regression (7.45), we either need to drop the first p observations or replace the unobserved lagged values of $\tilde{u}_t$ with zeros.

If we wish to test against an MA(q) process, it turns out that we can proceed exactly as if we were testing against an AR(q) process. The reason is that an autoregressive process of any order is **locally equivalent** to a moving-average process of the same order. Intuitively, this means that, for large samples, an AR(q) process and an MA(q) process look the same in the neighborhood of the null hypothesis of no serial correlation. Since tests based on the GNR use information on first derivatives only, it should not be surprising that the GNRs used for testing against both alternatives turn out to be identical; see Exercise 7.7.

The use of the GNR (7.43) for testing against AR(1) errors was first suggested by Durbin (1970). Breusch (1978) and Godfrey (1978a, 1978b) subsequently showed how to use GNRs to test against AR(p) and MA(q) errors. For a more detailed treatment of these and related procedures, see Godfrey (1988).

Older, Less Widely Applicable, Tests

Readers should be warned at once that the tests we are about to discuss are not recommended for general use. However, they still appear often enough in current literature and in current econometrics software for it to be necessary that practicing econometricians be familiar with them. Besides, studying them reveals some interesting aspects of models with serially correlated errors.

To begin with, consider the simple regression

$$\tilde{u}_t = b_\rho \tilde{u}_{t-1} + \text{residual}, \quad t = 1, \ldots, n, \tag{7.46}$$

where, as above, the $\tilde{u}_t$ are the residuals from regression (7.42). In order to be able to keep the first observation, we assume that $\tilde{u}_0 = 0$. This regression yields an estimate of b_ρ, which we will call $\tilde{\rho}$ because it is an estimate of ρ based on the residuals under the null. Explicitly, we have

$$\tilde{\rho} = \frac{n^{-1}\sum_{t=1}^{n} \tilde{u}_t \tilde{u}_{t-1}}{n^{-1}\sum_{t=1}^{n} \tilde{u}_{t-1}^2}, \tag{7.47}$$

where we have divided numerator and denominator by n for the purposes of the asymptotic analysis to follow. It turns out that, if the explanatory variables $\boldsymbol{X}$ in (7.42) are all exogenous, then $\tilde{\rho}$ is a consistent estimator of the parameter ρ in model (7.40), or, equivalently, (7.41), where it is not assumed that $\rho = 0$. This slightly surprising result depends crucially on the assumption of exogenous regressors. If one of the variables in $\boldsymbol{X}$ is a lagged dependent variable, the result no longer holds.

Asymptotically, it makes no difference if we replace the sum in the denominator by $n^{-1}\sum_{t=1}^{n} \tilde{u}_t^2$, because we are effectively just replacing the term $\tilde{u}_0^2$ by the term $\tilde{u}_n^2$. Then we can write the denominator of (7.47) as $n^{-1}\boldsymbol{u}^\top \boldsymbol{M_X} \boldsymbol{u}$,