hint: This is something of a 'trick' question. Why? You may find χ^2 tables helpful.

R Weber, January 22, 2008