18.05 Practice Exam 2b Solutions

Problem 1. Concept questions

(a) *A certain august journal publishes psychological research. They will only publish results that are statistically significant when tested at a significance level of 0.05.*

Could all of their published results be true? ***Yes*** ***No***

Could all of their published results be false? ***Yes*** ***No***

(b) *True or false: Setting the prior probability of a hypothesis to 0 means that no amount of data will make the posterior probability of that hypothesis the maximum over all hypotheses.*

True **False**

(c) *A researcher collected data that fit the criteria for a two-sided Z-test. He set the significance level at 0.05. He ran 80 trials and got a z-value of 1.7. This gave a p-value of 0.0892, so he could not reject the null hypothesis. Convinced that his alternative hypothesis was correct he ran 80 more trials. The combined data from the 160 trials now had a z-value of 2.1. He wrote a paper carefully describing his experiments and submitted it to the journal in part (a).*

Will the journal publish his results? ***Yes*** ***No***

(d) *Let θ be the probability of heads for a bent coin. Suppose your prior $f(\theta)$ is Beta$(6, 8)$. Also suppose you flip the coin 7 times, getting 2 heads and 5 tails. What is the posterior pdf $f(\theta|x)$?*

Answers. (a) Yes and yes. Frequentist statistics don't give the probability an hypothesis is true.

(b) True. Bayesian updating involves multiplying the likelihood and the prior. If the prior is 0 then this product will be 0.

(c) No. The actual experiment that was run would reject the null hypothesis if it were true, more than 5% of the time.

(d) Beta(8,13).

Problem 2. *The Pareto distribution with parameter α has range $[1, \infty)$ and pdf*

$$f(x) = \frac{\alpha}{x^{\alpha}}$$

Suppose the data

$$5,\ 2,\ 3$$

was drawn independently from such a distribution. Find the maximum likelihood estimate (MLE) of α.

Solution: likelihood $= \phi(\text{data} \mid \alpha) = \dfrac{\alpha}{5^{\alpha}} \cdot \dfrac{\alpha}{2^{\alpha}} \cdot \dfrac{\alpha}{3^{\alpha}} = \dfrac{\alpha^3}{30^{\alpha}}.$

Therefore, log likelihood $= \ln(\phi(\text{data} \mid \alpha)) = 3\ln(\alpha) - \alpha\ln(30)$. We find the maximum likelihood by setting the derivative equal to 0:

$$\frac{d}{d\alpha}\ln(\phi(\text{data} \mid \alpha)) = \frac{3}{\alpha} - \ln(30) = 0.$$

Solving we get $\boxed{\hat{\alpha} = \frac{3}{\ln(30)}}$.

Problem 3.

Your friend grabs a die at random from a drawer containing two 4-sided dice, one 8-sided die, and one 12-sided die. They roll the die once and report that the result is 5.

(a) *Make a discrete Bayes table showing the prior, likelihood, and posterior for the type of die rolled given the data.*

Solution: (We include the last column for part (d).)

hypoth. θ	prior	likelihood $P(x_1 = 5 \mid \theta)$	Bayes numer.	posterior	likelihood $P(x_2 = 7 \mid \theta)$
4-sided	1/2	0	0	0	0
8-sided	1/4	1/8	1/32	$\frac{3}{5}$	1/8
12-sided	1/4	1/12	1/48	$\frac{2}{5}$	1/12
			$T = \frac{1}{32} + \frac{1}{48} = \frac{5}{96}$		

Note, the posterior probabilites have to have the same ratio as the Bayes numerators, i.e. $0 : \frac{1}{32} : \frac{1}{48} = 0 : \frac{3}{5} : \frac{2}{5} = 0 : 3 : 2$.

(b) *What is the prior predictive probability of rolling a 5?*

Solution: $P(x_1 = 5) = T = \frac{1}{32} + \frac{1}{48} = \frac{5}{96} \approx 0.0521$. (Any of these expressions is a fine answer.)

(c) *What are your posterior* ***odds*** *that the die has 12 sides?*

Solution: $\text{Odds}(\theta = 12 \mid x_1 = 5) = \dfrac{P(\theta = 12 \mid x_1 = 5)}{P(\theta \neq 12 \mid x_1 = 5)} = \dfrac{2/5}{3/5} = 2/3$.

(d) *Given the data of the first roll, what is your probability that the next roll will be a 7?*

Solution: Here we use the posterior probabilities of the for each of the dice to compute the predictive probability

$$P(x_2 = 7 \mid x_1 = 5) = 0 \cdot 0 + \frac{3}{5} \cdot \frac{1}{8} + \frac{2}{5} \cdot \frac{1}{12} = \frac{13}{120} \approx 0.108$$

(Any of these answers is okay.)

Problem 4. *Everyone knows that giraffes are tall, but how much do they weigh? Let's suppose that the weight of male giraffes is normally distributed with mean 1200 kg and standard deviation 200 kg.*

I volunteered at the zoo and was given the task of weighing their male giraffe Beau. Now weighing a giraffe is not easy and the process produces random errors following a $N(0, 100^2)$ distribution. To compensate for the inaccuracy of the scale I weighed Beau three times and got the following measurements:

1250 kg, 1300 kg, 1350 kg .

What is the posterior expected value of Beau's weight?

Solution: This is a normal/normal conjugate prior pair, so we use the normal-normal update formulas.

μ = Beau's weight.

$n = 3, \quad \overline{x} = 1300$

Prior $\sim$ N(1200, 200^2), so $\mu_{\text{prior}} = 1200, \quad \sigma^2_{\text{prior}} = 200^2$.

Likelihood $\sim$ N($\mu \mid 100^2$), so $\sigma^2 = 100^2$.

$a = \frac{1}{\sigma^2_{pr}} = \frac{1}{200^2}, \qquad b = \frac{n}{\sigma^2} = \frac{3}{100^2}$.

$$\mu_{\text{posterior}} = \frac{a \cdot \mu_{\text{prior}} + b \cdot \overline{x}}{a+b} = \frac{\frac{1}{200^2} \cdot 1200 + \frac{3}{100^2} \cdot 1300}{1/200^2 + 3/100^2}$$

Problem 5. *Data is drawn from a binomial(5, θ) distribution, where θ is unknown. Here is the table of probabilities $p(x \mid \theta)$ for 3 values of θ:*

x	*0*	*1*	*2*	*3*	*4*	*5*
$\theta = 0.5$	*0.031*	*0.156*	*0.313*	*0.313*	*0.156*	*0.031*
$\theta = 0.6$	*0.010*	*0.077*	*0.230*	*0.346*	*0.259*	*0.078*
$\theta = 0.8$	*0.000*	*0.006*	*0.051*	*0.205*	*0.410*	*0.328*

You want to run a significance test on the value of θ. You have the following:

Null hypothesis: $\theta = 0.5$.
Alternate hypotheses: $\theta > 0.5$.
Significance level: $\alpha = 0.1$.

(a) *Find the rejection region.*

(b) *Compute the power of the test for each of the two hypotheses $\theta = 0.6$ and $\theta = 0.8$.*

(c) *Suppose you run an experiment and the data gives $x = 4$. Compute the p-value of this data.*

Answers. **(a)** Since the H_A is right-sided we use a right-sided rejection region: $\boxed{\text{rejection region } x = 5}$.

x	0	1	2	3	4	5
$\theta = 0.5$	0.031	0.156	0.313	0.313	0.156	0.031
$\theta = 0.6$	0.010	0.077	0.230	0.346	0.259	0.078
$\theta = 0.8$	0.000	0.006	0.051	0.205	0.410	0.328

(b) Power $= P(\text{reject} \mid \theta)$.
$\theta = 0.6$: power $= 0.078$
$\theta = 0.8$: power $= 0.328$

(c) $p = P(x \geq 4 \mid \theta = 0.5) = 0.156 + 0.031 = 0.187$.

Problem 6. *You have data drawn from a normal distribution with a known variance of 16. You set up the following NHST:*

- *H_0: data follows a $N(2, 4^2)$*

- *H_A: data follows a $N(\mu, 4^2)$ where $\mu \neq 2$.*
- *Test statistic: standardized sample mean z.*
- *Significance level set to $\alpha = 0.05$.*

You then collected $n = 16$ data points with sample mean 1.5.

***(a)** Find the rejection region. Draw a graph indicating the null distribution and the rejection region.*

***(b)** Find the z-value and add it to your picture in part (a).*

***(c)** Find the p-value for this data and decide whether or not to reject H_0 in favor of H_A.*

Answers. (a) The test statistic is z, so we need a Z-graph

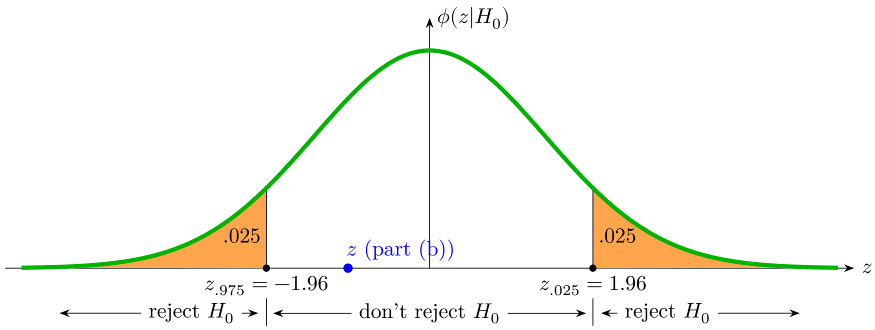

The rejection region is $z < -1.96$ or $z > 1.96$.

(b) We standardize $\overline{x}$ to get z: $z = \dfrac{\overline{x} - 2}{\sigma_{\overline{x}}} = \dfrac{1.5 - 2}{4/\sqrt{16}} = -0.5$

(c) $p = 2P(Z \leq -0.5) = 2 \cdot (0.3085) = 0.6170$. Since $p > 0.05$ we do not reject H_0.

Problem 7. *Someone claims to have found a long lost work by Jane Austen. She asks you to decide whether or not the book was actually written by Austen.*

You buy a copy of Sense and Sensibility *and count the frequencies of certain common words on some randomly selected pages. You do the same thing for the 'long lost work'. You get the following table of counts.*

Word	a	an	this	that
Sense and Sensibility	150	30	30	90
Long lost work	90	20	10	80

Using this data, set up and evaluate a significance test of the claim that the long lost book is by Jane Austen. Use a significance level of 0.1.

Solution: The null hypothesis H_0: For the 4 words counted the long lost book has the same relative frequencies as *Sense and Sensibility*

Total word count of both books combined is 500, so the maximum likelihood estimate of the relative frequencies assuming H_0 is simply the total count for each word divided by the total word count.

Word	a	an	this	that	Total count
Sense and Sensibility	150	30	30	90	300
Long lost work	90	20	10	80	200
totals	240	50	40	170	500
rel. frequencies under H_0	240/500	50/500	40/500	170/500	500/500

Now the expected counts for each book under H_0 are the total count for that book times the relative frequencies in the above table. The following table gives the counts: (observed, expected) for each book.

Word	a	an	this	that	Totals
Sense and Sensibility	(150, 144)	(30, 30)	(30, 24)	(90, 102)	(300, 300)
Long lost work	(90, 96)	(20, 20)	(10, 16)	(80, 68)	(200, 200)
Totals	(249, 240)	(50, 50)	(40, 40)	(170, 170)	(500, 500)

The chi-square statistic is

$$\begin{aligned} X^2 &= \sum \frac{(O_i - E_i)^2}{E_i} \\ &= \frac{6^2}{144} + \frac{0^2}{30} + \frac{6^2}{24} + \frac{12^2}{102} + \frac{6^2}{96} + \frac{0^2}{20} + \frac{6^2}{16} + \frac{12^2}{68} \\ &\approx 7.9 \end{aligned}$$

There are 8 cells and all the marginal counts are fixed, so we can freely set the values in 3 cells in the table, e.g. the 3 blue cells, then the rest of the cells are determined in order to make the marginal totals correct. Thus $df = 3$.

Looking in the $df = 3$ row of the chi-square table we see that $X^2 = 7.9$ gives p between 0.025 and 0.05. Since this is less than our significance level of 0.1 we reject the null hypothesis that the relative frequencies of the words are the same in both books. Based on the assumption that all her books have similar word frequencies (which is something we could check) we conclude that the book is probably not by Jane Austen.