6

Estimation, Smoothing, Interpolation, and Distribution for Structural Time-Series Models in Continuous Time

A. C. Harvey and James H. Stock

Introduction

Rex Bergstrom's work has stressed not just the technical aspects but also the philosophical basis for applying continuous time models to time-series data; see Bergstrom (1966, 1976, 1984). Because many economic variables are essentially continuous and decisions are made continuously, it is often more appealing to set up models in continuous time even though observations are made at discrete intervals. The dynamic structure of a model is then not dependent on the observation interval, something which may bear no relation to the underlying data generation process.

An application of continuous-time models emphasized in this paper is their use to estimate intermediate values of a discretely sampled time series. Adopting Litterman's (1983) terminology, estimation at points between observations will be termed interpolation for a stock variable (sampled at a point in time) and distribution for a flow (sampled as an integral over a time interval). Using a continuous-time model in this context is appealing for several reasons. First, as emphasized by the contributors to Bergstrom (1976), the continuous-time framework provides a logically consistent basis for the handling of stocks and flows. Second, it provides a natural conceptual framework, with considerable technical simplifications, for handling irregularly spaced observations. Third, it provides a well-defined framework for interpolation and distribution to arbitrary subintervals.

Historically, a key technical hurdle in applying continuous-time models to economic data has been the difficulty of evaluating the exact Gaussian likelihood for flow data and for mixed stock-flow systems. These problems have largely been solved for large classes of models by Bergstrom (1983, 1984, 1985, 1986) and by his students and collaborators. Here, we consider interpolation for stocks and distribution for flows. To simplify the discussion, we restrict attention to univariate series.

This article studies continuous-time formulations within the context of structural time-series models in the sense of Harvey (1989). Structural models are formulated directly in terms of components of interest, such as trends, seasonals, and cycles. These components are functions of time and it is natural to regard them as being continuous. The essence of a structural time-series model is that its components are stochastic rather than deterministic. A continuous-time model can be set up to parameterize these stochastic movements. It can then be shown that, for the principal structural time-series models, the implied discrete-time model is, apart from some minor differences, of the same form as a discrete-time model which one would set up without reference to the continuous-time formulation. This is true for both stocks and flows. Thus there is a logical consistency in the structural class.

Since the components of structural time-series models have a direct interpretation, these models can often be specified without a detailed initial analysis of the data. The appropriateness of a particular specification is then checked by various diagnostics, and the whole model selection exercise is much more akin to what it is in econometrics; see for example Harvey (1985). Thus although data are not available on a continuous basis, the greater emphasis on prior considerations in model specification means that it is just as easy to adopt a continuous-time model formulation as a discrete one.

The next two sections examine the exact discrete-time models implied by the underlying continuous-time structural models sampled at the observation timing interval. This is straightforward for stock variables, less so for flows. In each case we consider the following statistical problems: time domain estimation of the model parameters by maximum likelihood; prediction of future observations; estimation of the unobserved components at the observation points and at intermediate points; and estimation of what the observations themselves would have been at intermediate points.

Structural Time-series Models in Discrete Time

A structural time-series model is one which is set up in terms of components which have a direct interpretation. For an economic time series, these components will typically consist of a trend, a seasonal, an irregular, and perhaps even a cycle. Examples of the application of such models can be found in Engle (1978), Harvey (1985), and Kitagawa (1981). Other components can be brought into the model. For example, daily or weekly components can be included if appropriate data are available. A general review can be found in Harvey (1989). In the present article, attention is restricted primarily to trend, cycle, seasonal, and irregular components, defined as follows.

Trend The level, μ_t, and slope, β_t, are generated by the multivariate random walk process,

$$\mu_t = \mu_{t-1} + \beta_{t-1} + \eta_t, \tag{1a}$$

$$\beta_t = \beta_{t-1} + \zeta_t, \tag{1b}$$

where η_t and ζ_t are mutually uncorrelated white-noise processes with zero means and variances σ_η^2 and σ_ζ^2 respectively.

Cycle The cycle, ψ_t, is stationary and is centered on a frequency λ_c, which lies in the range $[0, \pi]$. Its statistical formulation is

$$\begin{bmatrix} \psi_t \\ \psi_t^* \end{bmatrix} = \rho \begin{bmatrix} \cos \lambda_c & \sin \lambda_c \\ -\sin \lambda_c & \cos \lambda_c \end{bmatrix} \begin{bmatrix} \psi_{t-1} \\ \psi_{t-1}^* \end{bmatrix} + \begin{bmatrix} \kappa_t \\ \kappa_t^* \end{bmatrix}, \tag{2}$$

where κ_t and κ_t^* are uncorrelated white-noise disturbances with a common variance σ_κ^2, and ρ is a damping factor which lies in the range $0 \leqslant \rho \leqslant 1$.

Seasonal The seasonal component, γ_t, is defined as the sum of an appropriate number of trigonometric terms, γ_{jt}, each having a specification of the form (2) with ρ equal to unity and λ_c equal to a given seasonal frequency, $\lambda_j = 2\pi j/s$. Thus $\gamma_t = \sum_{j=1}^{s/2} \gamma_{jt}$, where s is the number of "seasons" (assumed to be even) and where

$$\begin{bmatrix} \gamma_{jt} \\ \gamma_{jt}^* \end{bmatrix} = \begin{bmatrix} \cos \lambda_j & \sin \lambda_j \\ -\sin \lambda_j & \cos \lambda_j \end{bmatrix} \begin{bmatrix} \gamma_{j,t-1} \\ \gamma_{j,t-1}^* \end{bmatrix} + \begin{bmatrix} \omega_{jt} \\ \omega_{jt}^* \end{bmatrix} \tag{3}$$

with $\mathrm{Var}(\omega_{jt}) = \mathrm{Var}(\omega_{jt}^*) = \sigma_\omega^2$ for all j.

Irregular The irregular term, ε_t, is generally taken to be a white-noise process, with variance σ_ε^2, unless there are strong *a priori* grounds to assume otherwise, as in Hausman and Watson (1985).

These components – trend, cycle, seasonal, and irregular – combine in various ways to give the principal structural time-series models for an observed series, Y_t, $t = 1, \ldots, T$. These models are:

1. *Local Linear Trend* The discrete-time process obeys

$$Y_t = \mu_t + \varepsilon_t, \tag{4}$$

where μ_t is a stochastic trend of the form (1) and ε_t is a white-noise irregular term.

2. *Local Level* This is a special case of the local linear trend in which μ_t is just a random walk:

$$\mu_t = \mu_{t-1} + \eta_t. \tag{5}$$

3. *Basic Structural Model with Cycles* Both seasonal and cyclical components may be brought into the model by expanding (4) to give

$$Y_t = \mu_t + \gamma_t + \psi_t + \varepsilon_t. \tag{6}$$

Each of these models can be handled statistically by putting them in state space form:

$$\boldsymbol{\alpha}_t = T_t \boldsymbol{\alpha}_{t-1} + R_t \boldsymbol{\eta}_t, \ \mathrm{Var}(\boldsymbol{\eta}_t) \equiv Q_t \tag{7a}$$

$$Y_t = z_t' \boldsymbol{\alpha}_t + \varepsilon_t, \ \mathrm{Var}(\varepsilon_t) \equiv h_t \tag{7b}$$

where $\boldsymbol{\alpha}_t$ is an $m \times 1$ state vector and ε_t and $\boldsymbol{\eta}_t$ are respectively scalar and $g \times 1$ zero mean white-noise disturbances which are mutually uncorrelated. The matrices z_t, T_t, R_t, and Q_t are $m \times 1$, $m \times m$, $m \times g$ and $g \times g$ respectively. These matrices may depend on a number of parameters, known as hyperparameters. Thus, for example, in the local linear trend model, (4), the hyperparameters are the variances σ_η^2, σ_ζ^2 and σ_ε^2. In some circumstances z_t and R_t will depend on time, say if there are some missing observations. More often, z_t and R_t will be time-invariant and henceforth will be denoted by z and R.

The state vector may be estimated by the Kalman filter. Furthermore, if the disturbances are normally distributed, the unknown hyperparameters may be estimated by maximum likelihood via the prediction error decomposition; see Ansley and Kohn (1985), De Jong (1991) and Harvey (1989, chapter 4).

General State Form of Continuous Time Models with Stocks and Flows

This section summarizes some results for the general continuous-time model when the data are observed at T irregular observation times $\{t_\tau\}$, $\tau = 1, \ldots, T$. These observation times are separated by calendar time units δ_τ, so that $t_\tau \equiv t_{\tau-1} + \delta_\tau$. The continuous-time state vector is denoted by $\alpha(t)$; at the observation times, it is denoted by $\alpha_\tau \equiv \alpha(t_\tau)$. Thus, in the notational convention adopted here, $\alpha(t)$, $y(t)$, etc. denote continuous-time processes, and α_τ, y_τ, etc. denote these processes at the appropriate discrete-time sampling dates. The observable (discrete time) process is Y_τ, and the data are T observations on the (discrete time) time series $(Y_1, \ldots, Y_T)$.

The continuous-time analog of the time-invariant discrete-time transition equation in (7a) is

$$d\alpha(t) = A\alpha(t)\,dt + R\,d\eta(t) \tag{8}$$

where the matrices A and R are $m \times m$ and $m \times g$ respectively and may be functions of hyperparameters and $\eta(t)$ is a $g \times 1$ continuous-time multivariate Wiener process. For a discussion of the formal interpretation of linear stochastic differential equations, see Bergstrom (1983). The Wiener process has independent increments that are Gaussian with mean zero and covariance matrix

$$E\left[\int_r^s d\eta(t) \int_r^s d\eta(t)'\right] = (s - r)Q.$$

Suppose we have a univariate series of observations at time $\{t_\tau\}$ for $\tau = 1, \ldots, T$. For a stock variable the observations are defined by

$$Y_\tau = z'\alpha(t_\tau) + \varepsilon_\tau, \quad \tau = 1, \ldots, T, \tag{9}$$

where ε_τ is white-noise disturbance term with mean zero and variance σ_ε^2 which is uncorrelated with differences of $\eta(t)$ in all time periods. For a flow

$$Y_\tau = \int_{t_{\tau-1}}^{t_\tau} z'\alpha(r)\,dr + \int_{t_{\tau-1}}^{t_\tau} d\varepsilon(r), \quad \tau = 1, \ldots, T, \tag{10}$$

where $\varepsilon(t)$ is a continuous-time Gaussian process with uncorrelated increments, mean zero and variance σ_ε^2, which is uncorrelated with $\eta(t)$ in all time periods in that

$$E\left[\int_r^s d\eta(t) \int_p^q d\varepsilon(r)\right] = 0$$

for all $r < s$ and $p < q$.

The state space formulation of continuous-time models derives from the stochastic integral equations that properly define the stochastic differential equations (8). The relationship between the state vector at time t_τ and time $t_{\tau-1}$ is given by

$$\alpha(t_\tau) = e^{A\delta_\tau}\alpha(t_{\tau-1}) + \int_{t_{\tau-1}}^{t_\tau} e^{A(t_\tau - s)}R \, d\eta(s). \tag{11}$$

This yields the discrete-time transition equation,

$$\alpha_\tau = T_\tau \alpha_{\tau-1} + \eta_\tau, \; \tau = 1, \ldots, T, \tag{12}$$

where $T_\tau = e^{A\delta_\tau}$, $\alpha_\tau \equiv \alpha(t_\tau)$, and η_τ is a multivariate white-noise disturbance term with mean zero and covariance matrix

$$Q_\tau = \int_0^{\delta_\tau} e^{A(\delta_\tau - s)}RQR' \, e^{A'(\delta_\tau - s)} \, ds. \tag{13}$$

The condition for $\alpha(t)$ in (8) to be stationary is that the real parts of the characteristic roots of A are negative. Then

$$\alpha(t) = \int_{-\infty}^{t} e^{A(t-s)}R \, d\eta(s)$$

so that $E\alpha(t) = 0$ and

$$\text{Var}\,\{\alpha(t)\} = \int_0^\infty e^{As}RQR' \, e^{As'} \, ds \tag{14}$$

which provides initial conditions for $\alpha(t)$ in the Kalman filter.

Structural Components in Continuous Time

The main structural components have the following natural formulations in continuous time.

Trend The linear trend component is

$$d\begin{bmatrix}\mu(t)\\ \beta(t)\end{bmatrix} = \begin{bmatrix}0 & 1\\ 0 & 0\end{bmatrix}\begin{bmatrix}\mu(t)\\ \beta(t)\end{bmatrix}dt + \begin{bmatrix}d\eta(t)\\ d\zeta(t)\end{bmatrix}, \tag{15}$$

where the continuous time processes $\eta(t)$ and $\zeta(t)$ have mutually and serially uncorrelated increments and variances σ_η^2 and σ_ζ^2, respectively. The local level model obtains as a special case of (15) with $\beta(0) = 0$ and $\sigma_\zeta^2 = 0$.

Cycle The continuous-time cycle component is

$$\mathrm{d}\begin{bmatrix}\Psi(t)\\ \Psi^*(t)\end{bmatrix} = \begin{bmatrix}\log\rho & \lambda_c\\ -\lambda_c & \log\rho\end{bmatrix}\begin{bmatrix}\Psi(t)\\ \Psi^*(t)\end{bmatrix}\mathrm{d}t + \begin{bmatrix}\mathrm{d}\kappa(t)\\ \mathrm{d}\kappa^*(t)\end{bmatrix}, \tag{16}$$

where $\kappa(t)$ and $\kappa^*(t)$ have mutually and serially uncorrelated increments and the same variance, σ_κ^2, and ρ, and λ_c are parameters, the latter being the frequency of the cycle. The characteristic roots of the matrix containing these parameters are $(\log\rho) \pm \mathrm{i}\lambda_c$. Since the general condition for stationarity of a model of the form (8) is that the characteristic roots must have negative real parts, the condition for $\Psi(t)$ to be a stationary process is $\log\rho < 0$, which corresponds to $\rho < 1$.

Seasonal The continuous-time seasonal component is the sum of a suitable number of trigonometric components, $\gamma_j(t)$, generated by processes of the form (16) with ρ equal to unity and λ_c set equal to the appropriate seasonal frequency λ_j. That is, for $j = 1, \ldots, s/2$,

$$\mathrm{d}\begin{bmatrix}\gamma_j(t)\\ \gamma_j^*(t)\end{bmatrix} = \begin{bmatrix}0 & \lambda_j\\ -\lambda_j & 0\end{bmatrix}\begin{bmatrix}\gamma_j(t)\\ \gamma_j^*(t)\end{bmatrix}\mathrm{d}t + \begin{bmatrix}\mathrm{d}\omega_j(t)\\ \mathrm{d}\omega_j^*(t)\end{bmatrix}, \tag{17}$$

where $\omega_j(t)$ and $\omega_j^*(t)$ are processes with serially and mutually uncorrelated increments and with equal variance σ_ω^2.

Continuous-time structural models constructed from these components constitute special cases of the general process (8) and of more general continuous-time processes such as those studied by Phillips (1988). The aim of these structural models is to provide a practical framework for forecasting and – in the continuous-time setting – interpolation and distribution.

Stock Variables

The discrete state space form for a stock variable generated by a continuous-time process consists of the transition equation (12) together with the measurement equation (9). The Kalman filter can therefore be applied in a standard way. When the observations are equally spaced the implied discrete time model is time-invariant and typically it is convenient to set $\delta_\tau = 1$. One of the main practical advantages, however, of the continuous-time framework is the easy handling of irregularly spaced observations, so the case of general δ_τ is considered here. For related applications, see Jones (1984) and Kitagawa (1984).

Structural Models

The continuous-time components defined above can be combined to produce a continuous-time structural model. As in the discrete case, the components are usually assumed to be mutually uncorrelated. Hence the A and Q matrices in (13) are block diagonal and so the discrete-time components can be evaluated separately.

Trend For the local linear trend model (15), $T_\tau = e^{A\delta_\tau}$ is

$$T_\tau = \exp\left\{\begin{bmatrix} 0 & 1 \\ 0 & 0 \end{bmatrix}\delta_\tau\right\} = I + \begin{bmatrix} 0 & \delta_\tau \\ 0 & 0 \end{bmatrix} = \begin{bmatrix} 1 & \delta_\tau \\ 0 & 1 \end{bmatrix}.$$

Thus the exact discrete-time representation of (15) is

$$\begin{bmatrix} \mu_\tau \\ \beta_\tau \end{bmatrix} = \begin{bmatrix} 1 & \delta_\tau \\ 0 & 1 \end{bmatrix}\begin{bmatrix} \mu_{\tau-1} \\ \beta_{\tau-1} \end{bmatrix} + \begin{bmatrix} \eta_\tau \\ \zeta_\tau \end{bmatrix}. \tag{18a}$$

In view of the simple structure of this matrix exponential, the evaluation of the covariance matrix of the discrete-time disturbances can be carried out explicitly, yielding

$$\mathrm{Var}\begin{bmatrix} \eta_\tau \\ \zeta_\tau \end{bmatrix} = \delta_\tau \begin{bmatrix} \sigma_\eta^2 + \delta_\tau^2\sigma_\zeta^2/3 & \frac{1}{2}\delta_\tau\sigma_\zeta^2 \\ \frac{1}{2}\delta_\tau\sigma_\zeta^2 & \sigma_\zeta^2 \end{bmatrix}. \tag{18b}$$

When δ_τ is equal to unity, (18) reduces to the discrete-time local linear trend (1). However, while in (1) the disturbances are uncorrelated, (18b) shows that uncorrelatedness of the continuous-time disturbances implies that the corresponding discrete-time disturbances are correlated.

The discrete-time model for the local level trend obtains by setting $\beta_0 = 0$ and $\sigma_\zeta^2 = 0$ in (18), in which case μ_t evolves in discrete time as a random walk. Thus with only a local level trend and the irregular term ε_t, $Y_\tau = \mu_\tau + \varepsilon_\tau$, $\tau = 1, \ldots, T$, so that Y_τ evolves according to the familiar random-walk-plus-noise model.

Cycle For the cycle model (16), use of the matrix exponential definition together with the power series expansions for the cosine and sine functions gives the discrete-time model

$$\begin{bmatrix} \Psi_\tau \\ \Psi_\tau^* \end{bmatrix} = \rho^{\delta_\tau}\begin{bmatrix} \cos\lambda_c\delta_\tau & \sin\lambda_c\delta_\tau \\ -\sin\lambda_c\delta_\tau & \cos\lambda_c\delta_\tau \end{bmatrix}\begin{bmatrix} \Psi_{\tau-1} \\ \Psi_{\tau-1}^* \end{bmatrix} + \begin{bmatrix} \kappa_\tau \\ \kappa_\tau^* \end{bmatrix}. \tag{19}$$

When δ_τ is one, the transition matrix corresponds exactly to the transition matrix of the discrete-time cyclical component (2). As regards the properties of the disturbances, specifying that $\kappa(t)$ and $\kappa^*(t)$ be mutually uncorrelated with equal variances means that in the corresponding discrete time model κ_τ and κ_τ^* will also be uncorrelated with the same variance for any δ_τ. In fact, the covariance matrix of $(\kappa_\tau, \kappa_\tau^*)'$ is $(-\sigma_\kappa^2/2\log\rho)(1 - \rho^{2\delta_\tau})I$.

There are three noteworthy parallels between (19) and the cyclical process as originally defined in discrete time. First, as in the discrete analog, letting $\kappa(t)$ and $\kappa^*(t)$ be uncorrelated with equal variances imposes one more restriction than is necessary for identifiability. However, it ensures that the specification of the discrete-time model is consistent with the continuous-time model. Second, setting λ_c equal to zero in (19) means that Ψ_t collapses to a continuous-time AR(1) process exhibiting positive autocorrelation. Third, a pseudocyclical process is also obtained from a continuous time AR(2) or ARMA(2,1) model since, as shown in Phadke and Wu (1974), the roots of the corresponding discrete time AR(2) polynomial are typically complex. One attraction of working with the cycle model is that it is easier to handle as well as being set up in terms of the parameters of interest.

Seasonal For a trigonometric seasonal component, (17), the discrete time specification is similar to (19) with $\rho = 1$. The covariance matrix of the disturbance is $\sigma_\omega^2 \delta_\tau I$. When $\delta_\tau = 1$, this specification corresponds exactly to that given in (3).

As regards starting values for the Kalman filter for irregularly spaced observations, the considerations which arise are almost exactly as for a conventional discrete-time model. Thus, if the state vector contains d nonstationary components, d observations are needed before the elements of the state vector can be estimated with finite MSE. A diffuse prior can be used to initiate the nonstationary elements of the state vector. The starting values for the stationary components are provided by their respective unconditional distributions. Thus for the stationary cyclical component (16), the unconditional mean of $[\Psi(t), \Psi^*(t)]'$ is zero while the covariance matrix is obtained by evaluating (14). Thus, because $RQR' = \sigma_\kappa^2 I$, Var $(\alpha_0) = \sigma_\kappa^2 \int_0^\infty e^{(A+A')s}\, ds$. However $A + A' = (2 \log \rho)I$, so Var $(\alpha_0) = (-\sigma_\kappa^2/2 \log \rho)I$.

In summary, putting the various continuous-time components together yields discrete-time models which, for regularly spaced observations, are almost identical to the discrete-time models set up in equations (1) to (7). In fact, the only case where the implied discrete time model is different from the model originally formulated is the local linear trend and there the difference is minor. Thus, for a stock variable, the principal discrete-time structural models are consistent with the corresponding continuous-time models. The continuous-time models are more general, however, in that they can handle irregularly spaced observations and interpolate at any point.

Smoothing and Interpolation

Interpolation is the estimation of the series and/or its components at some point between observations. This may be carried out by defining the required points, constructing the appropriate transition equations, and treating the corresponding observations as missing. Thus, suppose interpolation is to be carried out at J points $t_\tau + r_1$, $t_\tau + r_2$, and so on where $0 < r_1 < r_2 < \cdots < r_J \leqslant \delta_{\tau+1}$. All that is required is the definition of the discrete-time transition equation (12) at $t_\tau + r_1, t_\tau + r_2, \ldots, t_{\tau+1}$ and the subsequent applications of the Kalman filter. The optimal estimators of $y(t)$ at these intermediate points are then obtained by applying a suitable smoothing algorithm; see for example Anderson and Moore (1979) or Harvey (1989, chapter 3). The estimators at time $t_\tau + r_j$ may be written as $y(t_\tau + r_j | T)$. This yields the minimum MSE linear estimator of an observation at time $t_\tau + r_j$, $y(t_\tau + r_j|T) = z'a(t_\tau + r_j|T)$, where $a(t|\tau)$ denote the optimal predictor of $\alpha(t)$ using data through the τth observation (i.e. the observable time series through calendar date t_τ). The corresponding MSE is MSE $[y(t_\tau + r_j|T)] = z'P(t_\tau + r_j|T)z + \sigma_\varepsilon^2$, $j = 1, 2, \ldots, J$, where $P(t_\tau + r_j|T)$ is the MSE matrix of $\alpha(t_\tau + r_j|T)$, the smoothed estimator of the state vector at time $t_\tau + r_j$.

Prediction

Let $a_\tau \equiv a(t_\tau|\tau)$. In the general model (8), the optimal predictor of the state vector for any positive lead time, l (i.e. at time $t_T + l$, given data through calendar time t_T), is given by the forecast function $a(t_T + l \mid T) = e^{Al}a_T$. The state vector at lead time l satisfies:

$$\alpha(t_T + l) = e^{Al}\alpha_T + \int_{t_T}^{t_T+l} e^{A(t_T+l-s)}R \, d\eta(s) \tag{20}$$

and so the MSE matrix associated with $a(t_T + l \mid T)$ is $P(t_T + l \mid T) = T_l P_T T_l' + Q_l$, where $T_l = e^{Al}$ and Q_l is given by (13) evaluated with $\delta_\tau = l$.

The forecast function for the systematic part of the series, $\bar{y}(t) \equiv z'\alpha(t)$, can also be expressed as a continuous function of l, namely $\bar{y}(t_T + l \mid T) = z' e^{Al}a_T$. Note that, if this is considered to be the forecast of an observation to be made at time $t_T + l$, then we can simply set $\delta_{T+1} \equiv l$, so that $Y_{T+1 \mid T} = \bar{y}(t_T + l \mid T)$. Thus $Y_{T+1 \mid T} = z' e^{Al}a_T$. Here, the observation to be forecast has arbitrarily been classified as the one indexed by $\tau = T + 1$. The MSE of this forecast obtains directly as $\text{MSE}\,(Y_{T+1 \mid T}) = E(Y_{T+1} - Y_{T+1 \mid T})^2 = z'P(t_T + l \mid T)z + \sigma_\varepsilon^2$.

The evaluation of forecast functions for the various structural models is relatively straightforward. For example, consider the local level model with measurement equation $Y_\tau = \mu(t_\tau) + \varepsilon_\tau$, $\tau = 1, \ldots, T$, where $\text{Var}\,(\varepsilon_\tau) = \sigma_\varepsilon^2$. This model has the forecast function, $y(t_T + l \mid T) = \mu(t_T + l \mid T) = \mu_{T \mid T}$, which is simply a horizontal straight line passing through the final estimate of the trend. The MSE of the forecast of the $(T + 1)$th observation (to be made at calendar time $t_T + l$) is $\text{MSE}\,(Y_{T+1 \mid T}) = P_T + l\sigma_\eta^2 + \sigma_\varepsilon^2$.

The forecast functions for more complicated components models obtain directly. For example, introducing a slope component into the trend – the local linear trend model – yields the straight line forecast function, $y(t_T + l \mid T) = \mu(t_T + l \mid T) = \mu_{T \mid T} + \beta_{T \mid T}l$. Also, the forecast function for a cyclical component takes the form of a damped cosine wave, $\Psi(t_T + l \mid T) = \rho^l[(\cos \lambda_c l)\Psi_{T \mid T} + (\sin \lambda_c l)\Psi_{T \mid T}^*]$. The forecast function for the seasonal component has the form of the cyclical component with $\rho = 1$ (no damping factor).

Estimation

The estimation of Gaussian continuous-time structural models poses no new problems when the observations are regularly spaced: algorithms already developed can be applied directly with only a minor modification needed to handle the covariance matrix of the disturbances of the local linear trend model in (18b). On the other hand, when the observations are irregularly spaced, the time domain estimation procedure must be adapted to account for the fact that the state space model is no longer time invariant. Once this has been done, the construction of the likelihood function can proceed via the prediction error decomposition as implemented by the Kalman filter; see for example Jones (1981) or Harvey and Stock (1985).

Flow Variables

Observed flow variables were defined in (10). To develop a state-space model of flow variables in continuous time, it is useful to introduce a continuous time cumulator (or integrator) variable, $y^f(t)$. This cumulator is defined as

$$y^f(t_\tau + l) = \int_{t_\tau}^{t_\tau + l} y(r)\,dr, \qquad 0 < l \leqslant \delta_\tau. \tag{21}$$

Thus $Y_\tau = y^f(t_\tau)$ for $\tau = 1, \ldots, T$. This definition, (10) and (11) imply that the cumulator at time t_τ can be written as

$$\begin{aligned} y^f(t_\tau) &= \int_0^{\delta_\tau} y(t_{\tau-1} + r)\,dr = z' \int_0^{\delta_\tau} \alpha(t_{\tau-1} + r)\,dr + \int_{t_{\tau-1}}^{t_\tau} d\varepsilon(r) \\ &= z' \int_0^{\delta_\tau} \left[e^{Ar}\alpha(t_{\tau-1}) + \int_{t_{\tau-1}}^{t_{\tau-1}+r} e^{A(t_{\tau-1}+r-s)} R\,d\eta(s) \right] dr + \int_{t_{\tau-1}}^{t_\tau} d\varepsilon(r) \\ &= z' \left[\int_0^{\delta_\tau} e^{Ar}\,dr \right] \alpha(t_{\tau-1}) + z' \int_0^{\delta_\tau} \int_{t_{\tau-1}}^{t_{\tau-1}+r} e^{A(t_{\tau-1}+r-s)} R\,d\eta(s)\,dr + \int_{t_{\tau-1}}^{t_\tau} d\varepsilon(r) \\ &= z' W(\delta_\tau)\alpha(t_{\tau-1}) + z'\eta^f(t_\tau) + \varepsilon^f(t_\tau), \end{aligned} \tag{22}$$

where $\eta^f(t_\tau) = \int_{t_{\tau-1}}^{t_\tau} W(t_\tau - s) R\,d\eta(s)$, $\varepsilon^f(t_\tau) = \int_{t_{\tau-1}}^{t_\tau} d\varepsilon(s)$, and $W(r) = \int_0^r e^{As}\,ds$. Now letting $\eta_\tau^f \equiv \eta^f(t_\tau)$, $\varepsilon_\tau^f \equiv \varepsilon^f(t_\tau)$, and $y_\tau^f \equiv y^f(t_\tau)$, and remembering that $y^f(t_\tau) = Y_\tau$, we have, on combining (12) with (22), the augmented state space form

$$\begin{bmatrix} \alpha_\tau \\ y_\tau^f \end{bmatrix} = \begin{bmatrix} e^{A\delta_\tau} & 0 \\ z'W(\delta_\tau) & 0 \end{bmatrix} \begin{bmatrix} \alpha_{\tau-1} \\ y_{\tau-1}^f \end{bmatrix} + \begin{bmatrix} I & 0 \\ 0 & z' \end{bmatrix} \begin{bmatrix} \eta_\tau \\ \eta_\tau^f \end{bmatrix} + \begin{bmatrix} 0 \\ \varepsilon_\tau^f \end{bmatrix}, \tag{23a}$$

$$Y_\tau = (0 \quad 1) \begin{bmatrix} \alpha_\tau \\ y_\tau^f \end{bmatrix} \tag{23b}$$

with $\mathrm{Var}(\varepsilon_\tau^f) = \delta_\tau \sigma_\varepsilon^2$ and

$$\mathrm{Var} \begin{bmatrix} \eta_\tau \\ \eta_\tau^f \end{bmatrix} = \int_0^{\delta_\tau} \begin{bmatrix} e^{Ar} RQR' e^{A'r} & e^{Ar} RQR' W(r)' \\ W(r) RQR' e^{A'r} & W(r) RQR' W(r)' \end{bmatrix} dr \equiv Q_\tau^\dagger. \tag{24}$$

Maximum likelihood estimators of the hyperparameters can be computed via the prediction error decomposition by running the Kalman filter on (23). No additional starting value problems are caused by bringing the cumulator variable into the state vector as $y^f(t_0)$ is zero by construction.

An alternative, numerically equivalent approach is to treat the second equation in (23a) as a measurement equation rather than a state equation. Define $\alpha_\tau^* \equiv \alpha_{\tau-1}$. Then (23) can be rewritten as

$$\alpha_{\tau+1}^* = T_{\tau+1}^* \alpha_\tau^* + \eta_\tau, \tag{25a}$$

$$Y_\tau = z_\tau^{*\prime} \alpha_\tau^* + \varepsilon_\tau^*, \tag{25b}$$

where $z_\tau^{*\prime} \equiv z' W(\delta_\tau)$, $\varepsilon_\tau^* \equiv z'\eta_\tau^f + \varepsilon_\tau^f$, and $T_{\tau+1}^* \equiv e^{A\delta_\tau}$. Taken together equations (25) are a state model in which the measurement equation disturbance, ε_τ, and the transition equation disturbance, η_τ, are correlated, in contrast to the standard formulation with uncorrelated disturbances. The covariance matrix of $(\boldsymbol{\eta}_\tau', \boldsymbol{\varepsilon}_\tau^*)'$ is given by

$$\mathrm{Var}\begin{bmatrix}\boldsymbol{\eta}_\tau \\ \boldsymbol{\varepsilon}_\tau^*\end{bmatrix} = \begin{bmatrix} Q_\tau & G_\tau \\ G_\tau' & H_\tau \end{bmatrix} = \begin{bmatrix} I & 0 \\ 0 & z' \end{bmatrix} Q_\tau^\dagger \begin{bmatrix} I & 0 \\ 0 & z \end{bmatrix} + \begin{bmatrix} 0 & 0 \\ 0 & \delta_\tau \sigma_\varepsilon^2 \end{bmatrix}. \tag{26}$$

The modified version of the Kalman filter needed to handle such systems is described in Jazwinski (1970, chapter 7) and Harvey (1989, chapter 3).

Structural Models

The various matrix exponential expressions which need to be computed for the flow variable are relatively easy to evaluate for trend and seasonal components. The formulas for a stationary cyclical component are rather more tedious to derive and so will not be given here explicitly. Because the components in a basic structural model are typically assumed to be independent of each other, the various blocks in (24) can be treated separately, as though there were only a single component in the model. This simplifies the development for general structural models. Because the top left-hand block in (24) is the same as the corresponding Q_τ matrix evaluated for a stock variable in the second section of this article, only the remaining two terms are derived here.

Trend For the local linear trend component (15),

$$W(r) = \int_0^r \begin{bmatrix} 1 & s \\ 0 & 1 \end{bmatrix} \mathrm{d}s = \begin{bmatrix} r & \frac{1}{2}r^2 \\ 0 & r \end{bmatrix}. \tag{27}$$

Thus, in (24), the lower right-hand matrix is

$$\mathrm{Var}\,(\boldsymbol{\eta}_\tau^{\mathrm{f}}) = \int_0^{\delta_\tau} W(r)RQR'W(r)'\,\mathrm{d}r = \begin{bmatrix} \delta_\tau^3\sigma_\eta^2/3 + \delta_\tau^5\sigma_\zeta^2/20 & \delta_\tau^4\sigma_\zeta^2/8 \\ \delta_\tau^4\sigma_\zeta^2/8 & \delta_\tau^3\sigma_\zeta^2/3 \end{bmatrix}, \tag{28}$$

where in this case $\boldsymbol{\eta}_\tau^{\mathrm{f}}$ is the 2×1 vector $(\eta_\tau^{\mathrm{f}} \quad \zeta_\tau^{\mathrm{f}})'$.

The off-diagonal blocks are derived in a similar way. Thus

$$\mathrm{Cov}\,(\eta_\tau^{\mathrm{f}}, \eta_\tau') = \int_0^{\delta_\tau} W(r)RQR'\,\mathrm{e}^{A'r}\,\mathrm{d}r = \begin{bmatrix} \frac{1}{2}\delta_\tau^2\sigma_\eta^2 + \delta_\tau^4\sigma_\zeta^2/8 & \delta_\tau^2\sigma_\zeta^2/6 \\ \delta_\tau^3\sigma_\zeta^2/3 & \frac{1}{2}\delta_\tau^2\sigma_\zeta^2 \end{bmatrix}. \tag{29}$$

The local level model is just a special case in which $\mathrm{Var}\,(\eta_\tau^{\mathrm{f}})$ and $\mathrm{Cov}\,(\eta_\tau^{\mathrm{f}}, \eta_\tau)$ are scalars consisting of the top left-hand elements of (27) and (28), respectively, with $\sigma_\zeta^2 = 0$.

Seasonal For a trigonometric component in the seasonal model (17),

$$W(r) = \int_0^r \begin{bmatrix} \cos\lambda s & \sin\lambda s \\ -\sin\lambda s & \cos\lambda s \end{bmatrix} \mathrm{d}s = \lambda^{-1} \begin{bmatrix} \sin\lambda r & 1-(\cos\lambda r) \\ (\cos\lambda r)-1 & \sin\lambda r \end{bmatrix}. \tag{30}$$

Thus, if η_τ^f in (23) relates to the disturbances in a trigonometric term,

$$\mathrm{Var}\,(\eta_\tau^f) = 2\delta_\tau\sigma_\omega^2/\lambda^2\begin{bmatrix} 1-(1/\lambda\delta_\tau)(\sin\lambda\delta_\tau) & 0 \\ 0 & 1-(1/\lambda\delta_\tau)(\sin\lambda\delta_\tau)\end{bmatrix}, \tag{31}$$

$$\mathrm{Cov}\,(\eta_\tau^f, \eta_\tau') = \sigma_\omega^2/\lambda^2\begin{bmatrix} 1-(\cos\lambda\delta_\tau) & (\sin\lambda\delta_\tau)-\lambda\delta_\tau \\ \lambda\delta_\tau-(\sin\lambda\delta_\tau) & 1-(\cos\lambda\delta_\tau)\end{bmatrix}. \tag{32}$$

The state space form with correlated disturbances in the measurement and transition equations, (25), shows that the structure of the discrete-time model corresponding to a particular continuous-time model is essentially the same for a flow as for a stock. For example, consider the continuous-time local level model $y(t) = \mu(t) + \varepsilon(t)$, with $\mu(t)$ given by the local level specialization of (15). In terms of (25), the state equation is $\mu_{\tau+1}^* = \mu_\tau^* + \eta_\tau$ and the measurement equation is $Y_\tau = \delta_\tau\mu_\tau^* + \varepsilon_\tau^*$, with

$$\mathrm{Var}\begin{bmatrix}\eta_\tau \\ \varepsilon_\tau^*\end{bmatrix} = \begin{bmatrix}\delta_\tau\sigma_\eta^2 & \tfrac{1}{2}\delta_\tau^2\sigma_\eta^2 \\ \tfrac{1}{2}\delta_\tau^2\sigma_\eta^2 & \delta_\tau^3\sigma_\eta^2/3 + \delta_\tau\sigma_\varepsilon^2\end{bmatrix}. \tag{33}$$

When the observations are evenly spaced, δ_τ can be set equal to unity. The forecasts formed by the steady-state Kalman filter for the local level model with flows are then equivalent to the exponentially weighted moving average,

$$Y_{\tau+1|\tau} = (1-\lambda)Y_{\tau|\tau-1} + \lambda Y_\tau, \tag{34}$$

where $0 < \lambda \leqslant 1.27$. Note that the smoothing constant has a maximum value of 1.27 (to two decimal places) rather than unity. The reason is that this model has a wider range of dynamic properties than the discrete-time model formulated in (4) and (5). Taking first differences of (34) with $\delta_\tau = 1$, and evaluating the autocorrelations, one obtains $\rho(1) = (q-6)/(4q+12)$, where $q = \sigma_\eta^2/\sigma_\varepsilon^2$ and, for $\tau \geqslant 2$, $\rho(\tau) = 0$. Whereas in (4) and (5), $\rho(1)$ is always negative, lying in the range $[-0.5, 0]$, in this case $\rho(1) \in [-0.5, 0.25]$.

One interesting consequence is that a series with a first-order autocorrelation of 0.25 in first differences can be modeled simply by a time-aggregated Brownian motion; compare Working (1960). A second point is that ΔY_τ follows a discrete-time random walk when $q = 6$. This means that a discrete-time random walk can be smoothed to a limited extent since the corresponding continuous-time model contains an additive disturbance term. Of course the same can be done when a discrete-time random-walk-plus-noise model is formulated at a finer timing interval than the observation interval; see Harvey (1989, chapter 6).

Smoothing and Distribution

Suppose that the observations are evenly spaced at intervals of δ and that one wishes to estimate certain integrals of linear combinations of the state vector at evenly spaced intervals Δ time periods apart where δ/Δ is a positive integer. For example, it might be desirable to distribute quarterly observations to a monthly level. The quantities to be estimated may be written in an $m^* \times 1$ vector as

$$\alpha^{\Delta}(t_i) \equiv \alpha_i^{\Delta} = \int_0^{\Delta} Z'\alpha(t_{i-1}+s)\,ds, \qquad t_i, i = 1, \ldots, (\delta/\Delta)T, \tag{35}$$

where Z is an $m^* \times m$ selection matrix. In the continuous-time basic structural model, the components of interest might be (for example) the level of the trend, the slope and the seasonal, in which case $Z'\alpha(t) = \{\mu(t), \beta(t), \gamma(t)\}'$. In addition it may be desirable to estimate the values of the series itself,

$$y_i^{\Delta} \equiv y(t_i) = \int_0^{\Delta} z'\alpha(t_{i-1}+s)\,ds + \int_{t_{i-1}}^{t_i} d\varepsilon(s). \tag{36}$$

Smoothed estimates of the quantities of interest may be obtained from an augmented discrete-time state space model. The transition equation is

$$\alpha_i^{\dagger} = \begin{bmatrix} \alpha_i \\ y_i^f \\ \alpha_i^{\Delta} \\ y_i^{\Delta} \end{bmatrix} = \begin{bmatrix} e^{A\Delta} & 0 & 0 & 0 \\ z'W(\Delta) & \phi_i & 0 & 0 \\ Z'W(\Delta) & 0 & 0 & 0 \\ z'W(\Delta) & 0 & 0 & 0 \end{bmatrix} \begin{bmatrix} \alpha_{i-1} \\ y_{i-1}^f \\ \alpha_{i-1}^{\Delta} \\ y_{i-1}^{\Delta} \end{bmatrix} + \begin{bmatrix} I & 0 \\ 0 & z' \\ 0 & Z' \\ 0 & z' \end{bmatrix} \begin{bmatrix} \eta_i \\ \eta_i^f \end{bmatrix} + \begin{bmatrix} 0 \\ 1 \\ 0 \\ 1 \end{bmatrix} \varepsilon_i^f, \tag{37}$$

with $\phi_i = 0$, $i = (\delta/\Delta)(\tau - 1) + 1$, $\tau = 1, \ldots, T$, and $\phi_i = 1$ otherwise. The covariance matrix of $(\eta_i', \eta_i^{f\prime})'$ is defined as in (24) with δ_τ replaced by Δ, while $\text{Var}(\varepsilon_i^f) = \Delta\sigma_\varepsilon^2$. The corresponding measurement equation is only defined at observation times and can be written as $Y_\tau = [0' \quad 1 \quad 0' \quad 0]\alpha_i^{\dagger}$ for $i = (\delta/\Delta)\tau$, $\tau = 1, \ldots, T$.

The distributed values of the series (the y_i^{Δ} terms) could alternatively be estimated by differencing the estimators of the y_i^f terms. The appearance of y_i^{Δ} in the state is really only necessary if the MSE of its estimator is required. Of course if $\Delta = \delta$, it becomes totally superfluous.

Predictions

In making predictions for a flow it is necessary to distinguish between the total accumulation from time t_τ to time $t_\tau + l$, which might include several unit time intervals, and the amount of the flow in a single time period ending at time $t_\tau + l$. The latter concept corresponds to the usual idea of prediction in a discrete model. Cumulative predictions are perhaps more natural in a continuous-time model, particularly when the observations are made at irregular intervals. Here, three types of predictions are discussed: cumulative predictions, predictions over the unit interval, and predictions over a variable lead time.

Cumulative Predictions Predictions for the cumulative effect of $y(t)$ are obtained by noting that the quantity required is $y^f(t_\tau + l)$ which in terms of the state space model (23) is y_{T+1} with $\delta_{T+1} = l$, where it is assumed that $l \geqslant 0$. The optimal predictor can therefore be obtained directly from the Kalman filter as can its MSE. Written out explicitly, $y^f(t_T + l \mid T) = Y_{T+1|T} = z'W(l)a_T$. Because the quantity to be estimated is

$y^f(t_T + l) = \int_0^l y(t_T + r)\,dr = z'W(l)\alpha_T + z'\eta^f_{T+1} + \varepsilon^f_{T+1}$, the prediction MSE is

$$\text{MSE}\,[y^f(t_T + l|T)] = z'W(l)P_T W(l)'z + z'\,\text{Var}\,(\eta^f_\tau)z + \text{Var}\,(\varepsilon^f_{T+1}). \tag{38}$$

Note that if the modified Kalman filter based on (25) is run, $\mathbf{a}_T = \mathbf{a}^*_{T+1|T}$ (because $\alpha^*_{T+1} = \alpha_T$) and so $\mathbf{a}^*_{T+1|T}$ is the optimal estimator of α_T based on all the observations.

As a simple example, consider the local level model. For $l > 0$,

$$y^f(t_T + l|T) = l\mu_{T|T}, \qquad \text{MSE}\,[y^f(t_T + l|T)] = l^2 P_T + l^3\sigma^2_\eta/3 + l\sigma^2_\varepsilon. \tag{39}$$

Corresponding expressions for the discrete-time models can be obtained; see for example Johnson and Harrison (1986). However, the derivation of (39) is both simpler and more elegant. For the local linear trend, (27) gives

$$y^f(t_T + l|T) = l\mu_{T|T} + \tfrac{1}{2}l^2\beta_{T|T}, \tag{40}$$

$$\text{MSE}\,[y^f(t_T + l|T)] = l^2 P_{11T} + l^3 P_{12T} + l^4 P_{22T}/4 + l^3\sigma^2_\eta/3 + l^5\sigma^2_\zeta/20 + l\sigma^2_\varepsilon, \tag{41}$$

where P_{ijT} is the (i,j) element of P_T.

Predictions over the Unit Interval Predictions over the unit interval emerge quite naturally from the state space form (23) as the predictions of $\mathbf{Y}_{T+l}$, $l = 1, 2, \ldots$ with δ_{T+l} set equal to unity for all l. The forecast function for the state vector, $\mathbf{a}_{T+l} = e^{Al}\mathbf{a}_T$, has the same form as in the corresponding stock variable model. The presence of the term $W(1)$ in (23a) leads to a slight modification when these forecasts are translated into a prediction for the series itself. Specifically,

$$\mathbf{Y}_{T+l|T} = z'W(1)\mathbf{a}_{T+l-1|T} = z'W(1)\,e^{A(l-1)}\mathbf{a}_T, \qquad l = 1, 2, \ldots. \tag{42}$$

As a special case, in the local linear trend model $\mathbf{a}_{T+l-1|T} = [\mu_{T|T} + (l-1)\beta_{T|T}, \beta_{T|T}]'$, so $\mathbf{Y}_{T+l|T} = \mu_{T|T} + (l - \tfrac{1}{2})\beta_{T|T}$ for $l = 1, 2, \ldots$. The one-half arises here because each observation is cumulated over the unit interval.

Predictions over a Variable Lead Time In some applications the lead time itself can be regarded as a random variable. This happens, for example, in inventory control problems where an order is put in to meet demand, but the delivery time is uncertain. In such situations it may be useful to determine the unconditional distribution of the cumulation of $y(t)$ from the current point in time, T. Assume the random lead time is independent of $\{y(s)\}$. The unconditional PDF of this cumulation from t_T to $t_T + l$ is

$$p(y^f(t_T + l)|M_T) = \int p(y^f(t_T + l)|M_T, l)\,dF(l), \tag{43}$$

where $F(l)$ is the distribution of lead times and $p(y^f(t_T + l)|M_T, l)$ is the predictive distribution of $y^f(t)$ at time $t_T + l$, that is the distribution of $y^f(t_T + l)$ conditional on l and on $M_T = \{Y_1, Y_2, \ldots, Y_T\}$, i.e. the information available at time t_T. In a

Gaussian model, the mean of $y^f(t_T + l)$ (conditional on l and M_T) is given by $y^f(r_T + l) = z'W(l)a_T$ and its conditional variance is the MSE given in (39). Although it can be difficult to derive the full PDF $p(y^f(t_T + l)|M_T)$, expressions for the mean and variance of this distribution may be obtained for the principal structural time series models; see Harvey and Snyder (1990). If the lead time distribution is taken to be discrete, the derivation of such expressions is much more tedious. Of course, in concrete applications the integral (43) can be evaluated numerically.

Conclusion

The formulas provided here for univariate stocks or flows are readily extended to multivariate mixed stock-flow systems; compare Harvey and Stock (1985) and Zadrozny (1988). Harvey and Stock (1988) develop this extension for a model in which the variables are cointegrated (so that the stochastic trend term is common among several multivariate time series); they also provide an empirical application to the estimation of the common stochastic trend among consumption and income using a mutivariate continuous-time components model.

An advantage of the continuous-time framework is that, in multivariate applications, the observational frequency need not be the same for all the series. As a concrete example, weekly observations on interest rates (a "stock" variable), observations on some of the components of investment that are available monthly (a "flow"), and quarterly observations on total investment could be used to distribute total quarterly investment to a monthly level. Some of the components – say, trend and cycle – could be modeled as common among these series, and some could be modeled as independent (or perhaps correlated) across series. It should be emphasized, however, that the distributed values (or, for stocks, the interpolated values) resulting from the procedures outlined in this article have unavoidable measurement error. Moreover, this article has not addressed issues of aliasing, which could pose additional difficulties for interpolation and distribution. Thus care must be taken in using these values in subsequent statistical analysis.

Structural time-series models are based on fitting stochastic functions of time to the observations. A continuous-time formulation of structural models both is intuitively appealing and provides a logical consistency for both stock and flow data. The form of the model does not depend on the observation timing interval and hence can be applied to irregular observations. From the technical point of view, estimation, prediction, interpolation, and distribution can all be based on state space algorithms.

Note

Stock thanks the Sloan Foundation and the National Science Foundation for financial support through grants SES-86-18984 and SES-89-10601.

References

Anderson, B. D. O., and J. B. Moore (1979). *Optimal Filtering.* Englewood Cliffs: Prentice-Hall.

Ansley, C. F., and R. Kohn (1985). "Estimation, filtering and smoothing in state space models with incompletely specified initial conditions". *Annals of Statistics,* 13, 1286–316.

Bergstrom, A. R. (1966). "Non-recursive models as discrete approximations to systems of stochastic differential equations," *Econometrica*, 34, 173–82.

Bergstrom, A. R. (1976) (ed.). *Statistical Inference in Continuous Time Economic Models.* Amsterdam: North-Holland.

Bergstrom, A. R. (1983). "Gaussian estimation of structural parameters in higher-order continuous time dynamic models," *Econometrica*, 51, 117–52.

Bergstrom, A. R. (1984). "Continuous time stochastic models and issues of aggregation over time," in Z. Griliches and M. D. Intrilligator (eds) *Handbook of Econometrics*, volume 2, pp. 1145–212. Amsterdam: North-Holland.

Bergstrom, A. R. (1985). "The estimation of parameters in non-stationary higher-order continuous time dynamic models," *Econometric Theory*, 1, 369–85.

Bergstrom, A. R. (1986). "The estimation of open higher-order continuous time dynamic models with mixed stock and flow data," *Econometric Theory*, 2, 350–73.

De Jong, P. (1991). "The Diffuse Kalman Filter," *Annals of Statistics*, 19, 1073–83.

Engle, R. F. (1978). "Estimating structural models of seasonality," in A. Zellner (ed.) *Seasonal Analysis of Economic Time Series*, pp. 281–308. Washington: Bureau of the Census.

Harvey, A. C. (1985). "Trends and cycles in macroeconomic time series," *Journal of Business and Economic Statistics*, 3, 216–27.

Harvey, A. C. (1989). *Forecasting, Structural Time Series Models, and the Kalman Filter.* Cambridge: Cambridge University Press.

Harvey, A. C., and R. D. Snyder (1990). "Structural time series models in inventory control," *International Journal of Forecasting*, 6, 187–98.

Harvey, A. C., and J. H. Stock (1985). "The estimation of higher order continuous time autoregressive processes," *Econometric Theory*, 1, 97–117.

Harvey, A. C., and J. H. Stock (1988). "Continuous time autoregressive models with common stochastic trends", *Journal of Economic Dynamics and Control*, 12, 365–84.

Hausman, J. A., and M. W. Watson (1985). "Errors in variable and seasonal adjustment procedures," *Journal of the American Statistical Association*, 80, 541–52.

Jazwinski, A. H. (1970). *Stochastic Processes and Filtering Theory.* New York: Academic Press.

Johnston, F. R., and P. J. Harrison (1986). "The variance of lead time demand," *Journal of the Operational Research Society*, 37, 303–8.

Jones, R. H. (1981). "Fitting a continuous time autoregression to discrete data," in D. F. Findley (ed.) *Applied Time Series Analysis II.* New York: Academic Press.

Jones, R. H. (1984). "Fitting multivariate models to unequally spaced data," in E. Parzen (ed.) *Time Series Analysis of Irregularly Observed Data*, pp. 158–88. New York: Springer-Verlag.

Kitagawa, G. (1981). "A nonstationary time series model and its fitting by a recursive filter," *Journal of Time Series Analysis*, 2, 103–16.

Kitagawa, G. (1984). "State space modelling of nonstationary time series and smoothing of unequally spaced data," in E. Parzen (ed.) *Time Series Analysis of Irregularly Observed Data*, pp. 189–210. New York: Springer-Verlag.

Litterman, R. B. (1983). "A random walk Markov model for the distribution of time series," *Journal of Business and Economic Statistics*, 1, 169–73.

Phadke, M. S., and S. M. Wu (1974). "Modelling of continuous stochastic processes from discrete observations with application to sunspots data," *Journal of the American Statistical Association*, 69, 325–9.

Phillips, P. C. B. (1988). "Error correction and long run equilibrium in continuous time," Cowles Foundation Discussion Paper 882, Yale University.

Working, H. (1960). "Note on the correlations of first differences of a random chain," *Econometrica*, 28, 916–18.

Zadrozny, P. (1988). "Gaussian likelihood of continuous-time ARMAX models when data are stocks and flows at different frequencies," *Econometric Theory*, 4, 108–24.