Seasonality and Trends in the Temperature Anomaly Data from Goddard Institute for Space Studies

Peter B Kenny
(PBK Research)

(For presentation at the Joint Statistical Meetings, Washington DC, August 2009)

Abstract

The GISS temperature anomaly data show the differences between monthly average temperatures and the average temperatures for the corresponding month in the reference period 1951-1980; hence in principle any stable seasonality should be eliminated. Using the analytical facilities of X-12-ARIMA, the published data for the entire globe, for the Northern and Southern hemispheres and for 8 latitude bands are analysed. It is found that there is identifiable seasonality in several of the series, particularly the Northern hemisphere and latitudes 44N to 64N. The trend chosen by the default X-12-ARIMA options is rather short (often the 13-term Henderson), and visually it seems implausible as a representation of the trend of the observed data. Using a longer Henderson (95 terms), a much more plausible trend is produced, but clearly showing an oscillation of about 8-10 years periodicity. Using a modified version of X-12-ARIMA, a much longer Henderson (301 terms) is applied; this removes the oscillation and gives a very convincing picture of the long-term trend over the entire history of the series (1880 to the present). There is some evidence that the oscillations in the 95-term trend may be correlated with sunspot numbers. The differences between the long-term trends for the different series are substantial, showing that global warming is far from uniform. The recent 95-term trend for some of the series shows a downward movement; using a method developed by the author, confidence limits for the latest trends are shown as a guide to whether this downward movement is likely to persist when the figures are revised in the light of later information.

Introduction

The Goddard Institute for Space Studies (GISS) is one of a number of organisations which publish data on temperature trends at various levels of detail and over a long period. This study looks at monthly data produced by GISS for what is known as the temperature anomaly, that is to say, the difference between current temperature and the average temperature for the corresponding month over some reference period. These figures are averages of all the available readings over specific geographical areas; the figures used here relate to the whole globe, the Northern and Southern Hemispheres and eight selected latitude bands. The reference period used for all the GISS data is 1951 -1980.

The initial question of interest was how to specify a trend value for some of the series, especially the whole Earth series, in order to compare with trends published by other authors. In some cases published trends seemed not to be following the course of the data very closely; in particular, some trends seemed to continue rising even when the most recent figures suggested a downward movement. For example, Figure 1 shows the crude monthly global series from GISS, while Figure 2 shows alternative trends suggested by Lockwood and Froehlich (2007) for the same data. A quick visual analysis suggests that the trends are ignoring certain local variations.

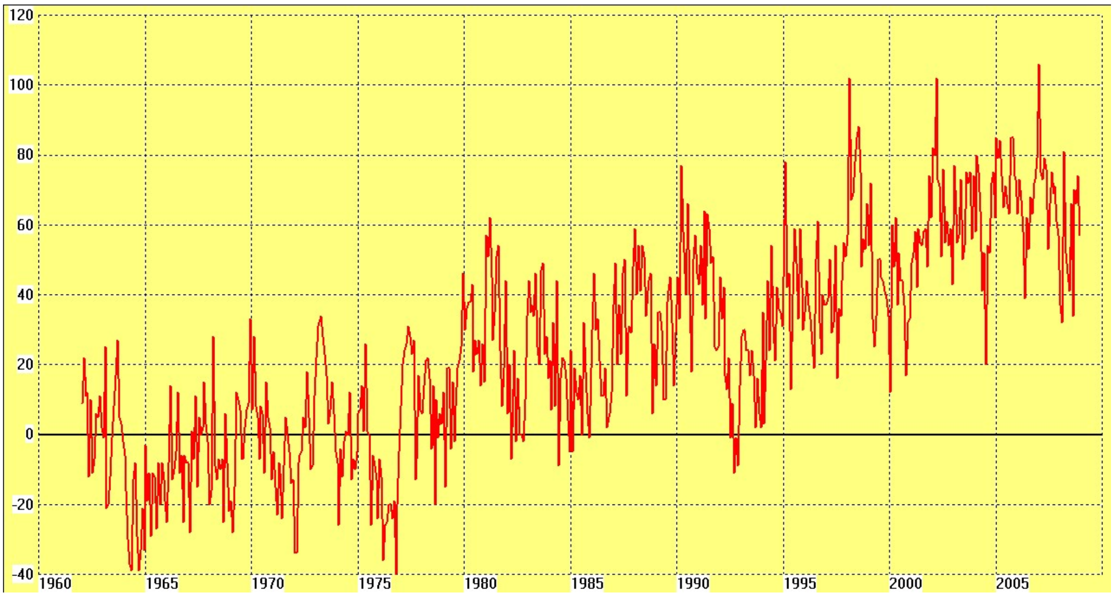

Figure 1: Whole earth raw data (vertical scale in hundredths of a degree Celsius)[1]

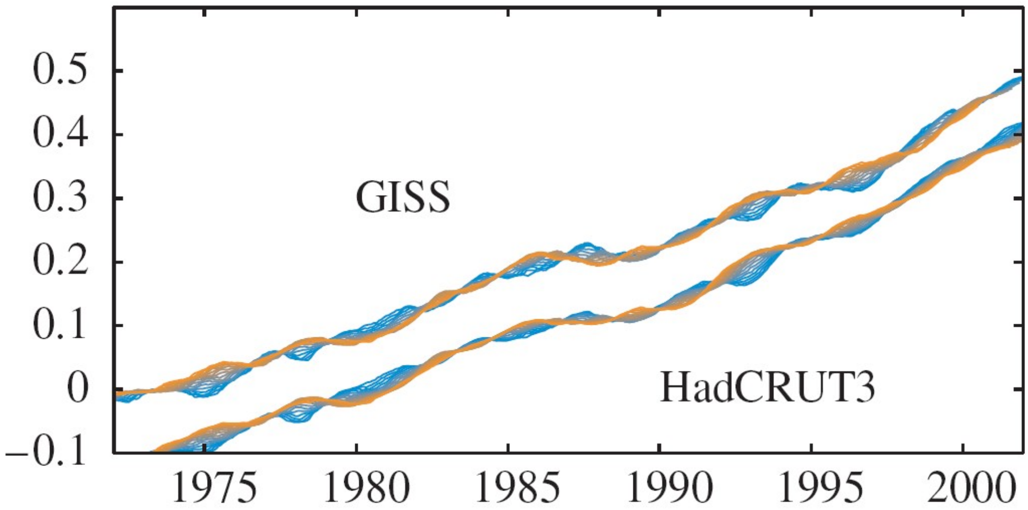

Figure 2: Lockwood and Froehlich trends (vertical scale in degrees Celsius)

Given a background in seasonal adjustment, it seemed natural to use the Henderson trend, as provided in the US Census Bureau X-12-ARIMA program, as an alternative definition of trend. Applying this program to a number of the GISS series with default options, the immediate and unexpected conclusion was that some of the series show marked seasonality. The trends selected automatically by the program tended to be very short, typically 13 or 23 months, and hence showed extensive local fluctuations. Experimentation with alternative lengths showed that a trend using a 95 term Henderson provided a good visual match to the underlying data, without introducing any obviously spurious oscillations.

---

[1] Note that the same vertical scale applies everywhere except in Figure 2.

This paper summarises the results of extensive comparisons of the GISS series, pointing out where seasonality is present and what effect it has on the interpretation of the trends in the temperature anomaly, and also showing comparisons which seem to justify the use of the 95 term Henderson to represent local trend. A much longer trend (301 terms) seems to represent the long-term trend well. As a final stage, the effect of revisions on the trend over recent years is considered; using a method described in Kenny (2006), 95% revision limits for the trend over the last three years are calculated, showing how much reliance can be placed on the apparent downward movement in recent years.

Seasonality in the data

There was no prior expectation that the GISS data would be seasonal; however, the default operation of X-12-ARIMA is to test for seasonality, and when one or two of the series had given positive results, the remainder were checked systematically. The method of construction of the anomaly series ensures that any stable seasonality in the reference period will be removed; therefore, it is to be expected that seasonal effects will show up in the more recent years, furthest from the reference period. Tests were done generally over the period 1996 to 2008, though other periods in recent years were sometimes tried. Over this period there were just two series which gave clear evidence of seasonality according to the usual X-12-ARIMA diagnostics; these were the series for the Northern Hemisphere and that for the latitude band from 44 N. to 64 N. Other series showed some evidence of seasonality, and the fitted seasonal patterns were similar to those identified in the series cited above, but the evidence is not conclusive.

In fact, there is some evidence that the nature of these series is such as to make it difficult for the standard X-12-ARIMA tests, especially the analysis of variance test in table D8A, to identify the stable seasonality. As a typical example, Figure 3 shows the seasonal components by month for the Northern Hemisphere series. In this plot, the red lines show extreme values which have been replaced. It can be seen that there is a large number of extremes in the winter months, which is the time of the seasonal peak, and very few in the months from May to October. Many series which give non-significant results in table D8 a have similar pattern to this, suggesting that a winter seasonal peak is obscured by the large number of extremes.

The series for the Southern Hemisphere is one which does not have identifiable seasonality according to the X-12-ARIMA diagnostics. If nevertheless we estimate seasonals, we obtain a pattern which is the inverse of that shown above for the Northern Hemisphere, that is to say high values in the months from June to September, with large numbers of extremes, and lower values in the other months. The seasonal peak in June and July reaches values somewhat lower than those for the Northern Hemisphere (around 20 units against around 30).

Consequently, it is not surprising that, when we analyse the series for the whole Earth, we obtain a seasonal pattern with two peaks. The whole Earth series is on the borderline for having identifiable seasonality, depending on the exact period over which the tests are carried out. The nearest to significance comes for the period 1990 to 2004, over which the test in table D8A gives "probably not present", but the stable seasonality test is significant that the 0.1% level.

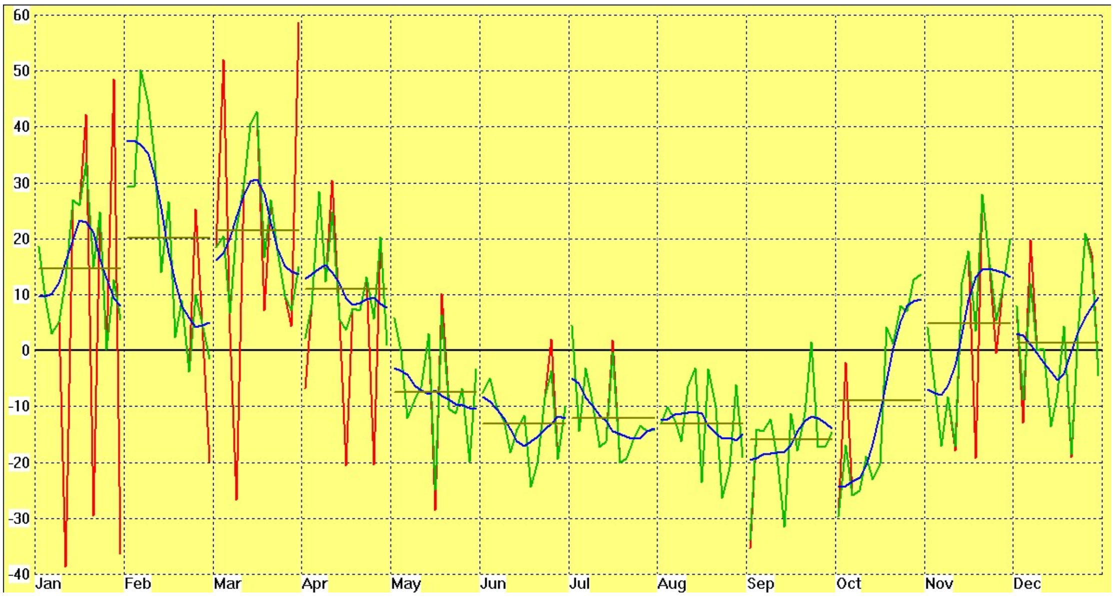

Figure 3: Northern Hemisphere, seasonal components, 1990-2008

(Red = S-I differences (D8). Green = Extreme-free differences (D9). Blue = seasonal factors (D10))

The most conspicuous seasonality is that shown by the series for the latitude band 44N to 64N. The seasonal effects can be large enough to completely offset the trend effects of warming. For example, consider Figure 4, which shows superimposed plots of the 95 term Henderson trend and the seasonal for this series analysed over the period 1990 to 2008; the same data are shown as the sum of trend and seasonal by month in Figure 4a. It will be seen that in 1995 the combined effect of seasonality and trend is to give warming of over 2° Celsius in winter and less than 0.5° in summer, which tells a substantially different story from that given just by the trend value of 0.7°. It is also clear from this plot that the amplitude of the seasonal has diminished sharply since about 2000, but it is still large enough to produce a substantial modulation of the trend.

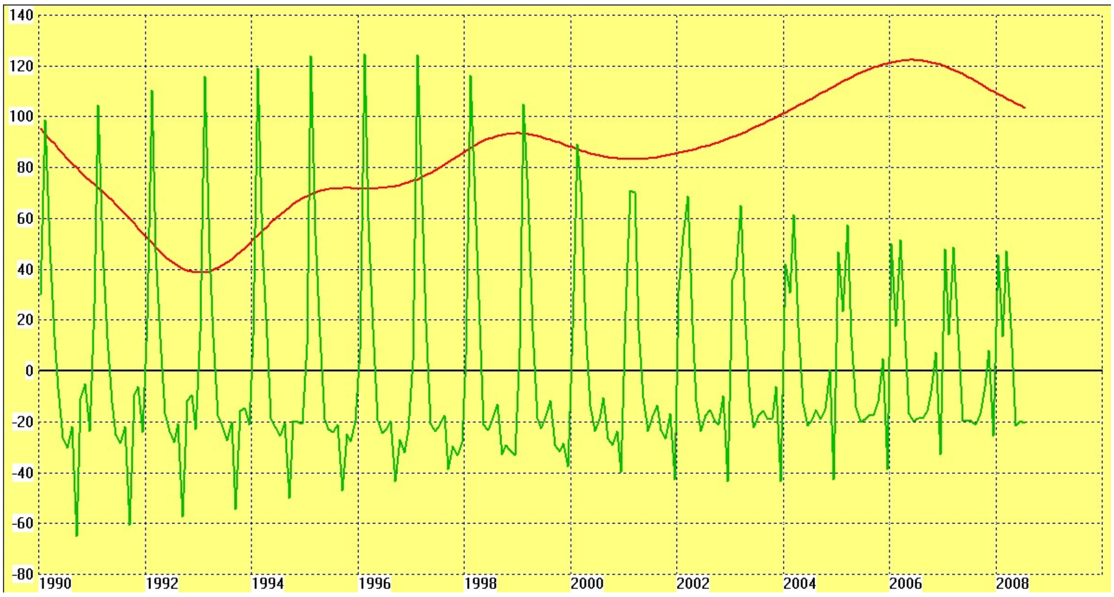

Figure 4: Latitude 44N-64N, final seasonal and 95-term trend

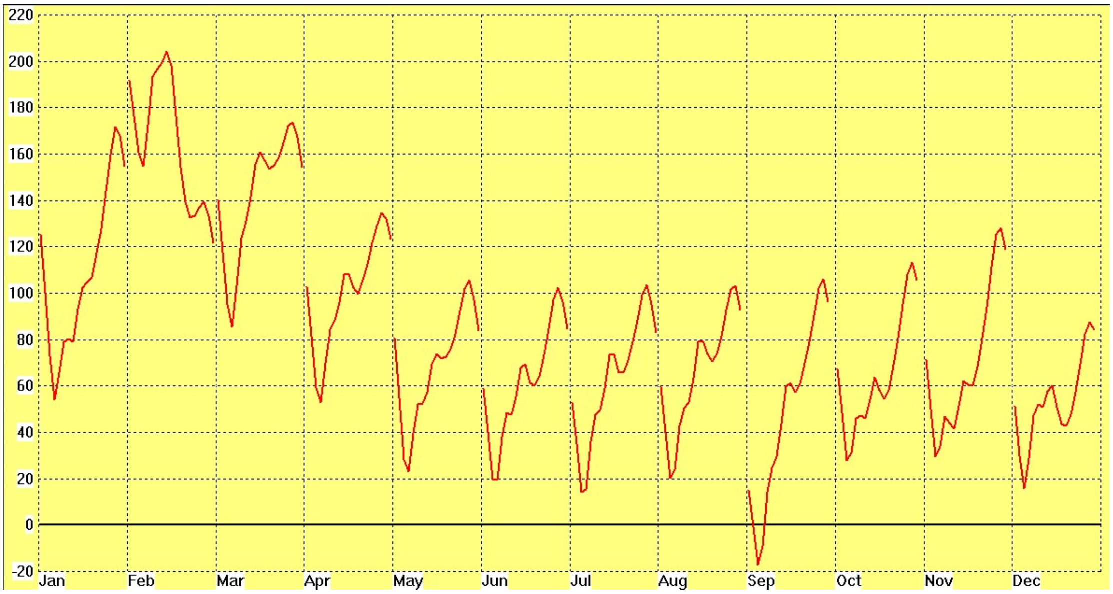

Figure 4a: Latitude 44N-64N, seasonal + 95-term trend by month, 1990-2008

Trends in warming

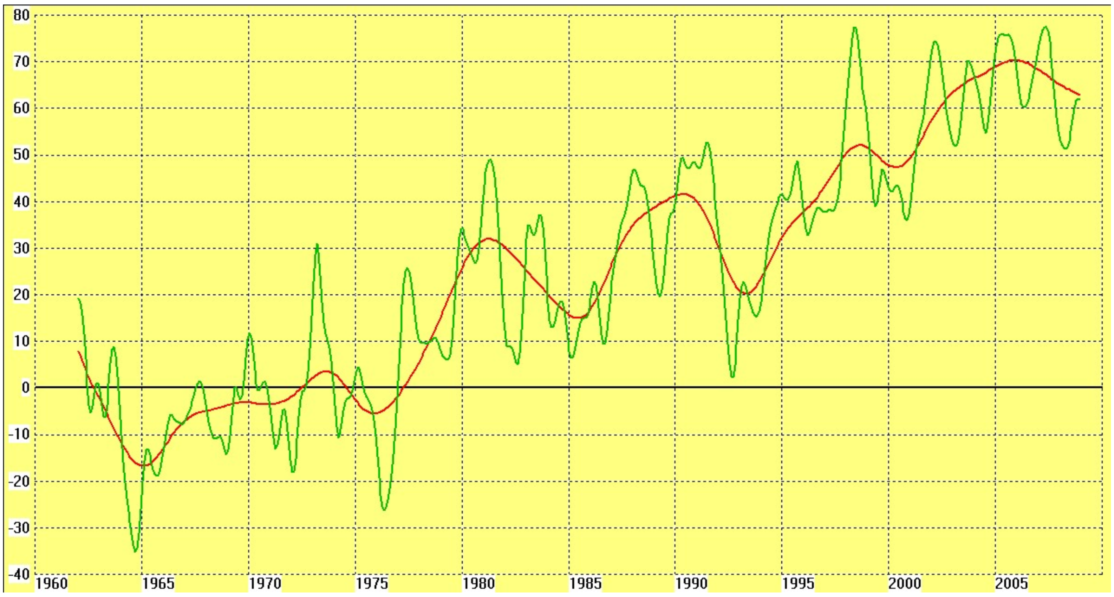

Figure 5: Whole earth, 13-term (default) trend and 95-term trend

Figure 5 above shows a comparison between the default 13 term Henderson for the whole Earth series and the 95 term Henderson. It is clear that for this series the 13 term does not provide a reasonable representation of the longer-term trend, since it has too many short-term oscillations. The 95 term curve seems much more plausible, since all its oscillations (such as the dips in 1985, 1993 and 2001) seem to match conspicuous excursions in the raw data (shown in Figure 1). There remains, however, the possibility that some of these might be artefacts of the trend filtering process. As a check on this, Figure 6 shows the 95 term trends for the Northern Hemisphere and Southern Hemisphere series separately. It is clear that major oscillations in the trends match between the two series. In other words, since the two series are quite independent, these oscillations must represent real effects applying to the whole Earth, rather than smoothing artefacts.

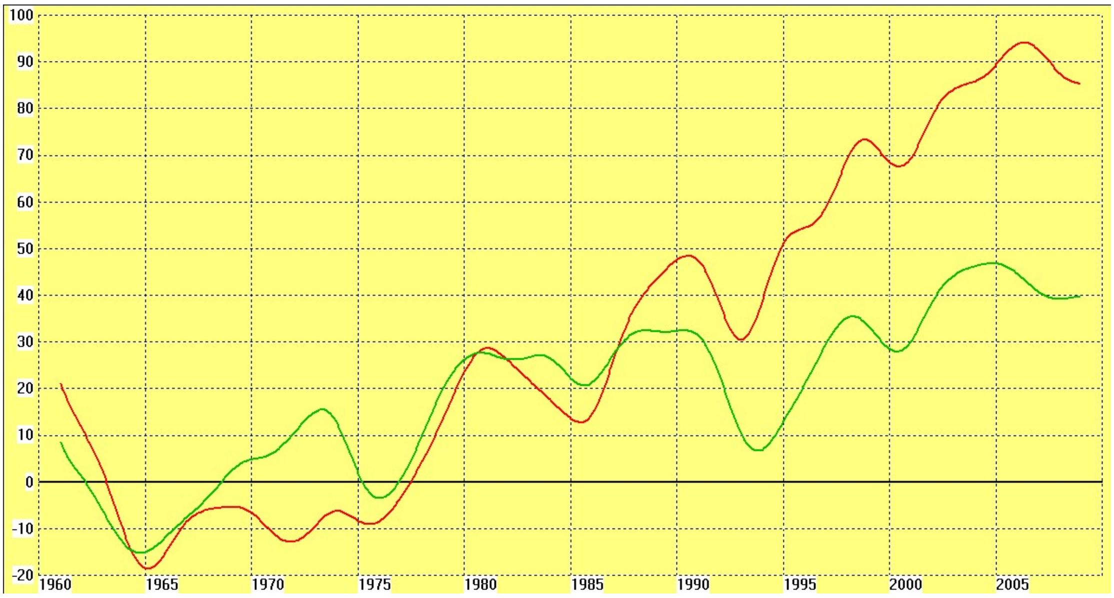

Figure 6: 95-term trends (Northern Hemisphere in red, Southern in green)

To a great extent, the choice of an appropriate trend filter must be a matter of judgment. It would in principle be possible to use a longer Henderson (though in fact standard X-12-ARIMA does not allow Hendersons longer than 101 terms, this is merely for convenience and can be overridden). Although this would smooth out oscillations which on this evidence represent real effects, they are clearly superimposed on a rising trend, and it is interesting to see that trend by itself. Using a modified version of X-12-ARIMA, kindly supplied by Brian Monsell of the US Census Bureau, a Henderson trend of 301 terms (about 25 years) was calculated. The results for the Northern Hemisphere, compared with the 95 term trend, are shown in Figure 7.

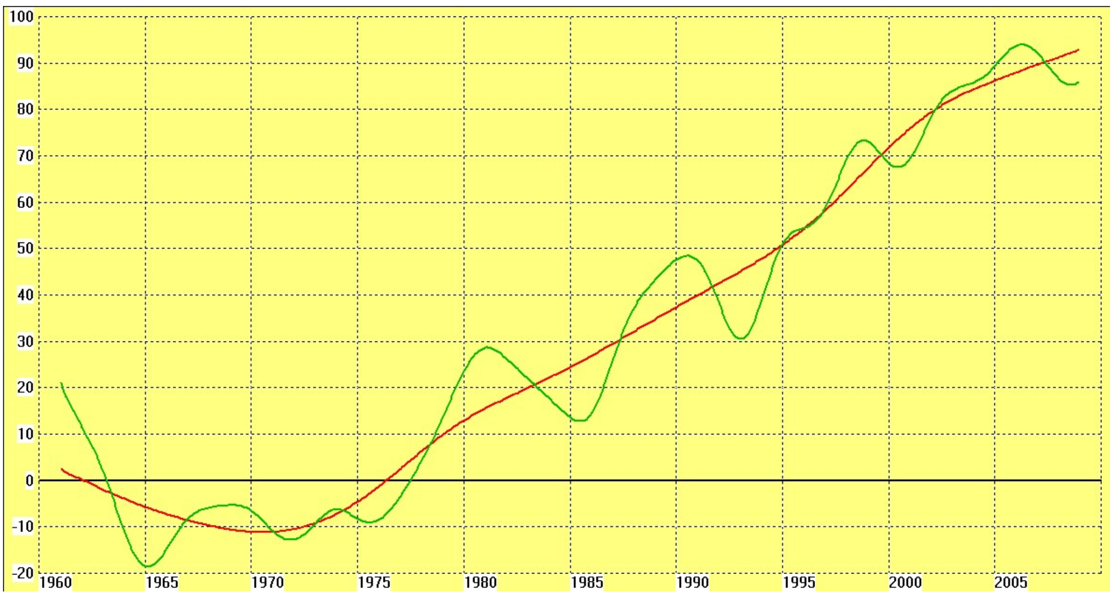

Figure 7: Northern Hemisphere, 95-term and 301-term trends

It is clear that the 301-term trend represents the long-term rising trend well, while the 95-term shows some form of medium-term oscillation about it.

Since the GISS data run from 1880 to the present, it would be interesting to have a view of the long-term trend over the whole of that period. Since the standard X-12-ARIMA can analyse only 600 observations (50 years), the 301-term trend was calculated over overlapping 50-year spans, and the combined trend produced by splicing together the middle 30 years of each span. The results for the Northern Hemisphere are shown in Figure 8.

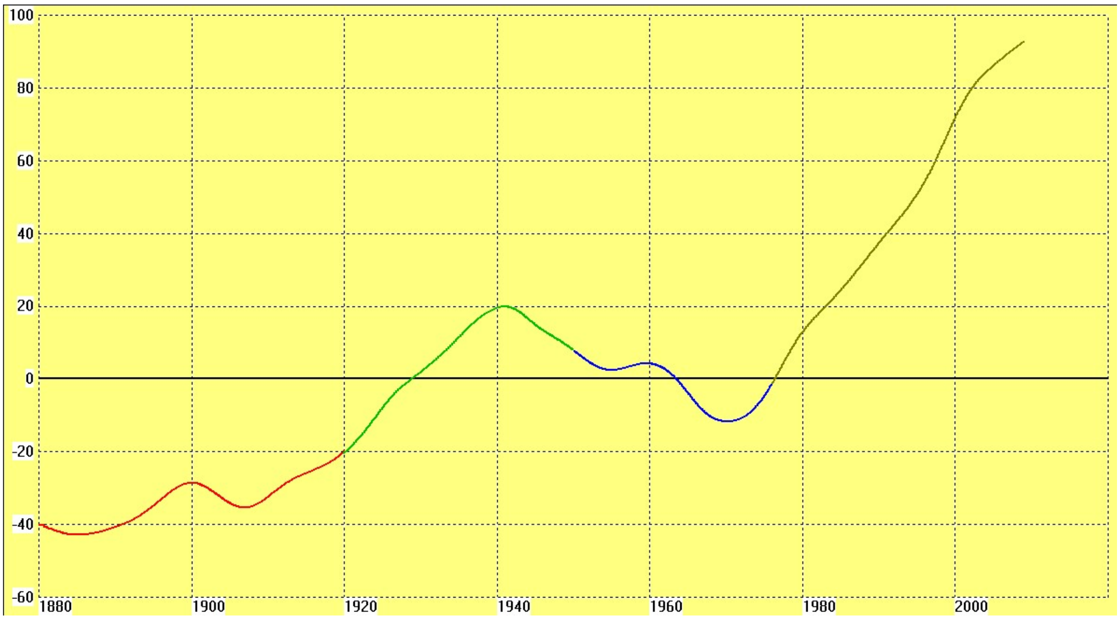

Figure 8: Northern Hemisphere, spliced 301-term trend, 1880-2008

The different colours in this graph show the separate spliced sections; clearly there is little or no discontinuity at the joins. This picture clearly shows the history of the series very well, while abstracting from the shorter-term movements. There is a steady rise up to 1940, then a sharp downturn – which is surprisingly sharp given that a 25-year trend can be expected to round off turning points – followed by an upturn in about 1970, then a steady rise up to the present. The rise in the Northern Hemisphere from 1970 to the present is just about 1° Celsius.

One question suggested by Figure 6 is what could explain the difference in overall growth between the two hemispheres -- it is clear that by 2005 the Northern Hemisphere temperature has grown by twice as much as the Southern. One possible "explanation", which simply transfers the question elsewhere, is to note that there are other data sets showing a marked difference in temperature growth between land and sea areas. A glance at the world map shows that the Northern Hemisphere has a substantially greater land area than the Southern, which has no land mass to match the area of Europe and Asia in the North. Given that the atmosphere circulates around the globe, it is not obvious why growth of air temperatures above land and sea should be so different, but it is clear that the two effects -- land versus sea and North versus South -- have the same explanation.

What could explain the oscillations in the 95-term trend?

The oscillations in the 95-term trend can be represented as the deviations of that trend from the 301-term trend. These are shown for the Northern and Southern Hemispheres in Figure 9.

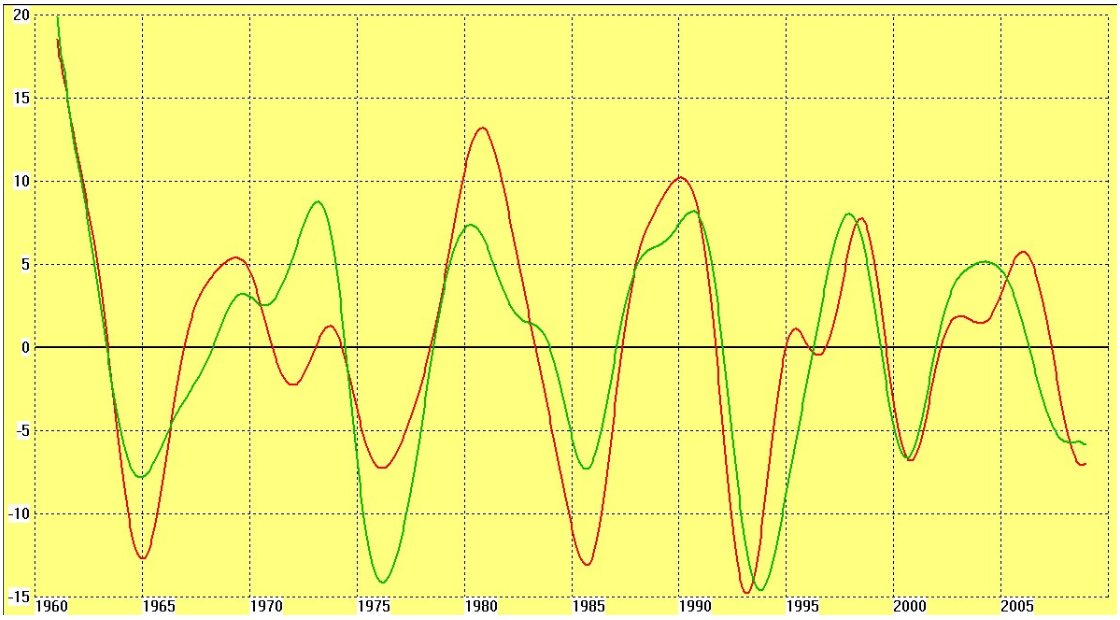

Figure 9: 95-term – 301 term trends (Northern Hemisphere in red, Southern in green)

Since these are derived from independent data from non-overlapping areas, we suspect that there must be some external effect with global range which influences them. One candidate which comes to mind is sunspot numbers. The smoothed sunspot numbers from the Solar Influences Data Centre in Brussels are plotted over the same range in Figure 10.

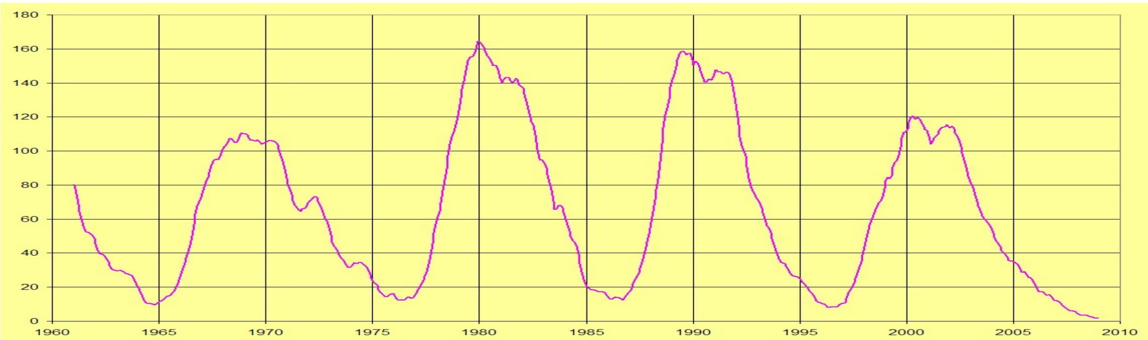

Figure 10: Smoothed sunspot numbers, 1961-2008 (from SIDC Brussels)

The agreement from 1961 up to the peak in 1990 is quite good, with turning points matching closely. After that things break down somewhat, so that in 2001 there is a trough in the temperature trends and a peak in the sunspot numbers. At the end of the series both are declining, so the correlation might appear better in the light of later data. At present, however, all we can say is that there is some appearance of correlation with sunspot numbers – though of course correlation does not necessarily imply causation.

Reliability of recent trend estimates

Figures 5 and 6 show that the 95-term trends for the two hemispheres and for the whole Earth have turned downwards since about 2005. As is well-known, the trends at the end of a series are subject to revision as more data are added; this must be particularly true with a trend filter as long as 95 terms. It is reasonable to ask, therefore, how certain we can be that these downturns will survive the revision process as more data are ended. In Kenny (2006), a method is described which makes it possible to calculate 95% revision limits for recent trends. In essence, the method assumes that the underlying unadjusted series may be projected using the Arima model fitted by X-12-ARIMA. The effects on the recent trend values of alternative realisations of the future series are calculated. As an illustration, this method has been applied to the 95 term trend of the Northern Hemisphere series, as shown in Figure 6, giving revision limits for the last three years of the trend. These results are shown in Figure 11.

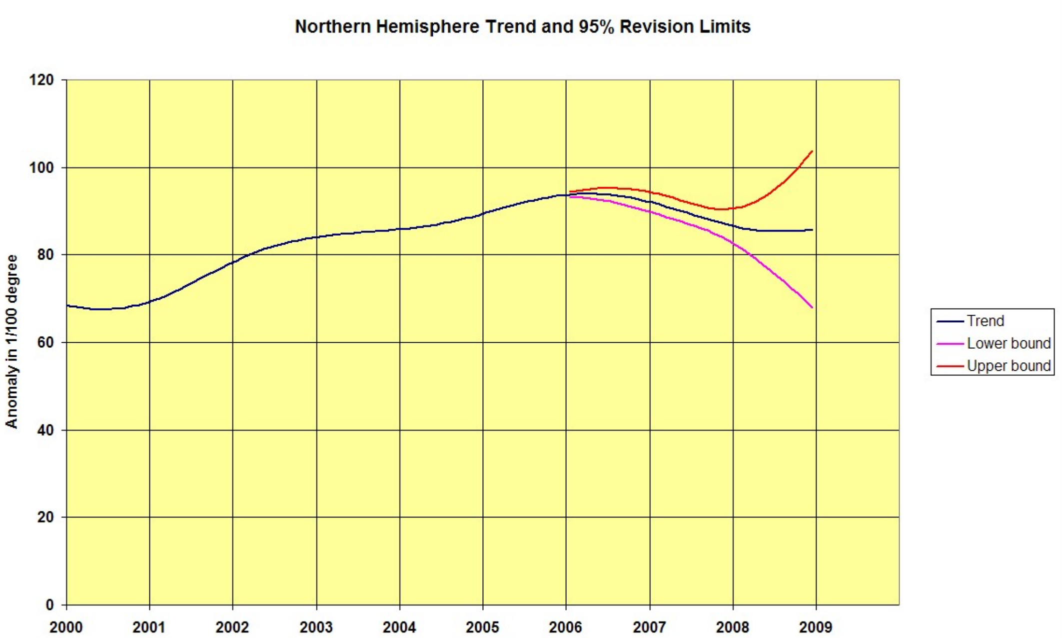

Figure 11

The main point of interest here is the peak in early 2006 and the subsequent decline. As can be seen clearly, the uncertainty is much greater for the last year of the trend. Prior to that, we can be fairly confident that there is a downturn which will not be removed in the course of revisions, but we cannot be sure about the last year. It is also clear that there is no need to calculate revision limits earlier than three years from the end of the series.

Conclusions

We can consider conclusions in two categories: whether X-12-ARIMA used in this way provides a useful tool which gives insights into the behaviour of the GISS figures, and what information is revealed about the data.

For the first category, the answer is that X-12-ARIMA has indeed revealed information which has not been widely noticed, in particular the extent of seasonal differences in global warming. The Henderson trend seems to provide a useful alternative to other choices, and its well-known optimality criteria might suggest that it is a preferable alternative. The revision limits for trends provide a useful qualification to any discussion about the current situation.

For the second category, the following points are of interest:

- Some of the series (especially 44N-64N) show strong seasonality, which means that trends differ considerably from month to month.
- The 95-term trend gives a good guide to local trends, and the 301-term for long-term trends.
- There are substantial differences in the trends in different geographical areas (e.g. between Northern and Southern Hemispheres).
- There is some indication that short-term variations world-wide may be correlated with sunspot numbers.
- The long-term rise is clear, but the fall from 1940 to 1970 is puzzling.

References

Kenny, P. B. (2006): “Measuring trend growth rates, with an application to consumer prices”. Proceedings of the 2006 Joint Statistical Meetings, Seattle.

Lockwood, M. and Froehlich, C. (2007): “Recent oppositely directed trends in solar climate forcings and the global mean surface air temperature”. Proceedings of the Royal Society, Series A.

P B Kenny

8 May 2009

Revised 29 July 2008