Solutions to Review Problems for Exam #2

1. In the book "Experimentation and Measurement," by W. J. Youden and published by the by the National Science Teachers Association in 1962, the author reported an experiment, performed by a high school student and a younger brother, which consisted of tossing five coins and recording the frequencies for the number of heads in the five coins. The data collected are shown in Table 1.

Number of Heads	0	1	2	3	4	5
Frequency	100	524	1080	1126	655	105

Table 1: Frequency Distribution for a Five–Coin Tossing Experiment

(a) Are the data in Table 1 consistent with the hypothesis that all the coins were fair? Justify your answer.

Solution: If we let X denote the number of heads observed in the five–coin toss, and all the coins are fair, then

$$X \sim \text{binomial}(5, 0.5).$$

Thus, the probability that we will see k coins out of the 5 showing heads is

$$p_{X}(k) = \binom{5}{k}\left(\frac{1}{2}\right)^{k}\left(\frac{1}{2}\right)^{5-k}, \quad \text{for} \quad k = 0, 1, 2, 3, 4, 5.$$

Thus, out of of the $n = 3590$ tosses of the five coins, on average, we expect to see

$$np_{X}(k)$$

of them showing k heads. These expected values are shown in Table 2. The table also shows the expected counts. We can therefore compute the value of the Pearson Chi–Square statistic to be

$$\widehat{Q} = 21.57.$$

In this case, the Pearson Chi-Square statistic has an approximate $\chi^2(5)$ distribution since there are 6 categories. The p–value of the goodness of fit test is then, approximately,

$$p\text{–value} = \mathrm{P}(Q > \widehat{Q}) \approx 0.0006,$$

Category (k)	p_k	Predicted Counts	Observed Counts
0	0.03125	112.1875	100
1	0.15625	560.9375	524
2	0.31250	1121.875	1080
3	0.31250	1121.875	1126
4	0.15625	560.9375	655
5	0.03125	112.1875	105

Table 2: Counts Predicted by the Binomial Model

which is very small. Thus, we may reject the null hypothesis that the data in Table 1 follows a binomial distribution at the 1% significance level. Therefore, we can say that the data do not support the assumption that the five coins are fair. □

(b) Assume now that the coins have the same probability, p, of turning up heads. Estimate p and perform a goodness of fit test of the model you used to do your estimation. What do you conclude?

Solution: Suppose now that the coins are not fair but they all have the same probability, p, of turning up head. We can estimate p from the data as follows:

$$5 \cdot \widehat{p} = \frac{0 \cdot 100 + 1 \cdot 524 + 2 \cdot 1080 + 3 \cdot 1126 + 4 \cdot 655 + 5 \cdot 105}{3590},$$

from which we get that

$$\widehat{p} \approx 0.5129.$$

We now test the null hypothesis

$$\mathrm{H}_o\colon X \sim \text{binomial}(5, \widehat{p}).$$

In this case we get the expected counts shown in Table 3 on page 3. The Pearson Chi–Square statistic, Q, has the value $\widehat{Q} \approx 8.75$, and the approximate p–value is

$$p\text{–value} = \mathrm{P}(Q > \widehat{Q}) \approx 0.068,$$

since Q has an approximate $\chi^2(4)$ statistic in this case because we estimated p from the data. Thus, we cannot reject the null

Category (k)	p_k	Predicted Counts	Observed Counts
0	0.02742	98.443	100
1	0.14437	518.286	524
2	0.30403	1091.476	1080
3	0.32014	1149.288	1126
4	0.16855	605.081	655
5	0.03549	127.426	105

Table 3: Counts Predicted by the binomial$(5, \widehat{p})$ Model

hypothesis at the 5% significance level, but we could reject at the 10% level of significance. Hence, the data gives moderate support to the hypothesis that the are slightly loaded towards yielding more heads on average. □

2. In 1,000 tosses of a coin, 560 yield heads and 440 turn up tails. Is it reasonable to assume that the coin if fair? Justify your answer.

***Solution**:* Test the hypothesis

$$\mathrm{H}_o\colon p = \frac{1}{2}$$

versus the alternative

$$\mathrm{H}_1\colon p > \frac{1}{2}.$$

We model the tosses by a sequence of $n = 1000$ independent Bernoulli(p) trials, $X_1, X_2, \ldots, X_n$ and form the test statistic

$$Y = \sum_{j=1}^{n} X_j.$$

We reject the null hypothesis if

$$Y > 500 + c,$$

for certain critical value c, determined by the level of significance, α, of the test. In this case,

$$\alpha = \mathrm{P}(Y + 500 > c) \quad \text{for } Y \sim \text{binomial}(1000, 0.5).$$

Using the Central Limit Theorem, we have that

$$\alpha \approx \mathrm{P}\left(Z > \frac{c}{\sqrt{1000 \cdot (0.5)(1-0.5)}}\right),$$

where $Z \sim \text{normal}(0,1)$. Thus, if we let z_α denote a value such that

$$\mathrm{P}(Z > z_\alpha) = \alpha,$$

we have that we can reject H_o at the α significance level if

$$Y > 500 + z_\alpha \sqrt{1000/4}.$$

if $\alpha = 0.05$, z_α is the value of z which yields

$$F_{_Z}(z) = 1 - \alpha.$$

Thus, $z_\alpha = F_{_Z}^{-1}(0.95) \approx 1.65$. We will then reject the null hypothesis if

$$Y > 526.$$

In this case, the observed value of Y is $\widehat{Y} = 560$. Hence, we may reject the null hypothesis at the 5% level of significance and conclude that the data lend evidence to hypothesis that the coin is biased towards more heads. □

3. In a random sample, $X_1, X_2, \ldots, X_n$, of Bernoulli(p) random variables, it is desired to test the hypotheses $\mathrm{H}_o\colon p = 0.49$ versus $\mathrm{H}_1\colon p = 0.51$ Use the Central Limit Theorem to determine, approximately, the sample size, n, needed to have the probabilities of Type I error and Type II error to be both about 0.01. Explain your reasoning.

***Solution**:* We use $Y = \displaystyle\sum_{i=1}^{n} X_i$ as a test statistic. Put $p_o = 0.49$ and $p_1 = 0.51$, and define the rejection region

$$R\colon \quad Y > c,$$

where c is some critical. We then have that the probability of a Type I error is

$$\alpha = \mathrm{P}(Y > c), \quad \text{given that } \; Y \sim \text{binomial}(n, p_o).$$

Similarly, the probability of a Type II error is

$$\beta = \mathrm{P}(Y \leqslant c), \quad \text{given that } \ Y \sim \text{binomial}(n, p_1).$$

We approximate these errors using the Central Limit Theorem as follows:

$$\begin{aligned}
\alpha &= \mathrm{P}\left(\frac{Y - np_o}{\sqrt{np_o(1-p_o)}} > \frac{c - np_o}{\sqrt{np_o(1-p_o)}}\right) \\
&\approx \mathrm{P}\left(Z > \frac{c - np_o}{\sqrt{np_o(1-p_o)}}\right),
\end{aligned}$$

where $Z \sim \text{normal}(0,1)$. Thus, we set

$$\frac{c - np_o}{\sqrt{np_o(1-p_o)}} = z_\alpha, \tag{1}$$

where z_α is the real value with the property that $\mathrm{P}(Z > z_\alpha) = \alpha$.
For the probability of a Type II error we get

$$\begin{aligned}
\beta &= \mathrm{P}\left(\frac{Y - np_1}{\sqrt{np_1(1-p_1)}} \leqslant \frac{c - np_1}{\sqrt{np_1(1-p_1)}}\right) \\
&\approx \mathrm{P}\left(Z \leqslant \frac{c - np_1}{\sqrt{np_1(1-p_1)}}\right).
\end{aligned}$$

Thus, we may set

$$\frac{c - np_1}{\sqrt{np_1(1-p_1)}} = z_\beta, \tag{2}$$

where z_β is the real value with the property that $F_{_Z}(z_\beta) = \beta$.
For the case in which $\alpha = \beta = 0.01$, we have $z_\alpha \approx 2.33$ and $z_\beta \approx -2.33$. Equations (1) and (2) then become

$$c = 2.33\sqrt{np_o(1-p_o)} + np_o \tag{3}$$

and

$$c - np_1 = -2.33\sqrt{np_1(1-p_1)}. \tag{4}$$

Subtracting (4) from (3) leads to

$$n(p_1 - p_o) = 2.33\sqrt{n}\left(\sqrt{p_o(1-p_o)} + \sqrt{p_1(1-p_1)}\right),$$

which leads to

$$\sqrt{n} = \frac{2.33}{p_1 - p_o}\left(\sqrt{p_o(1-p_o)} + \sqrt{p_1(1-p_1)}\right). \qquad (5)$$

Substituting the values for p_o and p_1 in (5) we obtain

$$\sqrt{n} \approx 116.5,$$

so that we want n to be at least $13,567$. $\square$

4. Let $X_1, X_2, \ldots, X_n$ be a random sample from a normal$(\theta, 1)$ distribution. Suppose you want to test $\mathrm{H}_o\colon \theta = \theta_o$ versus $\mathrm{H}_1\colon \theta \neq \theta_o$, with the rejection region defined by $\sqrt{n}|\overline{X}_n - \theta_o| > c$, for some critical value c.

   (a) Find and expression in terms of standard normal probabilities for the power function of this test.

   Solution: The power function of this test, $\gamma(\theta)$ is the probability that the the test will reject the null hypothesis when $\theta \neq \theta_o$; that is,

$$\gamma(\theta) = \mathrm{P}\left(\sqrt{n}|\overline{X}_n - \theta_o| > c\right) \quad \text{given that } \overline{X}_n \sim \text{normal}(\theta, 1/n),$$

for $\theta \neq \theta_o$. Thus, we can write $\gamma(\theta)$ as

$$\begin{aligned}
\gamma(\theta) &= 1 - \mathrm{P}\left(|\overline{X}_n - \theta_o| \leqslant \frac{c}{\sqrt{n}}\right) \\
&= 1 - \mathrm{P}\left(\theta_o - \frac{c}{\sqrt{n}} < \overline{X}_n \leqslant \theta_o + \frac{c}{\sqrt{n}}\right) \\
&= 1 - \mathrm{P}\left(\theta_o - \theta - \frac{c}{\sqrt{n}} < \overline{X}_n - \theta \leqslant \theta_o - \theta + \frac{c}{\sqrt{n}}\right) \\
&= 1 - \mathrm{P}\left(\sqrt{n}(\theta_o - \theta) - c < \frac{\overline{X}_n - \theta}{1/\sqrt{n}} \leqslant \sqrt{n}(\theta_o - \theta) + c\right) \\
&= 1 - \mathrm{P}\left(\sqrt{n}(\theta_o - \theta) - c < Z \leqslant \sqrt{n}(\theta_o - \theta) + c\right),
\end{aligned}$$

where $Z \sim \text{normal}(0,1)$. We therefore have that

$$\gamma(\theta) = 1 - \left(F_Z(\sqrt{n}(\theta_o - \theta) + c) - F_Z(\sqrt{n}(\theta_o - \theta) - c)\right), \qquad (6)$$

where F_Z denotes the cdf of the standard normal distribution. $\square$

(b) An experimenter desires a Type I error probability of 0.04 and a maximum Type II error probability of 0.25 at $\theta = \theta_o + 1$. Find the values of n and c for which these conditions can be achieved.

***Solution**:* The probability of a Type I error is $\gamma(\theta_o)$ where $\gamma(\theta)$ is given in Equation (6). Thus,

$$\alpha = \gamma(\theta_o) = 1 - (F_z(c) - F_z(-c)) = 2 - 2F_z(c).$$

Thus, if $\alpha = 0.04$, we need to set c so that

$$F_z(c) = 0.98,$$

which yields

$$c \approx 2.05.$$

The probability of a Type II error for $\theta = \theta_o + 1$ is

$$\begin{aligned}
\beta &= 1 - \gamma(\theta_o + 1) \\
&= 1 - (1 - (F_z(-\sqrt{n} + c) - F_z(-\sqrt{n} - c))) \\
&= F_z(-\sqrt{n} + c) - F_z(-\sqrt{n} - c) \\
&= \mathrm{P}(-\sqrt{n} - c < Z \leqslant -\sqrt{n} + c) \\
&\leqslant \mathrm{P}(-\infty < Z \leqslant -\sqrt{n} + c) \\
&= F_z(-\sqrt{n} + c).
\end{aligned}$$

Thus, in order to make $\beta \leqslant 0.25$, we require that

$$F_z(-\sqrt{n} + c) = 0.25.$$

This yields

$$-\sqrt{n} + c \approx -0.675.$$

Thus,

$$n \approx (c + 0.675)^2 \approx 7.43.$$

Thus, we may take n to be at least 8. □

5. Let $X_1, X_2, \ldots, X_n$ be a random sample from a normal(θ, σ^2) distribution. Suppose you want to test

$$\mathrm{H}_o\colon \theta \leqslant \theta_o$$

versus

$$\mathrm{H}_1 : \theta > \theta_1$$

with the rejection region defined by

$$T_n(\theta) > \frac{\sqrt{n}}{S_n}(\theta_o - \theta) + c,$$

for some critical value c. Here, $T_n(\theta)$ is the statistic

$$T_n(\theta) = \frac{\sqrt{n}(\overline{X}_n - \theta)}{S_n},$$

where $\overline{X}_n$ and S_n^2 are the sample mean and variance, respectively.

(a) If the significance level for the test is to be set at α, what should c be?

Solution: The power function of this test is

$$\gamma(\theta) = \mathrm{P}_\theta\left(T_n(\theta) > \frac{\sqrt{n}}{S_n}(\theta_o - \theta) + c\right),$$

where $T_n(\theta) \sim t(n-1)$; that is, $T_n(\theta)$ has a t distribution with $n-1$ degrees of freedom.
Observe that, if the null hypothesis is true, then $\theta \leqslant \theta_o$ and therefore

$$\frac{\sqrt{n}}{S_n}(\theta_o - \theta) + c \geqslant c$$

for all $\theta \leqslant \theta_o$. It then follows that

$$\mathrm{P}_\theta\left(T_n(\theta) > \frac{\sqrt{n}}{S_n}(\theta_o - \theta) + c\right) \leqslant \mathrm{P}(T_n(\theta) > c).$$

Thus,

$$\alpha = \sup_{\theta \leqslant \theta_o} \gamma(\theta) = \mathrm{P}(T_n(\theta) > c).$$

where $T_n(\theta) \sim t(n-1)$. Thus, to choose c, we find a real value, t, such that

$$\mathrm{P}(T > t) = \alpha, \quad \text{where } T \sim t(n-1).$$

Denoting that value by $t_{\alpha,n-1}$, we get that

$$c = t_{\alpha,n-1}.$$

□

(b) Express the rejection region in terms of the value c found in part (a), and the statistics $\overline{X}_n$ and S_n^2.

Solution: The rejection region is

$$\frac{\sqrt{n}(\overline{X}_n - \theta)}{S_n} > \frac{\sqrt{n}}{S_n}(\theta_o - \theta) + t_{\alpha,n-1},$$

which can be re-written as

$$\overline{X}_n > \theta_o + t_{\alpha,n-1}\,\frac{S_n}{\sqrt{n}}.$$

□

(c) Compute the power function, $\gamma(\theta)$, for the test.

Solution: From part (a) of this problem we have that

$$\begin{aligned}\gamma(\theta) &= \mathrm{P}_\theta\left(T_n(\theta) > \frac{\sqrt{n}}{S_n}(\theta_o - \theta) + t_{\alpha,n-1}\right)\\ &= 1 - \mathrm{P}_\theta\left(T \leqslant \frac{\sqrt{n}}{S_n}(\theta_o - \theta) + t_{\alpha,n-1}\right),\end{aligned}$$

where $T \sim t(n-1)$. Hence, the power function of the test is

$$\gamma(\theta) = 1 - F_{_T}\left(\frac{\sqrt{n}}{S_n}(\theta_o - \theta) + t_{\alpha,n-1}\right)$$

for $\theta > \theta_1$. □

6. A sample of 16 "10–ounce" cereal boxes has a mean weight of 10.4 oz and a standard deviation of 0.85 oz. Perform an appropriate test to determine whether, on average, the "10–ounce" cereal boxes weigh something other than 10 ounces at the $\alpha = 0.05$ significance level. Explain your reasoning.

Solution: We assume that the weight in each "10–ounce" cereal box follows a normal(μ, σ^2) distribution with mean μ and variance σ^2. We would like test the hypothesis

$$\mathrm{H}_o: \quad \mu = 10 \text{ oz}$$

against the alternative hypothesis

$$\mathrm{H}_1: \quad \mu \neq 10 \text{ oz}.$$

We consider the rejection region

$$R : |\overline{X}_n - \mu_o| > t_{\alpha/2,n-1} \frac{S_n}{\sqrt{n}}$$

where $\mu_o = 10$ oz, and $t_{\alpha/2,n-1}$ is chosen so that

$$\mathrm{P}(|T| > t_{\alpha/2,n-1}) = \alpha,$$

for $T \sim t(n-1)$. Then, if H_o is true, the statistic

$$T_n = \frac{\overline{X}_n - \mu_o}{S_n/\sqrt{n}},$$

where $n = 16$, has a $t(n-1)$ distribution, since we are assuming the the sample, $X_1, X_2, \ldots, X_n$, comes from a normal(μ_o, σ^2) distribution. Consequently, the test has significance level α. In the special case in which $\alpha = 0.05$, we get that $t_{\alpha/2,n-1} \approx 2.13$. Thus, the null hypothesis can be rejected at the 0.05 significance level if

$$|\overline{X}_n - \mu_o| > 2.13 \cdot \frac{S_n}{\sqrt{16}}.$$

In this problem, $\overline{X}_n = 10.4$, $S_n = 0.85$, and $n = 16$. We then have that

$$\frac{|\overline{X}_n - \mu_o|}{S_n/\sqrt{n}} \approx 1.88,$$

which is not bigger than 2.13, thus we cannot reject the null hypothesis at the 0.05 significance level. $\square$

7. Find the p–value of observed data consisting of 7 successes in 10 Bernoulli(θ) trials in a test of

$$\mathrm{H}_o\colon \theta = \frac{1}{2} \quad \text{versus} \quad \mathrm{H}_1\colon \theta > \frac{1}{2}.$$

***Solution**:* Let Y denote the number of successes in the $n = 10$ trials. Then $Y \sim \text{binomial}(10, \theta)$. This is the test statistic. The p–value is the probability that, if the null hypothesis is true, we will see the observed value of the statistic or more extreme ones. In this case, if the null hypothesis is true, $Y \sim \text{binomial}(10, 0.5)$ and the p–value is

$$\begin{aligned} p\text{–value} &= \mathrm{P}(Y \geqslant 7) \\ &= \sum_{k=7}^{10} \binom{10}{k} \frac{1}{2^{10}} \\ &\approx 0.1719 \end{aligned}$$

$\square$

8. Three independent observations from a Poisson(λ) distribution yield the values $x_1 = 3$, $x_2 = 5$ and $x_3 = 1$. Explain how you would use these data to test the hypothesis $\text{H}_o\colon \lambda = 1$ versus the alternative $\text{H}_1\colon \lambda > 1$. Come up with an appropriate statistic and rejection criterion and determine the p–value given by the data. What do you conclude?

***Solution**:* Denote the observations by X_1, X_2, X_3. Then, X_1, X_2 and X_3 are independent Poisson(λ) random variables. Define the test statistic

$$Y = X_1 + X_2 + X_3.$$

Then, $Y \sim \text{Poisson}(3\lambda)$. The p–value is the probability that the test statistic will take on the observed value, or more extreme ones, under the assumption that H_o is true; that is, $Y \sim \text{Poisson}(3)$. Thus,

$$\begin{aligned} p\text{–value} &= \text{P}(Y \geqslant 9) \\ &= 1 - \text{P}(Y \leqslant 8) \\ &= 1 - \sum_{k=0}^{8} \frac{3^k}{k!}\, e^{-3} \\ &\approx 0.0038. \end{aligned}$$

A rejection region is determined by the significance level that we set. For instance, if the significance level is α, then we can have the rejection criterion

$$p\text{–value} < \alpha \Rightarrow \text{Reject H}_o.$$

Thus, in this case, we can reject H_o at the $\alpha = 0.01$ significance level, and conclude that the data support the hypothesis that $\lambda > 1$. $\square$