INFORMATION POINT:

Chi-Squared Test

The chi-squared (χ^2) test is one of the most widely used statistical tests. It can be used in several different ways, but most commonly is used to test that there is no association between two categorical variables. By 'no association' we mean that for an individual, the response for variable 1 is not affected in any way by the response for variable 2, that is the two variables are independent.

The way the test works is to calculate the table of responses we would expect if it was true that there is no association between the two variables under consideration. The values in this table are called the expected values. The expected values are then compared to the values we actually observed in our data, the observed values. If the observed and expected values are similar, then the hypothesis of no association seems plausible and is retained. If the observed and expected values are not similar then we claim there is an association between the two variables.

The chi-squared test assumes that both variables are nominal: that is, the variable has categorical responses and the categories are unordered. In carrying out this test it is important to note that all of the expected values should be greater than 5. If this is not the situation, than categories of one or both variables need to be combined in order for the test to be valid. The authors of the preceding paper combined the category 'agreed very much' with 'agreed a little' and they combined 'disagreed very strongly' with 'disagreed a little' in order to make sure the expected values were large enough.

The chi-squared statistic is calculated directly from the difference between the observed and expected values and a characteristic called the degrees of freedom (d.f.) is also calculated. This relates to the number of rows and columns in the table we are considering. Whether the observed and expected values are similar enough to be able to claim no association is measured by the *P*-value, which is calculated from the chi-squared statistic and its degrees of freedom. A commonly adopted convention is to reject the hypothesis of no association between the two variables if the *P*-value is less than 0.05 ($P < 0.05$).

In the study above the authors consider whether there is an association between the type of unit attended (which has three categories) and whether the patient was less worried about their health after they had been seen (which has two categories). Comparison of the observed and expected values results in a chi-squared statistic of 6.24 with 2 degrees of freedom, leading to a *P*-value of 0.044. The *P*-value is less than 0.05, so the authors conclude that there is a significant association between the type of unit visited and whether the patient was less worried after the visit.

For a description of how the calculations are actually carried out see Altman (1992, pp. 242–248) or Bland (1995, pp. 225–231).

Further reading

Altman D.G. (1992) *Practical Statistics for Medical Research*. Chapman & Hall, London.
Bland M. (1995) *An Introduction to Medical Statistics* 2nd edn. Oxford University Press, Oxford.

NICOLA TURNER