A good trick to create a decline

During the review process for the IPCC's Fourth Assessment Report, Steve McIntyre, the retired mining consultant who is the scourge of the mainstream climatology community, suggested to the authors of the report's paleoclimate chapter that some of the key global temperature reconstructions were affected by a significant flaw, just one of the many criticisms he made of the work of the chapter authors.

Most of these temperature reconstructions are based upon tree ring measurements, and increasing ring widths in the twentieth century were being read as evidence of higher temperatures. This tallied nicely with measurements taken using thermometers, apparently giving these scientists and the public a great deal of confidence that the tree ring approach was valid and that the reconstructed temperatures in earlier centuries could therefore be relied upon. And if temperatures in the early sections of the reconstructions – the controversial Medieval Warm Period – were seen to be lower than those in the modern era then an important message was relayed to the public about how mankind was affecting the planet: current temperatures were unprecedented.

McIntyre's objection to this hypothesis was that the twentieth century growth spurt seen in one particular species of tree – the bristlecone pines – was not all it seemed. The bristlecones are nearly ubiquitous in temperature reconstructions and, in most, the hockey stick shape seen in the growth patterns over the centuries is critical to support a claim that twentieth century temperatures are unprecedented. McIntyre complained, however, that it was well established in the scientific literature that the uptick in bristlecone growth in the twentieth century was caused, not by higher temperatures, but some non-climatic effect. Among the possibilities was the presence of higher levels of carbon dioxide in the atmosphere, the so-called *carbon dioxide fertilization* effect.

A longer hockey stick

The IPCC author on this part of the report was Keith Briffa, a paleoclimatologist working at the Climatic Research Unit at the University of East Anglia. Briffa was unimpressed with McIntyre's arguments, saying in his response that carbon dioxide fertilization had been adjusted for in one of the key reconstructions used in the report.[1] The paper he was referring to was Mann *et al.* 1999 (MBH99), a follow-up to the famous Hockey Stick paper.

While the first version of the Hockey Stick reconstructed temperatures back to 1400 AD, MBH99 took the reconstruction right back to the start of the millennium, apparently

[1] Of course, the text of the report referred to a number of temperature reconstructions, all of which were said to support Mann's conclusions. So when Briffa referred to the correction in Mann's paper, he was neatly avoiding the knotty question of how carbon dioxide fertilization had been addressed in these other papers.

demonstrating that temperatures had declined slowly but steadily for almost nine hundred years. This decline, said Mann, was consistent with the Milankovitch cycle, the tiny changes in the Earth's orbit that are thought to cause the ice ages to come and go.

Then, at the start of the twentieth century, there had been a sudden uptick in temperatures coinciding with the onset of mass industrialisation. Since the long decline and the twentieth century uptick were apparent even after carbon dioxide fertilization had been adjusted for, it seemed that the findings of the original Hockey Stick paper and, by implication, the other proxy reconstructions, were confirmed and strengthened.

Researching the adjustment

After the appearance of the Fourth Assessment Report, McIntyre had explained this strange adjustment to his readers. In essence it was rather simple, although in practice the procedure adopted was rather complicated. Temperature reconstructions tend to use two different sets of tree ring records: those from high-elevation sites and those from sites at the northern limit of the geographical range of the species of tree concerned. Scientists believe that it is only these two groups where the tree rings change in line with temperature. However, they also think that it is only the high-elevation sites that are affected by carbon dioxide fertilization. This had given Mann a way to assess the magnitude of the effect and so to create an adjustment.

Mann had compared the high-elevation sites, including the bristlecones, to a northern treeline series created by another climatologist, named Gordon Jacoby. When overlaid, graphs of the two series showed that the two series tracked each other remarkably closely for most of the length of the record, until the nineteenth century, when the bristlecones started shooting upwards while the northern treeline series was apparently unaffected. The gap between the two was, in essence, the carbon dioxide fertilization effect. This can be seen in Figure 1.

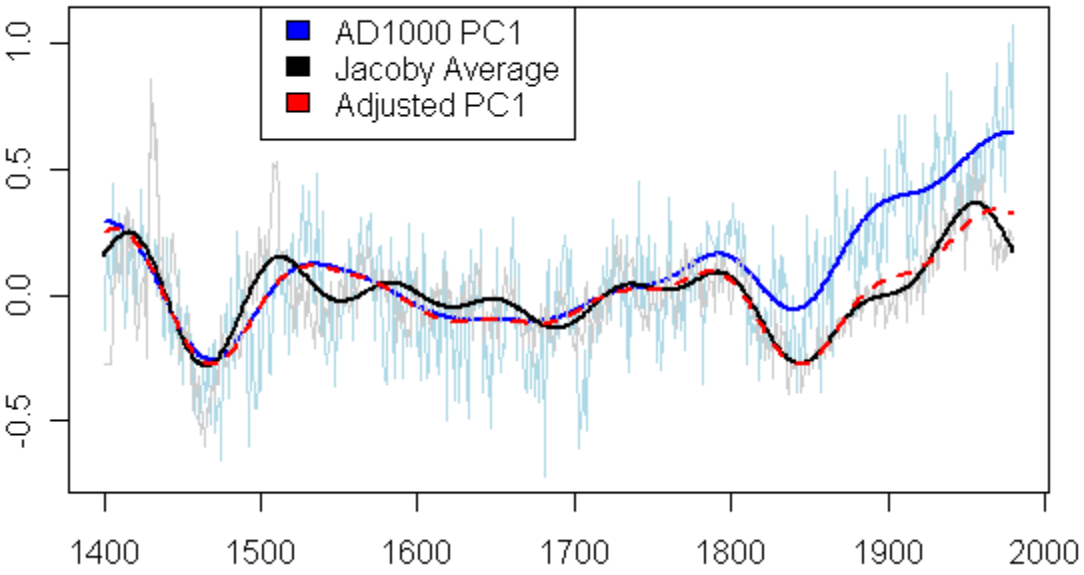

Figure 1 The Jacoby coercion
The bristlecones (in blue) are forced to match the Jacoby record (in black). The adjusted bristlecones are the red line.

Now of course, the size of the gap ("the residuals") would be expected to track carbon dioxide concentrations in the atmosphere and Mann had attempted to demonstrate that this was in fact the case. He had charted carbon dioxide concentration and the residuals on the same chart, as shown in Figure 2.

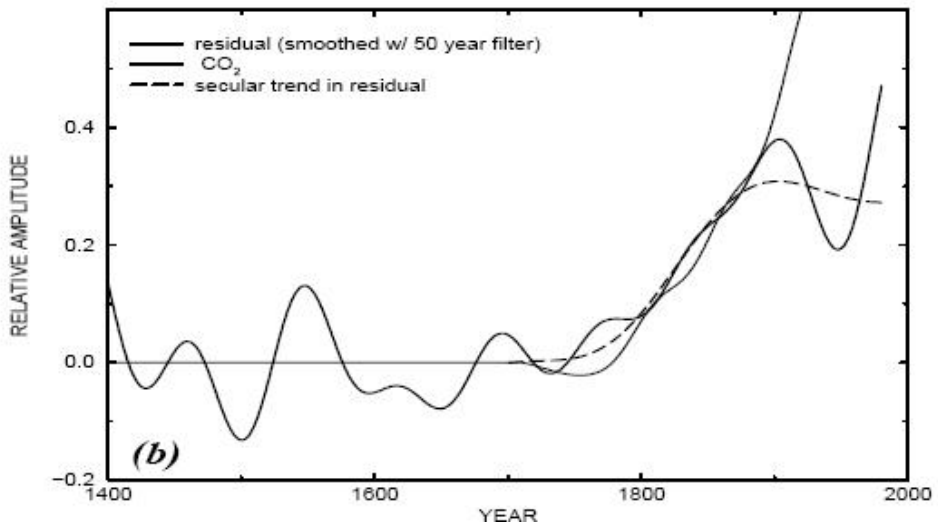

Figure 2 Correlation of the residuals to carbon dioxide

The wavy line is the residual – the size of gap between the two tree ring series – and as Mann showed it, this roughly tracked the carbon dioxide trend, which is the flat line at the left hand side, which then shoots off skywards at the top of the graph. However, while

there was a good correlation between the residual and carbon dioxide for most of the record, after around the start of the twentieth century the match seemed to fall apart, with carbon dioxide still trending upwards and the residuals falling away, only rising again after about 1950.

Mann's explanation for this phenomenon was as follows: after around 1900 the effect became "saturated", such that rises in carbon dioxide produced no further divergence between the two tree ring series. In support of his position, he presented what he called the "secular trend" in the residual. This was a rather inflated way of describing what was merely a version of the residuals that had been smoothed over the very long period of 75 years. This is the dotted line in the chart above, which, because of its long smoothing filter, does indeed seem to tail off at the end. Careful observers of his explanation might have pointed out that the underlying data, the wavy line, was actually shooting upwards at the end of the record, suggesting that this hypothesis was actually mistaken, but this point appears to have eluded the peer reviewers of Mann's paper and so, for the moment at least, the explanation stood.

With his explanation in hand, Mann had subtracted the secular trend – the smoothed difference between the two types of trees – from the bristlecone series, forcing it down to the levels of the Jacoby northern treeline series in the twentieth century. The result is the dotted line shown in Figure 1.

When McIntyre started to examine this part of Mann's calculations, however, things weren't quite as they might have seemed from the text of the original paper. Having obtained copies of the two tree ring series and the carbon dioxide record, McIntyre created his own version of Mann's comparison of the trend and obtained a result that was strikingly different (see Figure 3).

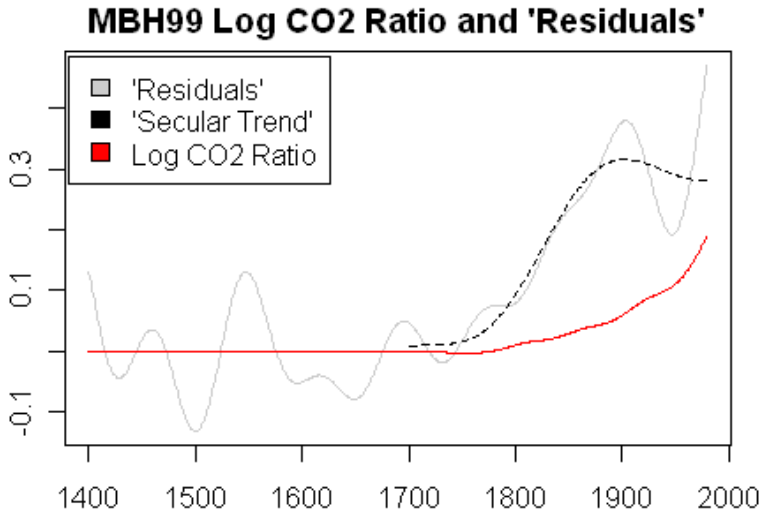

Figure 3 McIntyre's first attempt
The rise in tree growth (in grey) appears to precede the rise in carbon dioxide (in red).

In McIntyre's version of the comparison, there was virtually no match between the carbon dioxide trend and the residual series or the smoothed secular version of it. In fact it looked rather as if the growth spurt in the high altitude trees was *preceding* the rise in carbon dioxide.

It appeared then that Mann must have rescaled the carbon dioxide record to better match the residuals. The obvious way to do this was by means of regression analysis, a mathematical technique that would calculate a mathematical relationship between the two records and from this derive a transformed version of the carbon dioxide record that matched the tree ring residuals as closely as possible over its full length.

McIntyre's first attempt at the regression is shown in Figure 4:

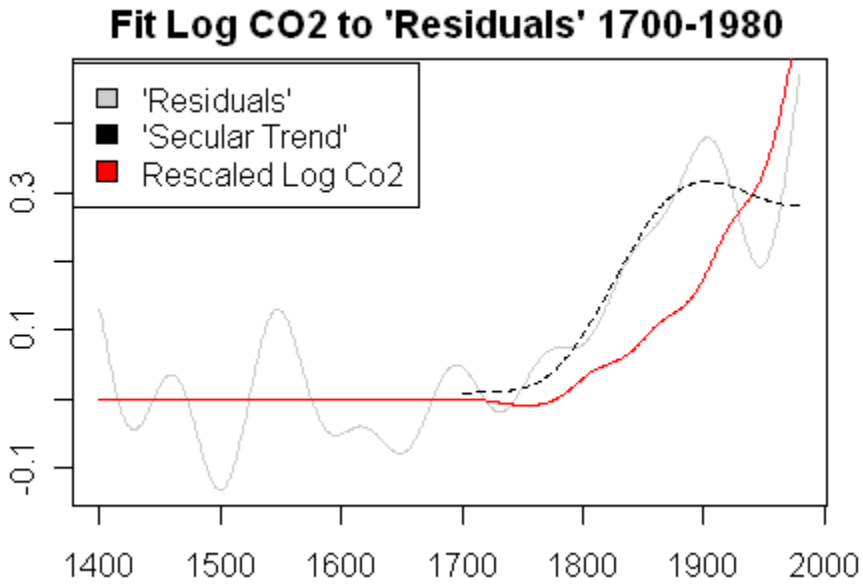

Figure 4 McIntyre's second attempt
By rescaling the carbon dioxide record, it can be made to track the rise in carbon dioxide better. This was still different to Mann's paper, however.

As can be seen, the rescaled carbon dioxide line now passes through the centre of all the waves in the residuals. This, of course, looked nothing like Mann's own version, as shown above, where the two lines parted company after 1900 AD. It looked almost as if Mann might have set the algorithm to match the two lines only as far as the end of the nineteenth century and so McIntyre performed a second version of this calculation forcing a fit only over this shorter period (see Figure 5).

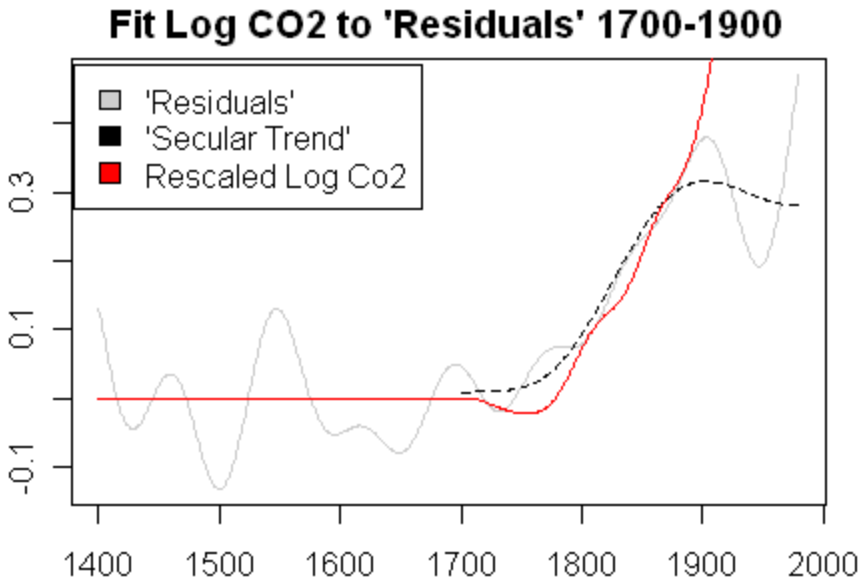

Figure 5 The restricted regression
By restricting the match to the period 1400–1900, McIntyre was able to get a close match to Mann's paper.

While not identical to Mann's results, this was certainly very close to it, suggesting that some slight variation on this calculation was what had been done in practice. If so, it raised a multitude of questions about all the different steps Mann had gone through. Why smooth the residuals over 75 years? Did such an adjustment have any meaning, particularly since the underlying tree ring series had already been smoothed? Wasn't the levelling off in this "secular trend" merely an artefact of the smoothing, which couldn't therefore be used to suggest a saturation of the carbon dioxide fertilization effect?

Then again, why restrict the fit of the carbon dioxide series to the residuals to the end of the nineteenth century? Did this have any meaning that was not better encapsulated by matching the series over their full length? There were many other questions too: why use Jacoby's northern treeline series, much of which was already obsolete by the time Mann performed his calculations? What is more, Jacoby's records had come from Northern Canada, thousands of miles from the bristlecones in the Western USA. Who was to say that there wasn't some regional effect distorting the relationship between the two series?

Many of these questions remain unanswered, but in time McIntyre and his Climate Auditors would get to the bottom of some of them and, as so often in the story of the Hockey Stick, there were some interesting tales uncovered.

A program

A few clues were unearthed among the files on Michael Mann's University of Virginia website – the same site where McIntyre had discovered the program Mann had used to

calculate tree ring principal components, in the process unravelling the main mystery of the Hockey Stick.

Among these files, McIntyre had discovered a computer program called co2detrend.f, which appeared to be the actual program used to perform the carbon dioxide adjustment – certainly the comments on the program suggested that this was the case:

```
c regress out co2-correlated trend (r=0.9 w/ co2)
c after 1800 from pc1 of ITRDB data

c remove co2-correlated portion (r=0.9) of 1800-1980
c corr= 0.9
```

The problem was that when McIntyre used co2detrend.f to process the tree ring and carbon dioxide data, it didn't turned out that it wasn't the adjustment used in the Hockey Stick paper. Quite what it did and where it was used remained a mystery.

Finlandia

The carbon dioxide adjustment remained a mystery until the sensational release of the Climategate emails at the end of 2009. Although public attention has been focused on the emails released from CRU, there were also large numbers of data and program files in the CRU archive. And while the media storm was raging, experienced eyes were examining these to see what they contained.

"Jean Sibelius" is the pseudonym for a professional statistician from Finland. A regular commenter and occasional guest poster at McIntyre's Climate Audit website, JeanS, as he is usually known, had been a valuable foil to McIntyre's work over the years, suggesting new lines of inquiry and challenging McIntyre's thinking. Over the years Sibelius had developed considerable expertise in the intricacies of the Hockey Stick paper, and was familiar with all of the remaining mysteries like the confidence interval calculations and of course the carbon dioxide adjustment.

One of the problems that Sibelius had been worrying about concerned an apparent discrepancy between the adjustment as described by Mann in the text of his paper and what actually appeared to have happened in practice. As we saw above, the narrative of the paper described smoothing of both the tree ring and the carbon dioxide records, the former then being smoothed *again* before the adjustment was calculated. But when Sibelius compared the original PC1 to the adjusted one, the difference, while very similar to the secular trend, was not identical. Rather than being a smooth line, it was stepped, or "piecewise linear" in the jargon (see Figure 6). This certainly ruled out co2detrend.f having been used to create the adjustment, and even suggested that, contrary to the narrative in the paper, the secular trend may not have been used in the adjustment at all.

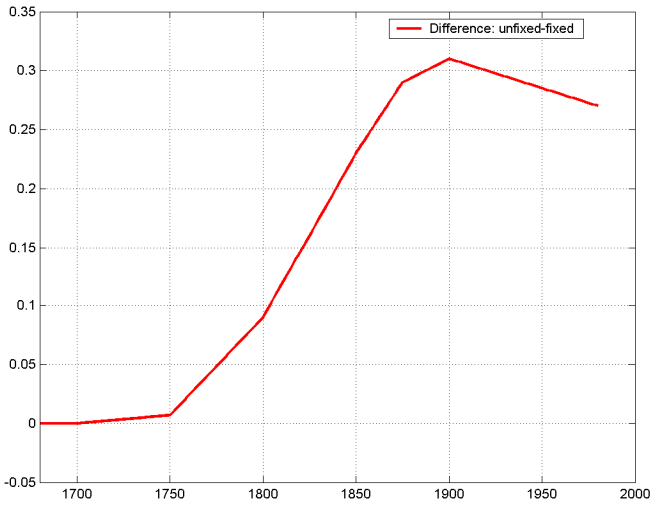

Figure 6 Mann's actual correction was stepwise linear

A new program

So there was a program that purported to do the adjustment but didn't. Moreover, the actual adjustment seemed to be only some sort of an approximation to the adjustment that was described in the paper. This was extremely strange. However, when he started examining the Climategate files, Sibelius stumbled across what turned out to be the solution to this conundrum. Trawling through the leaked data, he noticed that amongst the freed information was a directory of files relating to the Hockey Stick. This appeared to have been obtained by Tim Osborn, one of Briffa's colleagues at CRU and a regular co-author of Mann and the other key players in the Hockey Team. Examining these files, Sibelius noticed that as well as co2detrend.f there was another program called residualdetrend.f. This immediately stood out as being a file that had not previously been discussed and so he set to work to see what it contained.

To his surprise, almost the first thing he noticed was some remarkably familiar text:

```
c       regress out co2-correlated trend (r=0.9 w/ co2)
c       after 1800 from pc1 of ITRDB data
```

The comment was almost identical to the one from co2detrend.f. This was very strange. Why should Mann have created *two* programs to do the same adjustment? As he read further, he realized that he might have solved the problem. There, in the comments was a suggestion that this program might create an adjustment that was piecewise linear, as the final adjustment appeared to be.

```
c       linear segments describing approximate residuals
c       relative to fit with respect to secular trend
```

And as he checked over the code, he realized that it did just what the comments suggested, creating a piecewise linear adjustment that roughly matched up with the secular trend in carbon dioxide shown in the original paper. Why Mann should choose to

create this approximation to the secular trend, instead of simply using the secular trend itself, remained a mystery.

This then appeared to be the program Mann had used to create the adjustment, but this then raised the uncomfortable question of what the other program, co2detrend.f, was for?

Sibelius examined the two programs side-by-side and quickly noticed that they both output their results to the same file, pc01-fixed.dat. This, together with the similarity of the comments and the file names, certainly seemed to indicate that they were different attempts at the same adjustment, and the suggestion of a "trial and error" approach to the adjustment was disturbing, hinting at a search for a desired answer rather than a scientific consideration of how the issue could best be corrected from a physical point of view

And another

At this point, Sibelius had another surprise. In the same directory, there was *another* file that seemed relevant. This was a data file, the name of which – pc1-fixed-old.dat – suggested that it might hold the output from the rejected co2detrend.f program. However, knowing that surmises of this kind are often unwarranted when studying the Hockey Stick, Sibelius decided to make sure. To his surprise, it turned out that the data in the file was entirely different – it seemed to represent the results of a *third* attempt at creating a carbon dioxide adjustment.

The three adjustments are shown in Figure 7:

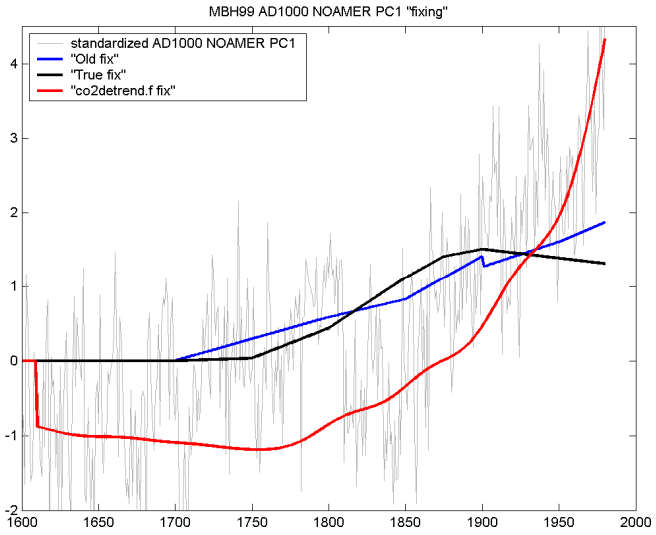

Figure 7 Three attempts to correct for carbon dioxide fertilization

Adjusting the Medieval Warm Period

Although we do not know what order the fixes were calculated in, it is probably fair to assume that Mann started from the unfixed graph and ended up with the version that he finally published. This analysis is instructive because it throws considerable light on the reasons for Mann's multiple smoothing of the records and the strange restriction of the residual–carbon dioxide match.

The first results of the MBH99 algorithm – before any fix for CO2 fertilization was applied – is shown in Figure 8:

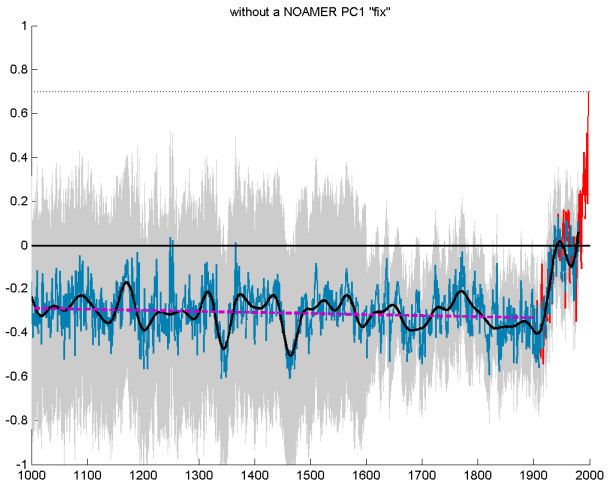

Figure 8 The Hockey Stick with no fix for carbon dioxide fertilization
The handle of the stick, from 1000 to 1900 is almost trendless.

To anyone used to looking at temperature reconstructions, this immediately looks odd. For 900 years, from the beginning of the chart to the end of the nineteenth century, there is essentially no trend in the reconstructed temperature at all, giving the chart a very artificial appearance. This would simply never be accepted by the climatological community where it was widely accepted that temperatures are affected by trends over much longer timescales than a century. The absence of any such long-term trend would immediately raise question marks over Mann's result.

To those who have followed the Hockey Stick story closely, the appearance of the uncorrected reconstruction is also reminiscent of the observation made by several skeptics that the effect of Mann's calibration algorithm was to pick out series with twentieth century upticks. These then gave a blade to the Hockey Stick graph. But in the rest of the series, these same series were essentially random and cancelled each other out. The long flat blade of the uncorrected reconstruction seemed to be strong confirmation that this was the case.

The stepwise calculation

At this point we need to remind ourselves how Mann put his temperature reconstructions together.

The essence of a temperature reconstruction is to find a mathematical relationship between the proxies, i.e. the tree rings, and temperature. Because so many of the tree rings are from one or two geographical areas though, these have first to be summarized using a technique called principal components analysis (PCA). Once this has been done a mathematical model can be worked out for the relationship between, on the one hand, this summary and all the other data series in the twentieth century and on the other hand the temperature records for the same period. These mathematical relationships can be used to work out temperatures in earlier periods from the tree ring data.

When summarizing using PCA it is important to have no gaps in the data and this should have been a problem for Mann because, as the bristlecone series went back in time, there were fewer and fewer trees and therefore more and more gaps in the data series: as the earliest date for each series was reached, any previous dates were seen as gaps. To get round this problem Mann had performed the calculation in a "stepwise" fashion. Essentially he prepared several reconstructions and spliced sections of them together to get his final graph.

Let's look at a simplified example which should help you understand how this process worked. In our example, all of the trees are at least six hundred years old. Some are even older than this, going back a further 250 years, and a smaller number again go right back to the start of the millennium. In Mann's stepwise calculation, a reconstruction back to 1400AD would have been prepared first (the top line in Figure 9). Then a second reconstruction was prepared on the subsample of the trees that extended back to 1150AD (the middle line). Finally the small number of trees that were 1000 years old would be used to prepare a reconstruction right back to the start of the millennium.

Now of course, the top line, having the most trees, is the most reliable reconstruction, so all of this graph was used in the final reconstruction. Between 1150 and 1400AD, however, the most reliable estimate of temperature available is from the middle graph, so the section of this line which is coloured red in Figure 9 would be extracted and spliced onto the top line. Finally the 1000–1150 section of the bottom chart would be used for the earliest section of the final Hockey Stick graph.

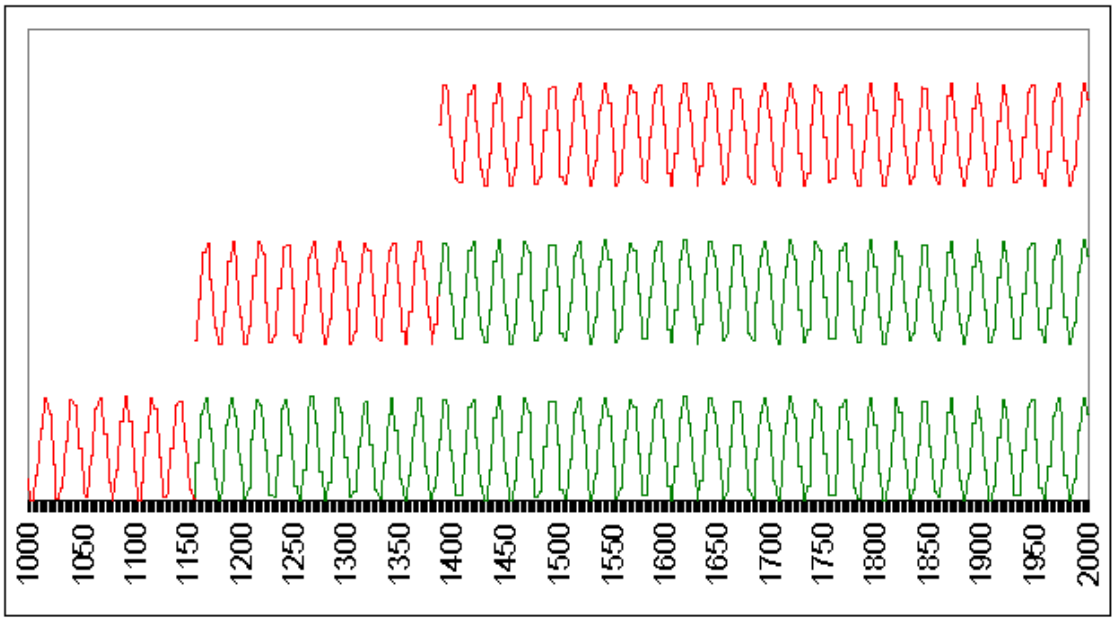

Figure 9 Example of a stepwise calculation
Reconstructions are calculated on different subsets of the tree ring data. The red sections of each reconstruction are spliced together to make the final graph.

Now this is a simplified example and in practice Mann used steps that were rather different. It is worth examining these as they are somewhat intriguing.

Table 1 Reconstruction steps in MBH99
These are equivalent to the red sections that would be spliced together in the final graph.

Step	Number of years
1000–1399	400
1400-1449	50
1450-1499	50
1500-1599	100
1600-1699	100
1700-1729	30
1730-1749	20
1750-1759	10
1760-1779	20
1780-1799	20
1800-1819	20
1820-1980	161

The varying length of the steps is very odd. Why should Mann choose to have one step that was only ten years long but another that was 400 years long? This remains one of the

methodological steps of the Hockey Stick that is still not understood by outsiders, but the impact of such a strange procedure is certainly worth questioning.

Applying the adjustment

When McIntyre and Sibelius had started examining how the adjustment worked, they had had another surprise. Mann had only applied the correction to the AD1000 step of his PC analysis. This meant that the correction only ended up affecting the very early sections of the final reconstruction – the period from 1000 to 1400 AD. This was extremely strange because of course he was trying to correct for twentieth century carbon dioxide fertilization.

It is somewhat counterintuitive that an adjustment for twentieth century carbon dioxide fertilization should affect only the shape of the graph in the first four centuries of the millennium, but that is how Mann chose to do it.

Assuming he prepared the co2detrend version first, he would have obtained a reconstruction that looked like Figure 10.

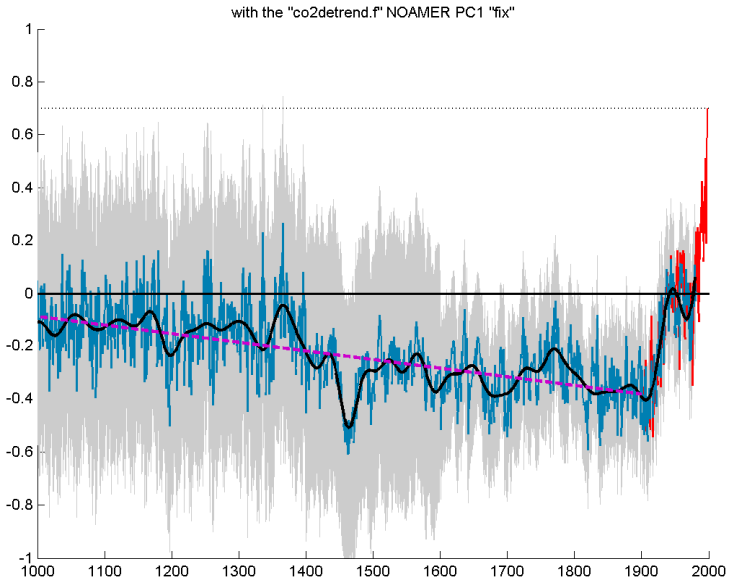

Figure 10 The Hockey Stick using the CO2detrend.f version of the fix
In this version, the Medieval Warm Period is relatively warm, reaching levels in line with the proxy record at the end of the twentieth century. It also had very poor verification statistics, rendering it unusable.

Although at first sight the charts are similar the new version had a much faster decline up to 1900. This would have presented something of a problem because the Medieval Warm Period now looked rather warm – in touching distance of the levels the proxies were

reaching at the end of the 1970s. A graph like this would have raised severe question marks over whether there really was anything unprecedented about recent temperatures.

What was worse was that the verification statistics for the revised version were dreadful and the results would therefore have been laughed off as unreliable.

The oldfix correction, meanwhile had better verification statistics (although still short of the values achieved by the unfixed version) but still had values in the early sections of the reconstruction that were too high for comfort (see Figure 11).

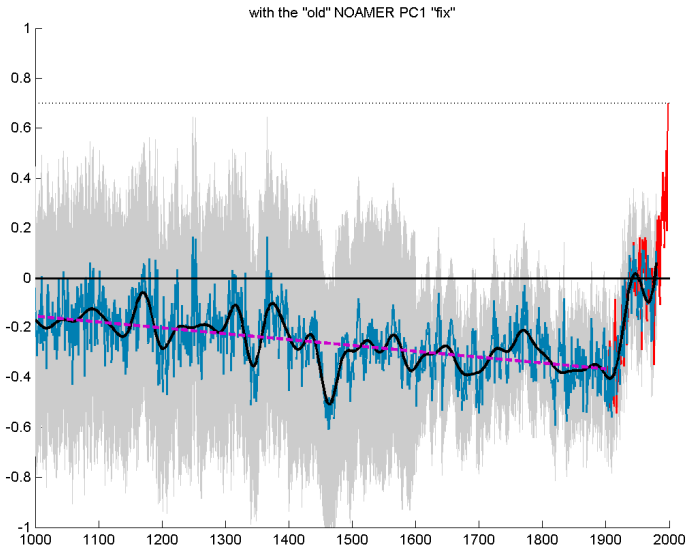

Figure 11 The Hockey Stick with the oldfix version of the fix
This version still suffers from very poor verification statistics and was therefore not usable.

This seems to have brought Mann on to the residualdetrend version he had finally settled on. With all the strange distortions of the graph – the repeated smoothing and the application only to the early section of the record – Mann appeared to get the result he wanted (see Figure 12). Here the slope was a gentle decline over 900 years that could be explained away as the results of the Milankovitch cycle; the Medieval Warm Period was satisfyingly cool, and the verification statistics were a whisker better than the unfixed version of the graph, the artificial appearance of which was now a thing of the past.

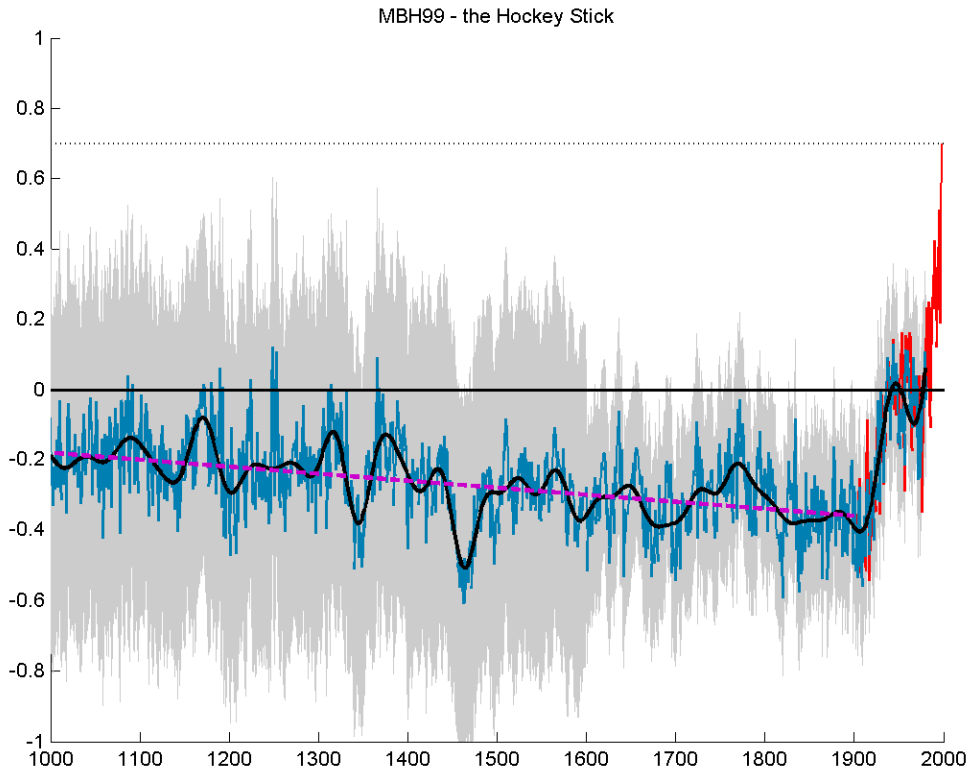

Figure 12 The Hockey Stick with the residualdetrend version of the fix
This was the version used in the final published article.