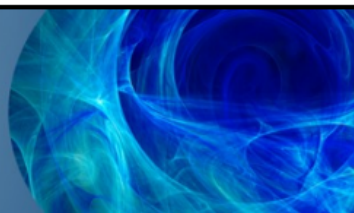# Climate scientists versus climate data

Posted on February 4, 2017 | 761 Comments

by John Bates

A look behind the curtain at NOAA's climate data center.

I read with great irony recently that scientists are "*frantically copying U.S. Climate data, fearing it might vanish under Trump*" (e.g., Washington Post 13 December 2016). As a climate scientist formerly responsible for NOAA's climate archive, the most critical issue in archival of climate data is actually scientists who are unwilling to formally archive and document their data. I spent the last decade cajoling climate scientists to archive their data and fully document the datasets. I established a climate data records program that was awarded a U.S. Department of Commerce Gold Medal in 2014 for visionary work in the acquisition, production, and preservation of climate data records (CDRs), which accurately describe the Earth's changing environment.

The most serious example of a climate scientist not archiving or documenting a critical climate dataset was the study of Tom Karl et al. 2015 (hereafter referred to as the Karl study or K15), purporting to show no 'hiatus' in global warming in the 2000s (Federal scientists say there never was any global warming "pause"). The study drew criticism from other climate scientists, who disagreed with K15's conclusion about the 'hiatus.' (Making sense of the early-2000s warming slowdown). The paper also drew the attention of the Chairman of the House Science Committee, Representative Lamar Smith, who questioned the timing of the report, which was issued just prior to the Obama Administration's Clean Power Plan submission to the Paris Climate Conference in 2015.

In the following sections, I provide the details of how Mr. Karl failed to disclose critical information to NOAA, *Science Magazine*, and Chairman Smith regarding the datasets used in K15. I have extensive documentation that provides independent verification of the story below. I also provide my suggestions for how we might keep such a flagrant manipulation of scientific integrity guidelines and scientific publication standards from happening in the future. Finally, I provide some links to examples of what well documented CDRs look like that readers might contrast and compare with what Mr. Karl has provided.

Background

In 2013, prior to the Karl study, the National Climatic Data Center [NCDC, now the NOAA National Centers for Environmental Information (NCEI)] had just adopted much improved processes for formal review of Climate Data Records, a process I formulated [link]. The land temperature dataset used in the Karl study had never been processed through the station adjustment software before, which led me to believe something was amiss. When I pressed the co-authors, they said they had decided not to archive the dataset, but did not defend the decision. One of the co-authors said there were 'some decisions [he was] not happy with'. The data used in the K15 paper were only made available through a web site, not in digital form, and lacking proper versioning and any notice that they were research and not operational data. I was dumbstruck that Tom Karl, the NCEI Director in charge of NOAA's climate data archive, would not follow the policy of his own Agency nor the guidelines in *Science* magazine for dataset archival and documentation.

I questioned another co-author about why they choose to use a 90% confidence threshold for evaluating the statistical significance of surface temperature trends, instead of the standard for significance of 95% — he also expressed reluctance and did not defend the decision. A NOAA NCEI supervisor remarked how it was eye-opening to watch Karl work the co-authors, mostly subtly but sometimes not, pushing choices to emphasize warming. Gradually, in the months after K15 came out, the evidence kept mounting that Tom Karl constantly had his 'thumb on the scale'—in the documentation, scientific choices, and release of datasets—in an effort to discredit the notion of a global warming hiatus and rush to time the publication of the paper to influence national and international deliberations on climate policy.

Defining an Operational Climate Data Record

For nearly two decades, I've advocated that if climate datasets are to be used in important policy decisions, they must be fully documented, subject to software engineering management and improvement processes, and be discoverable and accessible to the public with rigorous information preservation standards. I was able to implement such policies, with the help of many colleagues, through the NOAA Climate Data Record policies (CDR) [link].

Once the CDR program was funded, beginning in 2007, I was able to put together a team and pursue my goals of operational processing of important climate data records emphasizing the processes required to transition research datasets into operations (known as R2O). Figure 1 summarizes the steps required to accomplish this transition in the key elements of software code, documentation, and data.
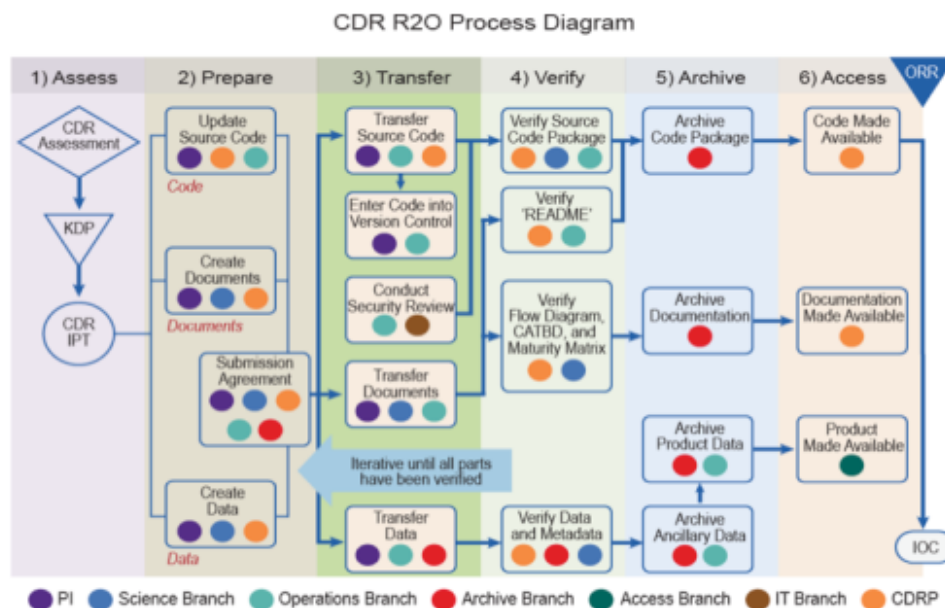

Figure 1. Research to operations transition process methodology from Bates et al. 2016.

Unfortunately, the NCDC/NCEI surface temperature processing group was split on whether to adopt this process, with scientist Dr. Thomas C. Peterson (a co-author on K15, now retired from NOAA) vigorously opposing it. Tom Karl never required the surface temperature group to use the rigor of the CDR methodology, although a document was prepared identifying what parts of the surface temperature processing had to be improved to qualify as an operational CDR.

Tom Karl liked the maturity matrix so much, he modified the matrix categories so that he could claim a number of NCEI products were "Examples of "Gold" standard NCEI Products (Data Set Maturity Matrix Model Level 6)." See his

NCEI overview presentation all NCEI employees [ncei-overview-2015nov-2] were told to use, even though there had never been any maturity assessment of any of the products.

NCDC/NCEI surface temperature processing and archival

In the fall of 2012, the monthly temperature products issued by NCDC were incorrect for 3 months in a row [link]. As a result, the press releases and datasets had to be withdrawn and reissued. Dr. Mary Kicza, then the NESDIS Associate Administrator (the parent organization of NCDC/NCEI in NOAA), noted that these repeated errors reflected poorly on NOAA and required NCDC/NCEI to improve its software management processes so that such mistakes would be minimized in the future. Over the next several years, NCDC/NCEI had an incident report conducted to trace these errors and recommend corrective actions.

Following those and other recommendations, NCDN/NCEI began to implement new software management and process management procedures, adopting some of the elements of the CDR R2O process. In 2014 a NCDC/NCEI Science Council was formed to review new science activities and to review and approve new science products for operational release. A draft operational readiness review (ORR) was prepared and used for approval of all operational product releases, which was finalized and formally adopted in January 2015. Along with this process, a contractor who had worked at the CMMI Institute (CMMI, Capability Maturity Model Integration, is a software engineering process level improvement training and appraisal program) was hired to improve software processes, with a focus on improvement and code rejuvenation of the surface temperature processing code, in particular the GHCN-M dataset.

The first NCDC/NCEI surface temperature software to be put through this rejuvenation was the pairwise homogeneity adjustment portion of processing for the GHCN-Mv4 beta release of October 2015. The incident report had found that there were unidentified coding errors in the GHCN-M processing that caused unpredictable results and different results every time code was run.

The generic flow of data used in processing of the NCDC/NCEI global temperature product suite is shown schematically in Figure 2. There are three steps to the processing, and two of the three steps are done separately for the ocean versus land data. Step 1 is the compilation of observations either from ocean sources or land stations. Step 2 involves applying various adjustments to the data, including bias adjustments, and provides as output the adjusted and unadjusted data on a standard grid. Step 3 involves application of a spatial analysis technique (empirical orthogonal teleconnections, EOTs) to merge and smooth the ocean and land surface temperature fields and provide these merged fields as anomaly fields for ocean, land and global temperatures. This is the product used in K15. Rigorous ORR for each of these steps in the global temperature processing began at NCDC in early 2014.
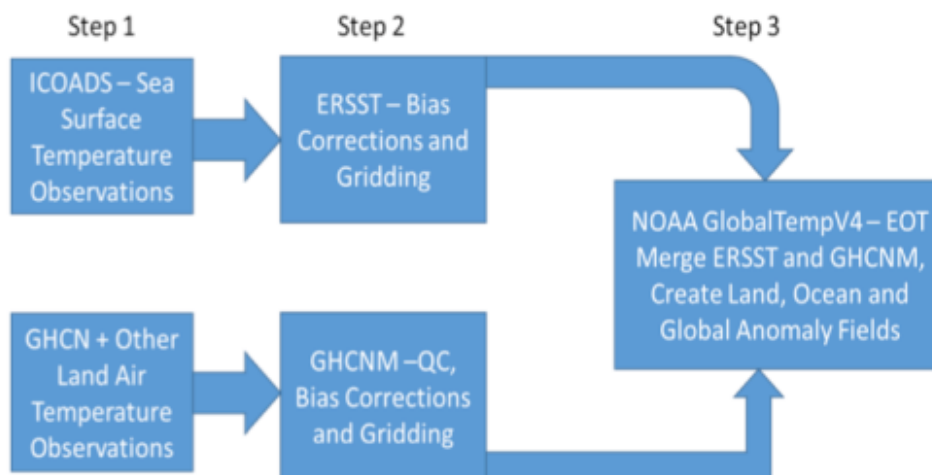
Figure 2. Generic data flow for NCDC/NCEI surface temperature products.

In K15, the authors describe that the land surface air temperature dataset included the GHCN-M station data and also the new ISTI (Integrated Surface Temperature Initiative) data that was run through the then operational GHCN-M bias correction and gridding program (i.e., Step 2 of land air temperature processing in Figure 2). They further indicated that this processing and subsequent corrections were 'essentially the same as those used in GHCN-Monthly version 3'. This may have been the case; however, doing so failed to follow the process that had been initiated to ensure the quality and integrity of datasets at NCDC/NCEI.

The GHCN-M V4 beta was put through an ORR in October 2015; the presentation made it clear that any GHCN-M version using the ISTI dataset should, and would, be called version 4. This is confirmed by parsing the file name actually used on the FTP site for the K15 dataset [link]; NOTE: placing a non-machine readable copy of a dataset on an FTP site does not constitute archiving a dataset). One file is named 'box.12.adj.4.a.1.20150119', where 'adj' indicates adjusted (passed through step 2 of the land processing) and '4.a.1' means version 4 alpha run 1; the entire name indicating GHCN-M version 4a run 1. That is, the folks who did the processing for K15 and saved the file actually used the correct naming and versioning, but K15 did not disclose this. Clearly labeling the dataset would have indicated this was a highly experimental early GHCN-M version 4 run rather than a routine, operational update. As such, according to NOAA scientific integrity guidelines, it would have required a disclaimer not to use the dataset for routine monitoring.

In August 2014, in response to the continuing software problems with GHCNMv3.2.2 (version of August 2013), the NCDC Science Council was briefed about a proposal to subject the GHCNMv3 software, and particularly the pairwise homogeneity analysis portion, to a rigorous software rejuvenation effort to bring it up to CMMI level 2 standards and resolve the lingering software errors. All software has errors and it is not surprising there were some, but the magnitude of the problem was significant and a rigorous process of software improvement like the one proposed was needed. However, this effort was just beginning when the K15 paper was submitted, and so K15 must have used data with some experimental processing that combined aspects of V3 and V4 with known flaws. The GHCNMv3.X used in K15 did not go through any ORR process, and so what precisely was done is not documented. The ORR package for GHCNMv4 beta (in October 2015) uses the rejuvenated software and also includes two additional quality checks versus version 3.

Which version of the GHCN-M software K15 used is further confounded by the fact that GHCNMv3.3.0, the upgrade from version 3.2.2, only went through an ORR in April 2015 (i.e., after the K15 paper was submitted and revised). The GHCN-Mv3.3.0 ORR presentation demonstrated that the GHCN-M version changes between V3.2.2 and V3.3.0 had impacts on rankings of warmest years and trends. The data flow that was operational in June 2015 is shown in figure 3.
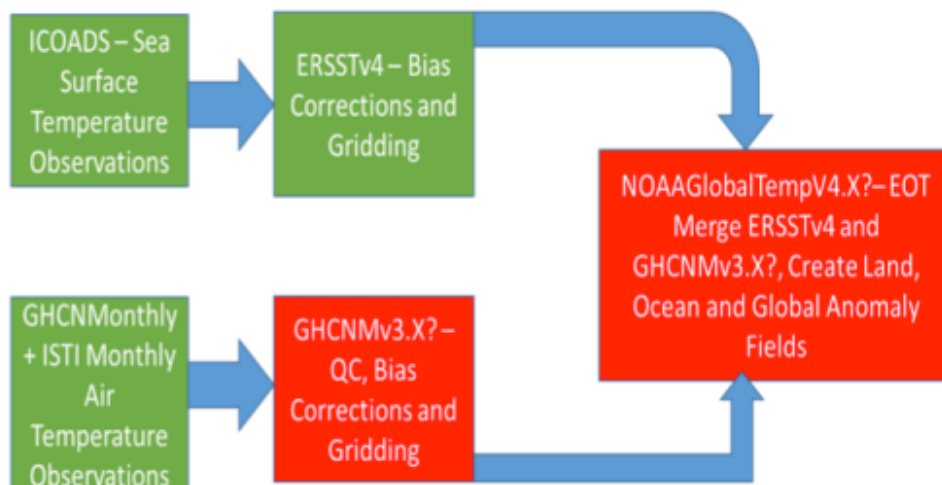

Figure 3. Data flow for surface temperature products described in K15 Science paper. Green indicates operational datasets having passed ORR and archived at time of publication. Red indicates experimental datasets never subject to ORR and never archived.

It is clear that the actual nearly-operational release of GHCN-Mv4 beta is significantly different from the version GHCNM3.X used in K15. Since the version GHCNM3.X never went through any ORR, the resulting dataset was also never archived, and it is virtually impossible to replicate the result in K15.

At the time of the publication of the K15, the final step in processing the NOAAGlobalTempV4 had been approved through an ORR, but not in the K15 configuration. It is significant that the current operational version of NOAAGlobalTempV4 uses GHCN-M V3.3.0 and does not include the ISTI dataset used in the *Science* paper. The K15 global merged dataset is also not archived nor is it available in machine-readable form. This is why the two boxes in figure 3 are colored red.

The lack of archival of the GHCN-M V3.X and the global merged product is also in violation of *Science* policy on making data available [link]. This policy states: "*Climate data. Data should be archived in the NOAA climate repository or other public databases*". Did Karl et al. disclose to *Science Magazine* that they would not be following the NOAA archive policy, would not archive the data, and would only provide access to a non-machine readable version only on an FTP server?

For ocean temperatures, the ERSST version 4 is used in the K15 paper and represents a major update from the previous version. The bias correction procedure was changed and this resulted in different SST anomalies and different trends during the last 15+ years relative to ERSST version 3. ERSSTV4 beta, a pre-operational release, was briefed to the NCDC Science Council and approved on 30 September 2014.

The ORR for ERSSTV4, the operational release, took place in the NCDC Science Council on 15 January 2015. The ORR focused on process and questions about some of the controversial scientific choices made in the production of that dataset will be discussed in a separate post. The review went well and there was only one point of discussion on process. One slide in the presentation indicated that operational release was to be delayed to coincide with Karl et al. 2015 *Science* paper release. Several Science Council members objected to this, noting the K15 paper did not contain any further methodological information—all of that had already been published and thus there was no rationale to delay the dataset release. After discussion, the Science Council voted to approve the ERSSTv4 ORR and recommend immediate release.

The Science Council reported this recommendation to the NCDC Executive Council, the highest NCDC management board. In the NCDC Executive Council meeting, Tom Karl did not approve the release of ERSSTv4, noting that he wanted its release to coincide with the release of the next version of GHCNM (GHCNMv3.3.0) and NOAAGlobalTemp. Those products each went through an ORR at NCDC Science Council on 9 April 2015, and were used in operations in May. The ERSSTv4 dataset, however, was still not released. NCEI used these new analyses, including ERSSTv4, in its operational global analysis even though it was not being operationally archived. The operational version of ERSSTv4 was only released to the public following publication of the K15 paper. The withholding of the operational version of this important update came in the middle of a major ENSO event, thereby depriving the public of an important source of updated information, apparently for the sole purpose of Mr. Karl using the data in his paper before making the data available to the public.

So, in every aspect of the preparation and release of the datasets leading into K15, we find Tom Karl's thumb on the scale pushing for, and often insisting on, decisions that maximize warming and minimize documentation. I finally decided to document what I had found using the [climate data record maturity matrix approach](#). I did this and sent my concerns to the NCEI Science Council in early February 2016 and asked to be added to the agenda of an upcoming meeting. I was asked to turn my concerns into a more general presentation on requirements for publishing and archiving. Some on the Science Council, particularly the younger scientists, indicated they had not known of the *Science* requirement to archive data and were not aware of the open data movement. They promised to begin an archive request for the K15 datasets that were not archived; however I have not been able to confirm they have been archived. I later learned that the computer used to process the software had suffered a complete failure, leading to a tongue-in-cheek joke by some who had worked on it that the failure was deliberate to ensure the result could never be replicated.

Where do we go from here?

I have wrestled for a long time about what to do about this incident. I finally decided that there needs to be systemic change both in the operation of government data centers and in scientific publishing, and I have decided to become an advocate for such change. First, Congress should re-introduce and pass the [OPEN Government Data Act](#). The Act states that federal datasets must be archived and made available in machine readable form, neither of which was done by K15. The Act was introduced in the last Congress and the Senate passed it unanimously in the lame duck session, but the House did not. This bodes well for re-introduction and passage in the new Congress.

However, the Act will be toothless without an enforcement mechanism. For that, there should be mandatory, independent certification of federal data centers. As I noted, the scientists working in the trenches would actually welcome this, as the problem has been one of upper management taking advantage of their position to thwart the existing executive orders and a lack of process adopted within Agencies at the upper levels. Only an independent, outside body can provide the needed oversight to ensure Agencies comply with the OPEN Government Data Act.

Similarly, scientific publishers have formed the Coalition on Publishing Data in the Earth and Space Sciences ([COPDESS](#)) with a signed statement of commitment to ensure open and documented datasets are part of the publication process. Unfortunately, they, too, lack any standard checklist that peer reviewers and editors can use to

ensure the statement of commitment is actually enforced. In this case, and for assessing archives, I would advocate a metric such as the data maturity model that I and colleagues have developed. This model has now been adopted and adapted by several different groups, applied to hundreds of datasets across the geophysical sciences, and has been found useful for ensuring information preservation, discovery, and accessibility.

Finally, there needs to be a renewed effort by scientists and scientific societies to provide training and conduct more meetings on ethics. Ethics needs to be a regular topic at major scientific meetings, in graduate classrooms, and in continuing professional education. Respectful discussion of different points of view should be encouraged. Fortunately, there is initial progress to report here, as scientific societies are now coming to grips with the need for discussion of and guidelines for scientific ethics.

There is much to do in each of these areas. Although I have retired from the federal government, I have not retired from being a scientist. I now have the luxury of spending more time on these things that I am most passionate about. I also appreciate the opportunity to contribute to Climate Etc. and work with my colleague and friend Judy on these important issues.

Postlude

A couple of examples of how the public can find and use CDR operational products, and what is lacking in a non-operational and non-archived product

1. NOAA CDR of total solar irradiance – this is the highest level quality. Start at web site – https://data.nodc.noaa.gov/cgi-bin/iso?id=gov.noaa.ncdc:C00828

Here you will see a fully documented CDR. At the top, we have the general description and how to cite the data. Then below, you have a set of tabs with extensive information. Click each tab to see how it's done. Note, for example, that in 'documentation' you have choices to get the general documentation, processing documents including source code, data flow diagram, and the algorithm theoretical basis document ATBD which includes all the info about how the product is generated, and then associated resources. This also includes a permanent digital object identifier (doi) to point uniquely to this dataset.

2. NOAA CDR of mean layer temperature – RSS – one generation behind in documentation but still quite good – https://www.ncdc.noaa.gov/cdr/fundamental/mean-layer-temperature-rss

Here on the left you will find the documents again that are required to pass the CDR operations and archival. Even though it's a slight cut below TSI in example 1, a user has all they need to use and understand this.

3. The Karl hiatus paper can be found on NCEI here – https://www.ncdc.noaa.gov/news/recent-global-surface-warming-hiatus

If you follow the quick link 'Download the Data via FTP' you go here – ftp://ftp.ncdc.noaa.gov/pub/data/scpub201506/

The contents of this FTP site were entered into the NCEI archive following my complaint to the NCEI Science Council. However, the artifacts for full archival of an operational CDR are not included, so this is not compliant with archival standards.

Biosketch:

John Bates received his Ph.D. in Meteorology from the University of Wisconsin-Madison in 1986. Post Ph.D., he spent his entire career at NOAA, until his retirement in 2016. He spent the last 14 years of his career at NOAA's National Climatic Data Center (now NCEI) as a Principal Scientist, where he served as a Supervisory Meteorologist until 2012.

Dr. Bates' technical expertise lies in atmospheric sciences, and his interests include satellite observations of the global water and energy cycle, air-sea interactions, and climate variability. His most highly cited papers are in observational studies of long term variability and trends in atmospheric water vapor and clouds.

NOAA Administrator's Award 2004 for "outstanding administration and leadership in developing a new division to meet the challenges to NOAA in the area of climate applications related to remotely sensed data". He was awarded a U.S. Department of Commerce Gold Medal in 2014 for visionary work in the acquisition, production, and preservation of climate data records (CDRs). He has held elected positions at the American Geophysical Union (AGU), including Member of the AGU Council and Member of the AGU Board. He has played a leadership role in data management for the AGU.

He is currently President of John Bates Consulting Inc., which puts his recent experience and leadership in data management to use in helping clients improve data management to improve their preservation, discovery, and exploitation of their and others data. He has developed and applied techniques for assessing both organizational and individual data management and applications. These techniques help identify how data can be managed more cost effectively and discovered and applied by more users.

David Rose in the Mail on Sunday

[David Rose](#) of the UK Mail on Sunday is working on a comprehensive expose of this issue [[link](#)].

Here are the comments that I provided to David Rose, some of which were included in his article:

Here is what I think the broader implications are. Following ClimateGate, I made a public plea for greater transparency in climate data sets, including documentation. In the U.S., John Bates has led the charge in developing these data standards and implementing them. So it is very disturbing to see the institution that is the main U.S. custodian of climate data treat this issue so cavalierly, violating its own policy. The other concern that I raised following ClimateGate was overconfidence and inadequate assessments of uncertainty. Large adjustments to the raw data, and substantial changes in successive data set versions, imply substantial uncertainties. The magnitude of these uncertainties influences how we interpret observed temperature trends, 'warmest year' claims, and how we interpret differences between observations and climate model simulations. I also raised concerns about bias; here we apparently see Tom Karl's thumb on the scale in terms of the methodologies and procedures used in this publication.

Apart from the above issues, how much difference do these issues make to our overall understanding of global temperature change? All of the global surface temperature data sets employ NOAA's GHCN land surface temperatures. The NASA GISS data set also employs the ERSST datasets for ocean surface temperatures. There are global surface temperature datasets, such as Berkeley Earth and HadCRUT that are relatively independent of the NOAA data sets, that agree qualitatively with the new NOAA data set. However, there remain large, unexplained regional discrepancies between the NOAA land surface temperatures and the raw data. Further, there are some very large uncertainties in ocean sea surface temperatures, even in recent decades. Efforts by the global numerical weather prediction centers to produce global reanalyses such as the [European Copernicus effort](#) is probably the best way forward for the most recent decades.

Regarding uncertainty, 'warmest year', etc. there is a good article in the WSJ: Change would be healthy at U.S. climate agencies ([hockeyshtick](#) has reproduced the full article).

I also found this recent essay in phys.org to be very germane: Certainty in complex scientific research an unachievable goal. Researchers do a good job of estimating the size of errors in measurements but underestimate chance of large errors.

Backstory

I have known John Bates for about 25 years, and he served on the Ph.D. committees of two of my graduate students. There is no one, anywhere, that is a greater champion for data integrity and transparency.

When I started Climate Etc., John was one of the few climate scientists that contacted me, sharing concerns about various ethical issues in our field.

Shortly after publication of K15, John and I began discussing our concerns about the paper. I encouraged him to come forward publicly with his concerns. Instead, he opted to try to work within the NOAA system to address the issues –to little effect. Upon his retirement from NOAA in November 2016, he decided to go public with his concerns.

He submitted an earlier, shorter version of this essay to the Washington Post, in response to the 13 December article (climate scientists frantically copying data). The WaPo rejected his op-ed, so he decided to publish at Climate Etc.

In the meantime, David Rose contacted me about a month ago, saying he would be in Atlanta covering a story about a person unjustly imprisoned [link]. He had an extra day in Atlanta, and wanted to get together. I told him I wasn't in Atlanta, but put him in contact with John Bates. David Rose and his editor were excited about what John had to say.

I have to wonder how this would have played out if we had issued a press release in the U.S., or if this story was given to pretty much any U.S. journalist working for the mainstream media. Under the Obama administration, I suspect that it would have been very difficult for this story to get any traction. Under the Trump administration, I have every confidence that this will be investigated (but still not sure how the MSM will react).

Well, it will be interesting to see how this story evolves, and most importantly, what policies can be put in place to prevent something like this from happening again.

I will have another post on this topic in a few days.

Being retired sure is liberating . . .

Moderation note: As with all guest posts, please keep your comments civil and relevant.

SHARE THIS:

| ✉ Email | 🐦 Twitter | f Facebook 3K+ | 🖨 Print | Reddit | t Tumblr |

★ Like

30 bloggers like this.

RELATED

This entry was posted in Data and observations, Ethics. Bookmark the permalink.

761 RESPONSES TO "CLIMATE SCIENTISTS VERSUS CLIMATE DATA"

Pingback: Climate scientists versus climate data – Enjeux énergies et environnement

beththeserf | February 4, 2017 at 6:34 pm |

Incredible ain't it, non-archiving of critical evidence -?
Faith-based cli-sci- not for general exhibition.

> **Nick Stokes** | February 5, 2017 at 2:58 am |
>
> *"Incredible ain't it, non-archiving of critical evidence -?"*
> And just not true. There is an extensive archive. Bates even linked to it. It is here.
>
> Bates complaints seem to be
> 1. The archiving wasn't complete until six months after the paper appeared
> 2. Data is in ascii format which is not "machine readable". Of course it is, it just requires a format statement.
>
> > **johnfpittman** | February 5, 2017 at 7:00 am |
> >
> > NIck you left out complaints:
> > 1. Karl made administrative decisions contrary to data integrity;
> > 2. Karl used 90% rather than 95% standard;
> > 3. The use of non standardized data set implied a greater uncertainty to the data that was not, could not be addressed;
> > 4. Karl made poor data integrity decisions in order to meet a publishing date;
> > 5. Karl made decisions that benefited his publication not the organization;
> > 6. etc., etc.
>
> > **Nick Stokes** | February 5, 2017 at 7:43 am |
> >
> > jfp,
> > *"NIck you left out complaints:"*
> > I dealt with the 90% issue here. He was comparing results with the AR5, which used 90%.
> >
> > All the rest are matters of opinion, and his opinion does seem to be soured by something. 1,4,5 are pure opinion, and unsubstantiated. The non-standardized issue is bunk. If you go to the readme file in the archive, they say exactly what files they use:
> > *"This directory contains the adjusted land station data and metadata used in the Old Analysis;*
> > *data from GHCN-Monthly version 3.2.2.*
> > *ghcnm.tavg.v3.2.2.20150116.qca.dat.gz; Adjusted station data*
> > *ghcnm.tavg.v3.2.2.20150116.qca.inv.gz; Inventory of adjusted stations"*
> > Those are standard issue, dated files. I download them every day. And the copies are in the archive.