WORKING PAPERS AND STUDIES

ISSN 1725-4825

2006 EDITION

Revision confidence limits for recent data on trend levels, trend growth rates and seasonally adjusted levels

Conference on seasonality, seasonal adjustment and their implications for short-term analysis and forecasting

Luxembourg, 10-12 May 2006

EUROPEAN COMMISSION

eurostat

THEME
General and regional statistics

Europe Direct is a service to help you find answers to your questions about the European Union

Freephone number (*):

00 800 6 7 8 9 10 11

(*) Certain mobile telephone operators do not allow access to 00-800 numbers or these calls may be billed.

A great deal of additional information on the European Union is available on the Internet.
It can be accessed through the Europa server (http://europa.eu).

Luxembourg: Office for Official Publications of the European Communities, 2006

ISBN 92-79-03408-1

ISSN 1725-4825

Catalogue number: KS-DT-06-009-EN-N

Revision confidence limits for recent data on trend levels, trend growth rates and seasonally adjusted levels

Peter B. Kenny (PBK Research)

Paper submitted for "Conference on Seasonality, Seasonal Adjustment and Their Implications for Short-Term Analysis and Forecasting", Luxembourg, May 2006

Abstract:

It is generally realised that recent seasonally adjusted or trend values are liable to be revised, even without changes to the unadjusted data, as further data points are added. However, there are no generally accepted indications of the likely scale of such revisions. This paper describes a method of using an ARIMA model of the unadjusted series, which is normally produced as the first stage of seasonal adjustment with either X-12-ARIMA or Seats, as the basis of calculating confidence limits for the revisions.

The method of calculation involves expressing the revision as a function of the future innovations of the ARIMA model. To a good approximation the relevant function is a linear combination; since the innovations are by definition independent, it is straightforward to calculate the appropriate revision variance. The adequacy of the linear approximation is confirmed in a sample case by a Monte Carlo simulation.

The method is illustrated by application to actual series, including to UK data on unemployment and prices. It is shown that the results can be useful, among other things, in judging the reliability of recent apparent turning points and in assessing the value of forecast extension in seasonal adjustment. It also gives interesting results in the comparison of adjustments using Seats and X-12.

Address for correspondence:
PBK Research, 53 Lindal Crescent, Enfield, Middlesex, EN2 7RP, United Kingdom
e-mail: kenny@pbkresearch.co.uk

Introduction

Seasonal adjustment practitioners are well aware that the most recent figures in the seasonally adjusted or trend series are liable to revision in the light of new unadjusted figures, even if the existing unadjusted figures are not revised. This is the well-known consequence of moving from the initial one-sided filters to the more symmetric filters which are usable as later data are available. However, this fact is seldom well advertised to users of the published data. Even practitioners are aware only in a qualitative sense of the revisions, and may have little or no quantitative idea of their extent. What is really needed, in this view, is an indication of the likely range of revisions, which should be presented, at least in graphical form, with the published data.

To be able to do this, it is necessary to have a measure of the standard error of revisions at each point in the series. In the case of SEATS, one of the outputs is a series of standard errors for the seasonally adjusted and trend series, so that there is no technical problem in presenting revision confidence intervals. X-12-ARIMA, however, has no such output. The main purpose of this paper is to present a method of calculation which will provide a standard error of revisions for X-12-ARIMA. It will be shown that this method provides a basis for standard errors for the seasonally adjusted series, the trend series and various measures of growth rates for both series.

SEATS is able to provide revision standard errors by reference to the properties of the underlying Arima model. The method proposed here for X-12-ARIMA is also based on an Arima model, and hence presupposes that such a model has been fitted. We also assume that the model is in some sense adequate for forecasting the series; in a sense, we ask what range of revisions is consistent with the uncertainty of the forecast future values. It would be rather risky to assume that the fitted Arima model is valid as a representation of the entire future of the series; the objective is therefore limited to calculating the extent of revisions over 36 months after first publication of the unadjusted data.

Method of calculation

Because of the multiple layers of averaging in X-12-ARIMA, it is not possible to provide a closed form algebraic expression for the revisions. Instead, we give a method of calculation or algorithm. Essentially, we use a variant of the well-known "black box" approach to analysing a complex filter, in which the response of the filter to an impulse input is calculated.

The standard forecast of the Arima model is based on the assumption that all future innovations are zero, and the initial seasonal adjustment is based on a decomposition of the published data extended by such a forecast. It will be assumed that the number of forecast values used is 36. We consider the effect of a nonzero innovation at each future date on the past trend and seasonally adjusted values. Considering each date in turn, we evaluate the effect of a given value for the innovation (normally some fraction of the standard error of innovations). Since the innovations are by definition statistically independent, we can easily evaluate their combined effect.

To put this in algebraic form, suppose that we have actual data from time $t = 1...T$, followed by 36 forecasts based on zero innovations. Consider the effect on the seasonally adjusted value at time t of assuming an innovation equal to the standard error of innovations at time $T + k$, all other innovations remaining zero; let this be $w_{t,k}$. Note that this effect is the combination of two components; firstly, the effect of the nonzero innovation on the forecast values, and then the effect of the modified forecast on the seasonal adjustment process. Now suppose that the actual realisation of the next 36 months of the series is defined by an innovation of ε_k times the standard error of innovations at time $T + k, k = 1..36$. Then the total revision after 36 months of the seasonally adjusted value at time t will be given by $\sum_{k=1}^{36} w_{t,k}\varepsilon_k$. Here the values $\{\varepsilon_k\}$ are independent with zero mean and unit variance by assumption, so that the variance of the total revision is given by $\sum_{k=1}^{36} w_{t,k}^2$.

This outline explanation ignores a number of details. In the case of an additive adjustment, calculations can proceed essentially as stated, but for a multiplicative adjustment we have to consider the effect on a logarithmic scale. This is normally the scale which is used for the Arima modelling in such case, so that this is not a major problem, but nevertheless we have to consider whether the effect in the untransformed scale can be calculated in the usual way in this case. To confirm the validity of this approach, the confidence limits obtained by the linear approach on a logarithmic scale are checked against a Monte Carlo simulation in which 1000 realisations of the future path of the series are generated using random normal deviates, and the extent of revision calculated for each. It will be shown that the 95% limits of the resulting revision distributions are in good agreement with those calculated from the formula above.

The method of calculation may be described in the following steps:

1. Carry out a full modelling and adjustment of the series, including identification and estimation of the Arima model and generation of forecasts for 36 months ahead. If prior adjustments for outliers, calendar effects, level shifts or other phenomena are necessary, they should be included in this analysis. From the outputs of this analysis, select the prior adjusted series in table B1 and the summary table of estimates of the Arima model.

2. Using the parameters of the Arima model from stage 1, generate forecasts for an artificial series consisting of a sequence of zero values followed by a single 1. The forecasts should extend at least 36 months ahead.

3. Carry out an X11 run on the B1 series extracted in the first stage. This run should have no Arima modelling or forecasting, but should include the 36 forecasts in table B1 as actual data. The sigma limits should be set wide enough to ensure that no extreme

values are identified. The seasonally adjusted series D11 and trend series D12 from this run will form the starting point for the comparisons in future stages.

4. Construct a set of 36 variants of the B1 series. For the first variant, modify the forecast part of the series (i.e. the last 36 values) using the forecasts generated in step 2 multiplied by the standard error of innovations, with the first nonzero value modifying the first forecast value in B1. For each later variant, lag the forecast series by one additional month. Carry out an X11 run on each of these 36 variants, similar to the run carried out in step 3, saving the D11 and D12 outputs.

5. Calculate the difference between each of the variants from step 4 and the starting point from step 3 for each time at which a revision margin is required. The sum of squares of these differences is the revision variance, and its square root is the revision standard error.

6. If revision margins are required for some function of the trend or seasonally adjusted value (e.g. some measure of growth rate), calculate that function for the output of the starting point run in step 3 and each of the variant runs in step 4 before calculating the differences in step 5.

These descriptions are deliberately vague in some respects, to cover the cases of additive and multiplicative adjustments. In step 4, for example, the modification of the B1 series is done by adding forecasts in the additive case, and by multiplying by the exponential of the forecasts in the multiplicative case. In step 5, the differences are calculated by subtraction in the additive case and by division (and then taking logarithms) in the multiplicative case.

One issue which needs further explanation is the decision, in steps 3 and 4 above, to run the X11 option with no extreme value modification. If the run is done with normal sigma limits, it will be found that the revisions are greater for some past observations which are close to the sigma limits. This is in a sense correct, since the addition of further data could well lead to changes in the amount of extreme modification. However, this is only valid if the revision limits are being calculated in a concurrent adjustment manner. In the probably more common situation, in which revision limits are updated annually, the wider limits may well be applied to values which are not subject to extreme modification. More seriously, limited experimentation suggests that the nonlinearity introduced by extreme modification invalidates the additivity assumption. If this is so, the only valid way of obtaining limits which are consistent with extreme modification is the Monte Carlo approach, which is much more computationally intensive.

Illustration and validation

The methods will first be illustrated on the well-known Box Jenkins airline passengers series. For this series, we present the results of various methods of calculation and choices of options, to demonstrate that the results are reasonably stable.

Figure 1: Box-Jenkins Airline Data
Original, SA and Trend

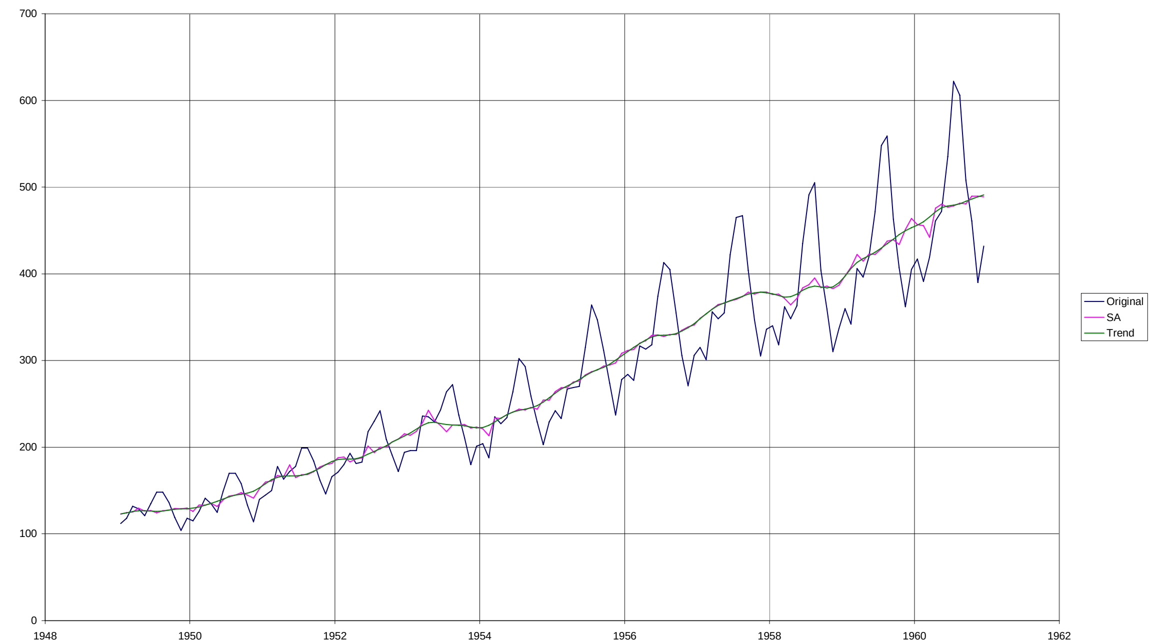

Figure 2: Box-Jenkins Airline Data
Revision Limits for X-11 SA

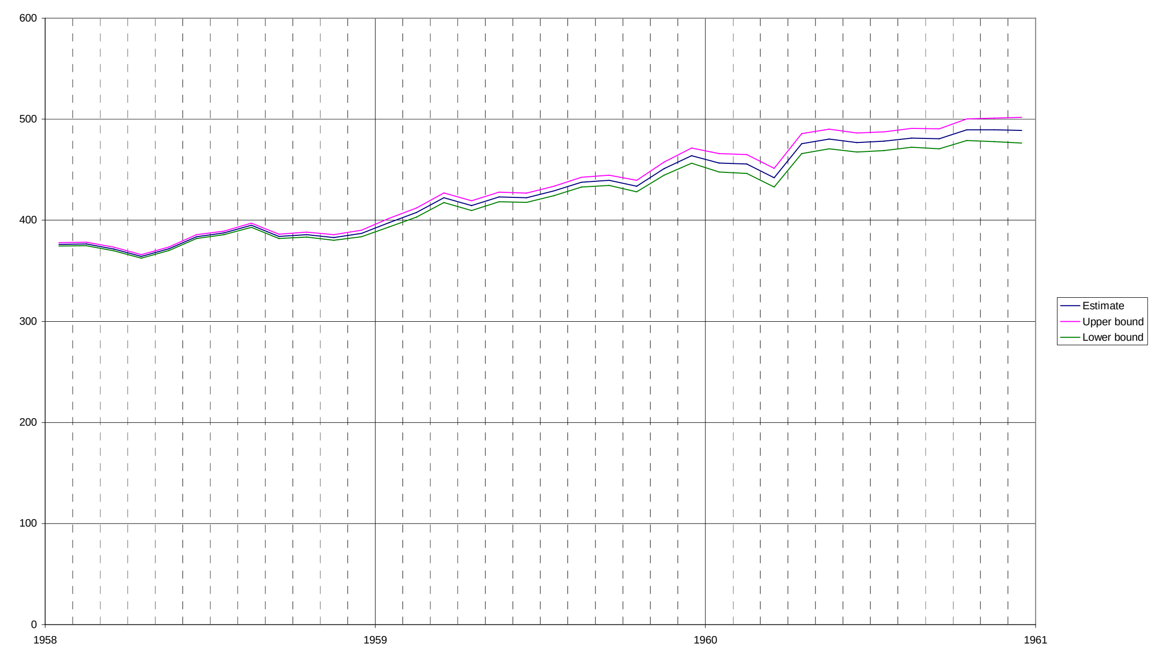

The first choice was to run with the default options of X11, using the automatic model choice. This conveniently gave the so-called airline model, but also gave a diagnostic indicating the presence of trading day variation. Inclusion of a trading day variable gave an intuitively obvious set of coefficients (a contrast between weekend and weekday numbers), but also gave a preference for a different Arima model (1 1 0)(0 1 1). The

Figure 3: Box-Jenkins Airline Data
Revision Limits for X-11 Trend

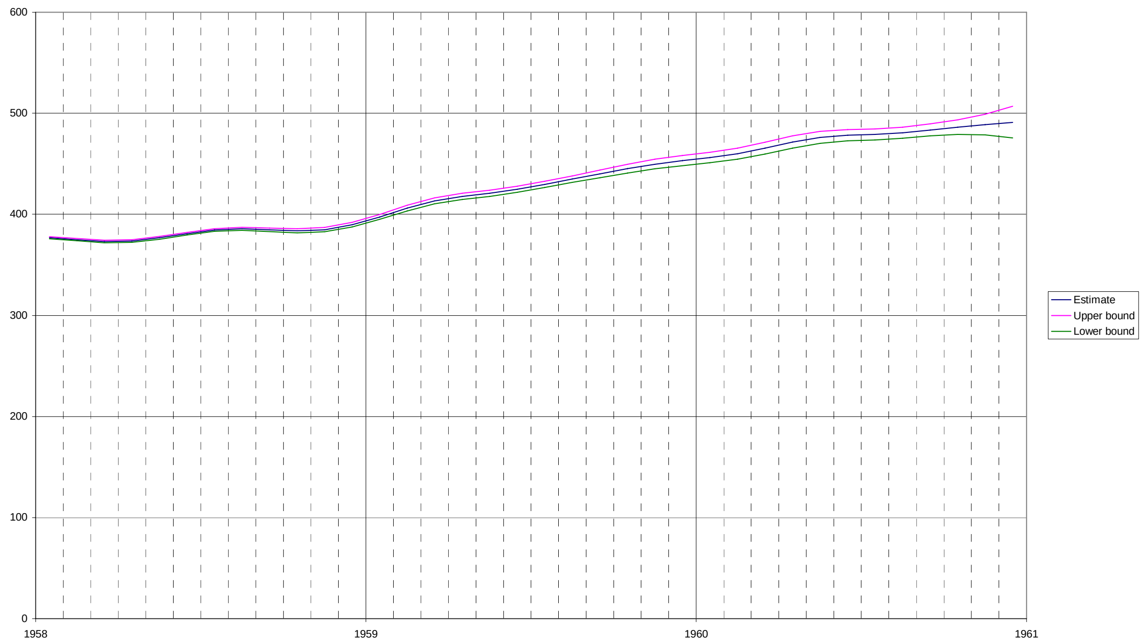

combination of this model and trading day regression was taken as the standard for the following runs.

Figure 1 shows the original data, seasonally adjusted series and trend series for this standard situation. Figures 2 and 3 show the 95% revision limits for seasonally adjusted and trend respectively from the last three years. It will be seen that both sets of revision limits show a kind of trumpet shape, widest at the most recent value. Closer observation shows that the trend trumpet is somewhat wider than the seasonally adjusted one for the latest observation, but for most earlier observations the seasonally adjusted trumpet is wider. This is a common situation for all the series which have been tried.

One question which needs to be asked is how stable are the revision limits against changes in the choice of Arima model. It is often found that the selection process does not give a clear-cut answer, since several models may seem almost equally valid. Changes in the choice of model may affect the most recent seasonally adjusted values, as is well-known, but we would like to be reassured that such changes will not have drastic effects on the calculation of revision limits. To check this, the X11 run was repeated with the standard airline model forced in. Figure 4 shows the revision limits for the seasonally adjusted series in this case; it can be seen that the limits are very similar to those in the standard case in Figure 2.

As mentioned earlier, it is possible to use the standard outputs of SEATS to provide revision limits. As a check on the validity of the method of calculating limits for X11, the same method was used to calculate limits for a SEATS adjustment; seasonally adjusted and trend limits are shown in Figures 5 and 6. Apart from the fact that SEATS

Figure 4: Box-Jenkins Airline Data
Revision Limits for X-11 Alternative SA

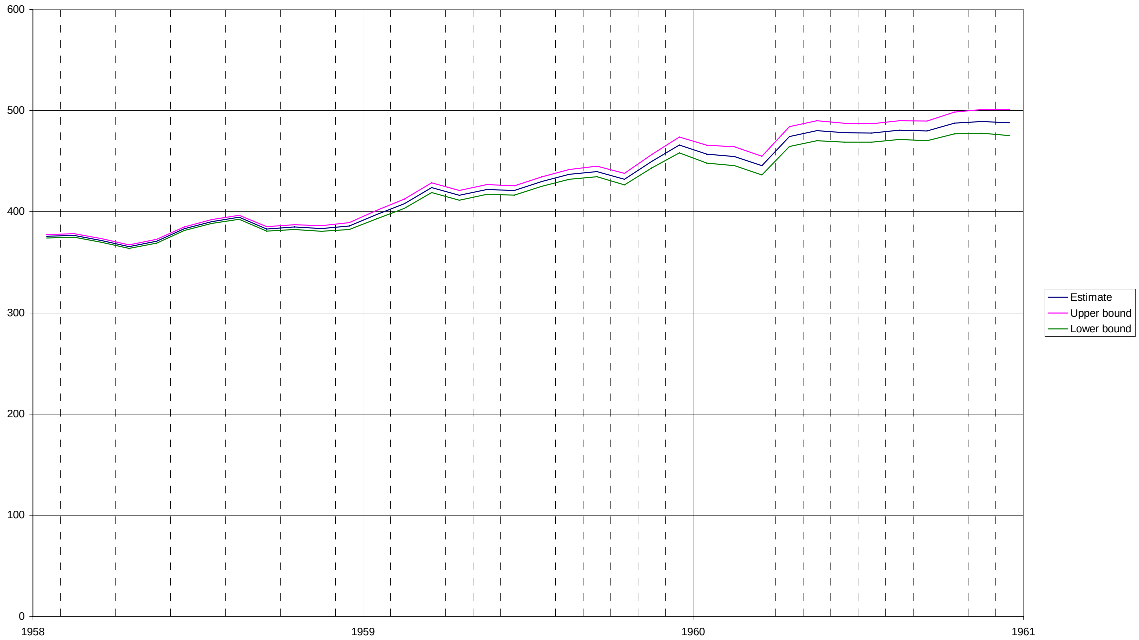

Figure 5: Box-Jenkins Airline Data
Revision Limits for Seats SA

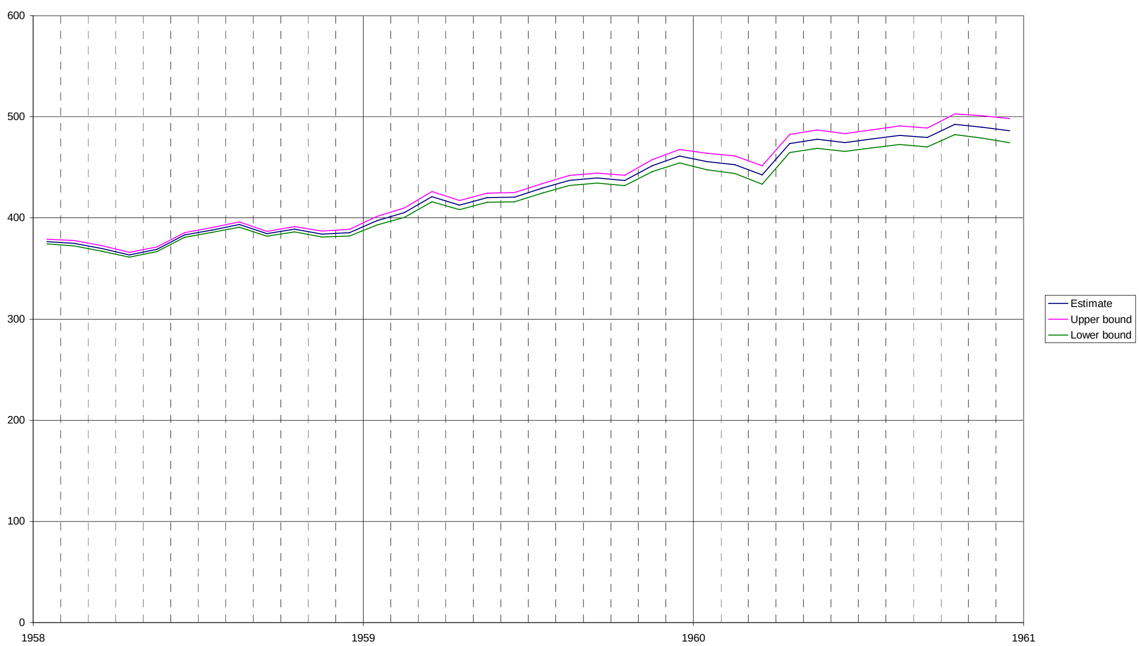

gives a slightly more flexible trend than X11, there is little difference in the graphs compared with Figures 2 and 3.

It is useful to be able to compare the limits from the method proposed here, both for X11 and for SEATS, with the limits obtained from the standard SEATS program output. If we show simply the difference between the estimate and the upper bound, it is easier to

Figure 6: Box-Jenkins Airline Data
Revision Limits for Seats Trend

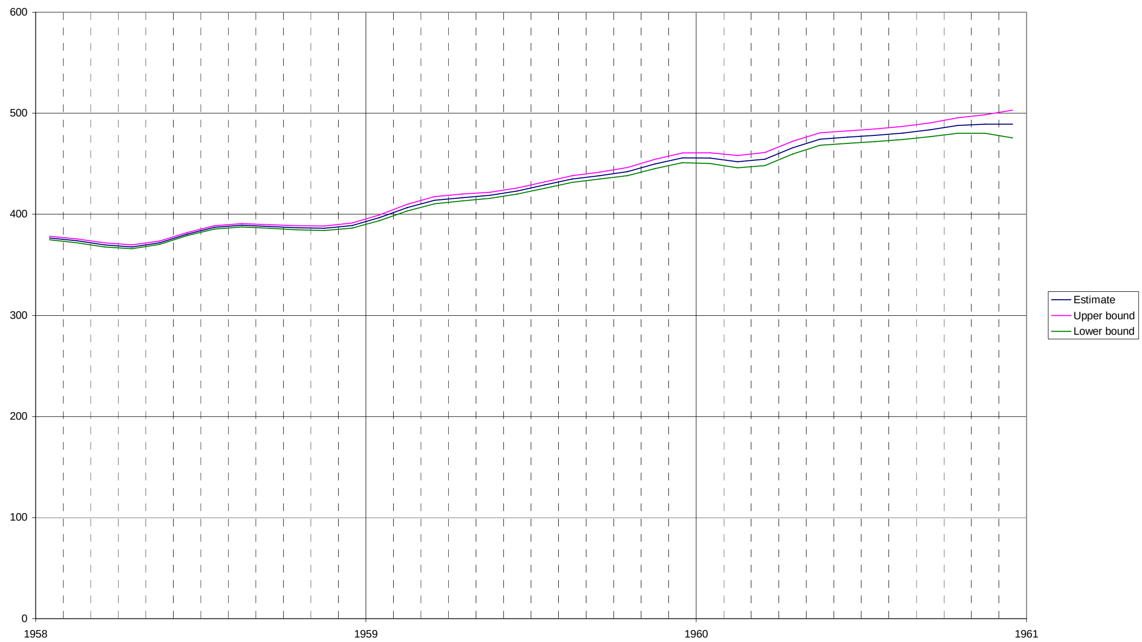

Figure 7: Box-Jenkins Airline Data
Comparison of Trend Revision Limits for Seats and X-11

see the differences between formulae. All three possibilities for the trend limits are shown in figure 7, in which the curve labelled "Maravall formula" represents the standard error from the SEATS program multiplied by 1.96. There are clearly differences between the formulae, but the general shape is very similar. One curiosity is that there are clearly systematic differences between the two curves based on SEATS. Over most of the span the curves are close together, but at around March of each year there is a kind of hump in

Figure 8: Box-Jenkins Airline Data
Comparison of Formula and Monte Carlo Trend Revision Limits

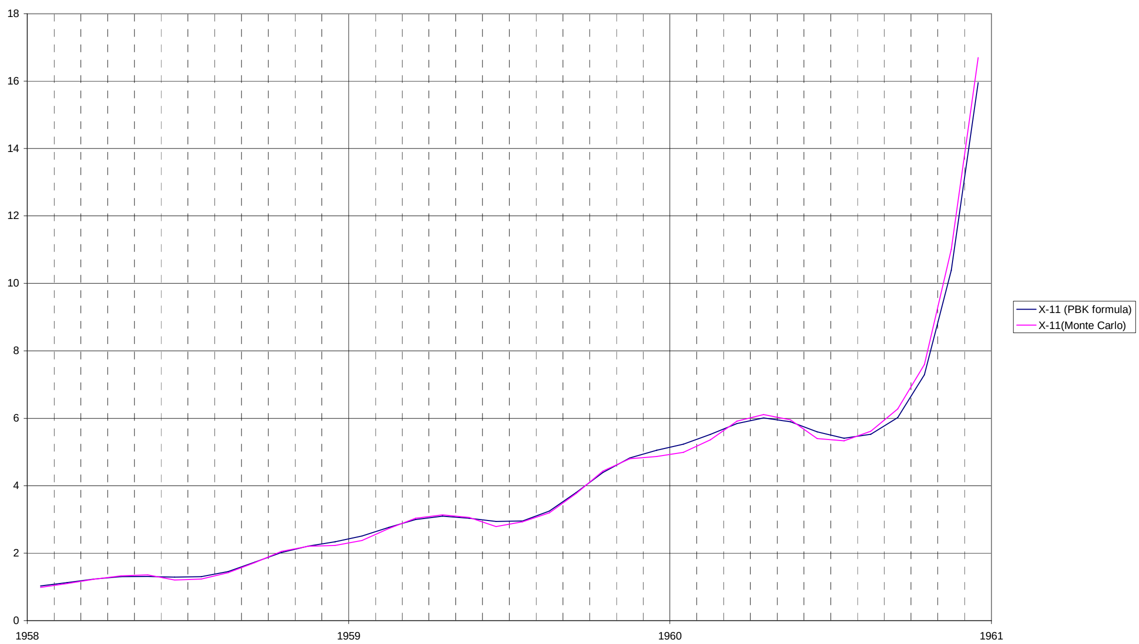

the results from the method of this paper. A similar hump appears in the X11 results, but mainly in April. At present these facts are unexplained. Nevertheless, the period of most interest is probably the last six months of the series, and over this period the curves are very close; the slightly wider limits for the X11 case are probably related to the different view of smoothness of trend between the two programs.

A further check on the validity of the limits is provided by the Monte Carlo calculation. This was carried out using X11 with the standard form of adjustment. Figure 8 shows the trend limits for the two cases, and clearly demonstrates that the results are close enough to be regarded as confirmation of the validity of the formula.

Further examples

Having shown on the rather antique data of the airline passenger series that the method proposed here gives sensible and stable results whether we adjust using X11 or SEATS, it is now interesting to see whether we can obtain useful results from a more recent set of data which are of some policy interest. The first series chosen is the UK Claimant Count unemployment series. This is no longer the preferred measure of unemployment in the UK, having been superseded by the harmonised count based on the Labour Force Survey, but it is nevertheless a useful economic indicator, being available quickly and accurately and apparently having a good cyclical relationship with the LFS series.

Figure 9 shows the seasonally adjusted series and the trend superimposed on the original data from 1994 to the latest available data (March 2006). The point of interest here is that the series showed a virtually monotonic decline up to the end of 2004, followed by a slight rise. At present there is an increase of about 10% from the low point. The question

Figure 9: UK Claimant Count Unemployment
Original, SA and Trend

Figure 10: UK Claimant Count Unemployment
Revision Limits for X-11 Trend

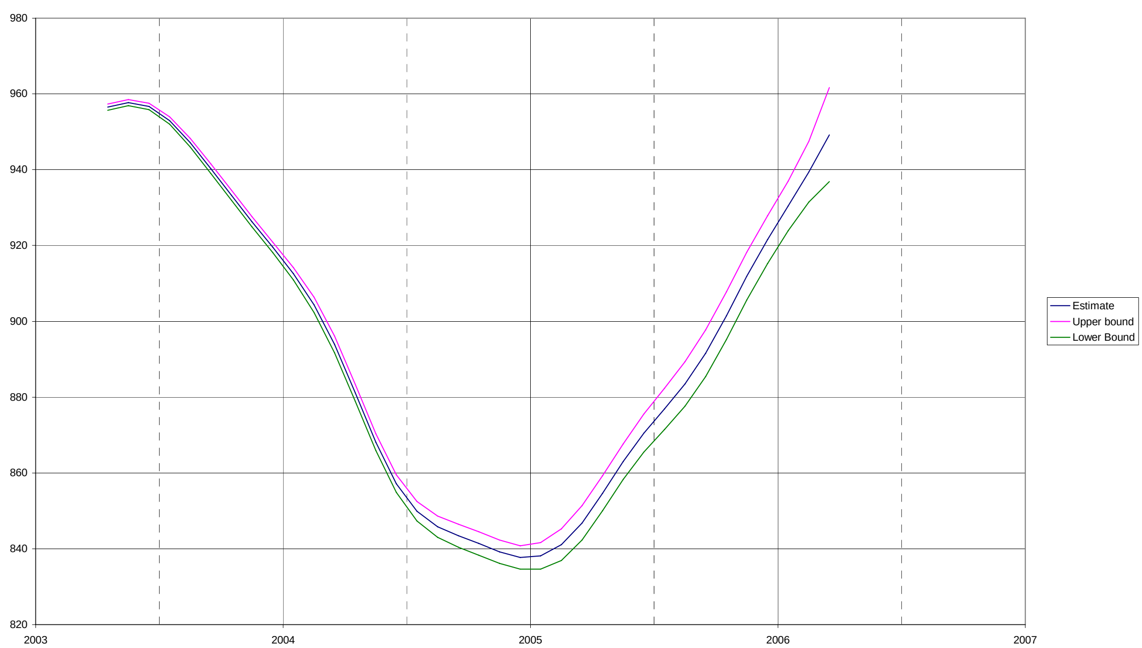

we ask is whether this rise is likely to be eliminated in the light of additional data; if not, we would like to know at what point it became clear that a turning point had been reached.

Figure 10 shows revision limits for the trend based on all data up to the present. It is absolutely clear that no change within the revision limits could possibly cancel out the

Figure 11: UK Claimant Count Unemployment
Revision Limits for X-11 Trend (data to March 2005)

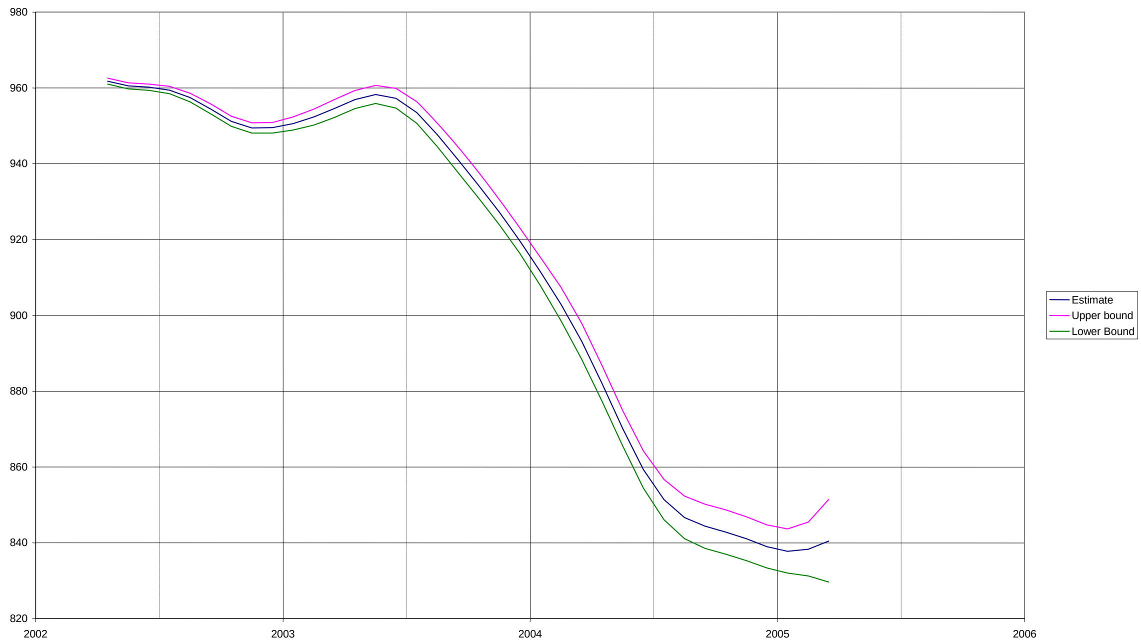

Figure 12: UK Claimant Count Unemployment
Revision Limits for X-11 Trend (data to April 2005)

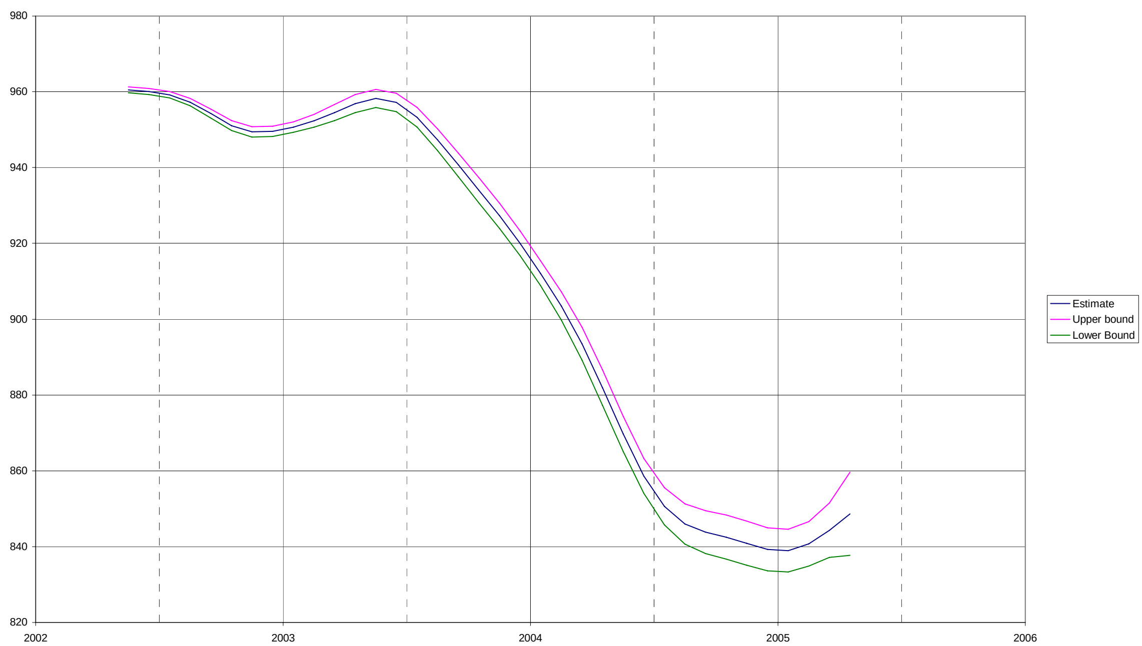

recent rise in trend level. Working backwards, figure 11 shows the picture we would have seen a year earlier, using data up to March 2005. At this point, the trend estimate shows an upturn, but the lower revision limit is still pointing downwards, so we cannot yet be sure that we have passed a turning point. Figure 12 shows the situation a month later, with data up to April 2005; clearly, we can now be fairly confident that the upturn in the trend is genuine.

This analysis provides a clear demonstration of the benefit of having revision limits in interpreting current data. Without the limits, it would be difficult to say at what point the upturn in trend could be taken as genuine; with the limits, especially in visual form, the task becomes a very easy one.

A final example considers the question of providing confidence limits for growth rates. The definition of trend growth rates raises a number of problems, which would require a more extensive discussion than can be given here. Some results are presented here without discussion of the background. The chosen series is the UK Consumer Price Index, whose growth rate is regarded as the main target variable of the Bank of England Monetary Policy Committee in its interest rate decisions. Figure 13 shows the index level, with seasonally adjusted and trend series. Note that no official seasonal adjustments of this series are published, although it is clear that there is small but stable seasonality present. The official inflation rate is obtained by comparing the current unadjusted index level with its value 12 months ago, which necessarily has a built-in lag of about six months. It is interesting to know whether a measure of trend growth, without this lag, could be provided from a seasonal adjustment of the index.

Figure 13: UK Consumer Price Index
Original, SA and Trend

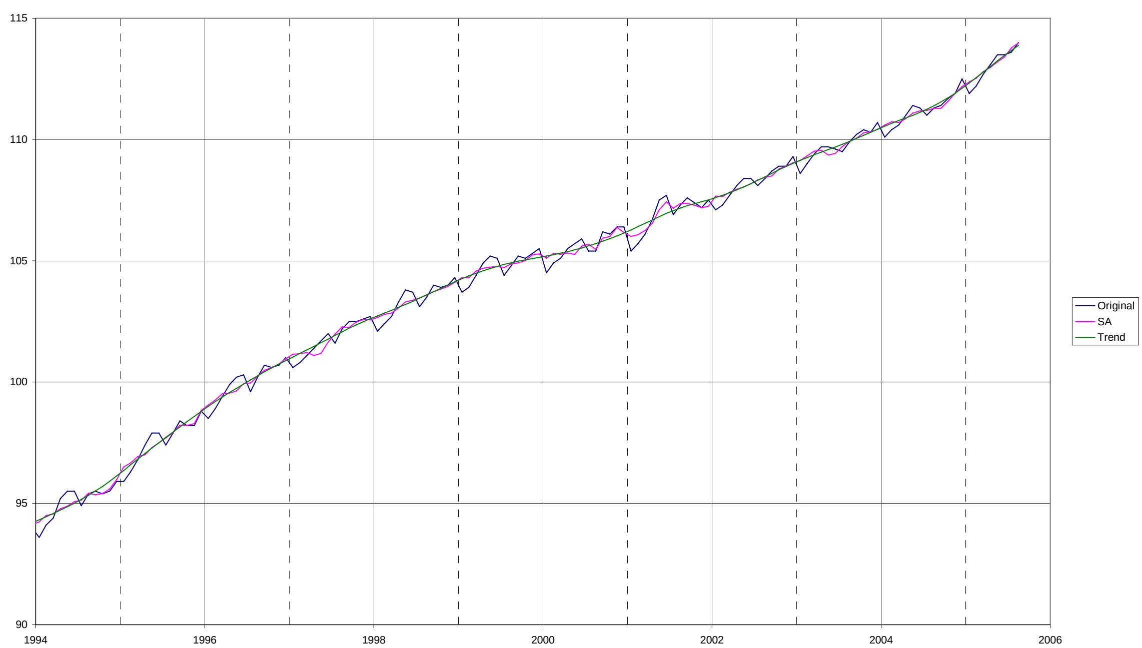

For reasons too complicated to explain here, although the automatic trend choice of X11 for this index series is the 9 term Henderson, the preferred choice for calculating growth rate is the 23 term Henderson. Figures 14 and 15 show revision limits for both these trends in the levels series, while figure 16 shows the revision limits for the annualised month to month growth rate of the 23 term Henderson. The curves for the levels series require little comment, other than the fact that the 23 term case gives limits which are wider at the end of the series but which taper more rapidly as we move backwards in

Figure 14: UK Consumer Price Index
Revision Limits for X-11 Trend (H9)

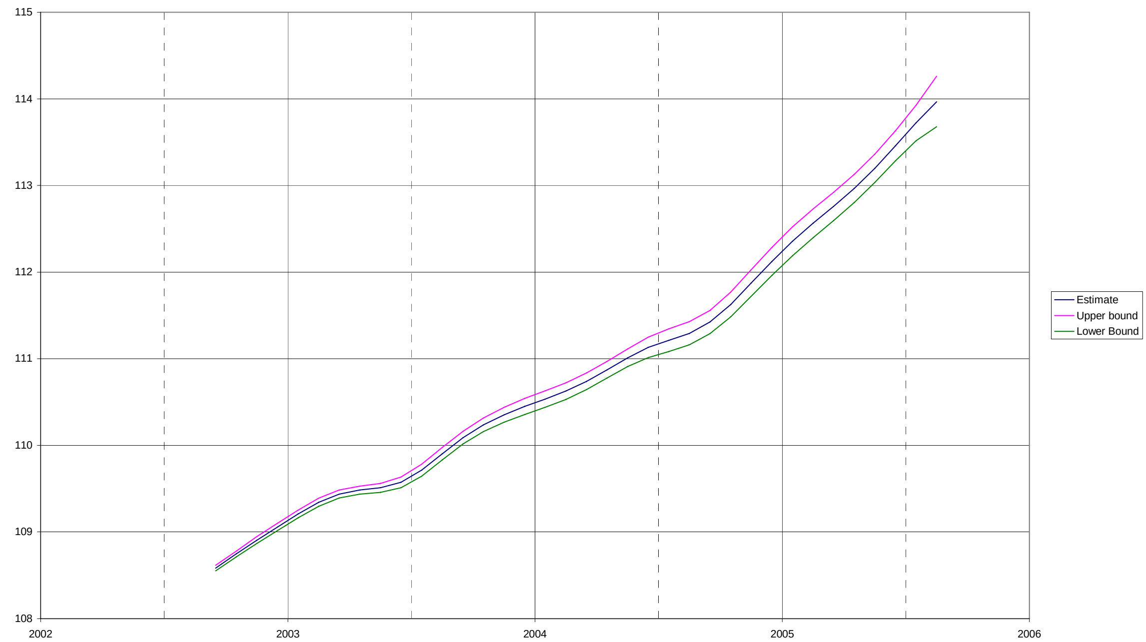

Figure 15: UK Consumer Price Index
Revision Limits for X-11 Trend (H23)

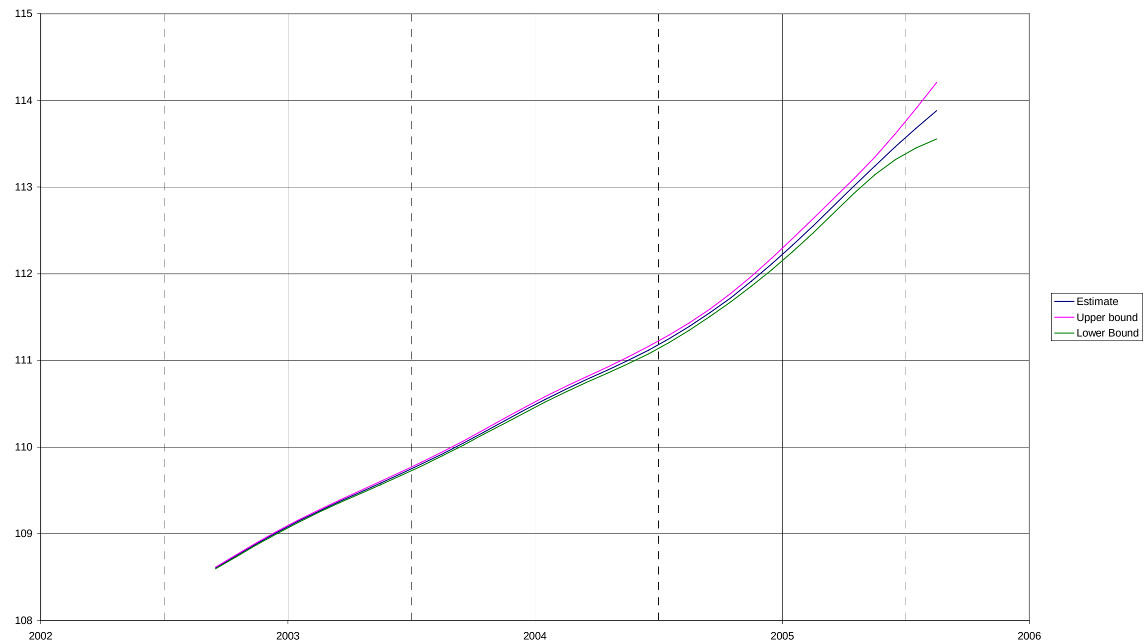

time. For the growth rate, however, we have a clear demonstration that, although the trend gives an up-to-date rate, the possible revisions are very large. This must raise questions about the usefulness of such a measure.

Figure 16: UK Consumer Price Index
Revision Limits for X-11 Trend Growth Rate (H23)

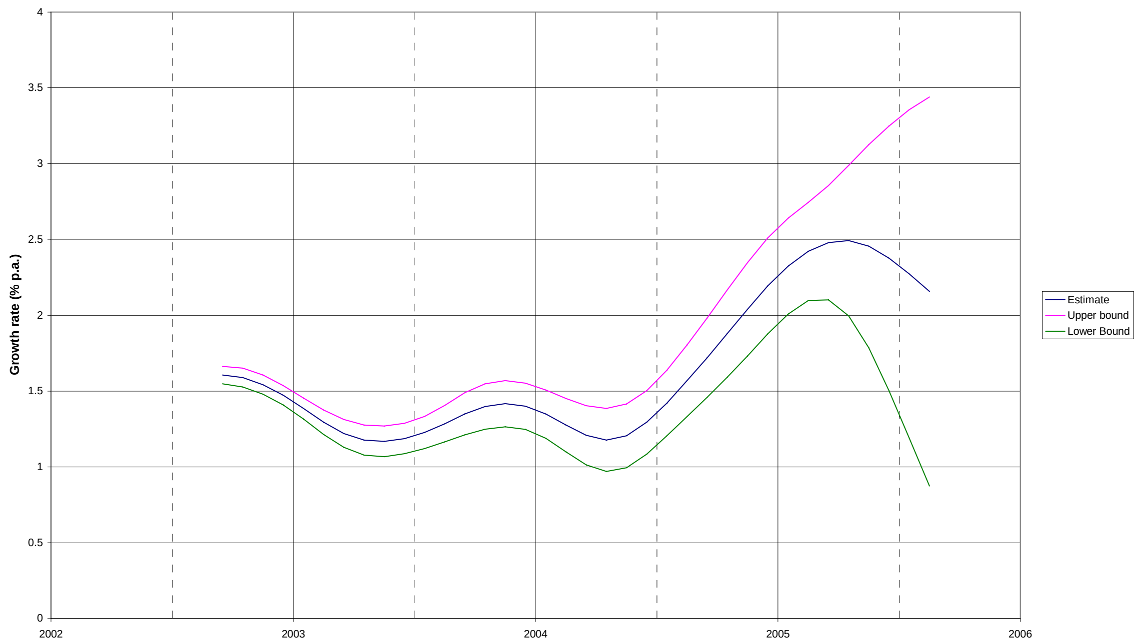

Summary and conclusions

The central argument of this paper is that users of seasonally adjusted data need to have an indication of how reliable the latest seasonally adjusted and trend values are. In the case of SEATS adjustment, the program already provides a way of calculating revision limits. The main result of the paper is a method of obtaining the revision limits for any adjustment method which is based on forecast extension using an Arima model. It is shown that the method is stable against plausible changes in the Arima model and gives similar results for X11 and SEATS adjustment. Illustrations of the use of the revision limits show that in appropriate circumstances the interpretation of current movements in the series can be made much easier.

Implementation of the method does not require heavy computation. For this study, a fairly simple Smalltalk program was written, which constructed appropriate specification files and data files, invoked X-12-Arima, collected the results and calculated the limits. The computation time for the 36 variants and the following calculation was about five seconds. If the method were incorporated as an extra option in X-12-Arima, it would involve about the same amount of computation as a history run involving 36 updates.

This research is to some extent work in progress. There should be further investigation of the question of using extreme value modification when calculating limits, and also an attempt to explain the minor but systematic differences between the limits obtained from this method applied to SEATS and the outputs of the SEATS program. Nevertheless, the results obtained are sufficiently encouraging to indicate that this is the basis of a useful addition to the outputs of any office publishing seasonally adjusted and trend values.