Applet #4 Activities

These activities accompany an applet located at

`http://www.zoology.ubc.ca/~whitlock/kingfisher/ContingencyAnalysis.htm`

The applet aims to assist understanding of the following concepts:

- Association between categorical variables.
- The chi-squared test for independence between categorical variables, including the sampling distribution of the test statistic.
- The P-value of a hypothesis test.
- Type I error in hypothesis testing.
- Type II error and the power of a hypothesis test.
- How the sample size can affect the power of the chi-squared test.
- How the population parameters can affect the power of the chi-squared test.

The applet is based on an imaginary population, which we will take to be the adult population of Statland. Suppose in Statland each infant may be vaccinated against an illness. Some adults were vaccinated as children, some were not, and so there are two subpopulations: those adults that were vaccinated and those that were not. We define an adult to be "healthy" if they have never had the illness which the vaccination targets; otherwise an adult is defined to be "sick" for our purposes. The question is, are the proportions of healthy and unhealthy adults the same within the two subpopulations?

Go to the applet linked above. The table in the top left corner gives the population parameters of interest which you will adjust later. For now, appreciate for example that within the adult population of Statland who were vaccinated as children, 40% are healthy in that they have never had the illness targeted by the vaccine whereas 60% have had the illness.

1. With the parameters set to their initial values, is there a dependence between vaccination status and health in this population?

   Yes   No   Not sure

2. The slider labelled "Choose the total sample size" allows us to select the sample size n chosen from the entire population sampling evenly from the two subpopulations (vaccinated and unvaccinated). Keeping $n = 120$, so that we sample 60 from the vaccinated and 60 from the unvaccinated, what sort of values would you expect to see in the contingency table below? Fill in values in the table below that you think are likely, making sure your column totals each add to 60.

		Vaccination Status	
		Vaccine	No vaccine
Health	Sick		
	Healthy		

3. Click the "Make one sample" button to see a sample table from this population. Compare the table with the one you predicted. Are they the same? If not, are they very different from each other?

4. The chi-squared test for independence in this context computes a measure of how different the observed values in a table are from what we would expect to see if the two categorical variables are independent. The test statistic χ^2 involves the differences in the counts observed and what we would expect under the null hypothesis of independence, H_0. A large value of χ^2 could lead us to reject the null hypothesis, but how large?
   Directly under the table you created there is the χ^2 value for your table, and also the probability of seeing a value at least as big if the null hypothesis is true. This probability is known as the *P-value* of the test. If the P-value is small, that means we have seen a surprisingly large value of χ^2 if the null hypothesis holds. But how small should the P-value be in order to reject H_0? We are free to choose this threshold, and here will take the common preference of 0.05, this defining the *significance level* of the test. In the table you observed, was the null hypothesis rejected?

   Yes No Not sure

5. Repeat the above, by clicking "Make one sample" again. Note the graphic below the table shows the sampling distribution of χ^2 if H_0 is true. The P-value is the area under the curve to the right of the observed test statistic. In the new table you observed, was the null hypothesis rejected?

   Yes No Not sure

6. Suppose you were to sample fifty tables like the one above, and compute the χ^2 statistic and the P-value for each one. About how many of the tables would you expect to lead to H_0 being rejected?

7. We can repeat this sampling procedure multiple times, to see how the test statistic and P-value vary. Click the "Simulate distribution" button until fifty sample tables have been taken. (The "Faster" button speeds up the rate of simulations, but do not use that for now so that you can look carefully at the simulations to ensure you understand what you are seeing.) In how many of these tables is the null hypothesis of independence rejected? That is, how many times is a *Type I error* (incorrectly rejecting the null hypothesis) observed?

8. Now change the parameters in the table to be as follows:

	Vaccine	No vaccine
P(Sick)	0.4	0.6
P(Healthy)	0.6	0.4

   Now is there an association between vaccination status and health?

   Yes No Not sure

   Fill in values in the table below that you think are likely now, making sure your column totals each add to 60.

		Vaccination Status	
		Vaccine	No vaccine
Health	Sick		
	Healthy		

9. Click the “Make one sample” button to see a sample table from this population. Compare the table with the one you predicted.

10. We can repeat this sampling procedure multiple times, to see how the test statistic and P-value vary. Click the “Simulate distribution” button until fifty sample tables have been taken. A *Type II* error occurs when a hypothesis test fails to reject a null hypothesis that is false. In how many of these fifty tables is the null hypothesis of independence rejected? That is, how many times is a *Type II error* avoided?

11. The probability that a test will correctly reject a null hypothesis that is false is the *power* of the test. We will explore how the power of the test here varies with the parameter values. Change the parameters in the table to be as follows:

	Vaccine	No vaccine
P(Sick)	0.4	0.7
P(Healthy)	0.6	0.3

How will the power change? Click the “Simulate distribution” button until fifty sample tables have been taken. What, approximately, is the power of the test now?

12. How if at all does the power change as n changes? Reduce the sample size from each subpopulation from 60 to 15 by changing n to 30, but keeping the parameter values unchanged. What impact will this have on the power of the test?

Increase Stay the same Decrease

13. Click the "Simulate distribution" button until fifty sample tables have been taken. What, approximately, is the power of the test now?

14. Keeping the sample size as 15, modify the parameter values to the following:

	Vaccine	No vaccine
P(Sick)	0.3	0.7
P(Healthy)	0.7	0.3

What impact will this have on the power of the test?

Increase Stay the same Decrease

15. Click the "Simulate distribution" button until fifty sample tables have been taken. What, approximately, is the power of the test now?

16. Now change the sample size n back to 120. What impact will this have on the power of the test?

Increase Stay the same Decrease