Mathematics 376 – Probability and Statistics 2
Solutions – Final Examination
May 6, 2006

I. Let $Y_1, \ldots, Y_n$ be a random sample from a distribution with probability density function $f(y|\theta) = \theta y^{\theta-1}$ if $0 < y < 1$ and 0 otherwise. We also assume $\theta > 0$.

A) (15) Find the method of moments estimator for θ.

Solution: To apply the method of moments, we first compute the expected value of Y with this pdf:

$$E(Y) = \int_0^1 y\theta y^{\theta-1}\ dy = \theta \int_0^1 y^{\theta}\ dy = \frac{\theta}{\theta+1}$$

Then we set the first sample moment (i.e. the sample mean $\overline{y}$) equal to this, and solve for θ:

$$\frac{\theta}{\theta+1} = \overline{y} \Rightarrow \widehat{\theta}_{MM} = \frac{\overline{y}}{1-\overline{y}}$$

B) (15) Find the maximum-likelihood estimator for θ.

Solution: We have

$$\begin{aligned}
\ln(L(y_1, \ldots, y_n|\theta)) &= \ln\left(\theta y_1^{\theta-1} \cdots \theta y_n^{\theta-1}\right) \\
&= \ln(\theta^n (y_1 \cdots y_n)^{\theta-1}) \\
&= n\ln(\theta) + (\theta - 1)\ln(y_1 \cdots y_n) \\
\Rightarrow \frac{d}{d\theta}\ln(L) &= \frac{n}{\theta} + \ln(y_1 \cdots y_n)
\end{aligned}$$

To find the maximum-likelihood estimator, we set this equal to zero and solve for θ:

$$\widehat{\theta}_{ML} = \frac{-n}{\ln(y_1 \cdots y_n)}$$

II. Let $X_1, \ldots, X_n$ and $Y_1, \ldots, Y_m$ be two random samples from normal distributions with common variance σ^2.

A) (10) Using the fact that the expected value of a χ^2 random variable with ν degrees of freedom is ν, show that $S_1^2 = \frac{1}{n-1}\sum_{i=1}^n (X_i - \overline{X})^2$ is an unbiased estimator for σ^2.

Solution: From discussions in class (or Theorem 7.3 in the text), we know that $\frac{(n-1)S_1^2}{\sigma^2} \sim \chi^2(n-1)$. Hence by the fact stated in the problem:

$$E\left(\frac{(n-1)S_1^2}{\sigma^2}\right) = n - 1$$

But the expected value is linear so this implies

$$\frac{(n-1)}{\sigma^2} E(S_1^2) = n - 1$$

so

$$E(S_1^2) = \sigma^2$$

which shows that S_1^2 is an unbiased estimator.

B) (10) Show that the pooled estimator S_p^2 using both the X_i and the Y_j is also unbiased for σ^2.

Solution: Using part A and linearity of expected value:

$$\begin{aligned}
E(S_p^2) &= E\left(\frac{(n-1)S_1^2 + (m-1)S_2^2}{n+m-2}\right) \\
&= \frac{(n-1)}{n+m-2}E(S_1^2) + \frac{(m-1)}{n+m-2}E(S_2^2) \\
&= \frac{(n-1)}{n+m-2}\sigma^2 + \frac{(m-1)}{n+m-2}\sigma^2 \\
&= \frac{(n-1)+(m-1)}{n+m-2}\sigma^2 \\
&= \sigma^2
\end{aligned}$$

This shows S_p^2 is an unbiased estimator for σ^2.

III.

A) (15) Let $Y_1, Y_2, \ldots, Y_{10}$ be a random sample from a normal distribution with mean $\mu = 2$ and variance $\sigma^2 = 81$. Let $U = \frac{1}{81}\sum_{i=1}^{9}(Y_i - 2)^2$. What is the distribution of $V = \frac{Y_{10}-2}{3\sqrt{U}}$?

Solution: First, we note that U is the sum of the $\left(\frac{Y_i-2}{9}\right)^2$. Each of these terms is the square of a standard normal random variable, so U has a χ^2 distribution with 9 d.f. (by Theorem 7.2 in the text). Hence

$$V = \frac{Y_{10}-2}{3\sqrt{U}} = \frac{(Y_{10}-2)/9}{\sqrt{U/9}}$$

has the form $Z/\sqrt{U/\nu}$ of a t-distributed random variable (Definition 7.2 in the text – note that Y_{10} is independent of $Y_1, \ldots, Y_9$ by assumption, hence Y_{10} and U are independent). It follows that V has a t-distribution with 9 d.f.

B) (15) If T has a t-distribution with ν degrees of freedom, what is the distribution of T^2? Explain.

Solution: Consider the form $T = Z/\sqrt{U/\nu}$ as in Definition 7.2. Then $T^2 = Z^2/(U/\nu)$. The numerator is the square of a standard normal, hence has a $\chi^2(1)$ distribution. The denominator is $\chi^2(\nu)$. Hence this has the form for a random variable with an F-distribution with 1 d.f. in the numerator, and ν d.f. in the denominator.

IV. Let p be the proportion of letters mailed in the Netherlands that are delivered the next day.

A) (15) A random sample of $n = 200$ letters are sent out and 142 are delivered the next day. Find an approximate 95% confidence interval for p based on this sample.

Solution: Our point estimate for p is $\widehat{p} = 142/200 = .71$ and using $z_{.025} = 1.96$ the confidence interval is

$$.71 \pm (1.96)\sqrt{\frac{(.71)(.29)}{200}} = .71 \pm .063$$

B) (15) ("Thought question") Note that part A says "approximate." What is the actual distribution of Y = the number of letters delivered the next day (out of a random sample of size $n = 200$)? Why does the method you used in part A give a reasonable interval estimate for p?

Solution: The actual distribution of Y is binomial with $n = 200$ and success probability p on one trial. The method using the normal distribution is based on the fact that, by the Central Limit Theorem, for large n and binomial Y, Y/n is close to a normal random variable with mean p and variance pq/n. The approximate interval is reasonable, since n is quite large here and the estimate $\sqrt{\frac{(.71)(.29)}{200}}$ is likely to be close to $\sqrt{pq/200}$.

V. A mathematics department wishes to evaluate a new method of teaching calculus with Maple labs. At the end of the course, 15 students who used the labs are given a standardized test. Their average score is 83, with standard deviation 9.

A) (10) Find a 95% confidence interval for the mean test score for students who are taught using the new method.

Solution: Since $n = 15$, we use the small-sample formulas. $t_{.025}(14) = 2.145$, so the interval is

$$83 \pm 2.145\frac{9}{\sqrt{15}} = 83 \pm 4.984.$$

B) (10) From departmental experience, students who are taught the course without the Maple labs average 79 on the same standardized test, and the standard deviation is also 9 for these students. Is there sufficient evidence to conclude that taking the course with the labs has an effect on students' performance, at the $\alpha = .05$ level?

Solution: No, since the value $\mu = 79$ is in the interval from part A: $[78.016, 87.984]$. (A different, but equivalent, method is to use a t-test of $H_0 : \mu = 79$ versus the alternative $H_a : \mu \neq 79$. The conclusion is the same, of course!)
Comment: A common solution on the exam papers you submitted used a small sample test of equality of means for this part, with $n_1 = n_2 = 15$. This is *not very appropriate* given the information as stated in the problem. Note the problem says "from departmental experience" the average test score of students who took the course without the

Maple labs is 79. This does not say the 79 is an average of 15 test scores the way the 83 is. In fact, the best way to interpret it is that that a *large number* of students who took the course without the Maple labs (probably over a period of years) have taken the standardized test, with mean score 79.

C) (10) What number n of students who took the course with the labs would have to be tested in order for the department to be "98% sure" that the sample mean test score is no farther than 1 away from the true population mean test score for the students taking the course with the labs. You may assume for this part that n will be significantly > 30.

Solution: Now, since $n > 30$, we go to the large-sample formula. To be "98% sure," we want the 98% confidence interval $\mu \pm z_{.01}\frac{s}{\sqrt{n}}$ to have width at most 1, i.e.

$$(2.33)\frac{9}{\sqrt{n}} \leq 1 \Rightarrow n \geq (2.33 \cdot 9)^2 = 439.74$$

$n \geq 440$ will do here.

VI. The fill weights of a random sample of $n_1 = 21$ 6-pound boxes of "Super-Sudsy" laundry soap produced at Plant 1 had a mean of 6.25 pounds and standard deviation $s = .095$ pounds. A similar sample of size $n_2 = 21$ produced at Plant 2 had mean weight 6.12 pounds and standard deviation $s = .065$.

A) (15) Is there sufficient evidence at the $\alpha = .01$ level to conclude that the standard deviations at the two plants are different?

Solution: Since $\sigma_i > 0$, $\sigma_1 = \sigma_2 \Leftrightarrow \sigma_1^2 = \sigma_2^2$, and we use the usual F-test for equality of variances. H_0 is that $\sigma_1^2 = \sigma_2^2$, and the alternative is $H_a : \sigma_1^2 \neq \sigma_2^2$. The test statistic is

$$F = s_1^2/s_2^2 = (.095)^2/(.065)^2 = 2.136.$$

(F-distribution with 20 numerator d.f. and 20 denominator d.f.) The two-tailed rejection region for a test with $\alpha = .01$ is

$$RR = \{F < F_{.995}(20, 20)\} \cup \{F > F_{.005}(20, 20)\} = \{F < 1/3.32 = .301\} \cup \{F > 3.32\}$$

Since 2.136 is not in RR we cannot reject H_0 based on this evidence.

B) (15) Is there sufficient evidence to conclude that the mean fill weights are different? Report the results by giving an estimate of the p-value of your test.

Solution: The results of part A indicate that it is reasonable to assume $\sigma_1^2 = \sigma_2^2$ and use the small-sample test for equality of means. We have test statistic

$$t = \frac{\overline{Y}_1 - \overline{Y}_2}{S_p\sqrt{1/21 + 1/21}}$$

The pooled estimator for the variance is

$$S_p^2 = \frac{(20)(.095)^2 + (20)(.065)^2}{40} = .006625$$

So our test statistic value is

$$t = \frac{6.25 - 6.12}{\sqrt{.006625}\sqrt{2/21}} = 5.18$$

From the t-table (with $n =$ inf.) we see that $5.18 > t_{.005}$, so the p-value of the test will be $p < 2 \cdot (.005) = .01$. (In fact, using the large-sample formulas here gives almost exactly the same results. You can also consult the z-table here to get a better estimate of p, and in fact p is much less than .01.) The fact that p is so small means that there is strong evidence to reject H_0 and conclude the mean fill weights are different.

VII. The following table gives measurements of the firmness of pickles stored in low-salt brine as a function of time:

x time (weeks)	y firmness (lb)
1	19.8
4	16.5
14	12.8
32	8.1
52	7.5

A) (15) Find the least-squares estimators for the coefficients in a model $Y = \beta_0 + \beta_1 x + \varepsilon$ for this data set.

Solution: Outline of the computation

$$\begin{aligned}
\overline{x} &= 20.6 \\
\overline{y} &= 12.94 \\
S_{xx} &= 1819.2 \\
S_{xy} &= -418.62 \\
S_{yy} &= 112.77 \\
\widehat{\beta_1} &= S_{xy}/S_{xx} = -.2301 \\
\widehat{\beta_0} &= \overline{y} - \widehat{\beta_1}\overline{x} = 17.68
\end{aligned}$$

B) (15) Is there sufficient evidence to say that $\beta_0 < 21$? Explain, using the p-value of an appropriate test.

Solution: We test $H_0 : \beta_0 = 21$ versus the alternative $\beta_9 < 21$. Using the standard formula the test statistic is

$$t = \frac{\widehat{\beta_0} - 21}{S\sqrt{c_{00}}}$$

(t-distribution with $n - 2 = 3$ d.f.) where

$$c_{00} = \frac{\sum x_i^2}{nS_{xx}} = .4333$$

and

$$S = \sqrt{\frac{SSE}{3}} = \sqrt{\frac{S_{yy} - \widehat{\beta_1} S_{xy}}{3}} = 2.3415$$

Then

$$t = \frac{17.68 - 21}{(2.3415)\sqrt{.4333}} = -2.154$$

From the t-table (with 3 d.f.), we see this is between $t_{.1}$ and $t_{.05}$. The p-value of the test is hence between .1 and .05. (We don't multiply by 2 here, since this is a lower-tail t-test.) This is relatively weak evidence to reject H_0.

Extra Credit (Short Essay – no more than a paragraph, please!) (10) Suppose we needed to choose between the estimators in question I, parts A and B. One criterion for comparing estimators we discussed in class used the relative efficiency $\text{eff}(\widehat{\theta_1}, \widehat{\theta_2}) = V(\widehat{\theta_2})/V(\widehat{\theta_1})$. But it shouldn't be clear how to compute either variance theoretically from your formulas. Propose a method to test which estimator is superior "experimentally," taking into account the possibility that which estimator is superior might depend on the true value of θ.

Solution: One possible method would be to simulate drawing samples from the distribution for various θ-values, use both estimators on the samples and compute the sample variances of the estimated θ-values. Ratios of those sample variances could be used to estimate the relative efficiency.