Commentary

Martin Fukač and Adrian Pagan

The paper by DiCecio and Nelson (2007; DCN) considers the estimation of the parameters of a dynamic stochastic general equilibrium (DSGE) model for the United Kingdom that is virtually the same as that which Christiano, Eichenbaum, and Evans (2005; CEE) estimated for the United States. The CEE model is much larger than existing academic DSGE models of the United Kingdom, such as Lubik and Schorfheide (2007). It is not as large as the Bank of England Quarterly Model (BEQM), which has both DSGE elements and data-imposed dynamics; however, because the BEQM is used for policymaking, there is a much greater imperative to match the data than found in most academic work. There are a few other DSGE models that have been applied to the United Kingdom—for example, Leitemo (2006)—but, in general, these are often used to examine some particular question and are also rather restricted in their mode of operation. Often they use a standard open-economy New Keynesian model rather than a straight DSGE model like CEE's. Moreover, the authors of these models are often not that familiar with the U.K. context and data; the current authors, however, are experts in this area, and it certainly shows in their discussion of alternative data sources. So, given the paucity of studies, any new one would be welcome.

Now, as the Chinese proverb says, a journey of a thousand miles starts with a single step. What we have here is more than single step but well short of a thousand miles. Reading it, one longed for a fully fleshed-out model along the lines of Smets and Wouters's (2003) work on the United States and the euro area (which is very similar to CEE's), where a complete set of shocks is described and estimated. Because the DCN model identifies only a money shock, there are few questions one can ask about the model. So it was disappointing that the authors were not a bit more adventuresome. But we presume that this will be part of a broader piece of research and look forward to seeing a more complex model that recognizes the open-economy characteristics of the United Kingdom. Of course, one does have to acknowledge that DSGE models have not had a good record of producing useful models of the open economy. One reason DCN point to is the prediction of stronger exchange effects than seen in the data. We agree with this, and it was a central conclusion about the mini-BEQM model that was calibrated to the U.K. economy in Kapetanios, Pagan, and Scott (2007). Moreover, as Justiniano and Preston (2006) argue, it has been very hard to find much of an influence of the foreign economy on a small open economy, and this is contrary to evidence we have from vector autoregressions (VAR). So there is quite a bit to be done both on the broader front of developing useful open-economy models and in getting a U.K. model that is in a more complete state than this one. Because the model is not fleshed out that much, we will restrict comments to what DCN do rather than alternatives that might have been tried.

ESTIMATION STRATEGY

The methodology used for estimation is that of CEE. It has four steps:

1. Identify monetary policy shocks using a structural equation for the interest rate and a VAR(2) to represent the rest of the system. The identification condition used is that monetary shocks have no contemporaneous effect on any variables of the system except the interest rate. The monetary policy rule depends on all variables in the VAR, and these enter the rule both with lags and contemporaneously.

2. Compute the monetary impulse responses C_j^D from this structural VAR (SVAR) (j indexes the jth impulse response).

3. Choose some values for the DSGE model parameters θ and use them to compute the DSGE model monetary impulse responses C_j^M.

4. Find the value of θ that minimizes

$$\sum_{j=1}^{M} \left(C_j^D - C_j^M \right)' W \left(C_j^D - C_j^M \right),$$

where W is a diagonal matrix of weights.

They apply this to U.K. data. The original DSGE model they employ has 15 variables, whereas the SVAR(2) has 6. Estimates of the parameters are presented, and some standard errors are given along with plots of the monetary impulses implied by the SVAR and the DSGE model calibrated with the estimated θ.

ESTIMATION PROBLEMS

What could go wrong with this methodology? We discuss three issues in the following subsections.

How Many Impulses To Use?

There is a maximum useful choice for M because the C_j^D are simply functions of the SVAR coefficients. Let there be n variables in the SVAR (in DCN $n = 6$ and it is an SVAR(2)). Then the total number of coefficients in the DCN SVAR(2)

is 77: Of these 77, 72 are from the 2 lags of the six variables in the six equations plus the 5 possible coefficients attached to contemporaneous coefficients in the interest rate rule. Because M seems to be 25, that would mean that 150 impulses are used. This is much larger than the number of parameters determining them. Hence there are many redundant impulse responses and the covariance matrix of $C_1^D, \ldots, C_{25}^D$ must be singular. This might be a problem when one uses the δ method to compute standard errors. Indeed, the standard errors of $\hat{\theta}$ found by moment matching in DCN seem to be incredibly small. Thus they have an estimate of the markup λ_f parameter of 2.27 with a standard error of 0.03. It's hard to believe that one could ever get that degree of precision with just 26 years of quarterly observations. Because the SVAR coefficients have a t ratio above 5 in only one case (lagged productivity), it's hard to see how we can end up with t ratios above 400 for $\hat{\theta}$, which are fundamentally derived from the SVAR coefficients.

Approximating the DSGE Model with an SVAR

There is a generic problem here in that the DSGE model often determines m variables and $m > n$; that is, the SVAR is fitted to a smaller number of variables than appear in the DSGE model. This is true of the DCN model, where it appears that $n = 6$, $m = 15$. Now the DSGE model will be an SVAR in the m variables but is unlikely to be an SVAR in the n variables. An old literature, due to Zellner and Palm (1974) and Wallis (1977), has noted that, when a system that is a VAR(p) in n variables is reduced to a smaller system with $m < n$ variables, the smaller system will generally be a vector autoregression moving-average (VARMA) process. Because the CEE procedure involves such compression of variables, it might be expected that a VARMA process is needed rather than a VAR; so, the use of a VAR could lead to specification bias. It might be expected that a VAR of very high order could compensate for this misspecification—and this is generally true—but the order of VAR needed to deliver a good approximation may in fact be far too high for the

data sets one is normally faced with. For example, Kapetanios, Pagan, and Scott (2007) find that reducing a model that is BEQM-like (but half the size) would require a VAR(50) to capture the effects of some shocks (and this with 30,000 observations). The problem has been analyzed in a DSGE context by Ravenna (forthcoming) and Fry and Pagan (2005). We adopt the exposition of the latter.

Suppose that the DSGE model followed a VAR(1) solution (assuming that u_t is i.i.d.):

$$z_t = Pz_{t-1} + Gu_t.$$

Now consider what happens if we model only a subset of the variables. We will call the modeled subset z_{1t} and the omitted variables z_{2t}. We can decompose the VAR above as

(1) $$z_{1t} = P_{11}z_{1t-1} + P_{12}z_{2t-1} + G_1u_t,$$

and we will assume that the following relation holds between z_{1t} and z_{2t}:

$$z_{2t} = \bar{D}_0 z_{1t} + \bar{D}_1 z_{2t-1} + \bar{D}_2 u_t.$$

Substituting this in we get

$$z_{1t} =$$
$$\left(P_{11} + P_{12}\bar{D}_0\right)z_{1t-1} + P_{12}\bar{D}_1 z_{2t-2} + G_1 u_t + P_{12}\bar{D}_2 u_{t-1},$$

so that the sufficient conditions for there to be a finite-order VAR in z_{1t} will be that either $P_{12} = 0$ (i.e., z_{2t} does not Granger cause z_{1t}; see Lütkepohl, 1993, p. 55, and Quenouille, 1957, p. 43-44) or $\bar{D}_1 = 0, \bar{D}_2 = 0$ (i.e., the variables to be eliminated must be connected to the retained variables through an identity and there can be no "own lag" in the omitted variables in the relation connecting z_{1t} and z_{2t}).

This observation looks trivial, but in fact many of the problems that have arisen where a finite-order VAR does not obtain occur because the omitted variables are connected with the retained variables through an identity, but one that contains an "own lag." The classic example is in the basic real business cycle (RBC) model where, after log linearization around the steady state, we would get

(2) $$l_t = y_t - c_t$$
$$C^* c_t + K^* k_t = Y^* y_t + (1-\delta)K^* k_{t-1}$$
(3) $$c_t = E_t\left(c_{t+1} - \alpha\gamma\left(y_{t+1} - k_{t+1}\right)\right)$$
$$y_t = a_t + \alpha k_{t-1} + (1-\alpha)l_t,$$

where c_t is the log of consumption, a_t is the log of the technology shock, k_t is the log of the capital stock, l_t is the log of labor input, and y_t is the log of output. An asterisk denotes steady-state values, and α is the steady-state share of capital in output. When a_t is an AR(1), the solution to this system can be made a VAR(1) in c_t, l_t, y_t, and k_t. It's clear that we could eliminate any of c_t, l_t, or y_t because none appear as a lagged variable in the system. Equally clearly, k_t cannot be eliminated unless we can find an identity that relates it to other variables but does not involve k_{t-1}. Thus, the identity (3) shows that this is not possible. Most of the literature that seeks to establish that a SVAR cannot approximate a DSGE model (Chari, Kehoe, and McGrattan, 2004; Erceg, Guerrieri, and Gust, 2005; Cooley and Dwyer, 1995) substitute out k_t and so end up with a non-finite-order VAR.

The implication of this for DCN's work is that the reduction of the system from 15 to 6 variables might necessitate a very long VAR and not the VAR(2) they adopted. They used statistical criteria to determine the order of the VAR. Kapetanios, Pagan, and Scott (2007) did this as well, and the tests produced a VAR of order six, far below what was needed (50th order) to produce the correct impulse responses. The reason is that the tests proceed on the assumption that the number of variables in the VAR is correct and it is only the order that needs to be found. So it seems that DCN might be matching impulses that are not strictly comparable. The appropriate procedure would be to (i) simulate a long history of data from the 15-variable DSGE model, incorporating just monetary shocks; (ii) fit a VAR(2) in just 6 of the variables; and then (iii) find the impulse responses from such an approximating VAR, being careful to note that some of the lagged values will be perfectly correlated with others and that it will be necessary to combine variables together to overcome that problem. These are then matched

up with the empirically observed VAR(2) impulse responses in the six variables. We have assessed this by examining a variant of their model, where information is dated at t rather than the combination of t and $t-1$ that is in their paper. However, we used the same model parameters as DCN. Although there are some differences between the true impulse responses and those delivered by a VAR(2), it seems that the approximations are quite good, except at longer horizons. So there does not seem to be an approximation issue here, although in any application one should check that there is no problem, as it is not very difficult to do. However, the δ method used by CEE to compute standard errors is correct only if the approximation is satisfactory. Basically, the estimator of the DSGE-model parameters is an indirect estimator, being derived from functions of the SVAR coefficients represented by the impulse responses. The covariance matrix of such an estimator requires that derivatives of the model-implied VAR impulse responses be computed with respect to the θ parameters, and not the derivatives with respect to the model impulses, as done by CEE. These are only the same if there is no approximation error.

Multiple Solutions

Ignoring the problem identified in the previous section, estimators such as the maximum likelihood estimator basically attempt to match the VAR coefficients from the data with those from the model, rather than attempting to match impulse responses. To see the problems you might encounter with the latter, let us look at the simple model

$$y_t = \beta E_{t-1}(x_{t+1}) + \varepsilon_{yt},$$
$$x_t = \rho x_{t-1} + \varepsilon_{xt}.$$

The VAR will be

$$y_t = a_1 x_{t-1} + \varepsilon_{yt}$$
$$x_t = a_2 x_{t-1} + \varepsilon_{xt},$$

where $a_1 = \beta\rho^2$, $a_2 = \rho$, and the impulse responses are

$$c_{1,y\varepsilon_x}^M = \beta\rho^2, c_{2,y\varepsilon_x}^M = \beta\rho^3, c_{1,x\varepsilon_x} = \rho, c_{2,x\varepsilon_x} = \rho^2.$$

If we would try to find β and ρ by matching the first two impulse responses, we would be minimizing (assuming that the weights in W are equal)

$$\left(c_{1,y\varepsilon_x}^D - \beta\rho^2\right)^2 + \left(c_{2,y\varepsilon_x}^D - \beta\rho^3\right)^2$$
$$+\left(c_{1,x\varepsilon_x}^D - \rho\right)^2 + \left(c_{2,x\varepsilon_x}^D - \rho^2\right)^2.$$

Clearly, such an approach has the problem of producing an order-six polynomial in ρ, so that we may get multiple solutions. This would not arise if we were matching to the VAR coefficients, because then $\hat{\rho} = \hat{a}_2$, $\hat{\beta} = \hat{a}_1/\hat{\rho}^2$. Bearing in mind the first point as well, it seems better to match the VAR to get estimates of θ and then to show the impulse response correspondence.

LOOKING AT SOME OF THE EULER EQUATIONS

Now it would seem useful to develop a method that uses the same information as impulse-response matching but that is a bit simpler, provides ready ways of computing standard errors, emphasizes the economics, and can be used to tell us something about the ability of the DSGE model to match the data. Basically the proposal is to work with the Euler equations and estimate the model parameters directly from them with a single-equation estimator. Of course this is an old idea, but it has fallen out of favor, possibly because of the literature claiming that systems estimators of parameters of the New Keynesian system performed better than the single-equation estimators because of weak instruments. But, in many DSGE models, enough parameters are prescribed that weak instrument problems are not present, and we will see this in the DCN context.

The Euler equations of DSGE models have the generic form[1]

$$E_{t-1}z_t = \eta_1 z_{t-1} + \eta_2 E_{t-1}(z_{t+1}) + \eta_3 E_{t-1}w_t.$$

In this equation, z_t is the endogenous variable whose coefficient is normalized to unity, w_t are

either exogenous or other endogenous variables, and the parameters η_j are functions of some of the DSGE model parameters θ. Now this can be written as

$$z_t = \eta_1 z_{t-1} + \eta_2 E_{t-1}(z_{t+1}) + \eta_3 E_{t-1} w_t + \varepsilon_t,$$

and the right-hand-side regressors are uncorrelated with the error ε_t. If we had these conditional expectations we could run a regression. We note that this equation holds for any subset of information used by the economic agents. Hence, let us define the information used in the estimation as that of the DCN VAR(2), that is, two lagged values of $y_t, c_t, i_t, y_t - h_t, r_t$, and Δp_t.[2] Call these the vector ζ_{t-1}. Then, if we can estimate $E_{t-1}(z_{t+1})$ and $E_{t-1} w_t$, we could simply fit a regression to this equation and thereby measure η_j. Because the model is linear, we can indeed estimate $E_{t-1}(z_{t+1})$ and $E_{t-1} w_t$ as the predictions from the regression of z_{t+1} and w_t against ζ_{t-1}. Basically, this estimation method uses the same information as moment matching, that is, the information contained in the VAR. Notice that standard errors are easily found from this by estimating the Euler equation parameters with an instrumental variables estimator. As we will see later, in most cases the instruments are very good and so there is no reason to doubt the standard errors of η_j found in this way.

This is a relatively simple way to estimate the η_j. Whether the DSGE model parameters θ can be estimated is a different question, because there may be a nonlinear mapping between the η and θ and so we may not be able to recover θ uniquely. This shouldn't concern us unduly because, fundamentally, the impact of monetary policy depends on the η_j; but there may be some cases where we want to think about changing θ and so would then need to identify it. Ma (2002) pointed out that there was an identification problem like this in strictly forward-looking New Keynesian Phillips curves, and we will see that it comes up in the CEE model as well.

Let us look at the above principles in the context of some of the equations in DCN. First we look at the Phillips curve. After normalizing on π_t, the Euler equation becomes

(4)
$$-E_{t-1}\pi_t + \frac{1}{1+\beta}\pi_{t-1} + \frac{\beta}{1+\beta}E_{t-1}\pi_{t+1}$$
$$+ \frac{(1-\beta\xi)(1-\xi)}{(1+\beta)\xi}E_{t-1}s_t.$$

We can write this as an equation of the form

$$\pi_t = \frac{1}{1+\beta}\pi_{t-1} + \frac{\beta}{1+\beta}\pi_{t+1}$$
$$+ \frac{(1-\beta\xi)(1-\xi)}{(1+\beta)\xi}E_{t-1}\left[\left(\alpha r_t^k + (1-\alpha)(R_t + w_t)\right)\right] + \varepsilon_t$$

or

$$\pi_t - \frac{1}{1+\beta}\pi_{t-1} - \frac{\beta}{1+\beta}\pi_{t+1}$$
$$= \eta_1 E_{t-1}\left[\left(\alpha r_t^k + (1-\alpha)(R_t + w_t)\right)\right] + \varepsilon_t,$$

where $E_{t-1}(\varepsilon_t) = 0$. We note that, because $\beta = 0.99$ is imposed, we are not trying to estimate the coefficients attached to π_t and π_{t+1}. Now, using data, one can form $\alpha r_t^k + (1-\alpha)(R_t + w_t)$. Because DCN pre-set α to 0.36, this can then be regressed against the information represented by the VAR lagged variables to get

$$E_{t-1}\left[\left(\alpha r_t^k + (1-\alpha)(R_t + w_t)\right)\right].[3]$$

The regression of this variable against ζ_{t-1} (the VAR(2) lagged variables) gives an R^2 of 0.99, so it is a very good instrument for $\alpha r_t^k + (1-\alpha)(R_t + w_t)$. If we fit a nonlinear regression to this equation, we get an estimated coefficient for ξ_p of 0.988, which is reasonably close to what is reported in the paper from impulse response matching. But the standard deviation is 0.092, which is nowhere near the 0.0004 given in the paper—although, if one makes it robust to serial correlation, it halves. Clearly, the estimate here implies very low frequency of price adjustment, as does DCN's.

---

[2] We work with data that are not deviations from steady-state values and so will have to include intercepts in any equation we estimate.

[3] Because DCN conclude that $\sigma_a = \infty$, there is no difference between the capital stock and services. We compute the capital stock recursively, but this means the estimates are inaccurate until the initial condition disappears. Because we start the recursion in 1955:Q2, but use data only after 1979:Q2, we feel that the effects of the initial condition have died away, as it will be multiplied by the term $(0.975)^{104}$.

Although this estimate seems implausible, the interpretation would seem to be that there are some problems in the specification of the Phillips curve (indeed the serial correlation in the residuals is consistent with that).

Another equation in DCN we consider estimating is the production function with the form

$$y_t = \lambda_f \left(\alpha k_t + (1 - \alpha) l_t \right).$$

Given that DCN prescribe α, we can form $\zeta_t = \alpha k_t + (1 - \alpha) l_t$ and treat this as a regressor to estimate λ_f. There will of course be technology in this relation and, by treating it as an AR(1) process, the equation will have an AR(1) error term. Because the regressor will generally be correlated with the white noise shock driving the AR in technology, we need instruments to estimate λ_f; for this we use y_{t-1} and ζ_{t-1}. We also include a constant to reflect the fact that we are not using variables that are deviations from a constant steady state and that technology should have a constant mean. Then we get an estimate of λ_f of 1.25, with a standard error of 0.05. This seems more reasonable than the value of 2.27 that DCN obtain, although they give a defense of it. Again, the standard errors are very different.

The interest rate rule parameter values are somewhat puzzling. Under the assumptions in force here, one should be able to fit this rule by ordinary least squares (OLS) regression because it is assumed that the regressors are all uncorrelated with the interest rate shock. If we run the regression, the fit we would get is

$$R_t = 0.883 R_{t-1} + (1 - 0.833)(0.0001 y_t + 1.28 \pi_{t-1})$$
$$+ 0.05 \Delta y_t + 0.10 \Delta \pi_t,$$

versus the estimated equation of the paper,

$$R_t = 0.872 R_{t-1} + (1 - 0.872)(0.352 y_t + 1.28 \pi_{t-1})$$
$$+ 0.43 \Delta y_t - 0.62 \Delta \pi_t.$$

The standard deviation on y_t from the OLS regression is very small, so these estimates are very different. DCN note that the rule they fit is not the one in the VAR(2), because that would include other lags in the variables. But if we just fit a VAR(1), then it should be comparable to what they claim the estimated money rule is. In fact,

there is not much difference if we add on extra lags. Notice also that the negative sign on $\Delta \pi_t$ that perturbed them has gone. Because this seems a logical way to estimate the money rule, given the assumptions made about the structure of the model, one is puzzled about the results that come from matching impulse responses.

What explains this? One possibility is that the DSGE model implies a particular value for the intercept of the equation, whereas we have just subsumed this into a constant term that is freely estimated. However, the steady-state values used in the model for variables seemed quite close to the sample means over the estimation horizon; so, it would seem that one would get much the same intercepts (provided of course the slope coefficients were correct).

There are some problems with multiple parameter values in both the Phillips curve and the wage equation. Because

$$\pi_t - \frac{1}{1+\beta} \pi_{t-1} - \frac{\beta}{1+\beta} \pi_{t+1} =$$
$$\eta_1 E_{t-1} \left[\left(\alpha r_t^k + (1 - \alpha)(R_t + w_t) \right) \right] + \varepsilon_t$$

and

$$\eta_1 = \frac{(1 - \beta \xi_p)(1 - \xi_p)}{(1 + \beta) \xi_p},$$

we see that the solution for ξ_p involves a quadratic. There are two estimates of ξ_p that produce exactly the same likelihood—the value of 0.988 given above and 1.02. A similar situation exists for the wage equation. Perhaps this is one reason why Bayesian methods might work better in these models—they would impose the restriction that ξ_p and ξ_w lie between 0 and 1.

STRUCTURAL CHANGE IN THE MODEL

The authors look at structural change in the SVAR and conclude that there was some back in the 1970s, but this was due to industrial issues and not monetary policy regime changes. But it's always difficult to learn something about the sta-

Figure 1

Coefficients of Lagged Inflation with Two Standard Errors on Each Side from a Rolling Regression (68 periods)
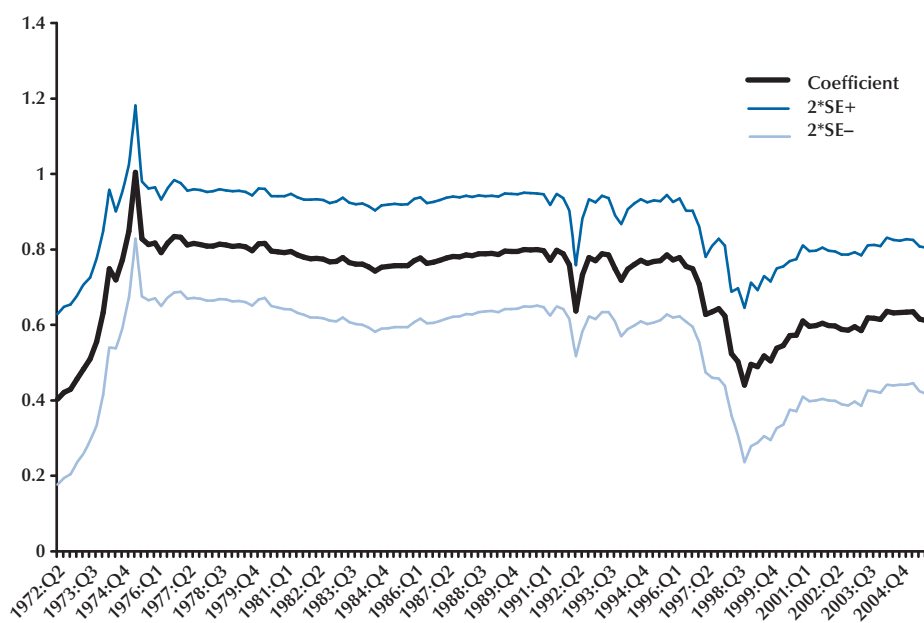
bility of the parameters in a VAR. One might also want to ask where to place a possible monetary policy regime change. Is it when Thatcher came in, when inflation targeting was adopted, or when there was a formal change to the institution with the formation of the Monetary Policy Committee (MPC)? Pagan (2003) argued that there had been a change in the level of persistence in inflation in the United Kingdom after the formation of the MPC. This is still evident in the data: See Figure 1, which gives estimates of the coefficient of π_{t-1} using a rolling horizon of 68 quarters.

So this looks like structural change in the dynamics, and perhaps the VAR stability tests should have focused more around the post-1997 point, although this means a very short post-break sample. At the end of the day, graphs like this have to make one wonder about applying a constant-parameter DSGE model to such data. It would seem one might need to use only the post-1997 period to estimate the DSGE model, although with such small samples one might need to use

some sort of Bayesian approach. Perhaps one could use the estimated values of this paper to produce priors.

REFERENCES

Chari, V.V.; Kehoe, Patrick J. and McGrattan, Ellen R. "A Critique of Structural VARs Using Real Business Cycle Theory." Working Paper No. 631, Federal Reserve Bank of Minneapolis, 2004.

Christiano, Lawrence J.; Eichenbaum, Martin and Evans, Charles L. "Nominal Rigidities and the Dynamic Effects of a Shock to Monetary Policy." *Journal of Political Economy*, February 2005, *113*(1), pp. 1-45.

Cooley, Thomas F. and Dwyer, Mark. "Business Cycles Without Much Theory: A Look at Structural VARs." *Journal of Econometrics*, March/April 1995, *83*(1-2), pp. 57-88.

DiCecio, Riccardo and Nelson, Edward. "An Estimated DSGE Model for the United Kingdom." Federal Reserve Bank of St. Louis *Review*, July/August 2007, *89*(4), pp. 215-31.

Erceg, Christopher J.; Guerrieri, Luca and Gust, Christopher. "Can Long-Run Restrictions Identify Technology Shocks?" *Journal of the European Economic Association*, December 2005, *3*(6), pp. 1237-78.

Fry, Renee and Pagan, Adrian R. "Some Issues in Using VARs for Macroeconometric Research." Centre for Applied Macroeconomic Analyses, CAMA Working Paper 2005/18, Australian National University, 2005.

Justiniano, Alejandro and Preston, Bruce. "Can Structural Small Open Economy Models Account for the Influence of Foreign Disturbances." Centre for Applied Macroeconomic Analyses, CAMA Working Paper, 2006/12, Australian National University, 2006.

Kapetanios, George; Pagan, Adrian R. and Scott, Alasdair. "Making a Match: Combining Theory and Evidence in Policy-Oriented Macroeconomic Modeling." *Journal of Econometrics*, February 2007, *136*(2), pp. 505-94.

Leitemo, Kai. "Targeting Inflation by Forecast Feedback Rules in Small Open Economies." *Journal of Economic Dynamics and Control*, March 2006, *30*(3), pp. 393-413.

Lubik, Thomas A. and Schorfheide, Frank. "Do Central Banks Respond to Exchange Rate Movements? A Structural Investigation." *Journal of Monetary Economics,* May 2007, *54*(4), pp. 1069-87.

Lütkepohl, Helmut. *An Introduction to Multiple Time Series.* Berlin: Springer-Verlag, 1993.

Ma, Adrian. "GMM Estimation of the New Phillips Curve." *Economics Letters*, August 2002, *76*(3), pp. 411-17.

Pagan, Adrian R. "Report on Modelling and Forecasting at the Bank of England." Bank of England *Quarterly Bulletin*, Spring 2003, *43*(1), pp. 1-29.

Quenouille, M.H. *The Analysis of Multiple Time Series.* Griffin's Statistical Monographs and Courses No. 1. London: Griffin, 1957.

Ravenna, Frederico. "Vector Autoregressions and Reduced Form Representations of Dynamic Stochastic General Equilibrium Models." *Journal of Monetary Economics* (forthcoming).

Smets, Frank and Wouters, Raf. "An Estimated Stochastic Dynamic General Equilibrium Model of the Euro Area." *Journal of the European Economic Association*, 2003, *1*(5), pp. 1123-75.

Wallis, Kenneth F. "Multiple Time Series and the Final Form of Econometric Models." *Econometrica*, September 1977, *45*(6), pp. 1481-97.

Zellner, Arnold and Palm, Franz. "Time Series Analysis and Simultaneous Equation Econometric Models." *Journal of Econometrics*, May 1974, *2*(1), pp. 17-54.