expression (7.78). In contrast, the regression function for the restricted model, H_1, has 6 parameters: β_1 through β_5, and ρ. Therefore, in this example, H_1 imposes just one restriction on H_2.

The phenomenon illustrated in this example arises, to a greater or lesser extent, for almost every model with common factor restrictions. Constant terms, many types of dummy variables (notably, seasonal dummies and time trends), lagged dependent variables, and independent variables that appear with more than one time subscript always lead to an unrestricted model H_2 with some parameters that cannot be identified. Therefore, the number of identifiable parameters is almost always less than $2k+1$, and, in consequence, the number of restrictions is almost always less than k.

Testing Common Factor Restrictions

Any of the techniques discussed in Sections 6.7 and 6.8 can be used to test common factor restrictions. In practice, if the error terms are believed to be homoskedastic, the easiest approach is probably to use an asymptotic F test. For the example of equations (7.73) and (7.74), the restricted sum of squared residuals, RSSR, is obtained from NLS estimation of H_1, and the unrestricted one, USSR, is obtained from OLS estimation of H_2. Then the test statistic is

$$\frac{(\text{RSSR} - \text{USSR})/r}{\text{USSR}/(n-k-r-2)} \overset{a}{\sim} F(r, n-k-r-2), \tag{7.80}$$

where r $(\leq k)$ is the number of restrictions. The number of degrees of freedom in the denominator reflects the fact that the unrestricted model has $k+r+1$ parameters and is estimated using the $n-1$ observations for $t = 2, \ldots, n$.

Of course, since both the null and alternative models involve lagged dependent variables, the test statistic (7.80) does not actually follow the $F(r, n-k-r-2)$ distribution in finite samples. Therefore, when the sample size is not large, it is a good idea to bootstrap the test. As Davidson and MacKinnon (1999a) have shown, highly reliable P values may be obtained in this way, even for very small sample sizes. The bootstrap samples are generated recursively from the restricted model, H_1, using the NLS estimates of that model. As with bootstrap tests for serial correlation, the bootstrap error terms may either be drawn from the normal distribution or obtained by resampling the rescaled NLS residuals; see the discussion in Sections 4.6 and 7.7.

Although this bootstrap procedure is conceptually simple, it may be quite expensive to compute, because the nonlinear model (7.73) must be estimated for every bootstrap sample. It may therefore be more attractive to follow the idea in Exercises 6.17 and 6.18 by bootstrapping a GNR-based test statistic that requires no nonlinear estimation at all. For the H_1 model (7.73), the corresponding GNR is (7.63), but now we wish to evaluate it, not at the NLS estimates from (7.73), but at the estimates $\acute{\boldsymbol{\beta}}$ and $\acute{\rho}$ obtained by estimating the linear H_2 model (7.74). These estimates are root-n consistent under H_2,