

COWLES FOUNDATION FOR RESEARCH IN ECONOMICS
AT YALE UNIVERSITY

Box 2125, Yale Station
New Haven, Connecticut 06520

COWLES FOUNDATION DISCUSSION PAPER NO. 702

Note: Cowles Foundation Discussion Papers are preliminary materials circulated to stimulate discussion and critical comment. Requests for single copies of a Paper will be filled by the Cowles Foundation within the limits of the supply. References in publications to Discussion Papers (other than mere acknowledgment by a writer that he has access to such unpublished material) should be cleared with the author to protect the tentative character of these papers.

CONDITIONAL PROJECTION BY MEANS OF KALMAN FILTERING

Richard H. Clarida

and

Diane Coyle

April 1984

CONDITIONAL PROJECTION BY MEANS OF KALMAN FILTERING

Richard H. Clarida

Yale University and the Cowles Foundation for Research in Economics

Diane Coyle

Harvard University

ABSTRACT

We establish that the recursive, state-space methods of Kalman filtering and smoothing can be used to implement the Doan, Litterman, and Sims (1983) approach to econometric forecast and policy evaluation. Compared with the methods outlined in Doan, Litterman, and Sims, the Kalman algorithms are more easily programmed and modified to incorporate different linear constraints, avoid cumbersome matrix inversions, and provide estimates of the full variance-covariance matrix of the constrained projection errors which can be used directly, under standard normality assumptions, to test statistically the likelihood and internal consistency of the forecast under study.

We are indebted to Mark Watson for invaluable discussions on this topic. Thanks also to John Campbell, Ben Friedman, Tom Sargent, and Chris Sims for their comments. We are of course responsible for all remaining errors.

CONDITIONAL PROJECTION BY MEANS OF KALMAN FILTERING

Richard H. Clarida and Diane Coyle

Recently, Sims (1982) and Doan, Litterman, and Sims (1983) have proposed and implemented methods for evaluating the plausibility and internal consistency of econometric forecasts. As Sims (1982) has convincingly argued, the usual method for constructing macro-econometric forecasts is potentially misleading since it assumes that the specified paths for a policy variable is generated by disturbances to the policy equation alone, with all other disturbances held to zero. By contrast, the methodology proposed by Doan, Litterman, and Sims explicitly takes into account policy endogeneity by generating true conditional projections given specified paths for the policy variables. More generally, the key insight of the Doan, Litterman, and Sims approach is that a constraint on the future path of one variable in the system carries information which can be used to help predict the most-likely current values of other variables in the system.

The Doan, Litterman, Sims methodology may be described as follows. The historical, dynamic correlations among the variables in some set of interest are summarized using a multivariate time series model such as a vector autoregression. The estimated parameters and innovation variance-covariance matrix of the model are then used to construct a projection of the most likely post-sample path for a given subset of variables subject to the constraint that (a linear combination of) the remaining variables attain the post-sample trajectory predicted by the forecast under study. The forecast is then compared against the derived conditional projections using a metric constructed from the estimated innovation variance-covariance matrix.

The principle behind the Doan, Litterman, and Sims methodology is that the estimated time series model provides a conditional joint density function for the post-sample trajectory of the system which can be used to find likelihood maximizing paths subject to linear restrictions on those paths. Doan, Litterman, and Sims use the following procedure to calculate these constrained conditional projections. A linear constraint upon future values of the system's variables is transformed, using the matrix of estimated moving average coefficients, into an equivalent constraint on the sequence of orthogonalized innovations. The least-squares estimate of the constrained innovations is computed and is used to obtain the least-squares projection of the system's variables subject to this constraint by constructing the path implied by the computed innovations. This classical approach requires first constructing a matrix of moving average coefficients whose dimension is equal to the product of the number of variables in the system and the number of post-sample periods in the forecast and then inverting, for each set of constraints studied, a transformation of this matrix of dimension equal to the product of the number of linear restrictions at each post-sample date and the number of post-sample periods in the forecast. For systems of even moderate size, this can be an imposing task.

Doan, Litterman, and Sims conjecture that these computations cannot be carried out recursively forward in time (as can unconstrained point forecasts) because of the aforementioned fact that a constraint on future values of a variable in the system carries information about the most likely current values of all variables in the system. The purpose of this note is to establish that in fact recursive, state-space methods can be used to implement the Doan, Litterman, and Sims approach to econometric forecast and policy evaluation. In particular, we show that the methods of Kalman filtering and smoothing can be

used to combine post-sample constraints on a linear combination of a system's variables with that system's historically estimated parameters to yield minimum mean square linear projections of the variables' future paths which optimally incorporate all of the post-sample constraints. Furthermore, the Kalman algorithms are easily programmed and modified to incorporate different linear constraints, avoid cumbersome matrix inversions, and provide estimates of the full variance-covariance matrix of the constrained projections which can be used directly, under standard normality assumptions, to test statistically the likelihood and internal consistency of the forecast under study.

The plan of the paper is as follows. In Section 1, we show that the methods of Kalman filtering and smoothing can be used to forecast the post-sample behavior of a linear dynamic system subject to linear restrictions on the future values attained by its variables. In Section 2, we show how differences between a particular forecast under study and the constrained projections generated by the Kalman filter can be tested for statistical significance. In Section 3, we discuss the empirical implementation of our state-space approach. In Section 4, we provide some concluding remarks.

1: State-Space Representation and the Kalman Filter

We consider the problem of forecasting the post-sample behavior of a linear dynamic system subject to linear restrictions on the future values attained by its variables. Harvey (1981) provides an excellent recent treatment of the necessary state-space tools.

Let x be a vector containing the system's r variables. The movement of x through time is assumed to be governed by a constant-coefficient, vector autoregression

$$x_t = \sum_{\tau=1}^p \phi(\tau) x_{t-\tau} + \varepsilon_t \quad (1)$$

which has been estimated from past observation on the x_t , $t = -S, \dots, 0$. Here $\phi(\tau)$ is an $r \cdot r$ matrix of autoregression coefficients and ε_t is an $r \cdot 1$ vector white noise with contemporary covariance matrix Q . This p 'th order vector autoregression can be written in the first-order form

$$\bar{x}_t = \begin{bmatrix} \phi(1) & \phi(2) & \dots & \phi(p) \\ & & & 0 \\ & & & \vdots \\ & I_{p-1} & & 0 \end{bmatrix} \bar{x}_{t-1} + \begin{bmatrix} I \\ 0^r \\ \vdots \\ 0 \end{bmatrix} \varepsilon_t$$

or

(2)

$$\bar{x}_t = \phi \bar{x}_{t-1} + G \varepsilon_t;$$

where \bar{x}_t is the $rp \cdot 1$ state vector $(x_t, x_{t-1}, \dots, x_{t-p+1})'$, ϕ is the $rp \cdot rp$ companion matrix, and G is constant $rp \cdot r$ matrix. Equation (2) constitutes the transition equation of the system.

Let y_t be an $m \cdot 1$ vector of of the restricted values attained by a linear combination of the system's variables in the post-sample period $t = 1, \dots, T$.

y_t is related to the state of the system \bar{x}_t by the measurement equation

$$y_t = S_t \bar{x}_t \quad (3)$$

where S_t is a fixed matrix of order m-rp. The measurement equation thus imposes m linear restrictions on the system at each post-sample date.

The Kalman filter consists of a set of recursive equations for calculating minimum mean square linear projections (MMSLP) of the variables contained in the \bar{x}_t given the information contained in measurements $Y_t = (y_t, y_{t-1}, \dots, y_1)'$. Let $\bar{x}_{t/t-1}$ be the best linear projection of \bar{x}_t given Y_{t-1} , $P_{t/t-1}$ the covariance matrix of the forecast error $\bar{x}_t - \bar{x}_{t/t-1}$, and H_t the covariance matrix of $y_t - y_{t/t-1} = S_t (\bar{x}_t - \bar{x}_{t/t-1})$. Note that since the initial state of the system x_0 is just the final observation in the sample of past \bar{x}_t $t = -S, \dots, 0$, $\bar{x}_{0/0} = x_0$ and $P_{0/0} = 0$. The set of recursions is given by the prediction equations:

$$\bar{x}_{t/t-1} = \Phi \bar{x}_{t-1/t-1}; \quad (4)$$

$$P_{t/t-1} = \Phi P_{t-1/t-1} \Phi' + GQG'; \quad (5)$$

$$H_t = S_t P_{t/t-1} S_t'; \quad (6)$$

and the updating equations:

$$\bar{x}_{t/t} = \bar{x}_{t/t-1} + P_{t/t-1} S_t H_t^{-1} (y_t - y_{t/t-1}); \quad (7)$$

$$P_{t/t} = P_{t/t-1} - P_{t/t-1} S_t H_t^{-1} S_t' P_{t/t-1}. \quad (8)$$

The prediction equations can be regarded as calculating recursively the conditional moments of the joint normal distribution of \bar{x}_t and y_t , updating the information set Y_{t-1} on each pass through the Kalman filter.

Each step in the Kalman filter yields the MMSLP of \bar{x}_t conditional on the information contained in Y_t . The only projection which incorporates all the post-sample restrictions is $\bar{x}_{T/T}$, the projection of the state in the final post-sample period. The method of Kalman smoothing may be used to calculate the MMSLP of each variable at each post-sample date which incorporates all the post-sample restrictions. These smoothed projections shall be denoted $x_{t/T}$.

There are several alternative methods of Kalman smoothing. We follow Anderson and Moore (1977) and stack the state vector so that only the Kalman filter recursions need to be used. Define the new $rT \cdot 1$ state vector

$\bar{x}_t \equiv (x_t, x_{t-1}, \dots, x_{t-T+1})'$. Then

$$\bar{x}_t = \begin{bmatrix} \phi(1) & \phi(2) & \dots & \phi(p) & \dots & 0 \\ I & 0 & \dots & \dots & \dots & 0 \\ 0 & I & \dots & \dots & \dots & 0 \\ \vdots & \vdots & \dots & \dots & \dots & \vdots \\ 0 & 0 & \dots & \dots & \dots & I & 0 \end{bmatrix} \bar{x}_{t-1} + \begin{bmatrix} I \\ 0^r \\ \vdots \\ \vdots \\ 0 \end{bmatrix} \epsilon_t$$

or

(9)

$$\bar{x}_t = \Phi \bar{x}_{t-1} + G \epsilon_t$$

where Φ is the $rT \cdot rT$ companion matrix and G is a fixed matrix of dimension $rT \cdot r$. Note that \bar{x}_T contains the entire post-sample trajectory of the forecast under study. With this augmented state vector, T passes through the Kalman filter will provide smoothed estimates of the variables in the system from 1 to T given the information in $Y_T = (\tilde{y}_T, \dots, \tilde{y}_1)'$. After the T 'th pass through, $\bar{x}_{T/T} = (x_{T/T}, x_{T-1/T}, \dots, x_{1/T})'$. It follows directly that the final step of the Kalman filter yields the MMSLP of the entire post-sample trajectory of the system subject to the restrictions on this path contained in Y_T .

There is a particularly simple relationship between the unconditional projection of \tilde{x}_T , $\tilde{x}_{T/0}$, and the projection which incorporates the information contained in Y_T , $\tilde{x}_{T/T}$. This is obtained by first setting $S_t = 0$, for $t = 1, \dots, T-1$. After $T-1$ passes through the filter, we obtain

$$\tilde{x}_{T/T-1} = \tilde{x}_{T/0} = \Phi^T x_0; \quad (10)$$

$$P_{T/T-1} = \Phi P_{T-1/T-1} \Phi' + GQG'; \quad (11)$$

For the final pass, we use the S_T which imposes the appropriate constraints on \tilde{x}_T so that

$$\tilde{y}_T = S_T \tilde{x}_T. \quad (12)$$

Then,

$$\tilde{x}_{T/T} = \Phi^T x_0 + K_T (\tilde{y}_T - S_T \Phi^T x_0) \quad (13)$$

where,

$$K_T = P_{T/T-1} S_T (S_T' P_{T/T-1} S_T)^{-1} \quad (14)$$

This implies (cf. Doan, Litterman, and Sims (1983) equation (22), p.37)

$$\tilde{x}_{T/T} = \tilde{x}_{T/0} + K_T (\tilde{y}_T - \tilde{y}_{T/0}). \quad (15)$$

Thus the difference between the unconditional projection $\tilde{x}_{T/0}$ and the projection which incorporates the post-sample restrictions $\tilde{x}_{T/T}$ is just an affine transformation of the difference between the post-sample constraints \tilde{y}_T and their best linear projection $\tilde{y}_{T/0}$.

2: Testing the Plausibility and Internal Consistency of a Forecast

As Doan, Litterman, and Sims emphasize, there is no unambiguously correct way to measure how likely it is that a particular condition on the projected post-sample path of a system will be realized. One possibility, for the case of a normal distribution over the post-sample paths, is to measure the plausibility of a set of linear restrictions directly by the significance level of the associated chi-squared statistic. To do this, we first construct the vector $\bar{v}_T = \bar{x}_T - \bar{x}_{T/T}$ where it should be recalled that $\bar{x}_{T/T}$ is the MMSLP of \bar{x}_T which incorporates all of the post-sample constraints $Y_T = (\bar{y}_T, \dots, \bar{y}_1)'$. This vector is multivariate normal with zero mean and covariance matrix $P_{T/T}$. It follows that the statistic

$$(\bar{x}_T - \bar{x}_{T/T})'(P_{T/T})^{-1}(\bar{x}_T - \bar{x}_{T/T}) \quad (16)$$

has the desired χ^2 distribution with mT degrees of freedom.

This procedure treats as the class of paths whose probability is to be measured all paths with lower likelihood than the most likely path satisfying the restrictions. An alternative "plausibility index" suggested by Doan, Litterman, and Sims instead looks only at paths lying on "one side" of the claimed path. This index uses the square-root of the chi-squared statistic given by (16) as if it were a normal random variable and measures plausibility by the probability in the upper tail of the normal p.d.f. at the level of this statistic. Whichever index is chosen, it is clear from (16) and equations (4) through (8) that it can be constructed directly from the output of the T'th pass through the Kalman algorithm.

3: Empirical Implementation

The Kalman recursions require as input an initial state vector \tilde{x}_0 , a companion matrix Φ , the variance-covariance matrix Q , and the matrix S which imposes the restrictions on the post-sample trajectory of the state vector. Short subroutines can easily be written to construct the Φ and Q matrices directly from the output files of the widely available Doan - Litterman RATS program. The S matrix is also easily constructed. For example, suppose that the values of x_t^1 , the first variable in x_t , are constrained to follow the path

$$Y_T = (x_T^1, \dots, x_1^1). \quad (17)$$

The appropriate S matrix is just the 1-rp vector $S = (1 \ 0 \ \dots \ 0)$.

Kalman filtering programs are readily available. The IMSL Fortran subroutine SFKALM has been used by the authors. For a six variable system and fifteen period forecast horizon, the computations used minutes of CPU time on a VAX/VMS system. The number of passes through the filter is equal to the number of forecast periods. However, only the output $\tilde{x}_{T/T}$ and $P_{T/T}$ from the last pass is required.

Computationally, there are several advantages to the state-space approach. In calculating the constrained conditional projections, the Kalman methods replace the cumbersome, if not intractable, matrix inversions required by the Doan, Litterman, and Sims method with a sequence of matrix multiplications. In addition, the Doan, Litterman, and Sims approach requires the non-trivial $rT \cdot rT$ matrix of moving average coefficients which does not appear, at least to us, as easily constructed from the RATS output files. In particular, the non-trivial block of the companion matrix Φ used in the Kalman filter is invariant to the number of post-sample periods under study.

4: Concluding Remarks

In this note we have established that the recursive, state-space methods of Kalman filtering and smoothing can be used to implement the Doan, Litterman, and Sims (1983) approach to econometric forecast and policy evaluation. Compared with the methods outlined in Doan, Litterman, and Sims, the Kalman algorithms are more easily programmed and modified to incorporate different linear constraints, avoid cumbersome matrix inversions, and provide estimates of the full variance-covariance matrix of the constrained projections which can be used directly, under standard normality assumptions, to test statistically the likelihood and internal consistency of the forecast under study. It is our hope that the wide understanding and ready availability of Kalman filtering technology will allow the Doan, Litterman, and Sims approach to be more easily, and thus, frequently employed. For a recent application of the methods outlined in this note, see Clarida and Friedman (1984).

REFERENCES

Anderson, R. and J. Moore, Optimal Filtering, Prentice-Hall: New York, 1979.

Clarida, R. and B. Friedman, "The Behavior of U.S. Short Term Interest Rates Since October 1979," forthcoming, The Journal of Finance, 1984.

Doan, T., R. Litterman and C. Sims, "Forecasting and Conditional Projection Using Realistic Prior Distributions," Mimeo, National Bureau of Economic Research, 1983.

Engle, R., "Wald, Likelihood Ratio and Lagrange Multiplier Statistics in Econometrics," The Handbook of Econometrics, Z. Grilliches and M. Intrilligator eds., North Holland: New York, 1982.

Harvey, A., Time Series Models, Halsted Press: New York, 1981.

Sims, C., "Policy Analysis with Econometric Models," The Brookings Papers on Economic Activity, I: 1982.