



UNIVERSITY OF
CAMBRIDGE

Cambridge Working Papers in Economics

Filtering with heavy tails

Andrew Harvey and Alessandra Luati

CWPE 1255

Filtering with heavy tails

Andrew Harvey (ach34@cam.ac.uk) and Alessandra Luati
Faculty of Economics, University of Cambridge
and Dept. of Statistics, Bologna

December 10, 2012

Abstract

An unobserved components model in which the signal is buried in noise that is non-Gaussian may throw up observations that, when judged by the Gaussian yardstick, are outliers. We describe an observation driven model, based on a conditional Student t-distribution, that is tractable and retains some of the desirable features of the linear Gaussian model. Letting the dynamics be driven by the score of the conditional distribution leads to a specification that is not only easy to implement, but which also facilitates the development of a comprehensive and relatively straightforward theory for the asymptotic distribution of the ML estimator. The methods are illustrated with an application to rail travel in the UK. The final part of the article shows how the model may be extended to include explanatory variables.

KEYWORDS: outlier; robustness; score; seasonal; t-distribution; trend.

JEL Classification: C22

1 Introduction

Linear Gaussian unobserved components models play an important role in time series modeling. The Kalman filter and associated smoother provide the basis for a comprehensive statistical treatment. The filtered and smoothed estimators of the signal are optimal, in the sense of minimizing the MSE, the likelihood is given as a by-product of one-step prediction errors produced by

the Kalman filter and the full multi-step predictive distribution has a known Gaussian distribution.

A model in which the signal is buried in noise that is non-Gaussian may throw up observations that, when judged by the Gaussian yardstick, are outliers. The purpose of this article is to investigate the practical value of an observation driven model that is tractable and retains some of the desirable features of the linear Gaussian model. The principal feature of the model is that the dynamics are driven by the score of the conditional distribution of the observations. As a result it is not only easy to implement, but its form also facilitates the development of a comprehensive and relatively straightforward theory for the asymptotic distribution of the maximum likelihood estimator. Models of this kind are called dynamic conditional score (DCS) models and they have already proved useful for modeling volatility; a full description can be found in the forthcoming monograph by Harvey (2013).

Modeling the additive noise with a Student t -distribution is effective and theoretically straightforward. Indeed the attractions of using the t -distribution to guard against outliers in static models is well-documented; see, for example, Lange, Little and Taylor (1989). Such an approach may be contrasted with the methods used in the robustness literature; see, for example, Muler, Peña and Yohai (2009). Section 2 sets out a simple unobserved components model and discusses the rationale for letting the dynamics depend on the conditional score. The first-order conditional score model for a Student t -distribution is described in Section 3. The asymptotic distribution of the maximum likelihood estimator is given in Section 4 and complemented by a Monte Carlo study on small sample properties. Section 5 then extends DCS models using the state space form and Section 6 discusses how to model trend and seasonality. The viability of a DCS model with trend and seasonal components is demonstrated with real data in Section 6. Explanatory variables are introduced into the model in Section 7.

2 Unobserved components and filters

A simple Gaussian signal plus noise model is

$$\begin{aligned} y_t &= \mu_t + \varepsilon_t, & \varepsilon_t &\sim NID(0, \sigma_\varepsilon^2) \\ \mu_{t+1} &= \phi\mu_t + \eta_t, & \eta_t &\sim NID(0, \sigma_\eta^2), \end{aligned} \tag{1}$$

for $t = 1, \dots, T$ and where the irregular and level disturbances, ε_t and η_t respectively, are mutually independent and the notation $NID(0, \sigma^2)$ denotes normally and independently distributed with mean zero and variance σ^2 . The autoregressive parameter is ϕ , while the signal-noise ratio, $q = \sigma_\eta^2 / \sigma_\varepsilon^2$, plays the key role in determining how observations should be weighted for prediction and signal extraction. The reduced form (RF) of (1) is an ARMA(1,1) process

$$y_t = \phi y_{t-1} + \xi_t - \theta \xi_{t-1}, \quad \xi_t \sim NID(0, \sigma^2), \quad t = 1, \dots, T \quad (2)$$

but with restrictions on θ . For example, when $\phi = 1$, $0 \leq \theta \leq 1$. The forecasts from the unobserved components (UC) model and RF are the same.

The UC model in (1) is effectively in state space form (SSF) and, as such, it may be handled by the Kalman filter (KF); see Harvey (1989). The parameters ϕ and q may be estimated by maximum likelihood, with the likelihood function constructed from the one-step ahead prediction errors. The KF can be expressed as a single equation which combines $\mu_{t|t-1}$, the optimal estimator of μ_t based on information at time $t - 1$, with y_t in order to produce the best estimator of μ_{t+1} . Writing this equation together with an equation that defines the one-step ahead prediction error, v_t , gives the innovations form of the KF:

$$\begin{aligned} y_t &= \mu_{t|t-1} + v_t, \\ \mu_{t+1|t} &= \phi \mu_{t|t-1} + k_t v_t. \end{aligned} \quad (3)$$

The Kalman gain, k_t , depends on ϕ and q . In the steady-state, k_t is constant. Setting it equal to κ in (3) and re-arranging gives the ARMA model (2) with $\xi_t = v_t$ and $\phi - \kappa = \theta$. A pure autoregressive model is a special case in which $\kappa = \phi$, so that $\mu_{t|t-1} = \phi y_{t-1}$.

Now suppose that the noise in (1) comes from a heavy tailed distribution, such as Student's t . Such a distribution can give rise to observations which, when judged against the yardstick of a Gaussian distribution, are considered to be outliers. The RF is still an ARMA(1,1) process, but with disturbances which, although they are serially uncorrelated, are not independent and identically distributed (IID). Allowing the ξ'_t s to have a heavy-tailed distribution does not deal with the problem as a large observation becomes incorporated into the level and takes time to work through the system. An ARMA model in which the disturbances are allowed to have a heavy-tailed distribution is designed to handle *innovations outliers*, as opposed to *additive outliers*.

There is a good deal of discussion of these issues in the robustness literature; see, for example the book by Maronna, Martin and Yohai (2006, ch. 8).

Simulation methods, such as Markov chain Monte Carlo (MCMC) and particle filtering, provide the basis for a direct attack on models that are nonlinear and/or non-Gaussian. The aim is to extend the Kalman filtering and smoothing algorithms that have proved so effective in handling linear Gaussian models. Considerable progress has been made in recent years; see Durbin and Koopman (2012). However, the fact remains that simulation-based estimation can be time-consuming and subject to a degree of uncertainty. In addition the statistical properties of the estimators are not easy to establish.

The DCS approach begins by writing down the distribution of the t -th observation, conditional on past observations. Time-varying parameters are then updated by a suitably defined filter. Such a model is called *observation driven*. In a linear Gaussian UC model, the KF, as in (3), is driven by the one step-ahead prediction error. The main ingredient in the filter developed here for non-Gaussian distributions is the replacement of v_t in the KF equation by a variable, u_t , that is proportional to the score of the conditional distribution. Thus the second equation in (3) becomes

$$\mu_{t+1|t} = \phi\mu_{t|t-1} + \kappa u_t \quad (4)$$

where κ is treated as an unknown parameter. This filter might be regarded as an approximation to the computer intensive solution for the *parameter driven* UC model. The attraction of regarding it as a model in its own right is that it becomes possible to derive the asymptotic distribution of the ML estimator and generalize in various directions. The same approach can be used to model scale, using an exponential link function; see, for example, Creal et al (2011).

3 Dynamic Student-t location model

When the location changes over time, it may be captured by a model in which, conditional on past observations, y_t has a t_ν -distribution

$$f_t(y_t|Y_{t-1}) = \frac{\Gamma\left(\frac{\nu+1}{2}\right)}{\sqrt{\pi}\Gamma\left(\frac{\nu}{2}\right)(\nu e^{2\lambda})^{\frac{1}{2}}} \left(1 + \frac{(y_t - \mu_{t|t-1})^2}{\nu e^{2\lambda}}\right)^{-\frac{\nu+1}{2}},$$

where $Y_{t-1} = \{y_{t-1}, y_{t-2}, \dots\}$, $\exp(\lambda)$ is the scale and the location, $\mu_{t|t-1}$, is generated by a linear function of

$$u_t = (1 + \nu^{-1} e^{-2\lambda} (y_t - \mu_{t|t-1})^2)^{-1} v_t, \quad t = 1, \dots, T, \quad (5)$$

where $v_t = y_t - \mu_{t|t-1}$ is the prediction error. Differentiating the log-density shows that u_t is proportional to the conditional score, $\partial \ln f_t / \partial \mu_{t|t-1} = (\nu + 1) \nu^{-1} \exp(-2\lambda) u_t$. No restriction is put on the degrees of freedom, ν , apart from requiring that it be positive: hence the reference to location rather than the mean. The scaling factor, $\exp(2\lambda)$, cancels out if the score is divided by the information quantity for the location.

The first-order model corresponds to the Gaussian innovations form, (3), and is

$$\begin{aligned} y_t &= \mu_{t|t-1} + v_t = \mu_{t|t-1} + \exp(\lambda) \varepsilon_t, & t = 1, \dots, T \\ \mu_{t+1|t} &= \delta + \phi \mu_{t|t-1} + \kappa u_t, \end{aligned} \quad (6)$$

where ε_t is serially independent, standard t -variate. More generally, a model of order (p, r) is

$$\mu_{t+1|t} = \delta + \phi_1 \mu_{t|t-1} + \dots + \phi_p \mu_{t-p+1|t-p} + \kappa_0 u_t + \kappa_1 u_{t-1} + \dots + \kappa_r u_{t-r}. \quad (7)$$

In the Gaussian case $u_t = v_t$. If q is defined as $\max(p, r + 1)$, y_t is an $ARMA(p, q)$ with MA coefficients $\theta_i = \phi_i - \kappa_{i-1}$, $i = 1, \dots, q$.

Re-parameterization in terms of the unconditional mean, ω , gives

$$\mu_{t|t-1} = \omega + \mu_{t|t-1}^\dagger, \quad t = 1, \dots, T, \quad (8)$$

where $\mu_{t|t-1}^\dagger$ is as in (7), but without δ , and $\omega = \delta / (1 - \phi_1 - \dots - \phi_p)$.

Figure 1 shows the impact, u , of a standardized ($\lambda = 0$) observation for different degrees of freedom. The Gaussian response is the 45 degree line. For low degrees of freedom, observations that would be seen as outliers for a Gaussian distribution are far less influential. As $|y| \rightarrow \infty$, the response tends to zero. Redescending M-estimators, which feature in the robustness literature, have the same property. On the other hand, the Huber M-estimator has a Gaussian response until a certain point, whereupon it is constant; see Maronna *et al* (2006, p 25-31). The implementation of M-estimates usually requires a (robust) estimate of scale to be pre-computed.

The variable u_t can be written

$$u_t = (1 - b_t)(y_t - \mu_{t|t-1}), \quad (9)$$

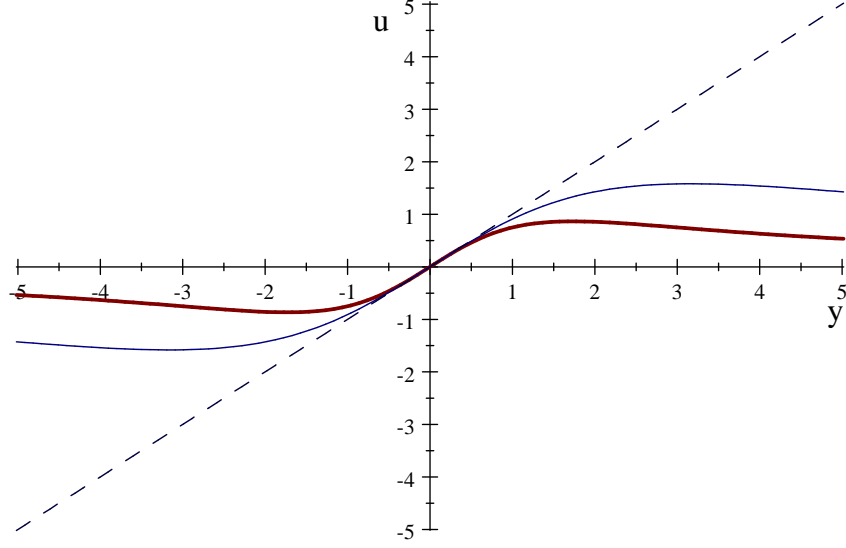


Figure 1: Impact of u_t for t_ν (with a scale of one) for $\nu = 3$ (thick), $\nu = 10$ (thin) and $\nu = \infty$ (dashed).

where

$$b_t = \frac{(y_t - \mu_{t|t-1})^2 / \nu \exp(2\lambda)}{1 + (y_t - \mu_{t|t-1})^2 / \nu \exp(2\lambda)}, \quad 0 \leq b_t \leq 1, \quad 0 < \nu < \infty, \quad (10)$$

is distributed as $beta(1/2, \nu/2)$; see Harvey (2013, Chapter 3). The u_t 's are $IID(0, \sigma_u^2)$ and symmetrically distributed. Even moments of all orders exist. In particular

$$Var(u_t) = \sigma_u^2 = \nu \exp(2\lambda) E(b_t(1 - b_t)) = \nu^2(\nu + 3)^{-1}(\nu + 1)^{-1} \exp(2\lambda),$$

Since the u_t 's are $IID(0, \sigma_u^2)$, $\mu_{t|t-1}$ is weakly and strictly stationary so long as $|\phi| < 1$. Although determining the statistical properties of $\mu_{t|t-1}$ requires assuming that it started in the infinite past, the filter needs to be initialized in practice and this may be done by setting $\mu_{1|0} = \omega$ or $\mu_{1|0}^\dagger = 0$ in (8).

The existence of moments of y_t is not affected by the dynamics. The autocorrelations can be found from the infinite MA representation; the patterns are as they would be for a Gaussian model.

4 Maximum likelihood estimation

The asymptotic distribution of the maximum likelihood estimator is given in the following sub-section. Sub-section 4.2 reports a set of Monte Carlo experiments to assess the small sample properties and Sub-section 4.3 fits a model to the growth rate in US GDP.

4.1 Asymptotic distribution of the maximum likelihood estimator

Proposition 1 *Let $y_t | Y_{t-1}$ have a t_ν -distribution with location, $\mu_{t|t-1}$, generated by (8) where $\mu_{t|t-1}^\dagger$ is a stationary first-order model, as in (4), and $|\phi| < 1$. Define*

$$a = \phi - \kappa \frac{\nu}{\nu + 3}, \quad b = \phi^2 - 2\phi\kappa \frac{\nu}{\nu + 3} + \kappa^2 \frac{\nu(\nu^3 + 10\nu^2 + 35\nu + 38)}{(\nu + 1)(\nu + 3)(\nu + 5)(\nu + 7)}$$

and let $\boldsymbol{\psi} = (\kappa, \phi, \omega)'$. Assuming that $b < 1$ and $\kappa \neq 0$, $(\tilde{\boldsymbol{\psi}}', \tilde{\lambda}, \tilde{\nu})'$, the ML estimator of $(\boldsymbol{\psi}', \lambda, \nu)'$, is consistent and the limiting distribution of $\sqrt{T}((\tilde{\boldsymbol{\psi}} - \boldsymbol{\psi})', \tilde{\lambda} - \lambda, \tilde{\nu} - \nu)'$ is multivariate normal with mean vector zero and covariance matrix given by $\text{Var}(\tilde{\boldsymbol{\psi}}, \tilde{\lambda}, \tilde{\nu}) = \mathbf{I}^{-1}(\boldsymbol{\psi}, \lambda, \nu)$, where the information matrix is

$$\mathbf{I}(\boldsymbol{\psi}, \lambda, \nu) = \begin{bmatrix} \frac{\nu+1}{\nu+3} \exp(-2\lambda) \mathbf{D}(\boldsymbol{\psi}) & 0 & 0 \\ 0 & \frac{2\nu}{\nu+3} & \frac{1}{(\nu+3)(\nu+1)} \\ 0 & \frac{1}{(\nu+3)(\nu+1)} & h(\nu)/2 \end{bmatrix} \quad (11)$$

with

$$h(\nu) = \frac{1}{2} \psi'(\nu/2) - \frac{1}{2} \psi'((\nu+1)/2) - \frac{\nu+5}{\nu(\nu+3)(\nu+1)},$$

where $\psi'(\cdot)$ is the digamma function, and

$$\mathbf{D}(\boldsymbol{\psi}) = \mathbf{D} \begin{pmatrix} \kappa \\ \phi \\ \omega \end{pmatrix} = \frac{1}{1-b} \begin{bmatrix} A & D & 0 \\ D & B & 0 \\ 0 & 0 & C \end{bmatrix} \quad (12)$$

with $A = \sigma_u^2$,

$$B = \frac{\kappa^2 \sigma_u^2 (1 + a\phi)}{(1 - \phi^2)(1 - a\phi)}, \quad C = \frac{(1 - \phi)^2 (1 + a)}{1 - a} \quad \text{and} \quad D = \frac{a\kappa \sigma_u^2}{1 - a\phi}.$$

The proof of the above result is given in Harvey (2013) and sketched out in the Appendix.

4.2 Monte Carlo experiments

A series of Monte Carlo experiments were carried out in order to get an indication of how good the asymptotic approximation is for small samples. The figures shown in Table 1 are based on computing RMSEs from 1000 replications¹ for $T = 500$ and 1000 observations from first-order models with $\nu = 6$ and a range of (realistic) values of κ and ϕ . The expression for the information matrix shows that the asymptotic standard errors (ASEs) are independent of ω and that λ only appears as a scaling factor. Hence setting $\omega = \lambda = 0$ should imply no loss in generality.

In most cases convergence was rapid and few computational problems were encountered. The ASEs were obtained from the square roots of the diagonal elements of the inverse of (11) divided by the sample size. The results show that the empirical RMSEs are somewhat larger than the ASEs, but not by a great deal. The numerical standard errors were computed from the Hessian matrix of the log-likelihood function. The figures in Table 1 are averaged over all replications. On the whole they are reasonably close to the empirical RMSEs. The row labelled ‘est. par’ is the sample average and its role is to give an indication of any bias. The most striking result is that ϕ seems to be underestimated when κ is small; see the rows at the top of the table for $\kappa = 0.1$.

4.3 Application to US GDP

A Gaussian AR(1) plus noise model, that is (1) with a constant, was fitted to the growth rate of US Real GDP, defined as the first difference of the logarithm, using the STAMP 8 package of Koopman *et al* (2009). The data were quarterly, from 1947(2) to 2012(1), and the parameter estimates were as follows:

$$\tilde{\phi} = 0.501, \quad \tilde{\sigma}_\eta^2 = 7.62 \times 10^{-5}, \quad \tilde{\sigma}_\varepsilon^2 = 2.30 \times 10^{-5}, \quad \tilde{\omega} = 0.0078.$$

There was little indication of residual serial correlation, but the Bowman-Shenton statistic is 30.04, which is clearly significant as the distribution under the null hypothesis of Gaussianity is χ_2^2 . The non-normality clearly comes from excess kurtosis, which is 1.9, rather than from skewness.

¹We carried out some simulations using 5000 and 1000 replications, but since the results were the same up to the third decimal, we concentrated on 1000 replications which was much faster (Matlab codes are available upon request).

A DCS-location- t model gave the results shown below. Given that the sample size is only 260, the numerical and analytic standard errors are reasonably close. The estimated degrees of freedom of 6.3 means that the DCS filter is less responsive to more extreme observations, such as the fall of 2009(1), than is the Gaussian filter.

Parameter	κ	ϕ	λ	ω	ν
Estimate	0.520	0.497	-4.878	0.0079	6.303
Num SE	0.098	0.102	0.073	0.0009	2.310
ASE	0.090	0.140	0.057	0.0009	1.807

5 Higher-order models and the state space form

The general statistical treatment of unobserved components models is based on the state space form. The corresponding innovations form facilitates the handling of higher-order DCS models.

5.1 Linear Gaussian models and the Kalman filter

For simplicity assume a time-invariant univariate time series model and exclude any deterministic components. The general case is set out in Harvey (1989, Chapter 3). The observation in the Gaussian state space model is related to an $m \times 1$ state vector, $\boldsymbol{\alpha}_t$, through a measurement equation, $y_t = \omega + \mathbf{z}'\boldsymbol{\alpha}_t + \varepsilon_t$, $t = 1, \dots, T$, where ω is a constant, \mathbf{z} is an $m \times 1$ vector and $\varepsilon_t \sim NID(0, \sigma_\varepsilon^2)$. The elements of $\boldsymbol{\alpha}_t$ are usually unobservable but are known to be generated by a transition equation $\boldsymbol{\alpha}_{t+1} = \boldsymbol{\delta} + \mathbf{T}\boldsymbol{\alpha}_t + \boldsymbol{\eta}_t$, $t = 1, \dots, T$, where $\boldsymbol{\delta}$ is a vector of constants and $\boldsymbol{\eta}_t \sim NID(\mathbf{0}, \mathbf{Q})$. The specification is completed by assuming that $E(\boldsymbol{\alpha}_1) = \boldsymbol{\alpha}_{1|0}$ and $Var(\boldsymbol{\alpha}_1) = \mathbf{P}_{1|0}$, where $\mathbf{P}_{1|0}$ is positive a semi-definite matrix, and that $E(\varepsilon_t \boldsymbol{\alpha}'_0) = \mathbf{0}$ and $E(\boldsymbol{\eta}_t \boldsymbol{\alpha}'_0) = \mathbf{0}$ for $t = 1, \dots, T$. It is usually assumed that the disturbances are uncorrelated with each other in all time periods, that is $E(\varepsilon_t \boldsymbol{\eta}'_s) = \mathbf{0}$ for all $s, t = 1, \dots, T$, though this assumption may be relaxed.

When the disturbances and initial state are normally distributed, the minimum mean square error estimates of the state and observation at time t , based on information at time $t-1$, are their conditional expectations. The Kalman filter is a recursive procedure for computing these estimates, given

$\mathbf{z}, \sigma_\varepsilon^2, \mathbf{T}$ and \mathbf{Q} together with the initial conditions, $\boldsymbol{\alpha}_{1|0}$ and $\mathbf{P}_{1|0}$. When the initial conditions are unknown, the filter may be started off as discussed in Durbin and Koopman (2012).

The Kalman filter can be written as a single set of recursions going directly from $\boldsymbol{\alpha}_{t|t-1}$ to $\boldsymbol{\alpha}_{t+1|t}$. The innovations form, generalizing (3), is

$$\begin{aligned} y_t &= \omega + \mathbf{z}'\boldsymbol{\alpha}_{t|t-1} + v_t, & t = 1, \dots, T, \\ \boldsymbol{\alpha}_{t+1|t} &= \boldsymbol{\delta} + \mathbf{T}\boldsymbol{\alpha}_{t|t-1} + \mathbf{k}_t v_t, \end{aligned} \quad (13)$$

where $v_t = y_t - \omega - \mathbf{z}'\boldsymbol{\alpha}_{t|t-1}$ is the innovation and $f_t = \mathbf{z}'\mathbf{P}_{t|t-1}\mathbf{z} + \sigma_\varepsilon^2$ is its variance. The gain vector is $\mathbf{k}_t = (1/f_t)\mathbf{T}\mathbf{P}_{t|t-1}\mathbf{z}$ and $\mathbf{P}_{t|t-1}$ is calculated by a matrix recursion. Since (13) contains only one disturbance term, it may be regarded as a reduced form model with \mathbf{k}_t subject to restrictions coming from the original structural form². In the steady-state, \mathbf{k}_t and f_t are time-invariant.

5.2 The DCS model

A general location DCS model may be set up in the same way as the innovations form of a Gaussian state space model. The model corresponding to the steady-state of (13) is

$$\begin{aligned} y_t &= \omega + \mathbf{z}'\boldsymbol{\alpha}_{t|t-1} + v_t, & t = 1, \dots, T, \\ \boldsymbol{\alpha}_{t+1|t} &= \boldsymbol{\delta} + \mathbf{T}\boldsymbol{\alpha}_{t|t-1} + \boldsymbol{\kappa}u_t. \end{aligned} \quad (14)$$

The \mathbf{z} vector and \mathbf{T} matrix may be specified in the same way as for the Gaussian UC models. The transition equation in (14) is stationary provided that the roots of the transition matrix \mathbf{T} have modulus less than one. When this is the case, $\boldsymbol{\delta}$ is superfluous and initialization is achieved by setting $\boldsymbol{\alpha}_{1|0} = \mathbf{0}$. If $\boldsymbol{\alpha}_{t|t-1}$ contains nonstationary elements, the best option seems to be to treat their initial values as unknown parameters.

There remains the question of how to specify the parameters in the vector $\boldsymbol{\kappa}$. More specifically, what restrictions should be imposed? The issues are explored for trend and seasonal components below.

²The single source of error (SSOE) models discussed in Ord et al (1997) are effectively in innovations form but if this is the starting point of model formulation some way of putting constraints on \mathbf{k}_t has to be found.

Remark 2 The general model, (7), of order (p, r) may be put in the state space form of (14) in a similar way to an $ARMA(p, r)$ plus noise unobserved components models.

6 Trend and seasonality

Stochastic trend and seasonal components may be introduced into UC models for location. These models, called structural time series models, are described in Harvey (1989) and implemented in the STAMP package of Koopman *et al* (2009). The way in which the innovations forms of structural time series models lead to corresponding DCS-t models is explored below.

6.1 Local level model

The Gaussian random walk plus noise or *local level* model is

$$y_t = \mu_t + \varepsilon_t, \quad \mu_t = \mu_{t-1} + \eta_t, \quad (15)$$

where $\varepsilon_t \sim NID(0, \sigma_\varepsilon^2)$, $\eta_t \sim NID(0, \sigma_\eta^2)$ and $E(\varepsilon_t \eta_s) = 0$ for all t and s . The signal noise ratio is $q = \sigma_\eta^2 / \sigma_\varepsilon^2$ and the parameter, θ in the $ARIMA(0, 1, 1)$ reduced form representation, (2), lies in the range $0 \leq \theta < 1$ when $\sigma_\varepsilon^2 > 0$. Since $\theta = 1 - \kappa$, the range of κ in the steady-state innovations form is $0 < \kappa \leq 1$. In this case $\mu_{t+1|t}$ is an exponentially weighted moving average in which the weights on current and past observations are non-negative.

The local level DCS- t model is

$$y_t = \mu_{t|t-1} + v_t, \quad \mu_{t+1|t} = \mu_{t|t-1} + \kappa u_t. \quad (16)$$

The initialization of the KF in (15) is best done using a diffuse prior; see Harvey (1989, pp 107-8). This is not an option for the DCS model. One possibility is to set $\mu_{2|1} = y_1$, but the filter could be adversely affected if the first observation is an outlier. An alternative approach is to treat the initial value, $\mu_{1|0}$, as an unknown parameter that must be estimated along with κ and ν . This is the technique used by Ord, Koehler and Snyder (1997) to initialize nonlinear single source of error models and by De Livera, Hyndman and Snyder (2010).

Since $u_t = (1 - b_t)(y_t - \mu_{t|t-1})$, re-arranging the dynamic equation in (16) gives

$$\mu_{t+1|t} = (1 - \kappa(1 - b_t))\mu_{t|t-1} + \kappa(1 - b_t)y_t. \quad (17)$$

A sufficient condition for the weights on current and past observations to be non-negative is that $\kappa(1 - b_t) < 1$ and, because $0 \leq b_t \leq 1$, this is guaranteed by $0 < \kappa \leq 1$. However, the restriction that $\kappa \leq 1$ is much stricter than is either necessary or desirable. Estimates of κ greater than one are not unusual and are entirely appropriate when the signal is strong relative to the noise.

As regards asymptotic properties, the following is a corollary to Proposition 1.

Corollary 3 *When $b < 1$ and μ_{10} is fixed and known or $\mu_{21} = y_1$, where y_1 is fixed, the ML estimator of κ in (16) is consistent and $\sqrt{T}(\tilde{\kappa} - \kappa)$ has a limiting normal distribution with mean zero and variance*

$$Var(\tilde{\kappa}) = \left(2\kappa \frac{\nu}{\nu + 3} - \kappa^2 \frac{\nu(\nu^3 + 10\nu^2 + 35\nu + 38)}{(\nu + 1)(\nu + 3)(\nu + 5)(\nu + 7)} \right) \left(\frac{\nu + 3}{\nu} \right)^2.$$

It can be seen that $\kappa > 0$ is a necessary condition for $b < 1$ and hence $Var(\tilde{\kappa}) > 0$. When the initial value, μ_{10} , is treated as a parameter to be estimated, it appears from some limited simulation evidence that the asymptotic distribution of the ML estimator of κ is unchanged.

The result extends to the random walk plus drift trend, that is

$$\mu_{t+1|t} = \beta + \mu_{t|t-1} + \kappa u_t, \quad (18)$$

where β is an unknown constant. The ML estimators of κ and β are asymptotically independent. Thus $Var(\tilde{\kappa})$ is unchanged and adapting expression (2.56) in Harvey (2013) gives

$$Var(\tilde{\beta}) = \left(2\kappa \frac{\nu}{\nu + 3} - \kappa^2 \frac{\nu(\nu^3 + 10\nu^2 + 35\nu + 38)}{(\nu + 1)(\nu + 3)(\nu + 5)(\nu + 7)} \right) \frac{\nu + 3}{\nu + 1} \frac{\nu}{(2 - \kappa)\nu + 6} \kappa e^{2\lambda}.$$

6.2 Application to hours worked by US employees

Fitting a local level DCS model (initialized with $\mu_{21} = y_1$) to seasonally adjusted monthly data on U.S. Average Weekly Hours of Production and Nonsupervisory Employees: Manufacturing (AWHMAN)³ from February 1992 to May 2010 (220 observations) gave

$$\tilde{\kappa} = 1.246 \quad \tilde{\lambda} = -3.625 \quad \tilde{\nu} = 6.35$$

³Source: U.S. Department of Labor: Bureau of Labor Statistics

with numerical (asymptotic) standard errors

$$SE(\tilde{\kappa}) = 0.161(0.090) \quad SE(\tilde{\lambda}) = 0.120(0.062) \quad SE(\tilde{\nu}) = 1.630(1.991)$$

A drift term was initially included but it was statistically insignificant. The value of b is 0.178. Figure 2 shows (part of) the series together with the contemporaneous filter, which for the random walk is $\mu_{t|t} = \mu_{t+1|t}$. As can be seen, the values are very close to the corresponding observations in the majority of cases because $\kappa(1 - b_t)$ is close to one. On the other hand, unusually large prediction errors result in a small value of $\kappa(1 - b_t)$ and most of the weight in (17) is assigned to $\mu_{t|t-1}$. When a Gaussian local level model is fitted, the outliers affect the filtered level more than is desirable.

Estimating the series without imposing the unit root gave

$$\tilde{\kappa} = 1.229 \quad \tilde{\lambda} = -3.652 \quad \tilde{\nu} = 6.20 \quad \tilde{\phi} = 0.978 \quad \tilde{\omega} = 40.18$$

The estimates of κ , λ and ν are close to those obtained with the local level model. Re-estimating both models with a longer series, starting in April 1947, gave very similar results, so the models seem to be stable.

6.3 Local linear trend

The DCS filter corresponding to the UC local linear trend model is

$$\begin{aligned} y_t &= \mu_{t|t-1} + v_t, \\ \mu_{t+1|t} &= \mu_{t|t-1} + \beta_{t|t-1} + \kappa_1 u_t, \quad \beta_{t+1|t} = \beta_{t|t-1} + \kappa_2 u_t. \end{aligned} \tag{19}$$

The initialization $\beta_{3|2} = y_2 - y_1$ and $\mu_{3|2} = y_2$ can be used, but, as in the local level model, initializing in this way is vulnerable to outliers at the beginning. Estimating the fixed starting values, $\mu_{1|0}$ and $\beta_{1|0}$, may be a better option.

An integrated random walk trend in the UC local linear trend model implies the constraint $\kappa_2 = \kappa_1^2 / (2 - \kappa_1)$, $0 < \kappa_1 < 1$, which may be found using formulae in Harvey (1989, p. 177). The restriction can be imposed on the DCS- t model by treating $\kappa_1 = \kappa$ as the unknown parameter, but without unity imposed as an upper bound.

6.4 Stochastic seasonal

A fixed seasonal pattern may be modeled as $\gamma_t = \sum_{j=1}^s \gamma_j z_{jt}$, where s is the number of seasons and the dummy variable z_{jt} is one in season j and zero

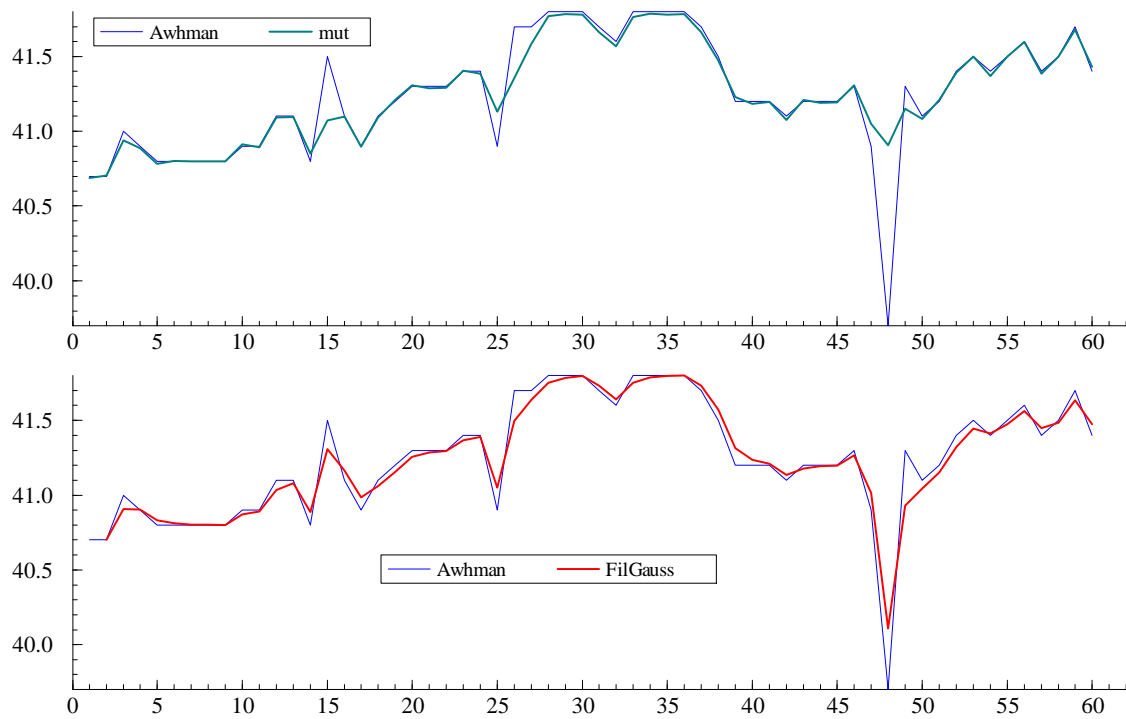


Figure 2: DCS and Gaussian (bottom panel) local level models fitted to US average weekly hours of production.

otherwise. In order not to confound trend with seasonality, the coefficients, γ_j , $j = 1, \dots, s$, are constrained to sum to zero. The seasonal pattern may be allowed to change over time by letting the coefficients evolve as random walks. If γ_{jt} denotes the effect of season j at time t , then

$$\gamma_{jt} = \gamma_{j,t-1} + \omega_{jt}, \quad \omega_t \sim NID(0, \sigma_\omega^2), \quad j = 1, \dots, s. \quad (20)$$

Although all s seasonal components are continually evolving, only one affects the observations at any particular point in time, that is $\gamma_t = \gamma_{jt}$ when season j is prevailing at time t . The requirement that the seasonal components evolve in such a way that they always sum to zero, that is $\sum_{j=1}^s \gamma_{jt} = 0$, is enforced by the restriction that the disturbances sum to zero at each point in time. This restriction is implemented by the correlation structure in $Var(\boldsymbol{\omega}_t) = \sigma_\omega^2 (\mathbf{I} - s^{-1}\mathbf{i}\mathbf{i}')$, where $\boldsymbol{\omega}_t = (\omega_{1t}, \dots, \omega_{st})'$, coupled with initial conditions requiring that the seasonals sum to zero at $t = 0$. It can be seen that $Var(\mathbf{i}'\boldsymbol{\omega}_t) = 0$.

In the state space form, the transition matrix is just the identity matrix, but the \mathbf{z} vector must change over time to accommodate the current season. Apart from replacing \mathbf{z} by \mathbf{z}_t , the form of the KF remains unchanged. Adapting the innovations form to the DCS observation driven framework, (14), gives

$$y_t = \mathbf{z}_t' \boldsymbol{\alpha}_{t|t-1} + v_t, \quad \boldsymbol{\alpha}_{t+1|t} = \boldsymbol{\alpha}_{t|t-1} + \boldsymbol{\kappa}_t u_t, \quad (21)$$

where \mathbf{z}_t picks out the current season, $\gamma_{t|t-1}$, that is $\gamma_{t|t-1} = \mathbf{z}_t' \boldsymbol{\alpha}_{t|t-1}$. The only question is how to parameterize $\boldsymbol{\kappa}_t$.

The seasonal components in the UC model are constrained to sum to zero and the same is true of their filtered estimates. Thus $\mathbf{i}'\boldsymbol{\kappa}_t = 0$ in the Kalman filter and this property should carry across to the DCS filter. If κ_{jt} , $j = 1, \dots, s$, denotes the j -th element of $\boldsymbol{\kappa}_t$ in (21), then in season j we set $\kappa_{jt} = \kappa_s$, where κ_s is a non-negative unknown parameter, while $\kappa_{it} = -\kappa_s/(s-1)$ for $i \neq j$. The amounts by which the seasonal effects change therefore sum to zero.

The seasonal recursions can be combined with the trend filtering equations of (19) in order to give a structure similar in form to that of the Kalman filter for the stochastic trend plus seasonal plus noise UC model, sometimes known as the ‘basic structural model’. Thus

$$y_t = \mu_{t|t-1} + \gamma_{t|t-1} + v_t, \quad (22)$$

where $\mu_{t|t-1}$ is as defined in (19). The filter can be initialized by regressing the first $s+1$ observations on a constant, time trend and seasonal dummies

constrained so that the coefficients sum to zero. Alternatively, the initial conditions at time $t = 0$ are estimated by treating them as parameters; there are $s - 1$ seasonal parameters because the remaining initial seasonal state is minus the sum of the others.

6.5 Application to rail travel

Figure 3 shows the logarithm of National Rail Travel, defined as the number of kilometres traveled by UK passengers. (Source: National Rail Trends). In a project carried out for the UK Department for Transport by one of the authors, an unobserved components model was fitted to this series using the STAMP 8 package of Koopman et al (2009). Trend, seasonal and irregular components were included but the model was augmented with intervention variables to take out the effects of observations that were known to be unrepresentative. The intervention dummies were: (i) the train drivers strikes in 1982(1,3); (ii) the Hatfield crash and its aftermath, 2000(4) and 2001(1); and (iii) the signallers strike in 1994(3).

Fitting a DCS model with trend and seasonal, that is (22), avoids the need to deal explicitly with the outliers. The ML estimates for the parameters in a model with a random walk plus drift trend, (18), are

$$\begin{aligned} \tilde{\kappa} &= 1.421(0.161) & \tilde{\kappa}_s &= 0.539 (0.070) & \tilde{\lambda} &= -3.787 (0.053) \\ \tilde{\nu} &= 2.564 (0.319) & \tilde{\beta} &= 0.003 (0.001) \end{aligned}$$

with initial values $\tilde{\mu} = 2.066(0.009)$, $\tilde{\gamma}_1 = -0.094(0.007)$, $\tilde{\gamma}_2 = -0.010(0.006)$ and $\tilde{\gamma}_3 = 0.086(0.006)$. The figures in parentheses are numerical standard errors. The last seasonal is $\tilde{\gamma}_4 = 0.018$; it has no SE as it was constructed from the others.

Figure 4 shows the filtered estimates of the trend from the UC and DCS-t models. The former includes intervention variables and so has gaps. The DCS-t trend appears not to be affected by the outliers. The same is true of the filtered seasonal shown in Figure 5. The pattern is very close to that of the seasonal in the UC model.

Figure 6 shows the residuals, that is the one-step ahead prediction errors, and score for the DCS model. The outliers, which were removed by dummies in the UC model, show up clearly in the residuals. (The Bowman-Shenton statistic is 137.82 indicating a massive rejection of the null hypothesis of

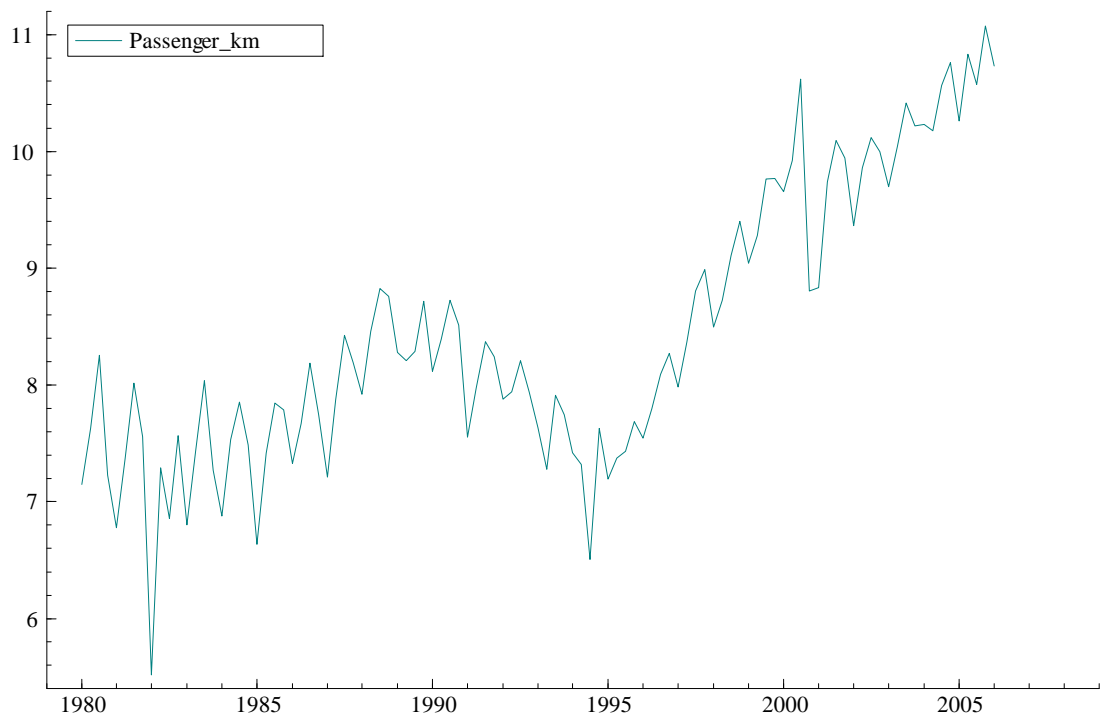


Figure 3: Logarithm of National Rail Travel in the UK (number of kilometres)

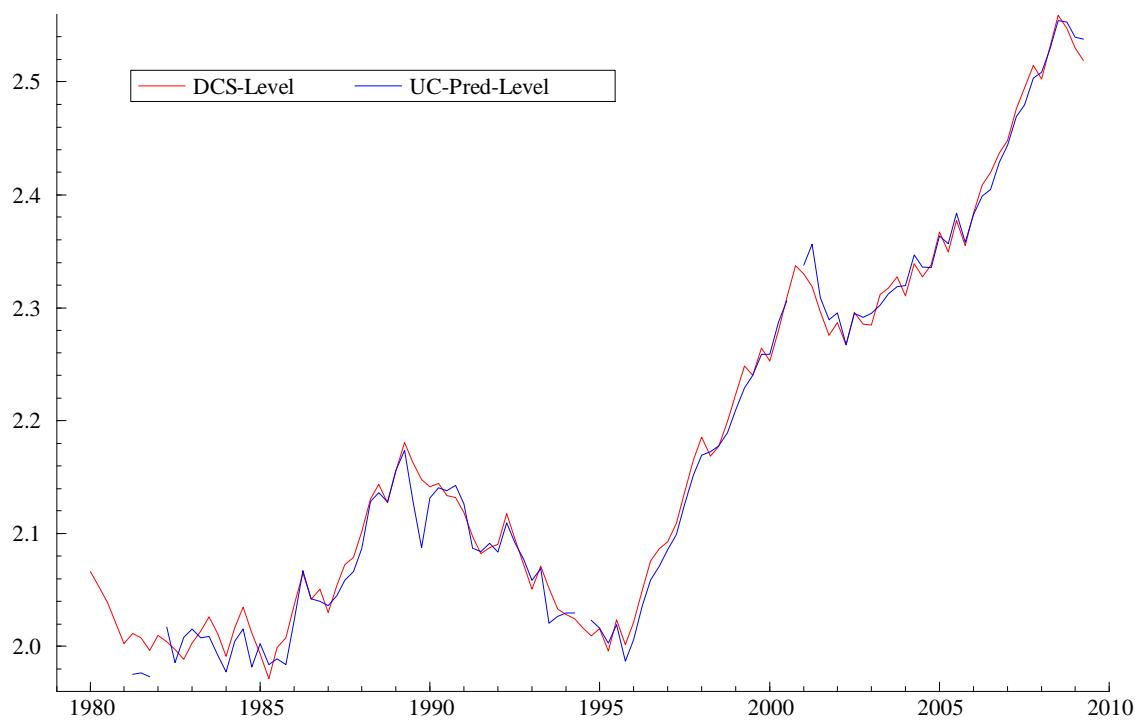


Figure 4: Trends in National Rail Travel from UC and DCS models

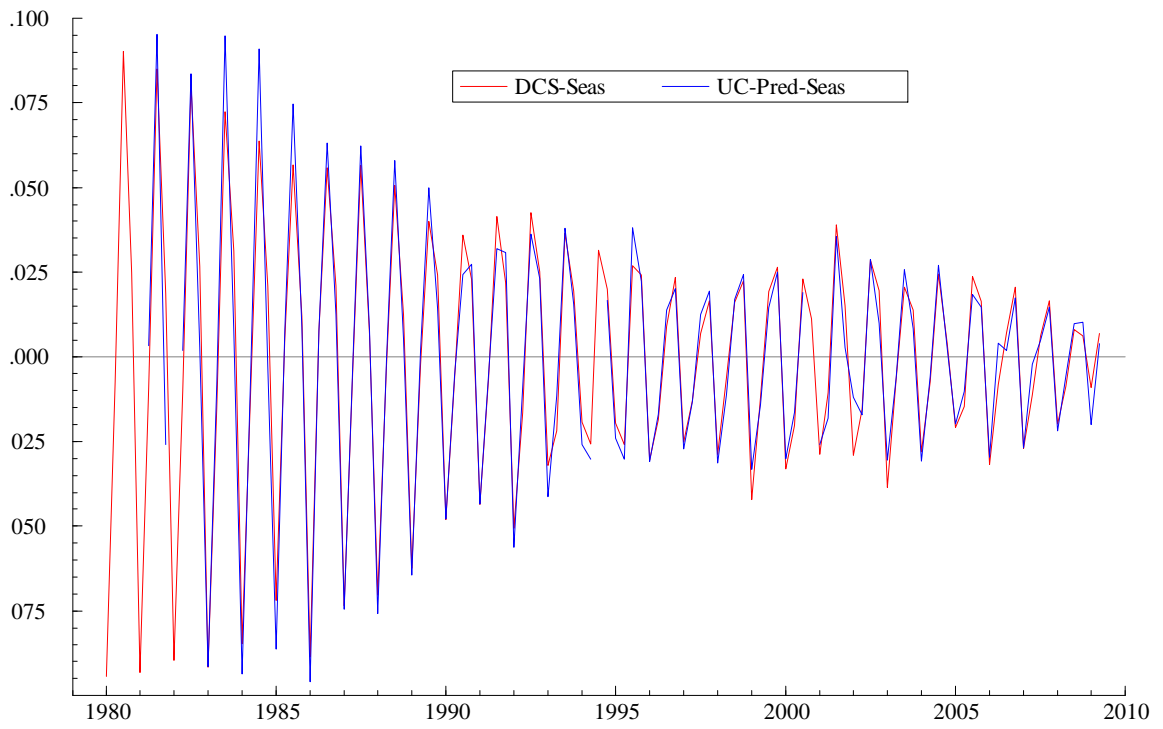


Figure 5: Seasonals in National Rail Travel from UC and DCS models

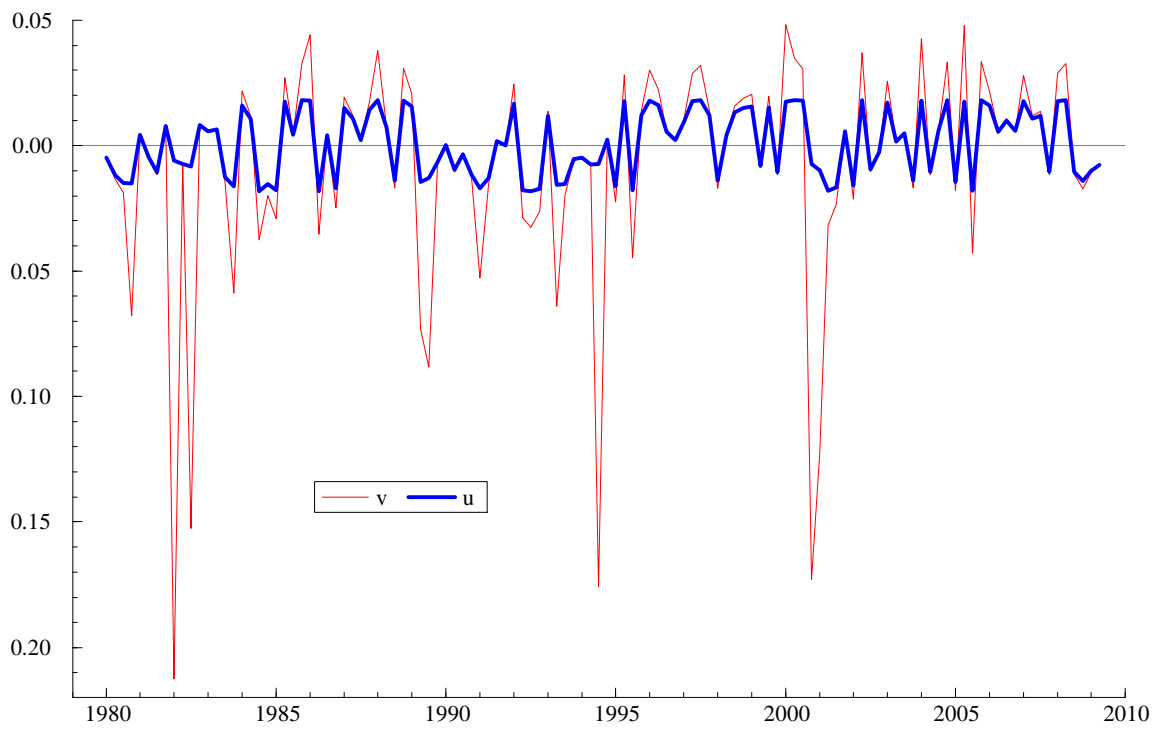


Figure 6: Residuals and scores from DCS-t model

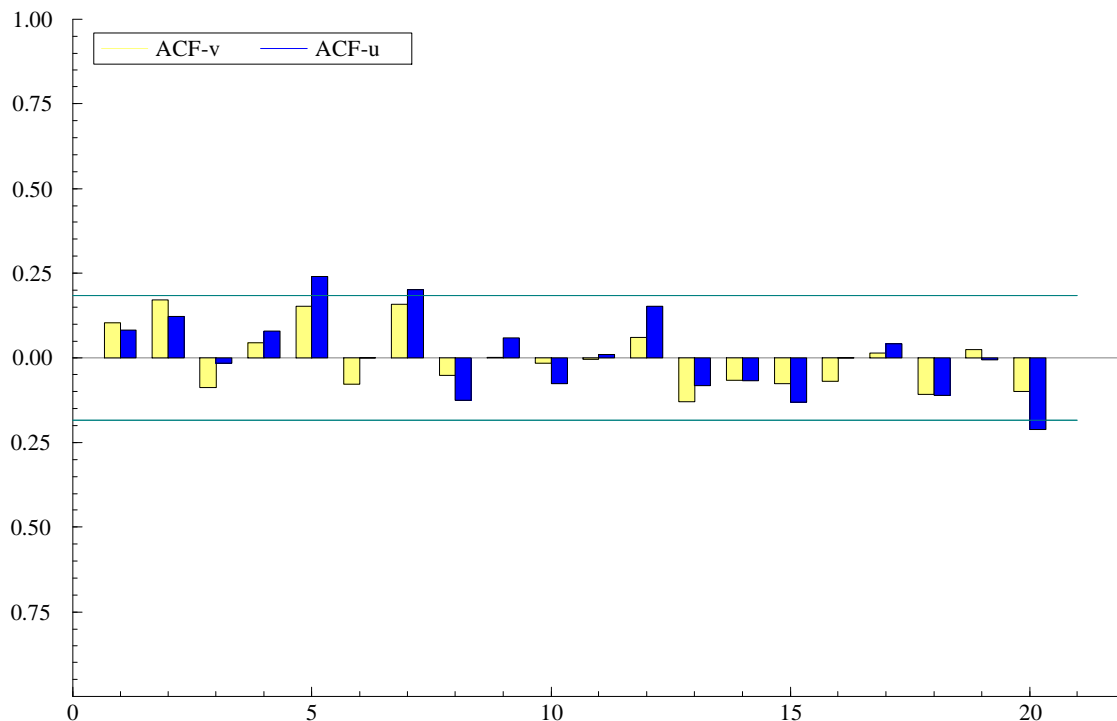


Figure 7: Residual correlograms for irregular and score residuals from DCS-t model fitted to National Rail Travel (lines are $\pm 1/\sqrt{T}$)

Gaussianity.) In the score series the outliers are downweighted and the autocorrelations for the score are slightly bigger than those of the residuals presumably because they are not weakened by aberrant values. The Box-Ljung $Q(12)$ statistic is 19.78 for the score and 12.40 for the residuals. If it can be assumed that only the number of fitted dynamic parameters affects the distribution of the Box-Ljung statistic, its distribution under the null hypothesis of correct model specification is χ_{10}^2 , which had a 5% critical value of 18.3. Thus the scores reject the null hypothesis, albeit only marginally, while the residuals do not. Having said that, the score autocorrelations do not exhibit any clear pattern and it is not clear how the dynamic specification might be improved.

7 Explanatory variables

The location parameter may depend on a set of observable explanatory variables, denoted by the $k \times 1$ vector \mathbf{w}_t , as well as on its own past values and the score. The model can be set up as

$$y_t = \mu_{t|t-1}^\dagger + \mathbf{w}_t' \boldsymbol{\gamma} + \varepsilon_t \exp(\lambda), \quad t = 1, \dots, T, \quad (23)$$

where $\mu_{t|t-1}^\dagger$ could be a stationary process, as in (8), or a stochastic trend such as (19). The model may be augmented by a seasonal component as in sub-section 6.4.

If it is possible to make a sensible guess of initial values of the explanatory variable coefficients, the degrees of freedom parameter, ν , and the dynamic parameters, ϕ and κ for a stationary first-order model or β and κ for a random walk with drift, can be estimated by fitting a univariate model to the residuals, $y_t - \mathbf{w}_t' \hat{\boldsymbol{\gamma}}$, $t = 1, \dots, T$. These values are then used to start off numerical optimization with respect to all the parameters in the model.

7.1 Asymptotic theory

The following result generalizes Proposition 1 by specializing a result in Harvey (2013, Section 2.6)

Proposition 4 *Consider model (23) with a stationary first-order component. Assume that the explanatory variables are weakly stationary with mean $\boldsymbol{\mu}_w$ and second moment $\boldsymbol{\Lambda}_w$ and are strictly exogenous in the sense that they are independent of the ε_t 's and therefore of the u_t 's. Assuming that $b < 1$ and $\kappa \neq 0$, the ML estimator of $(\kappa, \phi, \boldsymbol{\gamma}', \lambda, \nu)'$, is consistent and the limiting distribution of $\sqrt{T}(\tilde{\kappa} - \kappa, \tilde{\phi} - \phi, \tilde{\boldsymbol{\gamma}}' - \boldsymbol{\gamma}', \tilde{\lambda} - \lambda, \tilde{\nu} - \nu)'$ is multivariate normal with mean vector zero and covariance matrix given by the inverse of the information matrix in (11) but with $\boldsymbol{\psi}$ replaced by $(\kappa, \phi)'$ and $\mathbf{D}(\boldsymbol{\psi})$ replaced by*

$$\mathbf{D} \begin{pmatrix} \kappa \\ \phi \\ \boldsymbol{\gamma} \end{pmatrix} = \frac{1}{1-b} \begin{bmatrix} A & D & \mathbf{0}' \\ D & B & \mathbf{0}' \\ \mathbf{0} & \mathbf{0} & \mathbf{C}_w \end{bmatrix},$$

with A, B and D, E as in (12) while

$$\mathbf{C}_w = (1 + \phi^2) \boldsymbol{\Lambda}_w - 2\phi \boldsymbol{\Lambda}_w(1) + 2a(1-a)^{-1}(1-\phi)^2 \boldsymbol{\mu}_w \boldsymbol{\mu}_w',$$

with $\boldsymbol{\Lambda}_w(1) = E(\mathbf{w}_t \mathbf{w}_{t-1}') = E(\mathbf{w}_{t-1} \mathbf{w}_t')$.

An estimator of the asymptotic covariance matrix can be obtained by replacing $\mathbf{\Lambda}_w$ and $\mathbf{\Lambda}_w(1)$ by $T^{-1} \sum \mathbf{w}_t \mathbf{w}'_t$ and $T^{-1} \sum \mathbf{w}_t \mathbf{w}'_{t-1}$, respectively. The constant term, ω , will normally appear as an explanatory variable in which case the corresponding element in \mathbf{w}_t will be unity.

Corollary 5 *When $\mu_{t|t-1}^\dagger$ is known to be a random walk with drift, β , as in (18), and $\mu_{1|0}^\dagger$ is fixed and known, the information matrix is as in (11) but with*

$$\mathbf{D} \begin{pmatrix} \kappa \\ \gamma \\ \beta \end{pmatrix} = \frac{1}{1-b} \begin{bmatrix} \sigma_u^2 & \mathbf{0}' & \mathbf{0}' \\ \mathbf{0} & \mathbf{C}_{\Delta w} & \boldsymbol{\mu}_{\Delta w} \\ \mathbf{0} & \boldsymbol{\mu}'_{\Delta w} & 1 \end{bmatrix}, \quad b < 1,$$

where $\boldsymbol{\mu}_{\Delta w} = E(\Delta \mathbf{w}_t)$ and $\mathbf{C}_{\Delta w} = E(\Delta \mathbf{w}_t \Delta \mathbf{w}'_t)$.

The first differences of the explanatory variables must be weakly stationary but their levels may be nonstationary. It follows from the above result that the covariance matrix of the limiting distribution of $\sqrt{T} \tilde{\boldsymbol{\gamma}}$ is

$$\text{Var}(\tilde{\boldsymbol{\gamma}}) = \left(2\kappa \frac{\nu}{\nu+1} - \kappa^2 \frac{\nu(\nu^3 + 10\nu^2 + 35\nu + 38)}{(\nu+1)^2(\nu+5)(\nu+7)} \right) e^{2\lambda} (\mathbf{C}_{\Delta w} - \boldsymbol{\mu}_{\Delta w} \boldsymbol{\mu}'_{\Delta w})^{-1} \quad (24)$$

In principle, the above Corollary may be extended to models where seasonals are included.

7.2 Application to rail travel

Potential explanatory variables for the rail travel series of Sub-section 6.5 are: (i) Real GDP (in £2003 prices), (ii) Real Fares, obtained by dividing total revenue by the number of kilometres travelled and the retail price index (RPI), and (iii) Petrol and Oil index (POI), divided by RPI. See Figure 8. The fares series was smoothed by fitting a univariate UC model.

Fitting an unobserved components time series model using STAMP gave the following estimates for the coefficients of the logarithms of the explanatory variables: GDP was 0.716 (0.267), fares was -0.416 (0.245) and POI was 0.050 (0.065). Because the explanatory variables enter the model in logarithms, their coefficients are elasticities. All the estimates are all plausible. The coefficient of the petrol index is not statistically significant at any conventional level, but at least it has the right sign. The appendix shows the print-out of the full set of results.

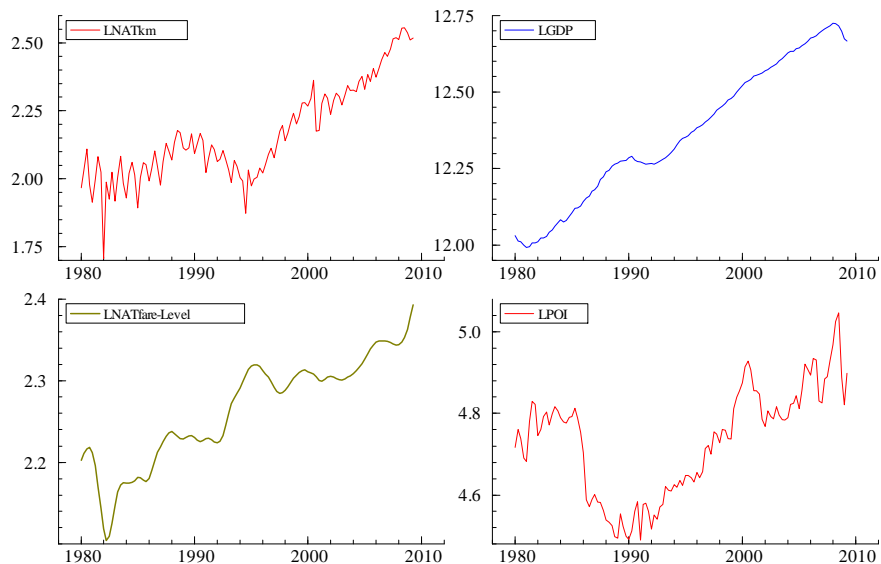


Figure 8: Rail travel in the UK and explanatory variables

Failure to deal with outliers in a time series regression can lead to serious distortions and this is well-illustrated by the rail series when the intervention variables are not included. In particular the fare estimate is *plus* 0.28.

When rail travel was seasonally adjusted by removing the seasonal component obtained from the univariate DCS- t model fitted in sub-section 6.5 and LPOI was also seasonally adjusted, estimating the DCS- t model without a seasonal component gave⁴

$$\tilde{\kappa} = 1.346(0.151) \quad \tilde{\lambda} = -3.879 (0.102) \quad \tilde{\nu} = 2.436 (0.534) \quad \tilde{\beta} = 0.001 (0.002),$$

where the figures in parentheses are asymptotic standard errors. The ASEs calculated for the coefficients of LGDP, Lfare (level) and LPOI (seasonally adjusted) using $Var(\tilde{\gamma})$ in (26) were 0.251, 0.246 and 0.050 respectively. These figures are close to the standard errors for the UC model (with seasonal component) reported in the first paragraph of this sub-section. (The estimated SEs obtained from a UC model fitted to seasonal adjusted data were similar).

⁴The value of b was 0.457.

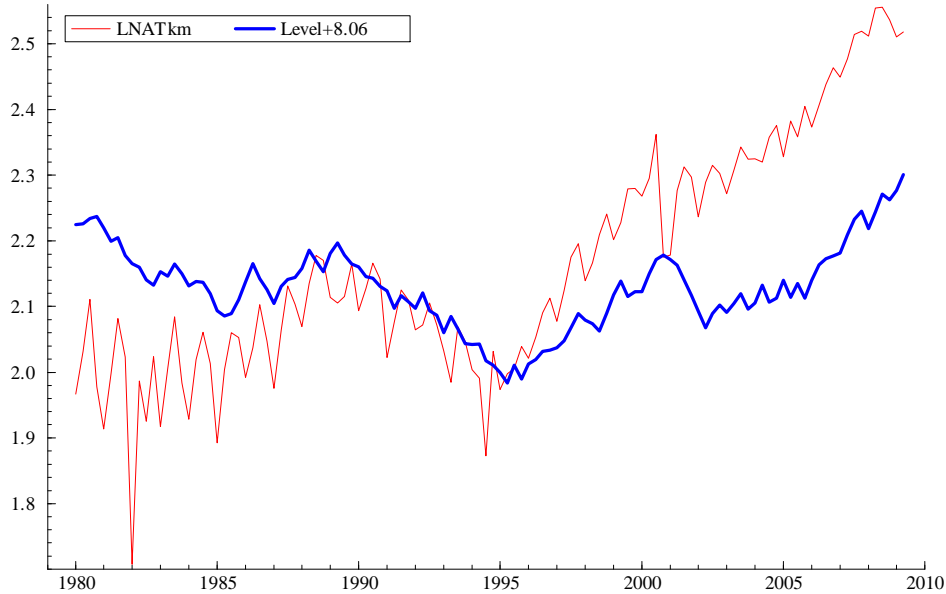


Figure 9: Trend in rail travel after explanatory variables have been taken into account. (A constant has been added to the trend so that it is at a level comparable with that of the series.)

Fitting the full DCS-t model with the seasonal gave $\tilde{\kappa} = 2.212$, $\tilde{\kappa}_s = 0.771$, $\tilde{\lambda} = -4.059$, $\tilde{\nu} = 2.070$ and $\tilde{\beta} = 0.0004$, with initial values $\tilde{\mu} = -6.162$, $\tilde{\gamma}_1 = -0.084$, $\tilde{\gamma}_2 = -0.007$ and $\tilde{\gamma}_3 = 0.070$. The coefficients of the explanatory variables were: $LGDP = 0.734$, $Lfare = -0.427$ and $LPOI = 0.056$. The Box-Ljung $Q(12)$ statistic is 5.30 for the score and 16.12 for the residuals. This result is a little surprising because in the univariate model the Q - statistic for the score was bigger than that of the residuals.

A good deal, but by no means all, of the growth in rail travel from the mid-nineties is due to the increase in GDP. The continued fall after the economy had moved out of the recession of the early nineties is partly explained by the fact that fares increased sharply in 1993 in anticipation of rail privatisation and continued to increase till 1995. Nevertheless, as is apparent from Figure 9, there remain long-term movements in rail travel that cannot be accounted for by the exogenous variables.

8 Conclusions

In this paper we develop, analyse and apply a filter that enables a time series model to be handled robustly. Our model-based approach, based on a t -distribution, is relatively simple, both conceptually and computationally, and is amenable to diagnostic checking and generalization. The article first considers stationary models and then moves on to include trend and seasonal components. The same techniques could be applied to robustify ARIMA and seasonal ARIMA models. Optimal forecasts can be computed recursively, either as in an ARMA model or by using the SSF, and multi-step conditional distributions can be easily constructed by simulation.

The viability of the techniques was illustrated with real data. The penultimate section introduced explanatory variables into the model. Other generalizations are possible. For example a skewed- t model may be adopted using the method used by Harvey and Sucarrat (2012) for a volatility model.

Acknowledgements. An earlier version of the paper was presented at the CFE conference in London in December, 2011. We would like to thank Phillip Andres, Ralph Snyder and Stephen Thiele for helpful comments on an earlier draft.

References

- [1] Creal, D., Koopman, S.J. and Lucas A. (2011), A Dynamic Multivariate Heavy-Tailed Model for Time-Varying Volatilities and Correlations, *Journal of Business and Economics Statistics*, 29, 4, 552–563.
- [2] De Livera, A. M., Hyndman, R. J. and R. D. Snyder (2011). Forecasting time series with complex seasonal patterns using exponential smoothing. *Journal of the American Statistical Association* **106**, 1513-1527.
- [3] Durbin, J. and S. J. Koopman (2012). *Time Series Analysis by State Space Methods*. 2nd edition. Oxford: Oxford Statistical Science Series.
- [4] Harvey A. C. (1989). *Forecasting, Structural Time Series Models and the Kalman Filter*. Cambridge: Cambridge University Press.
- [5] Harvey, A.C. (2013). Dynamic Models for Volatility and Heavy Tails: with applications to financial and economic time series. Econometric Society Monograph, Cambridge University Press (forthcoming).

- [6] Harvey, A. C. and G. Sucarrat (2012). EGARCH Models with Fat Tails, Skewness and Leverage. Cambridge Working paper in Economics, CWPE 1236.
- [7] Koopman, S. J. Harvey, A. C., Doornik, J. A. and N. Shephard (2009). *STAMP 8.2 Structural Time Series Analysis Modeller and Predictor*. London: Timberlake Consultants Ltd.
- [8] Lange, K. L., Little, R. J. A. and M. G. Taylor (1989). Robust statistical modeling using the t distribution. *Journal of the American Statistical Association* **84**, 881-896.
- [9] Maronna, R.A., Martin R.D. and Yohai, V.J. (2006), *Robust Statistics*, Wiley.
- [10] Muler, N., Peña D. and Yohai, V.J. (2009), Robust Estimation for ARMA Models, *Annals of Statistics*, **37**, 2, 816-840.
- [11] Ord, J. K., Koehler, A.B. and R.D. Snyder, (1997), Estimation and Prediction for a Class of Dynamic Nonlinear Statistical Models, *Journal of the American Statistical Association*, **92**, 1621-1629.

APPENDIX

A Consistency and asymptotic normality of the ML estimator

This appendix explains how to derive the information matrix of the ML estimator for the first-order model and outlines a proof for consistency and asymptotic normality. As noted in the text, if the model is to be identified, κ must not be zero and or such that the constraint $b < 1$ is violated. A more formal statement is that the parameters should be interior points of the compact parameter space which will be taken to be $|\phi| < 1$, $|\omega| < \infty$ and $0 < \kappa < \kappa_u$, $\kappa_L < \kappa < 0$ where κ_u and κ_L are values determined by the condition $b < 1$.

The first step is to decompose the derivatives of the log density wrt $\boldsymbol{\psi}$ into derivatives wrt $\mu_{t,t-1}$ and derivatives of $\mu_{t,t-1}$ wrt $\boldsymbol{\psi}$, that is

$$\frac{\partial \ln f_t}{\partial \boldsymbol{\psi}} = \frac{\partial \ln f_t}{\partial \mu_{t,t-1}} \frac{\partial \mu_{t,t-1}}{\partial \boldsymbol{\psi}}, \quad i = 1, 2, 3.$$

Since the scores $\partial \ln f_t / \partial \mu_{t,t-1}$ are $IID(0, \sigma_u^2)$ and so do not depend on $\mu_{t,t-1}$,

$$E_{t-1} \left[\left(\frac{\partial \ln f_t}{\partial \mu_{t,t-1}} \frac{\partial \mu_{t,t-1}}{\partial \boldsymbol{\psi}} \right) \left(\frac{\partial \ln f_t}{\partial \mu_{t,t-1}} \frac{\partial \mu_{t,t-1}}{\partial \boldsymbol{\psi}} \right)' \right] = \left[E \left(\frac{\partial \ln f_t}{\partial \mu} \right)^2 \right] \frac{\partial \mu_{t,t-1}}{\partial \boldsymbol{\psi}} \frac{\partial \mu_{t,t-1}}{\partial \boldsymbol{\psi}'} = \sigma_u^2 \frac{\partial \mu_{t,t-1}}{\partial \boldsymbol{\psi}} \frac{\partial \mu_{t,t-1}}{\partial \boldsymbol{\psi}'}$$

Thus the unconditional expectation requires evaluating the last term.

The derivative of $\mu_{t,t-1}$ wrt κ is

$$\frac{\partial \mu_{t,t-1}}{\partial \kappa} = \phi \frac{\partial \mu_{t-1,t-2}}{\partial \kappa} + \kappa \frac{\partial u_{t-1}}{\partial \kappa} + u_{t-1}, \quad t = 2, \dots, T.$$

However,

$$\frac{\partial u_t}{\partial \kappa} = \frac{\partial u_t}{\partial \mu_{t,t-1}} \frac{\partial \mu_{t,t-1}}{\partial \kappa},$$

Therefore

$$\frac{\partial \mu_{t,t-1}}{\partial \kappa} = x_{t-1} \frac{\partial \mu_{t-1,t-2}}{\partial \kappa} + u_{t-1} \tag{25}$$

where

$$x_t = \phi + \kappa \frac{\partial u_t}{\partial \mu_{t,t-1}}, \quad t = 1, \dots, T. \tag{26}$$

Define

$$a = E_{t-1}(x_t) = \phi + \kappa E_{t-1} \left(\frac{\partial u_t}{\partial \mu_{tt-1}} \right) = \phi + \kappa E \left(\frac{\partial u_t}{\partial \mu} \right)$$

Since $\partial u_t / \partial \mu_{tt-1}$ is IID, unconditional expectations can replace conditional ones. When the process for μ_{tt-1} starts in the infinite past and $|a| < 1$, taking conditional expectations of the derivatives at time $t - 2$, followed by unconditional expectations gives

$$E \left(\frac{\partial \mu_{tt-1}}{\partial \kappa} \right) = E \left(\frac{\partial \mu_{tt-1}}{\partial \phi} \right) = 0 \quad \text{and} \quad E \left(\frac{\partial \mu_{tt-1}}{\partial \omega} \right) = \frac{1 - \phi}{1 - a}.$$

To derive the information matrix, square both sides of (25) and take conditional expectations to give

$$\begin{aligned} E_{t-2} \left(\frac{\partial \mu_{tt-1}}{\partial \kappa} \right)^2 &= E_{t-2} \left(x_{t-1} \frac{\partial \mu_{t-1t-2}}{\partial \kappa} + u_{t-1} \right)^2 \\ &= b \left(\frac{\partial \mu_{t-1t-2}}{\partial \kappa} \right)^2 + 2c \frac{\partial \mu_{t-1t-2}}{\partial \kappa} + \sigma_u^2, \end{aligned} \quad (27)$$

where

$$\begin{aligned} b &= E_{t-1}(x_t^2) = \phi^2 + 2\phi\kappa E \left(\frac{\partial u_t}{\partial \mu} \right) + \kappa^2 E \left(\frac{\partial u_t}{\partial \mu} \right)^2 \geq 0, \quad \text{and} \\ c &= E_{t-1}(u_t x_t) = \kappa E \left(u_t \frac{\partial u_t}{\partial \mu} \right) \end{aligned}$$

Taking unconditional expectations gives

$$E \left(\frac{\partial \mu_{tt-1}}{\partial \kappa} \right)^2 = bE \left(\frac{\partial \mu_{t-1t-2}}{\partial \kappa} \right)^2 + 2cE \left(\frac{\partial \mu_{t-1t-2}}{\partial \kappa} \right) + \sigma_u^2$$

and so, provided that $b < 1$,

$$E \left(\frac{\partial \mu_{tt-1}}{\partial \kappa} \right)^2 = \frac{\sigma_u^2}{1 - b}.$$

Expressions for other elements in the information matrix may be similarly derived; see Harvey (2010, 2013). Fulfillment of the condition $b < 1$ implies

$|a| < 1$. That this is the case follows directly from the Cauchy-Schwartz inequality $E(x_t^2) \geq [E(x_t)]^2$.

The information matrix, (11), is given by noting that

$$\frac{\partial u_t}{\partial \mu} = 2(1 - b_t)b_t - (1 - b_t). \quad (28)$$

The distribution of (28) does not depend on μ and $E(\partial u_t / \partial \mu) = -\nu / (\nu + 3)$. Similarly

$$E\left(u_t \frac{\partial u_t}{\partial \mu}\right) = E(2(1 - b_t)b_t - (1 - b_t))(y_t - \mu_{t|t-1})(1 - b_t) = 0$$

because $E((y_t - \mu_{t|t-1}) | b_t) = 0$, and

$$E\left(\frac{\partial u_t}{\partial \mu}\right)^2 = E(2(1 - b_t)b_t - (1 - b_t))^2 = \frac{\nu(\nu^3 + 10\nu^2 + 35\nu + 38)}{(\nu + 1)(\nu + 3)(\nu + 5)(\nu + 7)} \leq 1.$$

Consistency and asymptotic normality can be proved by showing that the conditions for Lemma 1 in Jensen and Rahbek (2004, p 1206) hold. The main point to note is that the first three derivatives of $\mu_{t|t-1}$ wrt κ , ϕ and ω are stochastic recurrence equations (SREs); see Brandt (1986) and Straumann and Mikosch (2006, p 2450-1). The condition $b < 1$ is sufficient⁵ to ensure that they are strictly stationary and ergodic at the true parameter value. Similarly $b < 1$ is sufficient to ensure that the squares of the first derivatives are strictly stationary and ergodic.

Let $\boldsymbol{\psi}_0$ denote the true value of $\boldsymbol{\psi}$. Since the score and its derivatives wrt μ in the static model possess the required moments, it is straightforward to show that (i) as $T \rightarrow \infty$, $(1/\sqrt{T})\partial \ln L(\boldsymbol{\psi}_0)/\partial \boldsymbol{\psi} \rightarrow N(0, \mathbf{I}(\boldsymbol{\psi}_0))$, where $\mathbf{I}(\boldsymbol{\psi}_0)$ is p.d. and (ii) as $T \rightarrow \infty$, $(-1/T)\partial^2 \ln L(\boldsymbol{\psi}_0)/\partial \boldsymbol{\psi} \partial \boldsymbol{\psi}' \xrightarrow{P} \mathbf{I}(\boldsymbol{\psi}_0)$. The final condition in Jensen and Rahbek (2004) is concerned with boundedness of the third derivative of the log-likelihood function in the neighbourhood of $\boldsymbol{\psi}_0$. The first derivative of u_t , (28) is a linear function of terms of the form $b_t^* = b_t^h(1 - b_t)^k$, where h and k are non-negative integers, as is the second derivative. As regards u_t itself, since $u_t = (1 - b_t)(y_t - \mu_{t|t-1})$, it can be seen

⁵The necessary condition for strict stationarity is $E(\ln|x_t|) < 0$. This condition is satisfied at the true parameter value when $|a| < 1$ since, from Jensen's inequality, $E(\ln|x_t|) \leq \ln E(|x_t|) < 0$ and as already noted $b < 1$ implies $|a| < 1$.

that that $u_t = 0$ when $y_t = 0$ and $u_t \rightarrow 0$ as $|y_t| \rightarrow \infty$. Thus u_t , like its derivatives, is bounded for any admissible ψ . Since

$$b_t = h(y_t; \psi)/(1 + h(y_t; \psi)), \quad 0 \leq h(y_t; \psi) \leq \infty,$$

where $h(y_t; \psi)$ depends on y_t and ψ , it is clear that for any admissible ψ , $0 \leq b_t \leq 1$ and so $0 \leq b_t^* \leq 1$. Furthermore the derivatives of $\mu_{t|t-1}$ must be bounded at ψ_0 since they are stable SREs which are ultimately dependent on u_t and its derivatives. They must also be bounded in the neighbourhood of ψ_0 since the condition $b < 1$ is more than enough to guarantee the stability condition $E(\ln |x_t|) < 0$.

Unknown shape parameters, including degrees of freedom, pose no problem as the third derivatives (including cross-derivatives) associated with them are almost invariably non-stochastic.

B Explanatory variables: proof of Proposition 4.

Model (23) may be re-written as

$$\mu_{t+1|t} = (\mathbf{w}_{t+1} - \phi \mathbf{w}_t)' \gamma + \phi \mu_{t|t-1} + \kappa u_t. \quad (29)$$

Differentiating $\mu_{t+1|t}$ with respect to γ gives

$$\frac{\partial \mu_{t+1|t}}{\partial \gamma} = x_t \frac{\partial \mu_{t|t-1}}{\partial \gamma} + \mathbf{w}_{t+1} - \phi \mathbf{w}_t$$

and so

$$E_{t-1} \left(\frac{\partial \mu_{t+1|t}}{\partial \gamma} \right) = a \frac{\partial \mu_{t|t-1}}{\partial \gamma} + E_{t-1}(\mathbf{w}_{t+1} - \phi \mathbf{w}_t).$$

Taking unconditional expectations gives

$$E \left(\frac{\partial \mu_{t+1|t}}{\partial \gamma} \right) = \frac{1 - \phi}{1 - a} \boldsymbol{\mu}_w.$$

Now

$$\begin{aligned} E_{t-1} \left(\frac{\partial \mu_{t+1|t}}{\partial \gamma} \frac{\partial \mu_{t+1|t}}{\partial \gamma'} \right) &= b \frac{\partial \mu_{t|t-1}}{\partial \gamma} \frac{\partial \mu_{t|t-1}}{\partial \gamma'} + E_{t-1}((\mathbf{w}_{t+1} - \phi \mathbf{w}_t)(\mathbf{w}_{t+1} - \phi \mathbf{w}_t)') \\ &+ \frac{\partial \mu_{t|t-1}}{\partial \gamma} E_{t-1}(x_t (\mathbf{w}_{t+1} - \phi \mathbf{w}_t)') + E_{t-1}(x_t (\mathbf{w}_{t+1} - \phi \mathbf{w}_t)) \frac{\partial \mu_{t|t-1}}{\partial \gamma'}. \end{aligned} \quad (30)$$

Since \mathbf{w}_t is assumed to be strictly exogenous,

$$\frac{\partial \mu_{t|t-1}}{\partial \boldsymbol{\gamma}'} E_{t-1}(x_t(\mathbf{w}_{t+1} - \phi \mathbf{w}_t)) = a \frac{\partial \mu_{t|t-1}}{\partial \boldsymbol{\gamma}'} E_{t-1}((\mathbf{w}_{t+1} - \phi \mathbf{w}_t))$$

and when unconditional expectations are taken,

$$aE \left[\frac{\partial \mu_{t|t-1}}{\partial \boldsymbol{\gamma}'} (\mathbf{w}_{t+1} - \phi \mathbf{w}_t) \right] = aE \frac{\partial \mu_{t|t-1}}{\partial \boldsymbol{\gamma}} (1 - \phi) \boldsymbol{\mu}'_w = \frac{a(1 - \phi)^2}{1 - a} \boldsymbol{\mu}_w \boldsymbol{\mu}'_w.$$

Taking unconditional expectations in (30) then gives $(1 - b)^{-1} \mathbf{C}_w$.

Additional references

Brandt, A. (1986). The stochastic equation $Y_{n+1} = A_n Y_n + B_n$ with stationary coefficients. *Advances in Applied Probability* **18**, 211–220.

Harvey, A.C. (2010). Exponential Conditional Volatility Models (2010). Working Paper 10-36, Statistics and Econometrics Series 20, Carlos III University, Madrid.

Jensen, S.T. and A. Rahbek (2004). Asymptotic inference for nonstationary GARCH. *Econometric Theory*, 20, 1203–26.

Straumann, D. and T. Mikosch (2006). Quasi-maximum-likelihood estimation in conditionally heteroscedastic time series: a stochastic recurrence equations approach. *Annals of Statistics* **34**, 2449–2495.

Table 1: Simulation results for ML estimation of first-order DCS model based on 1000 replications

	$T = 500$					$T = 1000$				
	$\phi = 0.8$	$\nu = 6$	$\lambda = 0$	$\kappa = 0.1$	$\omega = 0$	$\phi = 0.8$	$\nu = 6$	$\lambda = 0$	$\kappa = 0.1$	$\omega = 0$
Est. par.	0.5618	6.6963	-0.0058	0.1068	-0.0012	0.6945	6.2203	-0.0040	0.1032	0.0002
Rmse	0.5137	4.3271	0.0534	0.0624	0.0666	0.3055	1.2869	0.0372	0.0411	0.0478
Num. se	0.2249	2.1046	0.0537	0.0539	0.0662	0.1582	1.2067	0.0377	0.0389	0.0475
Asy. se	0.1662	1.1932	0.0413	0.0543	0.0675	0.1175	0.8437	0.0292	0.0384	0.0478
Est. par.	$\phi = 0.8$	$\nu = 6$	$\lambda = 0$	$\kappa = 0.5$	$\omega = 0$	$\phi = 0.8$	$\nu = 6$	$\lambda = 0$	$\kappa = 0.5$	$\omega = 0$
Rmse	0.7843	6.3575	-0.0072	0.5013	0.0018	0.7928	6.1643	-0.0046	0.4994	0.0002
Num. se	0.0550	1.8529	0.0495	0.0763	0.1279	0.0367	1.1613	0.0348	0.0530	0.0932
Asy. se	0.0532	1.6744	0.0500	0.0746	0.1289	0.0363	1.0658	0.0352	0.0524	0.0925
Asy. se	0.0496	1.1932	0.0413	0.0612	0.1330	0.0371	0.8437	0.0292	0.0433	0.0940
Est. par.	$\phi = 0.8$	$\nu = 6$	$\lambda = 0$	$\kappa = 1$	$\omega = 0$	$\phi = 0.8$	$\nu = 6$	$\lambda = 0$	$\kappa = 1$	$\omega = 0$
Rmse	0.7910	6.1534	-0.0085	1.0043	0.0031	0.7956	6.0877	-0.0051	1.0009	-0.0000
Num. se	0.0363	1.3053	0.0445	0.0923	0.1958	0.0250	0.9203	0.0313	0.0673	0.1437
Asy. se	0.0353	1.2632	0.0443	0.0938	0.1996	0.0245	0.8549	0.0311	0.0657	0.1438
Asy. se	0.0340	1.1932	0.0413	0.0631	0.2077	0.0240	0.8437	0.0292	0.0446	0.1469
Est. par.	$\phi = 0.8$	$\nu = 6$	$\lambda = 0$	$\kappa = 1.3$	$\omega = 0$	$\phi = 0.8$	$\nu = 6$	$\lambda = 0$	$\kappa = 1.3$	$\omega = 0$
Rmse	0.7924	6.1488	-0.0072	1.3010	0.0115	0.7963	6.0538	-0.0051	1.3019	-0.0001
Num. se	0.0314	1.2080	0.0408	0.1025	0.2281	0.0217	0.7812	0.0292	0.0711	0.1670
Asy. se	0.0306	1.0903	0.0408	0.0981	0.2336	0.0213	0.7299	0.0287	0.0686	0.1690
Asy. se	0.0296	1.1932	0.0413	0.0608	0.2453	0.0209	0.8437	0.0292	0.0430	0.1735
Est. par.	$\phi = 0.95$	$\nu = 6$	$\lambda = 0$	$\kappa = 0.1$	$\omega = 0$	$\phi = 0.95$	$\nu = 6$	$\lambda = 0$	$\kappa = 0.1$	$\omega = 0$
Rmse	0.8504	6.6141	-0.0058	0.1062	0.0019	0.9026	6.2388	-0.0014	0.1006	-0.0014
Num. se	0.2931	2.9523	0.0530	0.0471	0.1207	0.2504	1.4187	0.0854	0.0292	0.0854
Asy. se	0.0822	2.0735	0.0535	0.0426	0.1059	0.0337	1.2099	0.0804	0.0275	0.0804
Asy. se	0.0336	1.1932	0.0413	0.0364	0.1180	0.0238	0.8437	0.0835	0.0257	0.0835
Est. par.	$\phi = 0.95$	$\nu = 6$	$\lambda = 0$	$\kappa = 0.5$	$\omega = 0$	$\phi = 0.95$	$\nu = 6$	$\lambda = 0$	$\kappa = 0.5$	$\omega = 0$
Rmse	0.9394	6.3371	-0.0072	0.5002	0.0022	0.9450	6.1401	-0.0042	0.4990	0.0127
Num. se	0.0228	1.7966	0.0488	0.0696	0.3248	0.0146	1.1003	0.0348	0.0478	0.2443
Asy. se	0.0198	1.6155	0.0492	0.0676	0.3188	0.0129	1.0301	0.0346	0.0475	0.2452
Asy. se	0.0170	1.1932	0.0413	0.0532	0.3805	0.0120	0.8437	0.0292	0.0376	0.2691
Est. par.	$\phi = 0.95$	$\nu = 6$	$\lambda = 0$	$\kappa = 1$	$\omega = 0$	$\phi = 0.95$	$\nu = 6$	$\lambda = 0$	$\kappa = 1$	$\omega = 0$
Rmse	0.9423	6.2434	-0.0066	0.9944	0.0189	0.9459	6.0660	-0.0047	1.0012	0.0231
Num. se	0.0188	1.3521	0.0449	0.0930	0.4861	0.0121	0.8815	0.0311	0.0641	0.3870
Asy. se	0.0161	1.2803	0.0439	0.0912	0.5306	0.0108	0.8357	0.0309	0.0642	0.4162
Asy. se	0.0144	1.1932	0.0413	0.0608	0.6844	0.0102	0.8437	0.0292	0.0430	0.4839
Est. par.	$\phi = 0.95$	$\nu = 6$	$\lambda = 0$	$\kappa = 1.3$	$\omega = 0$	$\phi = 0.95$	$\nu = 6$	$\lambda = 0$	$\kappa = 1.3$	$\omega = 0$
Rmse	0.9432	6.0804	-0.0085	1.3066	0.0141	0.9470	6.0482	-0.0042	1.3030	-0.0045
Num. se	0.0169	1.0811	0.0416	0.1072	0.5613	0.0108	0.7399	0.0292	0.0710	0.4448
Asy. se	0.0150	1.0612	0.0408	0.0959	0.6088	0.0100	0.7280	0.0287	0.0688	0.4945
Asy. se	0.0135	1.1932	0.0413	0.0606	0.8431	0.0095	0.8437	0.0292	0.0429	0.5962