Mathematics 376 – Mathematical Statistics
Solutions for Final Examination
May 16, 2012

I. Let X_1, X_2, X_3, X_4, X_5 be a random sample from a normal distribution with mean $\mu = 6$ and variance $\sigma^2 = 81$.

A) What is the distribution (give the type and all relevant parameter values) of $U = \frac{(X_1-6)^2+(X_2-6)^2+(X_3-6)^2+(X_4-6)^2}{81}$? Why?

Solution: For each i, $\frac{X_i-6}{9}$ has a standard normal distribution. Therefore

$$U = \left(\frac{X_1-6}{9}\right)^2 + \left(\frac{X_2-6}{9}\right)^2 + \left(\frac{X_3-6}{9}\right)^2 + \left(\frac{X_4-6}{9}\right)^2$$

is the sum of the squares of four independent standard normals. So U has a χ^2 distribution with 4 degrees of freedom.

B) What is the distribution (give the type and all relevant parameter values) of $V = \frac{2(X_5-6)}{9\sqrt{U}}$? Why?

Solution: We have

$$V = \frac{\frac{X_5-6}{9}}{\sqrt{U/4}}$$

This has the form of a standard normal, divided by the square root of a χ^2, divided by its number of degrees of freedom. So, by definition, V has a t-distribution with 4 degrees of freedom.

C) (10) What is the distribution (give the type and all relevant parameter values) of $W = (X_1-6)^2/(X_2-6)^2$? Why?

Solution: We have

$$W = \frac{\left(\frac{X_1-6}{9}\right)^2}{\left(\frac{X_2-6}{9}\right)^2}$$

This is a ratio of $\chi^2(1)$ random variables. Hence W has an F-distribution with 1 degree of freedom in the numerator and 1 degree of freedom in the denominator.

II. A random variable Y has pdf of the form

$$f(y) = \begin{cases} \frac{2y}{\alpha} e^{-y^2/\alpha} & \text{if } y \geq 0 \\ 0 & \text{otherwise} \end{cases}$$

(This is a particular case of a *Weibull* distribution; these are often used to model length of life of mechanical components and living organisms.) Let $Y_1, \ldots, Y_n$ be a random sample from this distribution with unknown α.

A) Find the maximum-likelihood estimator for α.

Solution: The likelihood function is

$$L(y_1, \dots, y_n \mid \alpha) = \frac{2^n \prod_i y_i}{\alpha^n} e^{-(\sum_i y_i^2)/\alpha}$$

so

$$\ln(L) = \ln(2^n \prod_i y_i) - n\ln(\alpha) - \frac{\sum_i y_i^2}{\alpha}.$$

First, let us see if $\ln(L)$ has any critical points as a function of α:

$$0 = \frac{d}{d\alpha}\ln(L) = \frac{-n}{\alpha} + \frac{\sum_i y_i^2}{\alpha^2}$$

when

$$\alpha = \frac{1}{n}\sum_{i=1}^{n} y_i^2$$

It is easy to check that the second derivative is negative here, so the maximum-likelihood estimator for α is

$$\widehat{\alpha} = \frac{1}{n}\sum_{i=1}^{n} y_i^2.$$

(Note that this is the second sample moment. In fact the *variance* of Y is a constant multiple of α.)

B) What is the pdf for the sample maximum, $Y_{(n)}$?

By integration (substitution with $u = y^2$), we see the cdf for this distribution is

$$F(y) = -e^{-y^2/\alpha} + c$$

for some constant of integration c. At $y = 0$ we should get $F(0) = 0$, so $c = 1$. Hence by the usual formula, the pdf for the sample maximum is

$$f_n(y) = n\left(1 - e^{-y^2/\alpha}\right)^{n-1} \frac{2y}{\alpha} e^{-y^2/\alpha}.$$

III. Let p be the proportion of letters mailed in Belgium that are delivered the next day.

A) A random sample of $n = 400$ letters is selected for tracking and 373 of those are delivered the next day. Find an approximate 98% confidence interval for p based on this sample.

Solution: We have $\widehat{p} = 373/400 = .9325$. Since $400 > 30$, to get a 98% confidence interval we can use the large sample two-sided formulas with $\alpha = .01$. The confidence interval is

$$\begin{aligned}\widehat{p} \pm z_{.01}\sqrt{\frac{\widehat{p}(1-\widehat{p})}{400}} &\doteq .9327 \pm 2.33 \cdot .0125 \\ &\doteq (.9033, .9617)\end{aligned}$$

All values in this interval are thought of as "reasonable estimates" for the proportion p. Since $p = .85$ is *not contained*, the answer to the last part is *no.*

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 $n = 400$)? Why does the method you used in part A give a reasonable interval estimate for p? Explain.

Solution: Y has a *binomial* distribution with $n = 400$, and p = the actual proportion of letters delivered the next day. An "exact" confidence interval would use information about that binomial distribution to develop the locations of the endpoints. The method we used in part A depends on the properties of the standard normal distribution. That is reasonable here since the Central Limit Theorem implies that $\frac{Y}{400}$ has an approximately normal distribution. (Recall that the estimator $\frac{Y}{400}$ for p can be thought of as a sample mean for 400 independent Bernoulli trials, all with success probability p, so the CLT applies in the form we discussed last semester.)

IV. The fill weights of a random sample of $n_1 = 13$ boxes of Twigs and Nuts Extra Crunchy Granola produced at Plant 1 had a mean of 20.5 ounces and standard deviation $s = .1$ ounce. A similar sample of size $n_2 = 11$ produced at Plant 2 had mean fill weight 20.3 ounces and standard deviation $s = .083$ ounce.

A) (20) Is there sufficient evidence at the $\alpha = .01$ level to conclude that the variances of the fill weights at the two plants are different?

Solution: To test for equality of variances, we want to use an F-test. The null hypothesis is $H_0 : \sigma_1^2 = \sigma_2^2 = \sigma^2$ and the alternative is $H_a : \sigma_1^2 \neq \sigma_2^2$. Under H_0, the test statistic

$$F = S_1^2/S_2^2 \quad \text{(which equals } \frac{(12S_1^2/\sigma^2)/12}{(10S_2^2/\sigma^2)/10}\text{)}$$

has an F-distribution with 12 degrees of freedom in the numerator and 10 degrees of freedom in the denominator. For a test with $\alpha = .01$, we want to set the rejection region to be

$$\{F \mid F > f_{.005}(12, 10)\} \cup \{F \mid 0 < F < f_{.995}(12, 10) = \frac{1}{f_{.005}(10, 12)}\},$$

which works out to the union

$$(5.66, +\infty) \cup (0, .1964)$$

Our test statistic value is

$$F = (.1)^2/(.083)^2 \doteq 1.45$$

This is not in the rejection region, so we *do not* reject H_0.

B) (20) 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: Because we did not find evidence to conclude the variances were different, we will use the basic small-sample t-test for equality of means (two-tailed version). The pooled estimator for the variance is

$$S^2 \doteq \frac{(12)(.1)^2 + (10)(.083)^2}{22} = .008586$$

This gives $S \doteq .0927$. Our test statistic is

$$t = \frac{20.5 - 20.3}{.0927\sqrt{1/13 + 1/11}} \doteq 5.27$$

From the t-table, for 22 degrees of freedom, we have $t_{.005} = 2.819$. We can say $p < 2 \cdot (.005) = .01$. So there is relatively strong evidence to conclude that *the mean fill weights are different*. (Using `R`, in fact the p value is quite a bit smaller even than that: $p \doteq .000028$.)

V. The following table gives measurements of water depth remaining in an evaporating reservoir as a function of time:

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

A) (20) Find the least-squares estimators for the coefficients β_0, β_1 in a model $Y = \beta_0 + \beta_1 x + \varepsilon$ for this data set.

Solution: Organizing the calculation the way we discussed in class,

$$\begin{aligned}
\overline{x} &= 20.6 \\
S_{xx} &= 1819.2 \\
\overline{y} &= 12.94 \\
S_{xy} &= -418.62 \\
\widehat{\beta_1} &= \frac{S_{xy}}{S_{xx}} \doteq -0.23 \\
\widehat{\beta_0} &= \overline{y} - \widehat{\beta_1}\overline{x} \\
&\doteq 17.68
\end{aligned}$$

So the estimated model would be

$$Y = 17.68 - .23x + \varepsilon$$

B) (20) Is there sufficient evidence to say that $\beta_1 < -0.1$? Test with $\alpha = .05$.

Solution: This is a lower-tail t-test with a null hypothesis saying $\beta_1 = -0.1$ (or something larger). The test statistic is

$$t = \frac{\widehat{\beta_1} - (-.1)}{S\sqrt{c_{11}}}$$

So we need to compute additionally

$$\begin{aligned} S_{yy} &= 112.772 \\ S^2 &= \frac{1}{5-2}\left(S_{yy} - \widehat{\beta_1} S_{xy}\right) \doteq 5.48 \\ S &= \sqrt{5.48} \doteq 2.34118 \\ c_{11} &= \frac{1}{S_{xx}} \doteq .0005497 \end{aligned}$$

So then

$$t = \frac{(-.23) + .1}{(2.34118)\sqrt{.0005497}} \doteq -2.368$$

For $n-2 = 3$ degrees of freedom, $-t_{.05} = -2.353$. We are (just barely!) in the rejection region, so we would say that there is enough evidence to conclude $\beta_1 < -0.1$.

VI. (20) Suppose we have measurements taken from a normal population with unknown μ and known σ^2. We test $H_0 : \mu = \mu_0$ versus $H_a : \mu > \mu_0$ using a Z-test and a rejection region designed to produce a Type I error probability α. Show that in order to make the Type II error probability equal a given β when μ is equal to some $\mu_a > \mu_0$, the sample size should be selected as follows:

$$n \geq \frac{(z_\alpha + z_\beta)^2 \sigma^2}{(\mu_a - \mu_0)^2}.$$

Solution: To get a given Type I error probability, α, we make the rejection region for the upper tail test

$$RR = \{Z \mid Z > z_\alpha\} \leftrightarrow \left\{\overline{Y} > \mu_0 + z_\alpha \cdot \frac{\sigma}{\sqrt{n}}\right\}$$

Now, assuming the actual population mean is μ_a, to compute the Type II error probability, we consider the case where we would fail to reject H_0, but where H_0 is actually false. That is we want the following:

$$\begin{aligned} \beta &= P\left(\overline{Y} \leq \mu_0 + z_\alpha \cdot \frac{\sigma}{\sqrt{n}}\right) \\ &= P\left(\frac{\overline{Y} - \mu_a}{\sigma/\sqrt{n}} \leq \frac{(\mu_0 - \mu_a)}{\sigma/\sqrt{n}} + z_\alpha\right) \end{aligned}$$

Since we have "restandardized" in the last equation, to get the probability to equal β, the right side of the inequality should be

$$\frac{(\mu_0 - \mu_a)}{\sigma/\sqrt{n}} + z_\alpha = -z_\beta$$

This rearranges to

$$\sqrt{n} = \frac{(z_\alpha + z_\beta)\sigma}{(\mu_a - \mu_0)}$$

and we get the desired equality after squaring.

Extra Credit

A New York Times/CBS poll conducted between April 5 and April 12, 2010 included the question "Do you approve or disapprove of the way your Representative in Congress is handling his or her job?" The poll was carried out in two stages. First a large national simple random sample was asked the question. The results were reported in three categories: 46% of respondents said they approved, 36% said they disapproved, and 18% said they had no opinion. Then the same question was asked of a separate sample of $n = 881$ people, all of whom had identified themselves as supporters of the "Tea Party" movement. The responses broke down like this in the same three categories:

	Approve	Disapprove	No opinion
Tea Party Sample	352	432	97

(Data is adapted from the New York Times web site.) It seems from the numbers that "Tea Party" supporters may differ from the general population when it comes to their opinions concerning their Congressional Representatives. But is this a real difference, or could it just be a product of chance variation in the sampling process?

We did not discuss this in class, but a standard statistical method in this case would be to use a χ^2-test. The test statistic

$$X^2 = \sum_{\text{categories}} \frac{(\text{observed} - \text{expected})^2}{\text{expected}}$$

has approximately a χ^2 distribution with degrees of freedom one less than the number of categories. For a test of H_0 : the distributions are really the same, versus H_a :, we would reject H_0 if X^2 fell in an upper tail rejection region.

A) (5) If the "Tea Party" sample did follow the national sample percentages, how many responses would we expect in each category?

Solution: If the "Tea Partiers" were really the same as the general population in this regard, then the expected numbers of people in the Approve, Disapprove, and No opinion categories would be:

$$.46 \times 881 \doteq 405.3, \quad .36 \times 881 \doteq 317.2, \quad .18 \times 881 \doteq 158.6.$$

B) (15) Carry out the χ^2 test and interpret the results.

Solution: We compute the X^2 statistic using the entries from the table above for the observed:

$$\chi^2 = \frac{(352 - 405.3)^2}{405.3} + \frac{(432 - 317.2)^2}{317.2} + \frac{(97 - 158.6)^2}{158.6} \doteq 72.5.$$

From the χ^2 table with 2 degrees of freedom, we see that $p < .005$. This is extremely strong evidence that "Tea Partiers" have a different set of opinions on this question. (They are much less likely to approve, and also much less likely not to have an opinion!)