Giant natural fluctuation models and anthropogenic warming

S. Lovejoy[1], L. del Rio Amador[1], R. Hébert[1], and I. de Lima[2]

[1]Physics, McGill University, Montreal, Quebec, Canada, [2]MARE–Marine and Environmental Sciences Centre, Department of Civil Engineering, University of Coimbra, Coimbra, Portugal

Abstract Explanations for the industrial epoch warming are polarized around the hypotheses of anthropogenic warming (AW) and giant natural fluctuations (GNFs). While climate sceptics have systematically attacked AW, up until now they have only invoked GNFs. This has now changed with the publication by D. Keenan of a sample of 1000 series from stochastic processes purporting to emulate the global annual temperature since 1880. While Keenan's objective was to criticize the International Panel on Climate Change's trend uncertainty analysis (their assumption that residuals are only weakly correlated), for the first time it is possible to compare a stochastic GNF model with real data. Using Haar fluctuations, probability distributions, and other techniques of time series analysis, we show that his model has unrealistically strong low-frequency variability so that even mild extrapolations imply ice ages every ≈1000 years. Helped by statistics, the GNF model can easily be *scientifically* rejected.

1. Introduction

Over the past decades, climate scientists have largely focused on proving the veracity of anthropogenic warming (AW) theory, developing ever more sophisticated numerical models and amassing more and more evidence. At the same time, the climate sceptics have rejected the models, have complained that the evidence is biased, and have promoted the theory that the industrial epoch warming is essentially a centennial-scale giant natural fluctuation (GNF). We are at a standstill. On the one hand, the sceptics are difficult to counter since up until now they have not proposed any explicit statistical model consistent with their claim. On the other hand, all scientific theories reach a point beyond which progress is only incremental and AW theory is no exception, whereas the International Panel on Climate Change (IPCC) Assessment Report (AR) 3 [*Intergovernmental Panel on Climate Change (IPCC)*, 2001] already stated that "There is new and stronger evidence that most of the warming observed over the last 50 years is attributable to human activities" [the AR4, *IPCC*, 2007, and AR5, *IPCC*, 2013] essentially upgraded their confidence levels to "likely" and then "extremely likely."

The GNF hypothesis was finally examined and statistically rejected with more than 99.9% confidence in *Lovejoy* [2014a] (hereafter L2014). This paper used some nonlinear geophysics theory combined with preindustrial multiproxy temperature reconstructions to estimate the probability distribution of centennial-scale global temperature GNFs. This statistical rejection of the GNF hypothesis has the advantage that it does not require any knowledge or assumptions about the relevant but complex physical processes [*Kondratyev and Varotsos*, 1995]. This is analogous to the case of medical testing, where without the need to understand any biology, properly designed statistical tests can reject ineffective medications or treatments.

A somewhat different approach was recently adopted in *Mann et al.* [2016], who—with the help of numerical models (and hence of extra climate physics assumptions)—estimated the likelihood that without anthropogenic warming, the year 2014 would be the hottest on record, concluding "that the recent record temperature years are roughly 600 to 130,000 times more likely to have occurred under conditions of anthropogenic than in its absence."

Although the conclusions of these two papers are mutually consistent, they differ in their statistical assumptions about the temperature correlation structure and in the nature of the probabilities of extreme fluctuations (the possibility and likelihood of "black swan" extreme fluctuations). While these differences do not alter their fundamental conclusions concerning AW and GNFs, they do have scientific implications for our understanding of natural climate variability. An important question is whether the residuals of temperature trends have weak or strong correlation structures; in mathematical terms, power laws versus exponential decorrelations (the latter include autoregressive and kindred processes). This issue is not academic since power law correlations are theoretically predicted as a consequence of the temporal-scale invariance of

the climate equations (see, e.g., the review in *Lovejoy and Schertzer* [2013]), and they imply that the atmosphere has a huge memory which can be exploited for monthly to decadal forecasts [*Lovejoy*, 2015; *Lovejoy et al.*, 2015].

Beyond potentially exciting scientific implications, there are also consequences for the statistical treatment of trend analysis uncertainties. Under the rubric "multidecadal oscillatory (or low-frequency) variations, (long-term) persistence, and/or secular trends" (AR5, chapter 2), the IPCC briefly discussed the issue, and several papers concerning the power law assumption were cited [*Koutsoyiannis and Montanari*, 2007; *Lennartz and Bunde*, 2009; *Mann*, 2011]. However, a little further (AR5, box 2.2), it was admitted that "The quantification and visualization of temporal changes are assessed in this chapter using a linear trend model that allows for first-order autocorrelation in the residuals"; i.e., they adopted the weak correlation assumption (also shared by *Mann et al.* [2016]). In order to attack the IPCC, the more sophisticated sceptics have latched onto this feeble justification (but this is not relevant to the attribution issue that is discussed in chapter 10). In particular, 2 years ago D. Keenan even made a submission to the UK House of Lords on the issue of uncertainty assumptions. In order to further attract public attention, on 18 November 2015, Keenan publically proposed a "climate contest" with a $100,000 prize (http://www.informath.org/Contest1000.htm). The contest aims to publicize trend uncertainties and to attack the IPCC.

Although at a general level, Keenan states that "The reliance on merely proclaimed assumptions, in statistical analyses of climatic data, implies that virtually all claims to have drawn statistical inferences from climatic data are untenable. In particular, there is no demonstrated observational evidence for significant global warming" (http://www.informath.org/AR5stat.pdf). However, since this statement follows pages of discussion of trend analysis and his contest is based entirely around trends, his (more precise) position seems to be that (a) trends are needed to establish AW, (b) that a trend without an uncertainty is meaningless, (c) that the IPCC failed to justify the weak correlation assumption in its uncertainty analysis; and (d) therefore that no trends have been adequately justified and hence that AW has not been proven. The focus on trends is ironic if only because L2014 shows that the GNF hypothesis can be rejected without any trend analysis so that point (a) and all that follow are irrelevant for the issue of AW (in May 2016, in a personal communication, Keenan acknowledged that although he was aware of L2014, he had not read it). Indeed, rather than trends, it suffices to analyze the probability distributions of changes—or (almost) equivalently—to determine their return periods [*Lovejoy*, 2014b].

We argue that Keenan concocted the first concrete stochastic model of how the warming since 1880 might have been generated by GNFs. Although his point about weak versus strong correlations may be valid, it is nearly inconsequential for AW; what his model reveals about sceptic misconceptions is much more significant. Thanks to Keenan's 1000 realizations of a "trendless statistical model, which was fit to a series of global temperatures" ($T_{init}(t)$ below), the GNF theory has finally been fleshed out and can be subjected to scientific criticism. This is the subject of this paper. For those wishing more information on the contest, see the supporting information.

2. Analysis of the GNF Models

According to Keenan's site, the $100,000 goes to the first person to correctly identify 900 of 1000 series of his random process. The exact details are in the supporting information and his site; in summary Keenan makes 1000 realizations of an initial trendless model of 135 points representing the global annual temperature anomaly since 1880, $T_{init}(t)$. He then randomly chooses an unspecified fraction of these and gives them trends of 1°C/century, i.e., either 0.01 or −0.01°C/yr. This second step yields $T(t) = T_{init}(t) + b(t)$, where $b(t) = B_n t$ and $B_n = -a$, 0, a (with $a = 0.01$°C/yr) is a three-state Bernoulli process with probability P_0 for 0 and $(1 - P_0)/2$ for each of the other two states (the equality of the probabilities for the ±0.01 trends is not stated but inferred; see the supporting information). We are told that "…the initial 1000 random series were obtained via a trendless statistical model, which was fit to a series of global temperatures." Although the contest started on 18 November 2015, there were various problems and Keenan relaunched the contest on 22 November and then—sometime after 30 November—quietly modified the winning criterion (see the supporting information).

Before evaluating Keenan's model, it is useful to recall the main points of L2014. The first part provided an estimate of the total warming from 1880 to 2004: 0.87 ± 0.11°C (very close to the AR5 estimate). The temperature

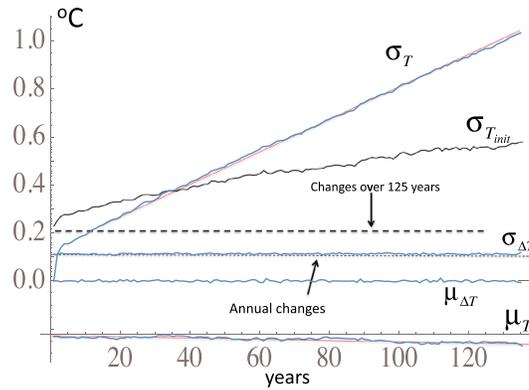**Figure 1.** The standard deviation of the model with added trends (σ_T) and the of the initial model ($\sigma_{T\text{init}}$), top two as a function of years since the beginning of the temperature series (1880), showing that Keenan's temperature processes, the initial and final processes (T_{init} and T), are nonstationary. Also shown are the mean temperatures, μ_T the mean increments $\mu_{\Delta T}$ (of the first difference), and the corresponding standard deviation $\sigma_{\Delta T}$ (third from the bottom); these are essentially indistinguishable for T and T_{init}. The observed value from both NASA's Goddard Institute for Space Studies (GISS) temperature series or using multiproxy preindustrial temperature reconstructions is shown as a fine dashed line (bottom). Finally, red reference lines (with value $\mu_T(t=0) \approx -0.23°C$) and slopes decreasing by 0.02°C/century (bottom) and the standard deviation increasing by 0.68°C per century (top) are shown. Also shown for reference (upper thick dashed line) is the typical temperature standard deviation from one 125 year period to the next as determined by preindustrial multiproxy temperature reconstructions (1500–1900 A.D.; see L2014).

was not regressed against time but rather against the logarithm of the CO_2 concentration ($\log CO_2$), which is proportional to the theoretical CO_2 radiative forcing. The only originality of this method was that $\log CO_2$ was used as a linear surrogate for all the anthropogenic forcings—including the difficult to quantify ones such as the effects of aerosol forcings. This was justified on the basis of the tight historical relationship between economic activity, CO_2 emissions, and anthropogenic effects. The issue of strong versus weak correlations in the residuals was avoided by the use of three different global temperature series. The estimated uncertainty was the *scientific* uncertainty involved in the different series constructions. This first part was only used to estimate the amplitude of the AW, no more.

The key second part of L2014 used preindustrial probability distributions of temperature changes as estimated from three disparate multiproxy reconstructions from 1500 to 1900 A.D. (see Figures 1–3 and the supporting information). With this probability distribution, it is possible to directly determine the probability of any given global temperature change over any time interval up to 125 years. This not only avoids any estimates of trends (and numerical models) but also avoids any attribution assumption since only preindustrial data are used.

Although Keenan seems to have used a stochastic model with some realism, his model is in fact quite unrealistic and this is easy to see. First note that although it may not be trivial to identify which realizations have added trends (i.e., which have the $b(t) \neq 0$ needed to win the contest), it is easy to deduce many of the statistics of the initial process. For example, by making histograms of the trends (estimated in various ways—for this purpose it makes little difference—see the supporting information), we find $P_0 = 0.54 \pm 0.016$, and squaring and averaging over the 1000 realizations (indicated by "<.>") gives $\langle T_{\text{init}}^2(t) \rangle = \langle T^2(t) \rangle - (1-P_0)a^2t^2$. It is now sufficient to determine the mean and the standard deviations of the model temperature as functions of time and of the first differences of the temperature (the increments), see Figure 1. The figure shows that although the average temperatures (μ_T) are roughly constant ($\mu_T = <T> \approx <T_{\text{init}}> \approx -0.23°C$), on the contrary, the standard deviations σ_T and $\sigma_{T\text{init}}$ *grow nearly linearly with time* (the top curve grows at a rate $a\sqrt{1-P_0} = 0.68 \pm 0.01°C/\text{century}$; see the supporting information). Before discussing this strongly unrealistic behavior, note that the increments of the processes (the first differences of Keenan's series) are on the contrary apparently stationary, with means $\mu_{\Delta T} \approx 0$ and with fairly realistic standard deviations $\sigma_{\Delta T} \approx \pm 0.11°C/\text{yr}$ corresponding to a "typical" year-to-year variation in the globally averaged temperature ($T(t)$ or $T_{\text{init}}(t)$) that give virtually identical results; NASA's Goddard Institute for Space Studies (GISS) temperature series gives $\pm 0.109°C$ [*Lovejoy et al.*, 2015], see Figure 1.

How do we know that σ_T and $\sigma_{T\text{init}}$ are aberrant? Over periods of 135 years, we find $\sigma_T \approx \pm 1.07°C$ and $\sigma_{T\text{init}} \approx \pm 0.58°C$ (the "±" is to remind us that this is a standard deviation, and it is a measure of the typical variation over a 135 year period; Figure 1). Imagine now n successive 135 year periods and consider the extreme case where the model correlations—which are strongly positive up to 135 year scales (see the supporting information)—completely disappear at larger scales so that successive 135 year periods are statistically
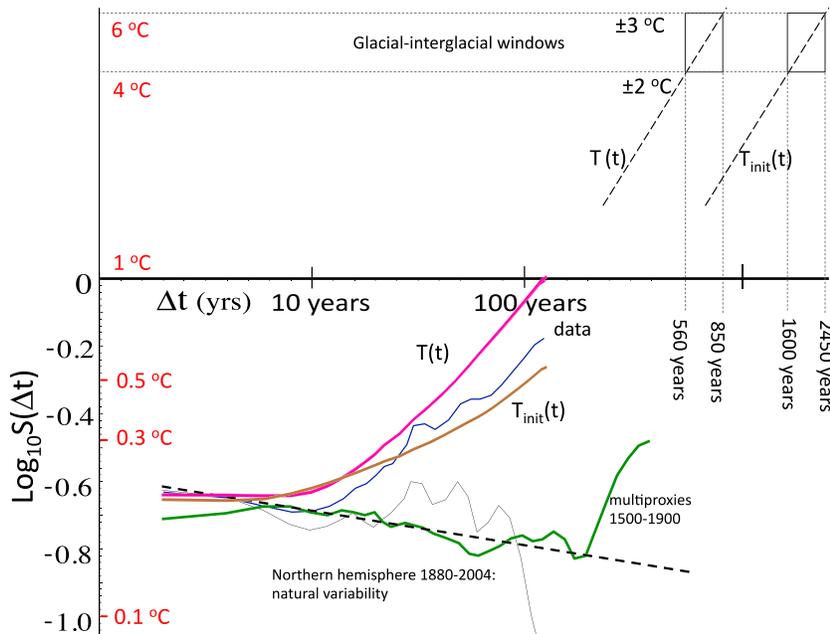
Figure 2. The RMS Haar structure functions of Keenan's model (top; magenta and brown for T and T_{init}, respectively) and of the average of three global surface data sets (the second from the top; blue, taken from L2014). Also shown for reference are $S(\Delta t)$ of the average of the three preindustrial temperature multiproxies analyzed in L2014 (green) along with the residuals with respect to a linear regression of the three 1880–2004 temperatures against $\log CO_2$ (thin black line). The fluctuations decrease roughly with exponent -0.1 (dashed line), for Gaussian processes, corresponding to a (statistically stationary) fractional Gaussian noise (fGn) process, not a (nonstationary) fractional Brownian motion (fBm) process as assumed by Keenan. All results were multiplied by a "canonical" factor of 2 for "calibration." This means that over the part of the curve that is increasing with Δt, the result is very close to the usual difference fluctuation. For example from the graph we see that typical (i.e., RMS) temperature differences at century scales are ≈1°C and 0.5°C (T and T_{init}) at the extreme large Δt. Also shown (top right) is the extrapolation of Keenan's models to longer time scales (see the supporting information). The model predicts typical temperature fluctuations of ±2°C at 560 and 1600 years and ±3°C at 850 and 2450 years (rectangles for T and T_{init}, respectively). Since going in and out of an ice age is a change of roughly this order, this is the models' prediction for the glacial-interglacial window. The time scales estimated from paleodata are roughly 50–100 times longer [*Lovejoy and Schertzer*, 1986].

independent of each other. In this case, over long time scales with n 135 year segments (periods $135 \times n$ years long), the temperature would perform a kind of random walk. For T and T_{init}, respectively, the total temperatures change would typically be $n^{0.5} \times 1.07$ and $n^{0.5} \times 0.58$. Going into and out of an ice age is a variation of roughly 4 to 6°C (corresponding to ±2 to ±3°C) so that for example changes of 4°C would typically occur every $135 \times (4/1.07)^2 = 1900$ years and $135 \times (4/0.58)^2 = 6400$ years for T and T_{init}, respectively, compared to 50–100 kyr as deduced from (much) paleoevidence. If instead of assuming that correlations in the fluctuations at scales larger than 135 years vanish, we assume that the same (positive scaling correlations) that holds at scales smaller than 135 year scales continues to hold at the longer time scales, we obtain the slightly lower estimates of 560 and 1600 years for a ±2°C change. This is shown graphically in Figure 2 as the "glacial-interglacial window"; it is 50–100 times shorter than the 50–100 kyr scales determined by paleodata.

Alternatively, the lack of model realism can be seen by noting that the typical empirical temperature difference at 125 year periods is about ±0.20 ± 0.01°C (L2014) so that at this time scale, the model T is roughly 5 standard deviations too variable and T_{init} is roughly 3 standard deviations too variable (for a Gaussian this has less than a 0.2% probability; see below).

Another way to understand this extreme low-frequency model variability is to consider the fluctuations in temperature $\Delta T(\Delta t)$ as a function of time scale Δt. While it is usual to define a fluctuation as a difference, for technical reasons, here we must use a generalization, the most convenient of which are called "Haar fluctuations" (the supporting information) [e.g., L2014]. These define $\Delta T(\Delta t)$ as the difference in the mean of $T(t)$ over the first and second half of an interval of length Δt (when the mean fluctuations increase with Δt,
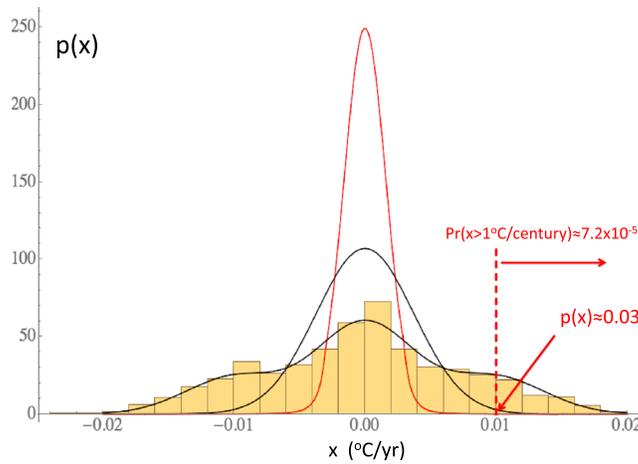
Figure 3. A comparison of the probability density of temperature changes (expressed as trends by dividing by the series length; in units of °C/yr) from $T(t)$ (yellow bars, with smooth bottom black curve being the three-Gaussian mixed distribution fit with normalization parameter $A = 0.00373$ and $P_0 = 0.553$, see supporting information). The top black line is the density for the T_{init} model, a pure Gaussian with standard deviation $A = 0.00373$. Also shown (red) is the empirical preindustrial probability density of temperature changes over 125 year periods (as empirically estimated in L2014). The dashed line indicates the 1°C/century imposed Bernoulli process trend, $p(x)$ is the empirical density, and $Pr(x > 1°C/century)$ is the cumulative probability of a fluctuation exceeding 1°C/century.

these are close to the usual differences and can be interpreted as such in most of analyses below). Figure 2 shows the result when the root-mean-square (RMS) fluctuations $(S(\Delta t) = \langle \Delta T(\Delta t)^2 \rangle^{1/2}$, are plotted for Keenan's model (for $T(t)$ and $T_{init}(t)$); for the NASA GISS global temperature series since 1880; for their residuals when regressed against $\log CO_2$; and finally, for the RMS fluctuations of three preindustrial reconstructions (multiproxies).

This figure helps clarify the misconceptions behind the GNF models. First, note that over the period of 1880–2014, the mean $S(\Delta t)$ of the T and T_{init} processes bound the data curve. The problem is thus not the comparison with the industrial epoch but rather with the preindustrial variability that according to the model should be a realization from the same stochastic process. However, the empirical preindustrial fluctuations (green) have much less centennial-scale variability. Indeed, the figure shows that at centennial scales, $S(\Delta t)$ from $T(t)$ and $T_{init}(t)$ are already, respectively, 6 and 4 times too strong. In addition, when one removes the $\log CO_2$ inferred anthropogenic part, the $S(\Delta t)$ (black curve) for the residual natural variability is indeed close to the preindustrial multiproxy $S(\Delta t)$, thereby statistically confirming the anthropogenic attribution.

Yet another way to judge the extreme lack of the model realism is to consider the probabilities of model trend x, here estimated as the difference between the first and last value of the series divided by the length $(x = (T(135) - T(1))/134)$. Figure 3 shows these temperature difference trend probabilities compared with the empirical probability densities $(p(x))$ of preindustrial differences over 125 year periods deduced in L2014 (red). As discussed in L2014 this $p(x)$ takes into account extreme black swan-type events; its tail (i.e., large x) falls off as $p(x) \approx x^{-(q_D+1)}$ (in the figure the estimated empirical value $q_D = 5$ was used; L2014 finds that the data are convincingly bounded between $q_D = 4$ and $q_D = 6$). When compared with the much more rapid (exponential) probability falloff of the usual Gaussian distribution, the empirical cumulative probability $Pr(x > 1°C/century)$ $= 7.2 \times 10^{-5}$ is more than 1000 times larger. The red curve thus implies such strong temperature changes that they would normally be considered outside the range of any "normal" models; [*Taleb*, 2010] popularized such outlier events under the term "black swans".

One way of quantifying the extreme nature of the GNF models is to consider what is the largest "acceptable" GNF that could be generated before the GNF hypothesis would be rejected at the (conventional) 2 standard deviation level (≈2.2%); smaller GNF's would be considered "typical". We find that for $T(t)$ and $T_{init}(t)$, this level corresponds to changes of, respectively, 1.99°C and 1.01°C over 135 years, whereas empirically, it is only 0.43°C. In this sense, Keenan's $T(t)$ and $T_{init}(t)$ models would both frequently generate GNFs comparable or larger than the industrial epoch warming even though the actual empirical probability of 1°C change in 135 years is ≈0.0003.

Keenan's stochastic model is the first to concretely implement the GNF hypothesis to explain the warming. In it, every 135 year period has potentially enormous temperature variations. If our only knowledge of the climate was the temperature series since 1880, the model could be plausible. However, we *do* have massive amounts of information about preindustrial temperatures and we *do* know about increasing CO_2 levels as well as about the relevant radiative transfer theory. Since a large body of climate knowledge exists we can confidently *scientifically* rule out Keenan's models as valid physical models.

3. Trend Uncertainties and GNFs

By analyzing the probabilities of temperature changes over centennial scales (or equivalently their return periods) L2014 avoided trend analysis, and the necessity of making weak/strong (exponential/power law) assumptions was mostly avoided. The qualifier "mostly" is only because L2014 did make mild use of the power law decorrelation hypothesis in order to extrapolate the empirical probabilities from 64 to 125 year intervals. In addition, nonlinear geophysics theory shows that power law correlations are generally associated with power law probability decays so that such black swan extreme events are theoretically expected (and can at least be used as bounds).

While trend uncertainties may be (almost) irrelevant to the issue of AW, Keenan's model does nicely illustrate the effect of various assumptions. For example, if one uses standard regressions (assuming weak correlations), one can only correctly guess 856 ± 9 correct trends, while using the alternative strong hypothesis, it is possible to obtain 893 ± 9 correct solutions (very near the 900 needed for winning the contest). The supporting information shows that this can be understood since to a good approximation $T_{init}(t)$ is a (long-range-dependent) fractional Brownian motion process. Significantly, such processes automatically generate random trends. Interestingly, if one simply uses differences to estimate the trends (as in Figure 3), one already obtains 877 ± 9 correct responses. The significance of the uncertainty assumptions is also brought out when estimating GNF trend exceedance probabilities $Pr(x > 1°C/century)$. In the T_{init} model with the strong correlation assumption, $Pr(x > 1°C/century) = 0.28\%$, whereas using the standard regression (weak) assumption, $Pr(x > 1°C/century) = 0.63\%$, a factor of 2.3 larger (ironically, this implies that the IPCC trends were slightly under - not over - estimated).

4. Conclusions

For decades, climate scientists have concentrated on proving the anthropogenic warming hypothesis, elaborating ever more complex numerical models. At the same time, the sceptics have simply rejected the models and continued loudly claiming that the warming is no more than a giant natural fluctuation (GNF). Although some have spent considerable effort attacking the statistical methods used by climate scientists, they have systematically refrained from proposing a concrete statistical model that might support their claims.

D. Keenan finally proposed such a model in November 2015, in the form of a challenging mathematical, statistical brain teaser with a $100,000 prize. By dressing it up as a problem in climate science, the sceptics' GNF theory has finally been made explicit in the form of 1000 realizations of a trendless statistical model, which was fit to a series of global temperatures. Irrespective of the details of the computer code and of any official solution that Keenan promises to unveil in November 2016, he has already allowed us to dismiss the sceptics' GNF hypothesis. Not only are his model's centennial-scale global temperature fluctuations far larger than the data allow (with probabilities hundreds of times too large) but in addition, mild model extrapolations imply that we should expect that ice age-scale temperature changes roughly every 500–1500 years rather than every 100,000.

Yet in spite of his model's lack of realism, Keenan has latched onto a kernel of truth with respect to one point: that the global temperature anomalies *do* have strong long-range statistical dependencies and this is important in climate *science* (they are also important in climate *statistics* although less so!). In particular, these long-range dependencies imply that the atmosphere has a huge memory that can be exploited for monthly, seasonal, and decadal ("macroweather") forecasting [*Lovejoy et al.*, 2015]. Unfortunately, for Keenan, this important fact—which is a consequence of the temporal-scale invariance of the dynamics (and hence is also a feature of the numerical general circulation models)—is not of much relevance to the scientific question: is the warming anthropogenic? Fortunately, for the rest of us, the giant natural fluctuation hypothesis has finally been subjected to scientific criticism.

References

Intergovernmental Panel on Climate Change (IPCC) (2001), *IPCC Third Assessment Report, vol. 1. Climate Change 2001. The Scientific Basis*, 881 pp., Cambridge Univ. Press, Cambridge, U. K., and New York.
Intergovernmental Panel on Climate Change (IPCC) (2007), Climate Change 2007: The Physical Science Basis, in *Contribution of Working Group I to the Fourth Assessment Report of the Intergovernmental Panel on Climate Change*, 996 pp., Cambridge Univ. Press, Cambridge, U. K., and New York.

Intergovernmental Panel on Climate Change (IPCC) (2013), Climate Change 2013: The Physical Science Basis, in *Contribution of Working Group I to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change*, Cambridge Univ. Press, Cambridge, U. K., and New York.

Kondratyev, K. Y., and C. A. Varotsos (1995), Volcanic eruptions and global ozone dynamics, *Int. J. Remote Sens.*, *16*(10), 1887–1895.

Koutsoyiannis, D., and A. Montanari (2007), Statistical analysis of hydroclimatic time series: Uncertainty and insights, *Water Resour. Res.*, *43* W05429, doi:10.1029/2006WR005592.

Lennartz, S., and A. Bunde (2009), Trend evaluation in records with long term memory: Application to global warming, *Geophys. Res. Lett.*, *36* L16706, doi:10.1029/2009GL039516.

Lovejoy, S. (2014a), Scaling fluctuation analysis and statistical hypothesis testing of anthropogenic warming, *Clim. Dyn.*, *42*, 2339–2351, doi:10.1007/s00382-014-2128-2.

Lovejoy, S. (2014b), Return periods of global climate fluctuations and the pause, *Geophys. Res. Lett.*, *41*, 4704–4710, doi:10.1002/2014GL060478.

Lovejoy, S. (2015), Using scaling for macroweather forecasting including the pause, *Geophys. Res. Lett.*, *42*, 7148–7155, doi:10.1002/2015GL065665.

Lovejoy, S., and D. Schertzer (1986), Scale invariance in climatological temperatures and the local spectral plateau, *Ann. Geophys.*, *4B*, 401–410.

Lovejoy, S., and D. Schertzer (2013), *The Weather and Climate: Emergent Laws and Multifractal Cascades*, 496 pp., Cambridge Univ. Press, Cambridge.

Lovejoy, S., L. del Rio Amador, and R. Hébert (2015), The ScaLIng Macroweather Model (SLIMM): using scaling to forecast global-scale macroweather from months to decades, *Earth Syst. Dynam.*, *6*, 1–22, doi:10.5194/esd-6-1-2015.

Mann, M. E. (2011), On long range dependence in global surface temperature series, *Clim. Change*, *107*, 267–276.

Mann, M. E., S. Rahmstorf, B. A. Steinman, M. Tingley, and S. K. Miller (2016), The likelihood of recent record warmth, *Sci. Rep.*, *6*, 6, doi:10.1038/srep19831.

Taleb, N. N. (2010), *The Black Swan: The Impact of the Highly Improbable*, 437 pp., Random House, New York.