5.4 The F test

The data for this example are arranged in a similar way to those for Example 5.3. There are 100 sets of two pairs of 10 samples, but the first set of each pair is drawn from a Gaussian of standard deviation unity and the second of each pair comes from a distribution with standard deviation 1.5. The first 200 data have contamination at the 10% level, the second at the 20% level, and so on.

Compute the empirical distribution of the F statistic, looking at what happens when the standard deviations are the same and the contamination is varied, and contrast with the sets where there is a difference in the population standard deviations. Work out the significance level for the hypothesis that the variance ratio is 1.5^2 or bigger. Take one-tailed values. It may be simplest to work out the empirical distribution of F and then sketch significance levels from tables onto the derived distributions. Use the contamination levels to see that comparing sets with, say, 10% contamination yields higher significance levels than 40%. The outliers weaken the ability of the test to detect intrinsic differences in width.