The Ljung Box test

It is used to test the linear dependence between the current return and past return values; in other words, to assess the presence of autocorrelation in the return series.

If we denote the lag k correlation coefficient between the current return (at T) and the return at T-k by r_k, the Ljung-Box test (also called Portmanteau test) is calculated as:

$$Q = T(T+2)\sum_{k=1}^{s} r_k^2/(T-k)$$

where s is the number of lags taken into consideration when testing for autocorrelation, and T is the number of observations. The statistic tests whether any of a group of s autocorrelations of a time series are different from zero. As a consequence, the choice of s (number of lags of interest) – will affect the result[1].

Hypothesis testing:

The Null Hypothesis is that *none* of the autocorrelation coefficients up to lag s are different from zero. If the sample value of Q exceeds the critical value of a chi-square distribution with s degrees of freedom, then at least one of the k autocorrelations is statistically different from zero at the specified significance level[2].

In practice:

For a chosen significance level of 5%, the critical values of the chi-square distribution are the following, as function of the number of lags:

No. of lags	Critical value
1	3.84
2	5.99
3	7.82
4	9.49
5	11.07
6	12.59
7	14.07
8	15.51
9	16.92
10	18.31

Interpretation: If the calculated statistic exceeds the tabulated critical value at any lag, we can infer that returns are predicted by their past values up to that lag, in 95% of the cases.

[1] Some authors suggest that the choice of $s = \ln(T)$ provides the best power of the test.
[2] This is a one-tailed test.