

Chi-Square Tests in R

R easily computes Chi-Square goodness of fit tests for any supplied matrix.

Reminder: R and matrices

Recall that the basic R command for a matrix is to start with a list of elements (`c(.,.,., .,)`) and the number of rows `nrow` in the matrix. For example, if we have the 2×4 contingency table

| | | | |
|---|----|----|----|
| 9 | 15 | 5 | 15 |
| 4 | 12 | 10 | 4 |

we express this as a matrix (which we will call `C`) with the command

```
> C<-matrix(c(9,4,15,12,5,10,15,4),nrow=2)
```

It's always a good idea to check to see that you have correctly enter the matrix. Typing

```
> C
```

R returns

```
  [,1] [,2] [,3] [,4]
[1,]  9 15  5 15
[2,]  4 12 10  4
```

To perform Pearson's Chi-squared test, simply put the contingency table data into a matrix (here we have used `C` and use the command

```
> chisq.test(C)
```

For the above table, R returns:

Pearson's Chi-squared test

```
data:  C
X-squared = 7.9266, df = 3, p-value = 0.04755
```

Note, however, that some of the cells of the table have values of 5 or less. In such cases, Fisher's exact test is better.

```
> fisher.test(C)
```

R returns

Fisher's Exact Test for Count Data

```
data:  C
p-value = 0.05133
alternative hypothesis: two.sided
```