both heteroskedastic and serially correlated that the HAC estimators really come into their own. Even in these circumstances, it is possible, with some patterns of heteroskedasticity, that the feasible GLS estimator, which takes no account of possible heteroskedasticity, can outperform the HAC estimators. But that is probably the exception rather than the rule, for Andrews finds other patterns of heteroskedasticity, which, in combination with serial correlation, require the use of HAC estimators for reasonably accurate inference.

Clearly, the last word on HAC estimators has by no means been said. For instance, in the usual implementation of the Newey-West estimator for linear IV models, we have that $\hat{\boldsymbol{\Gamma}}(0)$ is just $n^{-1}\boldsymbol{W}^\top\hat{\boldsymbol{\Omega}}\boldsymbol{W}$, with $\hat{\boldsymbol{\Omega}}$ the rather poor estimator associated with the HC_0 form of the HCCME. It would seem reasonable to suppose that it would be better to use other forms of $\boldsymbol{\Omega}$ in the Newey-West estimator, just as it is in HCCMEs, and to find similar ways of improving the estimators $\hat{\boldsymbol{\Gamma}}(j)$ for $j \neq 0$. At the time of writing, however, no evidence is available on whether these conjectures are justified. A quite different approach, which we do not have space to discuss, was recently suggested by Andrews and Monahan (1992).

In the next section, we will leave behind the "grubby details" of covariance matrix estimation, assume that a suitable covariance matrix estimator is available, and turn our attention to asymptotic tests of overidentifying restrictions and other aspects of specification testing in GMM models.

17.6 Inference with GMM Models

In this section, we undertake an investigation of how hypotheses may be tested in the context of GMM models. We begin by looking at tests of overidentifying restrictions and then move on to develop procedures akin to the classical tests studied in Chapter 13 for models estimated by maximum likelihood. The similarities to procedures we have already studied are striking. There is one important difference, however: We will not be able to make any great use of artificial linear regressions in order to implement the tests we discuss. The reason is simply that such artificial regressions have not yet been adequately developed. They exist only for some special cases, and their finite-sample properties are almost entirely unknown. However, there is every reason to hope and expect that in a few years it will be possible to perform inference on GMM models by means of artificial regressions still to be invented.

In the meantime, there are several testing procedures for GMM models that are not difficult to perform. The most important of these is a test of the overidentifying restrictions that are usually imposed. Suppose that we have estimated a vector $\boldsymbol{\theta}$ of k parameters by minimizing the criterion function

$$\boldsymbol{\iota}^\top\boldsymbol{F}(\boldsymbol{\theta})\hat{\boldsymbol{\Phi}}^{-1}\boldsymbol{F}^\top(\boldsymbol{\theta})\boldsymbol{\iota}, \tag{17.67}$$

in which the empirical moment matrix $\boldsymbol{F}(\boldsymbol{\theta})$ has $l > k$ columns. Observe that