BOX: Structural VARs with Long-Run Restrictions

I have not included a commonly used statistical representation known as the *structural vector autoregression* (SVAR) in this study. Although SVARs are widely used by business cycle analysts, there has been considerable debate recently about their usefulness.[1] One critique leveled by Chari, Kehoe, and McGrattan (2008) is discussed here and can be addressed by using a VARMA representation as is done in Section 3. However, given the wide use of SVARs, it may be helpful to review the substance of the critique.

The procedure.[2]

I'll focus attention on the *structural vector autoregression procedure with long-run restrictions* which is a simple time series technique that uses minimal economic theory to identify the pattern of responses of economic aggregates to possible shocks in the economy. Following this procedure, the analyst estimates a vector autoregression, inverts it to get a moving average (MA) representation, and imposes certain structural assumptions about the shocks hitting the economy.

To be more precise, let Y_t be a N-dimensional vector containing observations at time period t. The first step is to estimate a vector autoregression (VAR) by regressing Y_t on p lags,

$$Y_t = A_0 + A_1 Y_{t-1} + A_2 Y_{t-2} + \dots A_p Y_{t-p} + v_t \tag{1}$$

The second step is to invert this VAR to get the corresponding moving average,

$$Y_t = v_t + B_1 v_{t-1} + B_2 v_{t-2} + \dots \tag{2}$$

where v_t is the residual in (1) with variance-covariance matrix $E v_t v_t' = \Omega$. Mechanically, it is easy to recursively compute the B coefficients given estimates of the A coefficients in (1). An estimate of the matrix Ω is easily constructed from the VAR residuals.

One more step is needed to derive the structural MA, one that has interpretable shocks. In particular, the structural MA is given by

$$Y_t = C_0 e_t + C_1 e_{t-1} + C_2 e_{t-2} + \dots \tag{3}$$

where $E e_t e_t' = \Sigma$, $e_t = C_0^{-1} v_t$, and $C_j = B_j C_0$ for $j \geq 1$. The elements of v_t are simple residuals in a VAR, but the elements of e_t are the shocks of interest. To provide these shocks with an economic interpretation, we need to impose some restrictions on the elements of C_0 and Σ. Once we do that, we have an fully operational procedure: we can completely describe what happens to the economy (summarized by the time series $\{Y_t\}$) following some shock (summarized by a change in one of the elements of e_t). This is a very simple computer code.

An application.

To better understand the "structural" restrictions imposed on C_0 and Σ, it helps to describe a specific example. I'll use one that has been at the center of the recent debate concerned with the usefulness of SVARs.[3] In this case, Y_t is a two-dimensional vector containing the change in the log of labor productivity

[1] References.

[2] This is the general procedure developed by Blanchard and Quah (1989).

[3] See Chari, Kehoe, and McGrattan (2008) for details.

and the log of hours. The lag length p in (1) is set equal to four. In practice, analysts choose small values for p, typically between four and eight, because they have samples sizes of roughly 200 quarters. For the structural shocks, assume the first element of e_t is a "technology" shock and the second element is a "demand" shock.

Next we need identifying restrictions to determine the seven parameters in C_0 and Σ. Seven restrictions typically used in the SVAR literature are as follows. Three come from equating variance-covariance matrices ($C_0 \Sigma C_0' = \Omega$). Three come from assuming that the shocks are orthogonal ($\Sigma = I$). The last comes from the assumption that demand shocks have no long-run effect on labor productivity ($\sum_j C_j(1,2) = 0$). SVAR users call this last restriction a *long-run restriction.* Implicit in this restriction is the assumption that technology shocks have a permanent effect on the level of labor productivity whereas demand shocks do not.

The SVAR claims.

The main finding of the recent SVAR literature is that a positive technology shock (that is, an increase in the first element of e_t) leads to a fall in hours (the second element of Y_t). (See Galí and Rabanal (2005) for a review of this literature.) This has lead these researchers to conclude that a certain class of business cycle models, referred to as *real business cycle* or RBC models, are not promising for the study of business cycles since most RBC models predict a positive response in hours. They further claim that the SVAR evidence points to another class of business cycle models referred to as *sticky price* models as more promising because they can produce the fall in hours after a technology shock.

The CKM critique.

Users of the SVAR procedure claim that it is useful because it can confidently and correctly distinguish between promising and unpromising classes of models with minimal assumptions about the economic environment. Chari et al. (2008) evaluated this claim with a very simple test. They generated data from an RBC model, drawing a large number of samples with the same length as is available in U.S. data. With these artificial data, they repeatedly apply the SVAR procedure and find that it cannot confidently and correctly uncover the truth.

There are two problems associated with the SVAR procedure that lead to biased and uninformative results. The source of the first problem is *truncation bias.* Theoretical business cycle models currently in use cannot be represented by a finite-order VAR. In other words, p in (1) is infinity. But our datasets have finite length and, as a result, p must be finite. The source of the second problem is *small sample bias.* Most statistical procedures do poorly if sample sizes are not "sufficiently large." The ultimate question though is what is sufficient. Typically it depends on the details of the application.

Avoiding truncation bias.

One easy adaptation of the SVAR procedure is to use a VARMA representation that allows for moving average terms. This is what is done in this study. Using VARMAs takes care of the problem of truncation bias. Then, I can determine the extent of the problem of small sample sizes.