

STA 312f22 Assignment Three¹

Please bring your complete R input and output from Question 5 to the quiz. The non-computer questions are practice for the quiz on Friday Oct. 7th, and are not to be handed in. On the quiz, you will be given a copy of the [formula sheet](#), so please use the formula sheet while doing your homework, and let me know if there are problems. There is a link to the formula sheet from the [course home page](#), in case the one in this document does not work.

1. For a random sample of size n from a multinomial distribution $M(1, (\pi_1, \dots, \pi_c))$,
 - (a) What is the probability distribution of the frequency n_j ? It has a familiar name. Just write down the answer.
 - (b) What is the expected value of the frequency n_j ? Just write down the answer.
 - (c) Let $\hat{\pi}_1, \dots, \hat{\pi}_c$ be the estimated probabilities given some null hypothesis. What is the *estimated* expected value of the frequency n_j ? Just write down the answer.
2. The formula sheet has two formulas for the likelihood ratio test statistic G^2 . The one on the left is general, and the one on the right is specific to the multinomial model. Show how to get from $G^2 = -2 \log \left(\frac{\max_{\beta \in \mathcal{D}_0} \ell(\beta)}{\max_{\beta \in \mathcal{D}} \ell(\beta)} \right)$ to $G^2 = 2 \sum_{j=1}^c n_j \log \left(\frac{n_j}{\hat{\mu}_j} \right)$.
3. You must admit that the Bernoulli is a special case of the multinomial. Based on a random sample of size n from a Bernoulli, you want to test $H_0 : \pi = \pi_0$.
 - (a) Obtain a special-case formula for the likelihood ratio statistic G^2 . The answer is something you can compute based on π_0 and the frequencies.
 - (b) Obtain a special-case formula for the Pearson chi-squared statistic X^2 . The answer is something you can compute based on π_0 and the frequencies.
 - (c) What are the degrees of freedom? The answer is a number.
4. Here is a repeat of the question from the last assignment. This time, you are asked to do it with technology from the multinomial.

For years, brand awareness for Big Red chewing gum has been stuck at about 6%, meaning that about 6% of consumers who chew gum say they remember hearing about Big Red gum. The company started an advertising campaign to increase brand awareness, and after it had been running for a few weeks, they interviewed a random sample of 200 gum chewers, and found that twenty had heard of Big Red.

- (a) State the model and the null hypothesis, or look at your answer from last week.
- (b) The Pearson X^2 and large-sample likelihood test G^2 have the same critical value. What is it? The answer is a number.

¹Copyright information is at the end of the last page.

- (c) Carry out the large-sample likelihood test G^2 . Use the formula sheet and a calculator.
- What is the value of the test statistic? The answer is a number.
 - Do you reject H_0 at $\alpha = 0.05$? Answer Yes or No.
 - State your conclusion in plain, non-statistical language. For a directional conclusion, compare the observed and (estimated) expected frequencies.
- (d) Carry out the Pearson chi-squared test X^2 . Use the formula sheet and a calculator.
- What is the value of the test statistic? The answer is a number.
 - Do you reject H_0 at $\alpha = 0.05$? Answer Yes or No.
 - State your conclusion in plain, non-statistical language.
5. Under carefully controlled conditions, 120 beer drinkers tasted 6 beers and indicated which one they liked best. Here are the numbers preferring each beer.

	Preferred Beer					
	1	2	3	4	5	6
Frequency	30	24	22	28	9	7

The main question is whether preference for the 6 beers is different in the population from which this sample was taken. Use R whenever possible. *You may be asked to hand in your printout of the R parts, so please print this R session on a separate sheet of paper.*

- State a reasonable model for these data.
- What is the parameter space \mathcal{B} ?
- State the null hypothesis in symbols. It is a statement about the π_j s. Please be specific. The research question allows you to give a specific numerical value for each π_j under H_0 .
- What is the restricted parameter space \mathcal{B}_0 ? How many points are in \mathcal{B}_0 ?
- What are the degrees of freedom of the test? The answer is a number.
- What are the expected frequencies under H_0 ? Your answer is a set of 6 numbers. Are these estimated expected values, or exact expected values?
- Calculate the likelihood ratio test statistic G^2 using R.
- Obtain the critical value at $\alpha = 0.05$? with R. Compare the number on the formula sheet.
- Calculate the p -value using R. Print out all the R input and output, and bring it to the quiz.
- Do you reject the null hypothesis at $\alpha = 0.05$? Answer Yes or No.

- (k) It is tempting to ask you to state your conclusion in words. But all you can conclude without further testing is that preference for all the beers is not equal. It *looks* like preference for beers 1 through 4 is greater than preference for 5 and 6, and this is what you would tell your management or client in a job situation.
- (l) Calculate the Pearson chi-square statistic X^2 for these data, using R .
- (m) Do you reject the null hypothesis at $\alpha = 0.05$? Answer Yes or No. Just so you can check your answer to this question, my p -value for X^2 is 0.0002479085.
6. In Question 5, you found that preference for the 6 beers was not the same. But it seems that the first 4 beers are lagers and the last two are ales. No one would expect preference for lagers and ales to be the same.² So let's test whether preference for the 4 lagers is different, and at the same time, whether preference for the 2 ales is different.
- (a) State the null hypothesis in symbols. It is a statement about the π_j s.
- (b) What are the degrees of freedom of the test? The answer is a number.
- (c) Differentiating the log likelihood function, obtain the maximum likelihood estimator of the parameter $\boldsymbol{\pi}$, under the null hypothesis. Show all your work. The answer is a symbolic expression, a vector of length 6. If it is not absolutely the most natural thing you can imagine, then either it's wrong or you have not simplified enough.
- (d) Give the maximum likelihood estimate of $\boldsymbol{\pi}$ under the null hypothesis for this particular data set. The answer is a set of 6 numbers. Note the difference between an estimator and an estimate.
- (e) What are the expected frequencies under H_0 ? Your answer is a set of 6 numbers. Are these estimated expected values, or exact expected values?
- (f) Using the formula sheet, calculate the likelihood ratio test statistic G^2 . Show your work. Your answer is a number. This is something you should be able to do with a calculator if necessary on the quiz. You would be silly not to check it with R .
- (g) Calculate the Pearson X^2 statistic for these data. The answer is a number.
- (h) What is the critical value for this test at $\alpha = 0.05$? The answer is a number.
- (i) Do you reject the null hypothesis at $\alpha = 0.05$ based on the likelihood ratio test? Answer Yes or No.
- (j) Do you reject the null hypothesis at $\alpha = 0.05$ based on the Pearson Chisquare test? Answer Yes or No.
- (k) Based on this analysis, are you able to conclude that there is any difference in preference among the 4 lagers or between the 2 ales? Answer Yes or No. (In applied situations, beware of concluding that there is no effect — that is, *absolutely no difference at all* in the population.)

²Actually, I am making all this up with only a vague idea of what these terms mean.

Just so you can check your answer to this question, my p -value for the likelihood ratio test is 0.7735209.

7. A sample of 150 students each try to solve two difficult logic problems. The problems are complicated, but it's multiple choice, so the answers are either right or wrong. For each student, the data file indicates whether he or she got Question 1 right, and whether he or she got Question 2 right. It is expected that students who get one question correct will also tend to get the other correct; this is not the issue. The issue is whether the one question is more difficult than the other.

At first glance, this seems like just a problem of comparing two proportions, but there is a twist. Each student answers *both* questions. In a multinomial model, there are n *independent* observations, so each case (person or whatever) must contribute only one frequency to the table. This experiment is an example of a within-cases (repeated measures) design, but with categorical data.

We will set up the problem as a multinomial with four categories, as follows.

	Question 1 Correct	Question 1 Incorrect
Question 2 Correct	π_1	π_2
Question 2 Incorrect	π_3	π_4

The observed frequencies are:

	Question 1 Correct	Question 1 Incorrect
Question 2 Correct	66	41
Question 2 Incorrect	30	13

Note that even though these look like 2-dimensional tables, we have not reached contingency tables yet in the assignments. This is a one-dimensional multinomial.

- State the null hypothesis in symbols. It is a statement about the π_j s. Simplify.
- What are the degrees of freedom for this test?
- Differentiating the log likelihood function, obtain the maximum likelihood estimator of the parameter $\boldsymbol{\pi}$, under the null hypothesis. Show all your work. The answer is a symbolic expression, a vector of length 4.
- Give the maximum likelihood estimate of $\boldsymbol{\pi}$ under the null hypothesis for this particular data set. The answer is a set of 4 numbers.
- What are the expected frequencies under H_0 ? Your answer is a set of 4 numbers. Are these estimated expected values, or exact expected values?
- Using the formula sheet, calculate the likelihood ratio test statistic G^2 . Show your work. Your answer is a number. This is something you should be able to do with a calculator if necessary on the quiz.

- (g) Calculate the Pearson X^2 statistic for these data. The answer is a number. This is something you should be able to do with a calculator if necessary on the quiz.
- (h) What is the critical value for this test at $\alpha = 0.05$? The answer is a number.
- (i) Do you reject the null hypothesis at $\alpha = 0.05$ based on the likelihood ratio test? Answer Yes or No.
- (j) Do you reject the null hypothesis at $\alpha = 0.05$ based on the Pearson Chisquare test? Answer Yes or No.
- (k) What percent of students answered Question 1 correctly? The answer is a number.
- (l) What percent of students answered Question 2 correctly? The answer is a number.
- (m) Does this study provide solid evidence that the two questions differ in their difficulty? Answer Yes or No.

Please bring hard copy of your R input and output for Question 5 to the quiz.

This assignment was prepared by [Jerry Brunner](#), Department of Statistics, University of Toronto. It is licensed under a [Creative Commons Attribution - ShareAlike 3.0 Unported License](#). Use any part of it as you like and share the result freely. The L^AT_EX source code is available from the course website: <http://www.utstat.toronto.edu/~brunner/oldclass/312f22>