

4.2 Permutation Tests

We start with some samples from a bivariate t-distribution, where the variables are uncorrelated and the number of degrees of freedom is three.

The data in the file are 20 uncorrelated (x, y) pairs, followed by 20 correlated pairs.

In Figure 1 we see the data; this is a clear case where we need to be careful of an outlier driving a correlation.




Figure 1: Uncorrelated (x, y) pairs from a t-distribution.

For these data, values for the Kendall and Spearman rank correlation coefficients and the Pearson / Fisher r are -0.06, -0.08 and 0.49 respectively. The parametric test is going to be affected by the outlier.

We can do a permutation test by randomly assigning x s to y s to make new pairs; this should give the range of values of the test statistic which are consistent with there being no correlation. There are $20!$ distinct permutations for even this little dataset, so the results that follow picked only 1000 of these at random. This is achieved by sampling without replacement from the set of 20 x s and assigning each one in order to the set of y s.

From the results some cumulants can be estimated, as shown in Figures 2, 3 and 4.




Figure 2: Distribution of the Kendall statistic.

From these we can see that the r -statistic would be interpreted as having quite a high degree of significance (90%).

Now we turn our attention to a correlated set of data, again with outliers – Figure 5.




Figure 3: Distribution of the Spearman statistic.




Figure 4: Distribution of the Fisher / Pearson statistic.

We will just look at Spearman's statistic, which takes the value 0.77 for these data. (Kendall's statistic is 0.6 and $r = 0.96$. A bootstrap, with 1000 sampling-with-replacement operations, gives the cumulant in Figure 6.

Clearly values of the statistic near zero are very unlikely indeed and so the correlation is convincing. Finally, we can estimate the standard deviation on the statistic using the jackknife. Using Equation 6.20, we find 0.11 – with a mean of 0.77, to the extent that normality applies to the jackknifed estimates, we can see again that values near zero are very unlikely. The jackknife is quick, but for a sample this size a bootstrap is virtually instantaneous as well and gives more information. In this case, for instance, we can see that the distribution of the Spearman statistic is not very Gaussian – there is a pronounced tail to small values.




Figure 5: Correlated (x, y) pairs from a t-distribution.




Figure 6: Bootstrap distribution for the Spearman coefficient.