Fitting trends to time series data

Graeme Chamberlin
Office for National Statistics

Trend estimates are derived from seasonally adjusted data via an averaging process which attempts to remove the irregular component of the time series. This allows the underlying direction of a time series to be identified.

This article describes two of the most important and commonly used methods for constructing a trend estimate. These are the Henderson and Kalman filters.

The article also outlines several of the issues involved in constructing and interpreting trend estimates. These include phase shifting, the end point problem, revisions and the identification of turning points, and assessing the degree of trend smoothness.

Introduction

Many statistical organisations around the world present trend estimates as part of their data releases. This is done to different extents and with a variety of different methods, but in all cases the overarching aim is to remove the short-term irregular movements in the data so that users are left with a better idea of its true underlying path. For many time series, especially at monthly and quarterly frequencies, focusing on the latest release of data can be misleading because the data are volatile. Instead, it is more expedient to take a slightly longer-term view and place the latest release in the context of other recent figures. This is what a trend estimate sets out to achieve.

At present the Office for National Statistics (ONS) makes limited use of trends.[1] When it comes to putting trend estimates in the National Accounts, the Australian Bureau of Statistics (ABS) has gone the furthest down this line.[2] Not only does it fit trends to all its major time series, it actually headlines with the trend estimate in its data releases. This article considers two widely accepted methodologies used to fit trends, the Henderson and Kalman filters.

Constructing a trend estimate – smoothing

Once a data series (Y_t) has been seasonally adjusted, it is thought to consist of two components. The first is the trend (T_t), and the second is short-term irregular movements around this trend (I_t). Trend estimates can therefore be calculated by filtering out the irregular components from a series.

$$Y_t = T_t + I_t \qquad (1)$$

Most filters are based on a moving-average procedure. Given that the irregular movements often tend to be in opposite directions, averaging will effectively allow them to cancel each other out. Any remaining irregularity is then smoothed by having its impact spread over a number of observations. Generally, the longer the length of a filter, the smoother the resulting trend series because irregularities are spread over a greater number of observations.

Figure 1 plots the quarterly percentage change in UK GDP along with a four-period moving average. Here the current trend estimate reflects an average of the current and three previous observations. This produces a smoother series but also introduces the problem of phase shifting – that is, the trend lags behind the actual data. This is particularly evident where the trend estimate attempts to place turning points in the actual data such as the end of 1999 and the beginning of 2000. This is not surprising; being entirely backward-looking, the filter can only respond with a lag to new innovations or directions in the data.

Figure 1
GDP percentage change quarter on quarter (GDP) and 4 quarter moving average (MA4)

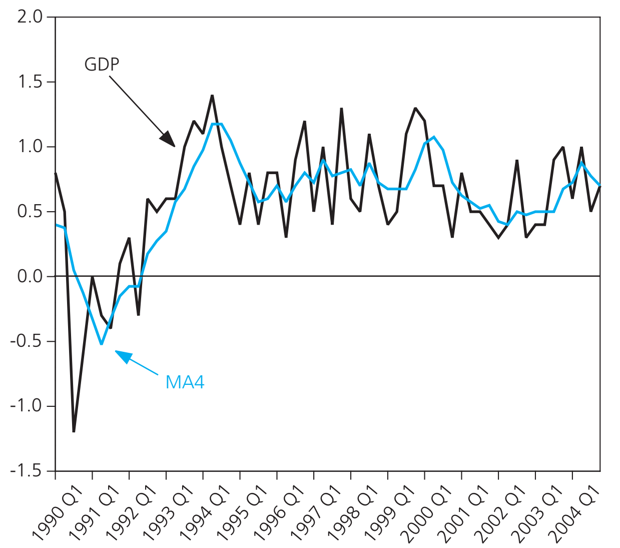

This is a feature of all non-symmetric filters. With many of its monthly data releases, ONS recognises the volatility inherent in high frequency data and presents a smoother estimate based on percentage changes in three-month averages.[3] As a backward-looking procedure this, too, is liable to phase shifting. The price for smoothing the data is to delay its reaction to new directions.

This problem can be avoided by using a symmetric filter. This takes a moving average of the data before and after the observation for which a trend estimate is required. However, because a symmetric filter by definition is centred in the middle of the data, applying a symmetric $2m+1$ term filter would leave m observations unaccounted for at the end of the sample. This is referred to as the end point problem. With data releases, it is of particular significance, as most interest will be in the recent data for which trend estimates are missing. Producing trend estimates up to the end of the sample will therefore require a compromise between satisfying the end point problem and limiting the degree of phase shift.

Henderson filters

The Henderson filter[4] is the traditional workhorse of trend estimation and is also the method of choice applied by the ABS. The Henderson filter weights are designed to strike a compromise between two characteristics expected of trends. The first is that the trend can reproduce a variety of different curvatures, and the second is that they should be as smooth as possible. The first condition is satisfied by designing the filter so that the trends it produces can follow a local cubic polynomial without distortion. This would enable the trend to track curves, peaks and troughs fairly well. The smoothness of the trend reflects the smoothness of the weighting pattern.

These two conditions specify a unique weighting pattern and hence a unique moving average for each possible length of filter considered. The general form of a Henderson filter of term $2m+1$ is given by:

$$w_j = \frac{315[(m+1)^2-j^2][(m+2)^2-j^2][(m+3)^2-j^2][3(m+2)^2-11j^2-16]}{8(m+2)[(m+2)^2-1][4(m+2)^2-1][4(m+2)^2-9][4(m+2)^2-25]} \quad (2)$$

for $j=-m,\ldots,m$

The ABS produces trend estimates for quarterly data using a 7-term Henderson filter. These weights shown in Table 1 can be produced using (2) and setting $m=3$.

Table 1
The Henderson 7-term filter

Period ($N+j$)	$N-3$	$N-2$	$N-1$	N	$N+1$	$N+2$	$N+3$
Weight	–0.059	0.059	0.294	0.413	0.294	0.059	–0.059

In Figure 2 a trend estimate of UK GDP is constructed using the Henderson 7-term weights and is compared with the estimate produced by using a simple[5] 7-term moving-average filter. There are two things to note. The first is that both are symmetric filters, so avoid phase shifting, but suffer from the end point problem. The second is that the simple 7-term moving average tends to flatten out the trend line relative to the Henderson 7-term trend estimate. This can be clearly seen in the period 1999 to 2000 where the Henderson trend reflects the jump in measured GDP growth to a larger extent.

Figure 2
GDP percentage change quarter on quarter (GDP), 7-term simple moving average (MA7) and 7-term Henderson (HF7) filters

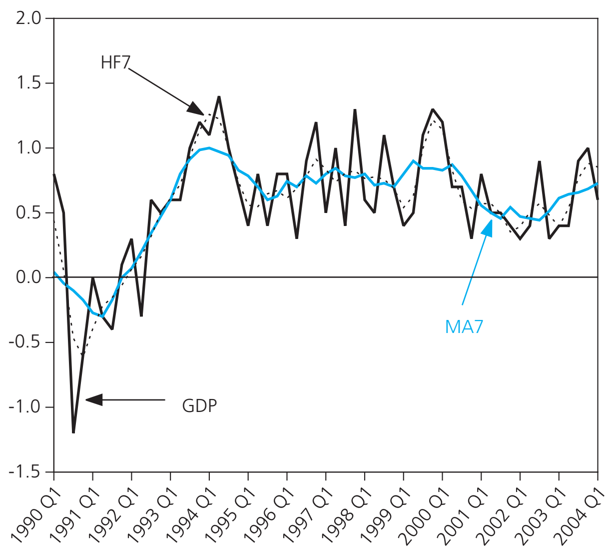

This reflects one of the advantages that non-simple filters like the Henderson have over simple filters. A simple filter can at best only reproduce straight line segments to the data for which it is applied. More generally, this feature arises when the weights are restricted to being non-negative values. Only with the use of negative weights can moving averages be constructed to track various curvatures. The weights in the Henderson filter are designed not only to track linear segments, but also quadratic and cubic segments, so will more accurately reflect curves and points of inflection in the data. In contrast, the simple moving average will suffer from several problems which are evident in Figure 2:

- they underestimate the height of peaks and depths of troughs
- they distort the shape of turning points and as a consequence extend the period over which a trough or a peak appears to exist
- they flatten out points of inflection in the series often resulting in their elimination

Dealing with the end point problem

There are essentially two solutions to this problem. The first is to forecast the original data m periods into the future and then apply the $2m+1$ Henderson filter as before. However, the accuracy of the trend estimate will depend on the accuracy of the forecasts produced. Alternatively, a set of surrogate filters can be applied to the data towards the end of the sample, but because these will be asymmetric filters the problem of phase shifting is reintroduced.

Extrapolating the GDP data forward three periods provides sufficient information to fit the conventional 7-term Henderson filter. This is plotted in Figure 3 where the end of sample trend estimate is represented by the dashed segment.[6]

Figure 3
GDP percentage change quarter on quarter (GDP) and Henderson 7-term filter trend estimate with Autoregressive Integrated Moving Average (ARIMA) extrapolation (ARIMA HF7)

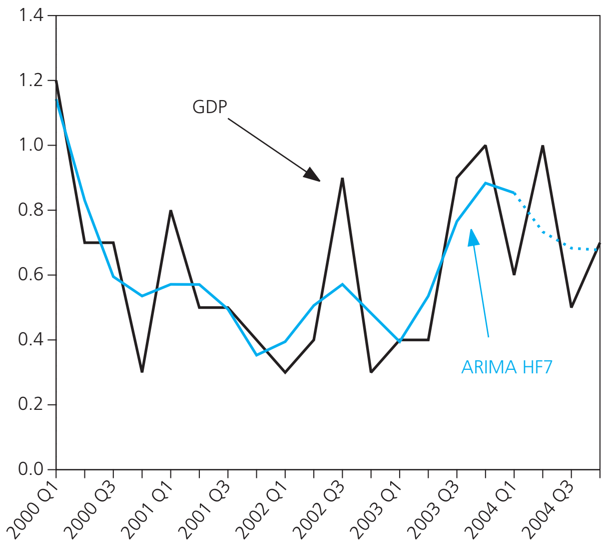

In order to develop a collection of surrogate Henderson filters designed to calculate trend estimates all the way to the end of the sample, it is necessary to consider non-symmetric moving averages. The preferred class of surrogate filter will be those that exhibit similar cycle dampening properties of the main Henderson filter but introduce a minimal amount of phase shift. The ABS uses a procedure outlined in Doherty (2001).[7] This essentially aims to minimise the mean square revision between the trend estimates of the surrogate and the main Henderson filters. Kenny and Durbin (1982)[8] observed that these surrogate Henderson weights are roughly what you achieve if you extrapolate the series by a simple linear model prior to smoothing. Therefore, it is not surprising that the surrogate filters produce a similar end of sample trend to that in Figure 3.

Trend revisions and the identification of turning points

Trend estimates towards the end of samples are subject to revision as forecasts are replaced with actual data or asymmetric filters are replaced with symmetric ones. This can be seen from Figure 4, where trend estimates are derived recursively using the type of ARIMA forecasts shown in Figure 3. The surrogate filter approach produces a similar looking pattern. Trend revisions are usually greater around turning points in the data.

Figure 4
Revisions to the ARIMA 7-term Henderson filter trend estimates

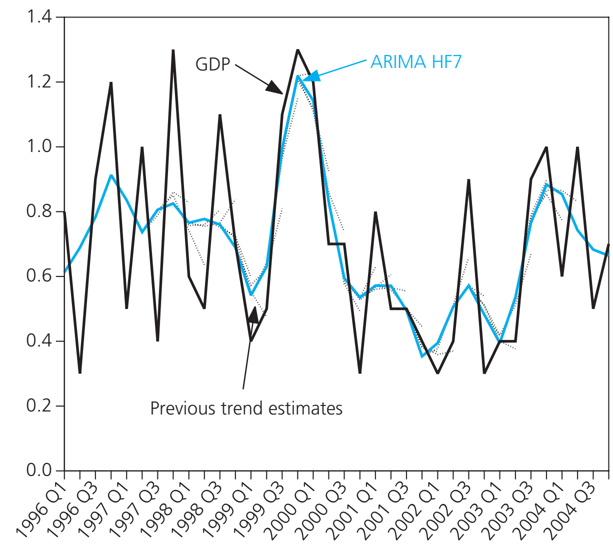

Both techniques are inherently backward-looking. Forecasts are generally produced by extrapolating recent trends, and a surrogate filter takes a moving average of past data. As a result, focusing on trend estimates may hamper the accurate and speedy detection of turning points in the data. This is one of the central issues concerning the use of trend estimates and is often cited as a justification for not using them.

However, this argument can be overstated, for it is wrong to suggest that it is otherwise easy to detect turning points in the data.[9] It must also be acknowledged that in real time data, the definitive existence of turning points might not emerge until later vintages of data are published, by which point the sample has moved on sufficiently to make trend estimates more reliable. As later vintages of the data are based on more information, it is a fair proposition that they will be more accurate in placing turning points. However, these later data vintages are only available with a lag. For example, Blue Book two estimates of National Accounts data are released around 18 to 24 months after the preliminary estimates, so a further five to eight quarters of information will be available before we have a mature estimate of the data at any point in time.[10]

Different types of Henderson filter

As filter lengths become longer, it is expected that the resulting trend estimates become smoother, because any idiosyncratic movement in the data will be averaged out over a larger number of observations. This is shown in Figures 5 and 6 where the central weights of the Henderson 7-, 9- and 11-term filters along with the respective trend estimates are plotted.

Figure 5
Central weights for the Henderson 7-, 9- and 11-term filters (HF7, HF9, HF11)

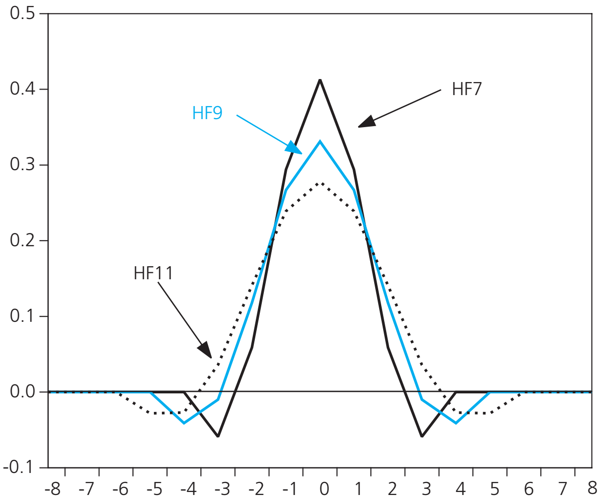

A common criticism levelled at the use of trends is arbitrariness. For example, by selecting different lengths of filter it is possible to produce different trends, and who is to say which is the most appropriate, and why? This takes another dimension when there are several possible ways of producing trends – with Henderson and Kalman filters just two examples. Again, different methods will likely produce different trend estimates, so the practitioner has fair scope to make the trend look how they want.

However, spectral analysis can go some way to justifying the use of different types of filter. A cycle refers to a repeating pattern of behaviour and a time series can be thought of as representing the complex interaction of many different cycles of different behaviours, strengths and frequencies. For example, in GDP we might expect there to be very short-run cycles accounted for by shocks, medium-run cycles accounted for by business cycles and longer-run cycles reflecting the long-term growth path of the economy.

Spectral analysis plots the power of these cycles at different frequencies. The frequency of a cycle refers to the number of times that a particular cycle repeats itself in a given time period. So a cycle with a length of six months will have an annual frequency of two and a ten-year cycle will have an annual frequency of just 0.1. The trend contributes to the behaviour of the spectrum around cycles of longer duration, so will have larger power at smaller frequencies, while the irregular fluctuations dominate in the shorter duration cycles, so will have higher power at bigger frequencies.

Figure 6
7-, 9-, and 11-term Henderson Filter trend estimates (HF7, HF9 and HF11)

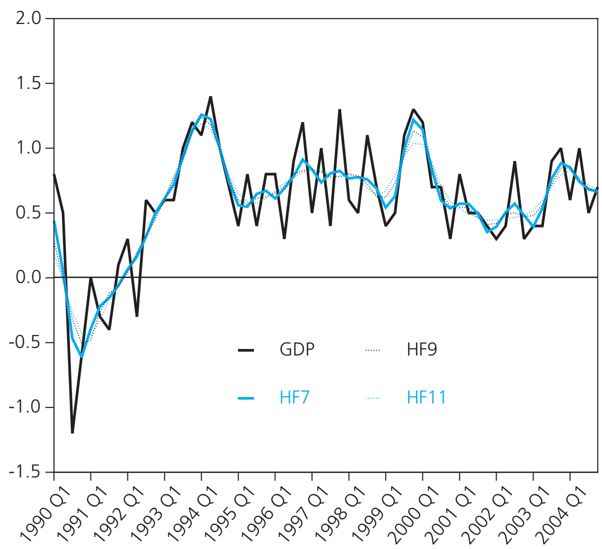

The job of the filter is therefore to identify and isolate the short-term cycles while maintaining the medium- to longer-term cycles. Spectral analysis of the original and trended data can therefore give an indication of what effects a filter has. This effect is summed up in the gain function which defines the length of each cycle that a filter allows to remain in the data.[11] Table 2 reports the gain of different types of Henderson filter.

When applied to quarterly data, a 7-term Henderson moving average preserves 50 per cent or more of the strength of cycles at least 4.63 quarters long; cycles shorter than this are therefore reduced to less than 50 per cent of their strength in the filtered series. When the 7-term Henderson filter is applied to quarterly series, cycles shorter than a year will be largely removed from the data. For example, any cycle with a frequency of less than 3.49 quarters will be reduced to less

Table 2
Impact of Henderson moving averages on cycles[12]

Quarters

No. of terms	10%	25%	50%	75%	90%
5	2.60	2.84	3.34	4.21	5.51
7	3.49	3.88	4.63	5.88	7.74
9	4.33	4.84	5.81	7.41	9.78
11	5.15	5.78	6.95	8.89	11.73
13	5.95	6.69	8.06	10.32	13.64
23	9.89	11.16	13.49	17.31	22.90
33	13.77	15.56	18.84	24.18	31.99

than 10 per cent of its strength. Any cycle longer than 7.74 periods (approximately 2 years) will largely remain with at least 90 per cent of its strength. Consequently, the 7-term Henderson filter removes a substantial part of the irregularity in a quarterly series while retaining medium-term business cycles and longer cycles associated with the secular trend.

As seen in Table 2, the filters with longer terms will smooth out cycles of increasing lengths. The choice of filter will therefore depend on what length of cycles the trend estimate is to observe. Also, it will differ for the frequency of the data, where longer-term filters will be applied to higher frequency data. For example, the ABS uses 7-term filters with quarterly data, but 13-term filters on monthly data.

Kalman filters

Deriving a trend estimate is an example of a signal extraction problem. Signal extraction is concerned with finding the optimal estimate of an unobserved component (UC) of a data series at some point in the sample. In this case, from (1), the trend estimate (T_t) can be viewed as an unobserved component of the measured series (Y_t). These are common problems in economics where the important information (signal) may be embedded in more noisy measured data.[13]

The Kalman filter[14] is an extremely useful algorithm for these types of problems. It is highly flexible and enables the practitioner ample scope to design trend estimates with different properties. It also offers the further advantage of producing the optimal set of filter weights as we approach the end of the sample.

The Kalman filter works in two parts: prediction and updating. Prediction is simply the attempt to make a best guess at the state variables given our knowledge of the system and historical data and estimates. Updating is the process of combining our initial estimate of the state variable with the information contained in the current observed estimates.

The modelling procedure identifies two types of equations. The measurement equation(s) describes the process taken by the observed data, whereas the state equation(s) defines the process taken by the unobserved component.

Measurement equation $Y_t = \gamma' T_t + \varepsilon_t \qquad \varepsilon_t \sim N(0,\sigma^2_\varepsilon)$ (3)

State equation $T_t = \delta' T_{t-1} + \psi_t \qquad \psi_t \sim N(0,\alpha^2\sigma^2_\varepsilon)$ (4)

This is a typical framework for estimating UC models – where in this case the UC is the trend estimate. The measurement equation describes the observed data as the sum of a trend and an error component which is a direct analogy to (1), and the state equation models the trend relative to its own past behaviour.[15] The key factor here is the hyperparameter α^2 which corresponds to the ratio of variances between the measurement and state equations. This is often referred to as the signal to noise ratio.

When the signal to noise ratio is low, then it implies that the variance of the state variable (trend) is low relative to the variance of the measured series. As a result, the variance of the measured series is considered to be largely accounted for by the irregular component and the trend estimate will be relatively smooth. However, as the hyperparameter increases, more of the variance in the measured data is attributed to the trend component. In this case the trend estimate will become less smooth and more reflective of the actual data. These parameters can either be estimated using maximum likelihood or imposed by the user depending on the degree of smoothing required.

Stochastic trends model

This is the basic workhorse of UC and trend modelling consisting of one measurement and two state equations:

$$Y_t = T_t + \varepsilon_t \qquad \varepsilon_t \sim N(0,\sigma^2_\varepsilon)$$

$$T_t = T_{t-1} + \mu_{t-1} + \eta_t \qquad \eta_t \sim N(0,\alpha_1^2\sigma^2_\varepsilon)$$

$$\mu_t = \mu_{t-1} + \varsigma_t \qquad \varsigma \sim N(0,\alpha_2^2\sigma^2_\varepsilon)$$

Here the trend follows a unit root with drift or a stochastic trend. The effect of η_t is to allow the trend line to shift up and down, whereas ς_t enables its slope to change. Therefore, the larger the two hyperparameters α_1 and α_2, the more volatile the trend estimate. If the two hyperparameters are set to zero ($\alpha_1=\alpha_2=0$), then these equations will produce a trend series that increases by the fixed component μ every period. This will simply be a linear deterministic trend (a straight line) exhibiting maximum smoothness. However, as the hyperparameters are cranked up, the UC or trend estimate will form a closer fit to the measured data.

A specific form of the Kalman filter that is widely used is the Hodrick-Prescott[16] or HP filter. The HP filter is normally defined in terms of a smoothness parameter λ. The trended series becomes smoother as this parameter becomes larger, with famously λ = 1,600 being the filter proposed by Hodrick and Prescott (1997) for quarterly observations of the level of US GNP.

The HP filter is simply the stochastic trends model where $\alpha_1=0$. The relationship between the HP smoothing parameter λ and the hyperparameter α_2 is given by:

$$\alpha_2 = \sqrt{(1/\lambda)}$$

Therefore, if λ = 1,600 then γ = 0.025. More specifically, as the smoothness parameter in the HP filter increases, the hyperparameter in the stochastic trends model falls; in each case a smoother trend estimate results. It should be clear that by varying the hyperparameters then the estimated trend can be as smooth as we want it to be.

Figures 7 and 8 plot various Kalman trends where $\alpha_1^2=0$ and α_2^2 can take on various degrees of smoothing behaviour. These hyperparameters were derived so that these stochastic trend models behave in a similar way to the Henderson 7-, 9- and 11-term filters in Figures 5 and 6.[17] The gain function can also be calculated for Kalman filters in order to identify their cycle dampening characteristics.

Other Kalman filter models

The flexibility of the Kalman filter approach does not just lie in the choice of smoothing parameters, but also in the

Figure 7

Kalman filter central weights with hyperparameter 1.5657 (KF (1.5657)), 0.7559 (KF (0.7559)) and 0.4094 (KF (0.4094))

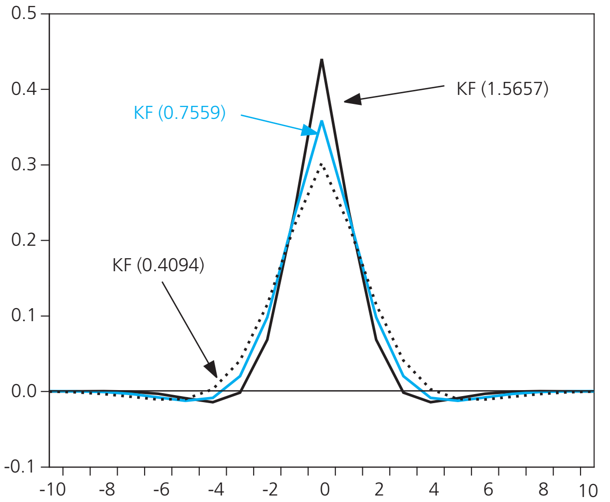

Figure 8

GDP percentage change quarter on quarter (GDP) and Kalman filter trend estimates with hyperparameter 1.5657 (KF 1.5657), 0.7559 (KF 0.7559) and 0.4094 (KF 0.4094)

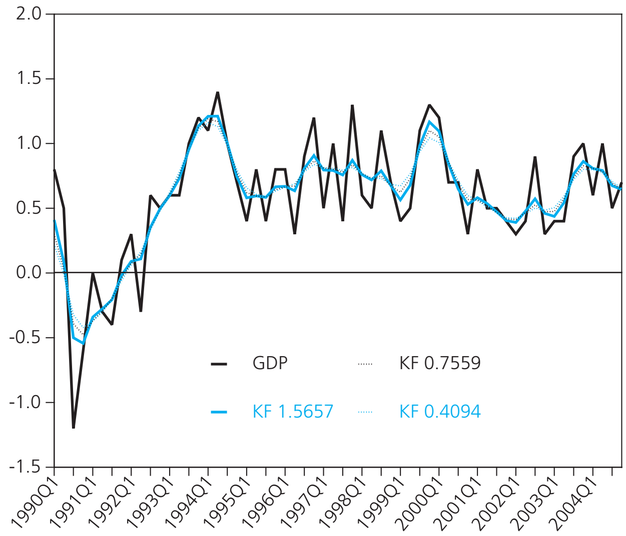

underlying models that constitute the measurement and state equations. By contrast, the Henderson filter is limited to just fitting trends that follow a cubic process. This filter might therefore over-fit a trend line if the data are best described by a linear or quadratic process, and likewise will under-fit higher order models. On the other hand, the Kalman filter state equations can be designed to fit a large number of different underlying processes.[18]

Kalman filters can also be modelled explicitly to achieve certain objectives. If the measured data have been affected by outliers, then dummy variables can be added to the measurement equation in order to limit their effect on the trend estimate. Also, other deterministic variables can be incorporated into the measurement and state equations in order to form richer models. In each case, the statistical significance of the model parameters can actually be tested by maximum likelihood techniques.

Conclusions

There are various ways in which users can fit a trend to time series data. This article has described two of the most common in the Henderson and Kalman filters. Both approaches allow the user a degree of flexibility. Although Henderson filters are widely used, Kalman filters arguably offer the practitioner a more comprehensive tool for fitting trends to time series data.

Notes

1. Currently trend estimates are used in some labour market statistics, UK trade, and the index of production.
2. See Trewin (2003).
3. For example, retail sales, index of production, consumer price indices and UK trade.
4. See Henderson (1916). Henderson filters are also commonly used to seasonally adjust data and are incorporated into the US Bureau Census X-11 and Statistics Canada X-11 ARIMA programs.
5. A simple filter is where each term is given the same weight (w_j=1/7 for j=–3,...,3).
6. In this case the GDP data were extrapolated using an ARIMA (0,1,1) process where $y_t=y_{t-1}+\varepsilon_t+\theta\varepsilon_{t-1}$.
7. See Doherty (2001) which formalises the methodology of Musgrave (1964).
8. See Kenny and Durbin (1982).
9. The Business Cycle Dating Committee at the National Bureau of Economic Research waits a relatively long time before dating peaks and troughs in the United States. This is because it wishes to avoid making premature judgements based on relatively immature data that are likely to be revised.
10. Chamberlin (2005) argues that trend revisions are relatively minor compared with revisions to the underlying data.
11. For a moving-average filter, the filter gain can be calculated using the formula

    $$G(\theta)=\left|\sum_{i=-m}^{m} w_j\cos(\theta j)\right|$$
12. This table is taken from Trewin (2003).
13. Such as the natural rate of unemployment or the level of permanent income.
14. See Kalman (1960).
15. For example, any ARIMA type model could be used here.
16. See Hodrick and Prescott (1997).
17. Koopman and Harvey (2003) explain how the central weights of the Kalman filter can be calculated. A numeric method can then be used to find α_2 so that the squared difference between the Kalman filter weights and that of the respective Henderson filter is minimised.
18. The stochastic trends model is generally a good representation of I(1) data such as log(GDP) but a local trends model might be more appropriate for I(0) data such as the GDP growth rate

    $Y_t=T_t+\varepsilon_t$ $\qquad \varepsilon_t\sim N(0,\sigma_\varepsilon)$

    $T_t=T_{t-1}+\eta_t$ $\qquad \eta_t\sim N(0,\alpha^2\sigma_\varepsilon)$

    For I(2) processes, like nominal GDP levels, you might like to consider an accelerationist model where the rate of change in the slope of the trend can vary so as to track accelerations or decelerations in the data. This form of Kalman filter is similar to the cubic processes mapped by the Henderson filter.

$$Y_t = T_t + \varepsilon_t \qquad \varepsilon_t \sim N(0, \sigma_\varepsilon)$$

$$T_t = T_{t-1} + \mu_{t-1} + \eta_t \qquad \eta_t \sim N(0, \alpha_1 \sigma_\varepsilon)$$

$$u_t = u_{t-1} + \phi_{t-1} + \varsigma_t \qquad \varsigma \sim N(0, \alpha_2 \sigma_\varepsilon)$$

$$\phi_t = \phi_{t-1} + v_t \qquad v \sim N(0, \alpha_3 \sigma_\varepsilon)$$

References

Chamberlin G (2005) *Fitting Trends to GDP*, ONS mimeo.

Doherty M (2001) *Surrogate Henderson filters in X-11, Australian and New Zealand Journal of Statistics*.

Henderson R (1916) *Note on graduation by adjusted average, American Society of Actuaries* No. 17, pp 43–48.

Hodrick R J and Prescott E C (1997) Post-war US Business Cycles: An Empirical Investigation, *Journal of Money Credit and Banking*, volume 29, pp 1–16.

Kalman R E (1960) *A new approach to linear filtering and prediction problems, Journal of Basic Engineering Series D*, volume 82, pp 35–45.

Kenny P B and Durbin J (1982) *Local trend estimation and seasonal adjustment of economic and social time series, Journal of the Royal Statistical Society*, Series A, volume 145, pp 1–41.

Koopman J S and Harvey A (2003) *Computing observation weights for signal extraction and filtering, Journal of Economic Dynamics and Control*, volume 27, pp 1317–33.

Musgrave J C (1964) *A theory of end weights to end all end weights*, US Bureau of the Census mimeo.

Trewin D (2003) *A Guide to Interpreting Times Series – Monitoring Trends*, Australian Bureau of Statistics Information Paper 1349.0.