

THE END-TO-END ATTRIBUTION PROBLEM: FROM EMISSIONS TO IMPACTS

DÁITHÍ A. STONE¹ and MYLES R. ALLEN²

¹*Departments of Physics (AOPP) and Zoology, Clarendon Laboratory, University of Oxford, Parks Road, Oxford OX2 6BD, U.K.; E-mail: stoned@atm.ox.ac.uk*

²*Department of Physics (AOPP), Clarendon Laboratory, University of Oxford, Parks Road, Oxford OX2 6BD, U.K.*

Abstract. When a damaging extreme meteorological event occurs, the question often arises as to whether that event was caused by anthropogenic greenhouse gas emissions. The question is more than academic, since people affected by the event will be interested in recurring damages if they find that someone is at fault. However, since this extreme event could have occurred by chance in an unperturbed climate, we are currently unable to properly respond to this question. A solution lies in recognising the similarity with the cause-effect issue in the epidemiological field. The approach there is to consider the changes in the risk of the event occurring as attributable, as against the occurrence of the event itself. Inherent in this approach is a recognition that knowledge of the change in risk as well as the amplitude of the forcing itself are uncertain. Consequently, the fraction of the risk attributable to the external forcing is a probabilistic quantity. Here we develop and demonstrate this methodology in the context of the climate change problem.

1. Introduction

Human and environmental infrastructures are designed to handle a range of weather events. For logistical reasons, however, these infrastructures cannot handle all possible events, and so there exists a non-negligible probability of an 'extreme' event occurring beyond the tolerated range, possibly with large impacts. For instance, in August 2003 the main concern in Western Europe was not the mean climatological temperature for that region during that month but that in this particular August the temperature was well above normal and in fact above the tolerance levels of some national health systems (Institut de veille sanitaire, 2003).

The usual reaction to an extreme weather event is to review the response measures in place, including the ability to predict the event. The goal is to use such review to improve the preparedness for the next similar extreme event. With the prospect of a changing climate responding to anthropogenic activities, however, there is another plausible reaction: we may be able to alter the likelihood of the event occurring again in the future. Of more direct relevance, however, is the point that past anthropogenic influence may have altered the likelihood of the event occurring now. Consequently, anthropogenic activities may be partly at fault for aiding the occurrence of the extreme event, perhaps even in a legal context (Grossman, 2003).

In a climate changing in response to some external forcing, the frequency and/or intensity of extreme events may change. Consequently, after an extreme event we may be able to attribute its occurrence to the external cause. For instance, with respect to anthropogenic emissions of greenhouse gases, we may ask the question “Was this event caused by that external forcing?” Unfortunately, this question is unanswerable under the current climate change attribution methods outlined in Mitchell et al. (2001). The problem lies in the fact that these methods assume a deterministic link between the forcing and the response in question. While this may be appropriate for the trend in global mean temperature, it is not so clearly applicable to the occurrence of a single local event that could easily have occurred anyway in the absence of the forcing. Thus we cannot say that some long term forcing ‘caused’ a short transient event. In the end, though, if the probability of that event occurring has changed as a result of the forcing, we should still be able to consider the attribution question.

A solution to this problem is to look at the risk of the event occurring, rather than the occurrence of the event itself (Barsugli et al., 1999; Palmer and Räisänen, 2002). Since the risk represents statistical information on the occurrence of the event, it can be viewed as being related to the forcing in a deterministic sense. In a way, we are defining risk to be the equivalent of the ‘climate’ while the actual occurrence of the event is the equivalent of the ‘weather’. The implication of this is that risk, just like climate, can be quite predictable since it depends only on knowledge of future external forcing, while event occurrence and weather are rather less predictable since they depend strongly on the initial conditions. Thus, we may ask the question instead “How has the risk of this event occurring changed as a consequence of that external forcing?” (Allen, 2003). We should point out here that strictly speaking we are actually considering the probability of a harm occurring and not the risk itself. We will refer to it as the risk however, under the assumption that the harm is independent of the external forcing.

The purpose of this paper is to transfer the measure of attributable risk from the epidemiological field to the climate change problem. This can also be viewed as a formal expansion of the probabilistic approaches developed in Barsugli et al. (1999) and Palmer and Räisänen (2002) to include an assessment of the change in risk and to include uncertainty as part of that assessment. In these earlier approaches, the probabilities are assumed to be known precisely, whereas here we explicitly recognise that the probabilities are uncertain since they are estimated from finite samples and the individual samples themselves are only approximations of the true system. In Section 2 we develop the method of attributing changes in risk, first in a case where we assume that the statistical properties of the climate system are known perfectly and then in a case where our knowledge is imperfect. We proceed in Section 3 with a test of the method in a simple nonlinear dynamical system exhibiting chaotic behaviour. Section 4 then contains an example exercise of the full methodology required for attributing weather events. We discuss the results in Section 5.

2. Method

2.1. THE PERFECT WORLD CASE

First, we consider the simple case in which the properties of the system are known exactly. We may then define a scalar, $p_0 = \psi(q_0)$, to be the probability of some event occurring in our system ψ in some reference state forced constantly with q_0 . p_1 is the corresponding value in some new state produced through a constant external forcing q_1 . Implicit in this formulation is the assumption of the existence of a well-defined climate attractor, such that p_0 and p_1 have single defined values. In practice these probabilities will usually be estimated from simulations with numerical models starting from a specified historical state, implying that this assumption can be relaxed somewhat to the condition that the evolution of the system from that historical state to the current reference and perturbed distribution of possible states is well defined.

In this definition we are also assuming a step function for our damage function, i.e. either the event does or does not occur. While this is a simplification, it is nevertheless a useful view point for the attribution problem. For instance, while a larger flood may cause more damage than a smaller flood, the concern for the attribution problem is the fact that the house was flooded in the first place.

An increase in the risk of the damaging event occurring in the externally forced state, as compared to the reference state, can then be attributed to the change in external forcing $\Delta q = q_1 - q_0$. This is given by taking the fraction of the new risk that cannot be attributed to the random probability (p_0) of the event occurring in the reference state. That is,

$$\text{far} = \frac{p_1 - p_0}{p_1} \quad (1a)$$

$$= 1 - \frac{p_0}{p_1} \quad (1b)$$

where the far is the change in risk attributable to the external forcing. This quantity is used extensively in the epidemiological field (Greenland and Rothman, 1998), for example in determining the relation between smoking and lung cancer. It is variously referred to in that field as the ‘excess fraction’, ‘attributable proportion’, ‘attributable risk’, and ‘etiologic fraction’. We adopt the term ‘fraction attributable risk’ (far).

Strictly speaking, this quantity is only valid if the events are repeatable in real time, that is if occurrence of an event does not preclude re-occurrence. While this is clear for meteorological events, it may not follow for their impacts where the damage may remove the victim from further possibility of damage. Accounting for this factor involves modifying the p_0 term such that $\text{far} = 1 - \frac{p_0}{p_1}(1 - c)$. Here $c = \frac{p_1 - p_0}{1 - p_0}$ is the fraction of the events that would have occurred normally in the absence of forcing but which will now be pre-empted and instead will be caused

by the new forcing. In the context of this paper, however, we mainly have rarer extreme events in mind, and so we will consider this c term to be negligible or that the events are repeatable in real time.

This fraction of attributable risk can be interpreted as follows. If the event never occurred in the reference state, but does occur in the new forced state, then $\text{far} = 1$, implying that all of the risk of this event occurring can be attributed to the forcing. If the probability of the event occurring does not change between states, then $\text{far} = 0$: there was no change, therefore none can be attributed to the forcing. If the risk of the event occurring decreases, then $\text{far} < 0$. In this case the far can range down to $-\infty$ (as $p \rightarrow 0$), making it difficult to interpret the result, other than recognising that no attribution can be made to an *increase* in the risk of the event occurring, since the risk did not increase. This measure of course could be redefined in this case to measure the attribution of a decrease in the risk by exchanging p_0 and p_1 in Equation (1), as is done in the epidemiological field in the definition of ‘preventable fractions’.

2.2. ACCOUNTING FOR UNCERTAINTY

Of course, in the real world, we cannot know either p_0 or p_1 exactly. Incomplete and finite observations, uncertainty in our knowledge of the external forcing, and limitations in our understanding of the system all limit the accuracy of our result. To account for this, we must consider these variables as being probabilistic, with some uncertainty in their estimation. In the epidemiological case it is possible to estimate both the value and uncertainty in these quantities through observational studies employing large cohorts. In the climatological case, however, we only have one observable sample for the forced state and none for the reference state. Thus we must turn to models as a proxy for the climate system, employing large ensembles of experiments with these models spanning the ranges of our uncertainties in order to estimate the relevant probability density functions (PDFs). While statistical models may be useful in some cases, various nonlinear relationships between observations, modelling, and forcing imply that the general approach should be with dynamical models.

The first source of uncertainty arises from finite sampling in our model experiments. Thus we must reformulate our expressions in terms of random variables, which we denote as uppercase characters. With this modification, the estimated probabilities of an event occurring, given the information from a finite sample, become

$$P_0 = \psi(q_0) + E_{\psi}(q_0) \tag{2a}$$

$$P_1 = \psi(q_1) + E_{\psi}(q_1). \tag{2b}$$

$\psi(q_0)$ and $\psi(q_1)$ are now the statistical mean values and $E_{\psi}(q_0)$ and $E_{\psi}(q_1)$ are random variables with zero mean which sample the uncertainty arising from the

finite samples of model experiments in the reference and perturbed states. The distributions of $E_\psi(q_0)$ and $E_\psi(q_1)$ may take any form and may even differ from each other if the sampling strategies differ or if model variability is a function of the forcing. Because of this latter possibility it is important to write this component as a function of the forcing.

Limitations in our knowledge of the forcing on the system can be another source of uncertainty. In many cases this is uncertain even in the reference state (e.g. current natural greenhouse gas forcing). Therefore we must include an ensemble of forcing scenarios for both states. We can express this as a random variable Q :

$$Q_0 = q_0 + E_{q_0} \quad (3a)$$

$$Q_1 = q_1 + E_{q_1}. \quad (3b)$$

q_0 and q_1 become the expected scenarios while E_{q_0} and E_{q_1} are random variables with zero mean which sample about these mean scenarios. Once again, no restrictions exist on the distributions of these random variables in either the reference or forced states. These forcing scenarios can then be used as input for the model, from which we can estimate event probabilities as in Equation (2), with Q substituted for q .

Finally, the model is by nature always imperfect. The arising uncertainty can in principle be quantified by using a family of models which span the range of our uncertainty in the operation of the climate system (Allen and Stainforth, 2002). With this addition, the model ψ is interpreted instead as a random variable Ψ . This results in the final formulation of the risk in the reference and perturbed states:

$$P_0 = \Psi(Q_0) + E_\Psi(Q_0) \quad (4a)$$

$$P_1 = \Psi(Q_1) + E_\Psi(Q_1). \quad (4b)$$

In this formulation, the model Ψ , the forcing Q , and the consequence of limited sampling E_Ψ are all treated as random variables.

With this full probabilistic expression of the risk, we can then calculate the fraction attributable risk (FAR) as in Equation (1):

$$\text{FAR} = 1 - \frac{P_0}{P_1}. \quad (5)$$

Now, however, we are expressing the risk in a probabilistic form: the FAR is a distribution rather than a scalar quantity.

3. Example: The Lorenz System

A good test study for this technique is a variant on the Lorenz (1963) system proposed by Palmer (1999). Like the climate system it is a nonlinear dynamical system producing chaotic behaviour. The nonlinearity of this behaviour appears

much stronger than that of the climate system though, which makes the Lorenz (1963) system a good test for robustness of a method against the most extreme nonlinear behaviour that we may expect in the climate system. Furthermore, it is much simpler and this allows larger test ensembles than is possible with a dynamical climate model. This system is described by the three equations:

$$\frac{\partial x}{\partial t} = s(y - x) + q \cos a \quad (6a)$$

$$\frac{\partial y}{\partial t} = (r - z)x - y + q \sin a \quad (6b)$$

$$\frac{\partial z}{\partial t} = xy - bz \quad (6c)$$

in the space spanned by x , y , and z . s , r , and b are parameters. The last terms in the first two equations represent an external forcing on the system along the x , y plane of magnitude q and at angle a from the positive x -axis. When a is 140° and common values are used for the other three parameters, the response of the system to this forcing is counter-intuitive. Rather than respond along the direction of forcing, the centre of gravity of the system responds by moving almost perpendicularly to the forcing, that is along the $x = y$ axis (Palmer, 1999).

This behaviour is demonstrated in Figure 1. The projection of the PDF of the system along the $\frac{x+y}{\sqrt{2}}$ axis is shown. The averaging interval is one unit of time in the Lorenz (1963) system. The lag correlation function of this system in time is similar to that of many surface weather variables upon removal of the annual cycle, so these ‘daily’ units can be considered analogous to conventional days in the climate system. The plot is produced from a 10,000-day simulation of the model, comparable to 30 years of daily weather data such as the conventional 1961–1990 period. In the unforced case there is a 4% chance of a daily event exceeding 10 on this axis, but when the forcing increases to $q = 3$ the risk of the event doubles.

For the tests we assume that the ‘real’ Lorenz (1963) system is described exactly by Equation (6) with $s \equiv 10$, $r \equiv 28$, $b \equiv \frac{8}{3}$, and $a \equiv 140^\circ$. Furthermore, in the reference ‘climate’ $q_0 \equiv 0$ and in the perturbed ‘climate’ $q_1 \equiv 3$. We also suppose that the dynamics of the system, as described by Equation (6), are




Figure 1. The PDF of the projection of the Lorenz (1963) system along the $\frac{x+y}{\sqrt{2}}$ axis for unforced and forced cases.

known exactly. However, uncertainties from sampling, in forcing magnitude, and in model parameters will be introduced progressively. For convenience, impacts are considered to occur on days with $\frac{x+y}{\sqrt{2}} \geq 10$.

3.1. UNCERTAINTY FROM LIMITED SAMPLING

For the first test, we include only the uncertainty arising from a limited sampling. The various parameter values used are listed in Table I. We take $n_{\text{sample}} = 1000$ samples of $n_{\text{days}} = 10,000$ days each of the modelled system with varying initial conditions.

The PDFs of the risk of the magnitude 10 event occurring in both the unforced and forced states (P_0 and P_1) are shown in Figure 2. In both states, uncertainty arising from internally generated variability is small due to the use of reasonably long simulations. In the forced state this uncertainty is broader, consistent with the behaviour of a binomial distribution with an increased mean, but P_1 is still well defined. Consequently, the resulting plot of attributable risk is also well defined at FAR ~ 0.5 . This value implies that half of the risk of the magnitude 10 event occurring in the forced state is a consequence of the applied forcing. This is necessarily consistent with the evident doubling of the risk (indicated on the top axis of the lower panel of Figure 2 as $\frac{P_1}{P_0} = \frac{1}{1-\text{FAR}}$, which is the quantity examined by Palmer and Räisänen (2002)). Within the resolution of our sampling, there is a $<1\%$ chance that none of the change in risk is attributable to the forcing, while there is a 13% chance that the FAR is in fact less than one half. Use of shorter samples (n_{day}) would widen the estimated distributions such that we would be less confident in attributing the change to the forcing; on the other hand use of fewer samples (n_{sample}) would only have made the estimated distributions more bumpy and more poorly defined.

TABLE I
Parameter values for the various tests with the Lorenz (1963) system

| Parameter | Section 3.1 | Section 3.2 | Section 3.3 | Section 3.4 |
|---------------------|---------------|----------------------|-------------------------------------|-------------------------------------|
| s, S | 10 | 10 | $10(1 + \frac{1}{250}E_G)$ | $10(1 + \frac{1}{250}E_G)$ |
| r, R | 28 | 28 | $28(1 + \frac{1}{250}E_G)$ | $28(1 + \frac{1}{250}E_G)$ |
| b, B | $\frac{8}{3}$ | $\frac{8}{3}$ | $\frac{8}{3}(1 + \frac{1}{250}E_G)$ | $\frac{8}{3}(1 + \frac{1}{250}E_G)$ |
| q_0, Q_0 | 0 | $0 + E_G$ | 0 | $0 + E_G$ |
| q_1, Q_1 | 3 | $3 + \frac{1}{4}E_G$ | 3 | $3 + \frac{1}{4}E_G$ |
| a | 140° | 140° | 140° | 140° |
| n_{sample} | 1000 | 1000 | 1000 | 1000 |
| n_{day} | 10000 | 10000 | 10000 | 10000 |

E_G is a Gaussian random variable of zero mean and unit standard deviation.

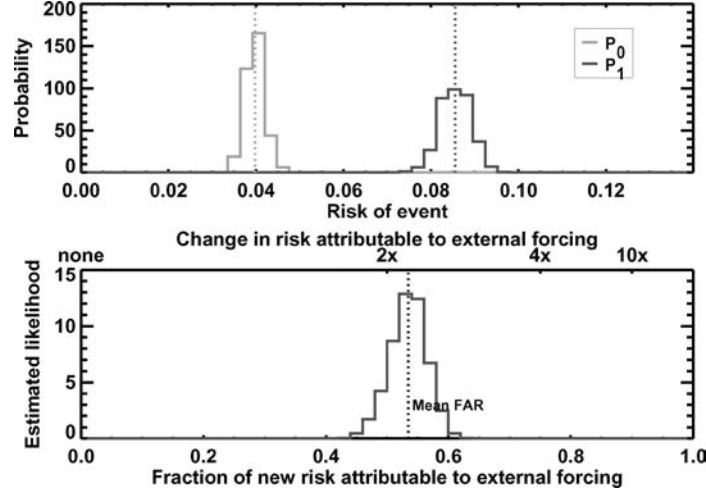


Figure 2. Top: the histogram estimated PDF of the risk of a $\frac{x+y}{\sqrt{2}} \geq 10$ event in the Lorenz (1963) system in unperturbed (P_0) and perturbed (P_1) cases including only uncertainty arising from limited sampling. Bottom: the fraction of the risk in the perturbed case attributable to the imposed external forcing (FAR). The top axis of the lower panel is the ratio between the risks in the perturbed and unperturbed cases, given as $\frac{P_1}{P_0} = \frac{1}{1-\text{FAR}}$. All values of FAR < 0 are pooled in the zero bin since there is nothing to attribute in these cases. See Table I for the parameter values used in these simulations.

3.2. INCLUDING FORCING UNCERTAINTY

For the second test, we add uncertainty in the forcing to the internally generated variability while still assuming a perfect model. The forcing for the reference state is still taken to be zero, but with a Gaussian distributed uncertainty of unit standard deviation. Thus we are replacing the constant q_0 with the random variable $Q_0 = E_G$, where E_G is a Gaussian distributed random variable of zero mean and unit standard deviation. For the forced state the forcing is taken to be both translated and better defined, as would be the case if we were comparing present day with non-industrial conditions; thus $q_1 = 3$ is replaced by $Q_1 = 3 + \frac{1}{4}E_G$.

The PDFs of the risk of the magnitude 10 event occurring in both of these states are shown in Figure 3. As expected the PDFs are noticeably wider than in the first test due to the addition of uncertainty in the forcing, especially in the reference state. The PDF of the FAR is correspondingly wider too. There is still a $< 1\%$ chance that none of the change in risk is attributable to the forcing, but there is now a 45% chance that the FAR is less than one half. In this example, recognising the uncertainty in the forcing is clearly important for the final result.

3.3. INCLUDING MODELLING UNCERTAINTY

For the third test, we once again ignore uncertainty in the forcing but add uncertainty in the model parameters to the internally generated variability. For simplicity, we

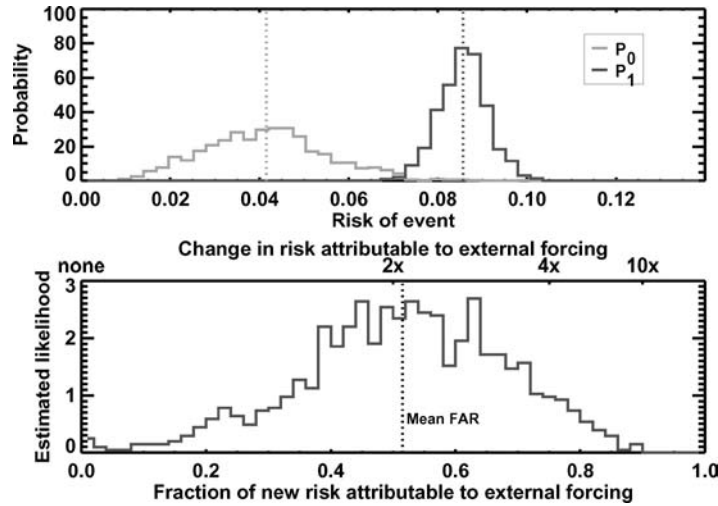


Figure 3. As in Figure 2 but including uncertainty in the forcing in both the perturbed and unperturbed cases as outlined in Table I.

assume the same uncertainty for all three model parameters. This uncertainty is characterised by a Gaussian distribution of standard deviation $\frac{1}{250}$ of the mean value.

The PDFs of the risk of the magnitude 10 event occurring in both the reference and forced states are shown in Figure 4. Since the model uncertainty is independent of the forcing, it has affected the two states equally, such that the PDF for the reference state is narrower than in Section 3.2 but the PDF for the forced state is




Figure 4. As in Figure 2 but including uncertainty in the model parameters in both the perturbed and unperturbed cases as outlined in Table I.

wider. The result is a narrower PDF of the FAR, although it is still considerably wider than when assuming a perfectly known model. In this case there is still a <1% chance that none of the change in risk is attributable to the forcing but a 41% chance that less than half is attributable.

3.4. INCLUDING FORCING AND MODELLING UNCERTAINTY

Finally, we add uncertainty in the three model parameters on top of the uncertainty in the forcing and internally generated variability. The PDFs of the risk of the magnitude 10 event occurring in both of these cases are shown in Figure 5 and are now even wider. Comparable to the previous two tests, there is a 43% chance that less than half of the change in risk is attributable to the forcing. However, there is now a non-negligible 2% chance that none of the change in risk is attributable. Accounting for all sources of uncertainty has stretched the PDF of the FAR considerably from Figure 2 to the extent that any statement concerning the attribution will have to be substantially more cautious.

As a probabilistic quantity, it may sometimes be more useful to view the FAR through its cumulative distribution function (CDF), as in Figure 6. From this plot we can see that if uncertainty from limited understanding of the dynamics and forcing were non-existent or neglected, we would say that a doubling of the risk is ‘likely’ according to the Intergovernmental Panel on Climate Change convention of a better than 2 in 3 chance. However, given some uncertainties in our knowledge, as assumed in preceding subsections, that possibility could only be stated as being of ‘medium likelihood’, that is between a 1 in 3 and 2 and 3 chance.




Figure 5. As in Figure 2 but including uncertainty in both the forcing and the model parameters in both the perturbed and unperturbed cases as outlined in Table I.

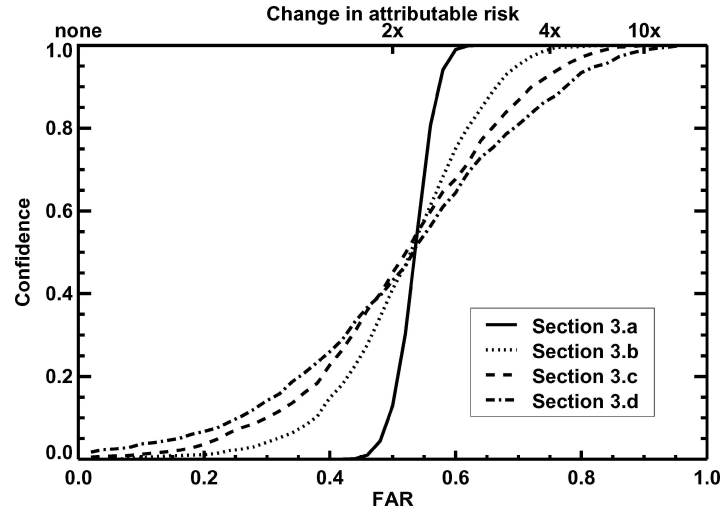


Figure 6. The cumulative distribution function of the FAR as estimated in Section 3. The horizontal axes are identical to those in Figure 2.

4. The End-to-End Procedure

The probabilistic attribution procedure outlined in the preceding sections assumes that we have full knowledge of the nature of the external forcings on our system and that the only uncertainty we have about them is in their magnitude and representation in our model. This cannot be assumed in the case of the climate system, considering the large array of potential forcings acting on it. Consequently, we must establish that our understanding of the forcings is consistent with observed climate change and not consistent with other plausible scenarios. In this section we present a methodology for a full probabilistic end-to-end attribution investigation, again using the Lorenz (1963) system.

First we need to establish that our understanding of the external forcing and system parameters is consistent with the observed change. It would be preferable to do this by looking at the measure of the extremes of the variable, with which we are ultimately concerned. However, beyond restrictions in modelling efficiency, the observational measurements of climate variables are generally too sparse to examine extremes in this manner. Thus the only option is to look at more robust measures of the physical quantity, like its mean and variance, and assume that verification of our understanding of the physics and forcings relevant to these measures applies to its extremes as well.

For this experiment we modify the experimental design of Allen and Stott (2003) to the following.

1. We generate a time history of pseudo observations of the system with the parameters used in Section 3.1. The forcing follows a scenario where it increases

- linearly from the reference value ($q_0 = 0$) to the perturbed value ($q_1 = 3$). This history covers a period of 30,000 Lorenz days, analogous to 100 years in the climate system, which are averaged to analogous decadal means. In addition, our observations of this change may be incomplete, so we have to include uncertainties, E_{obs} arising from this restriction. We choose a Gaussian distribution of zero mean and standard deviation 0.5 for this uncertainty, which we apply to the decadal averages of all three system variables.
2. We estimate a spatio-temporal response pattern using ensemble simulations with a model of the Lorenz (1963) system, but with the forcing and parameters as used in Section 3.4. The main effect of the inclusion of these uncertainties will be a larger uncertainty in our result. Only the single scenario described above (linear increase) is examined in this experiment, but our knowledge of this scenario is considered uncertain as in Section 3.4. The response pattern to this forcing scenario is estimated from 100 simulations.
 3. We estimate the internal variability of the system using a long ‘control’ simulation of the model, in this case 100 times longer than our observational ‘record’. The system parameters used are as in Section 3.1, and the forcing is maintained at the reference $q_0 = 0$. In this construction, it is implicitly assumed that the internal variability does not depend strongly on the system parameters and forcing, within the plausible range of values.

Time series for the observed history and the average scenario are shown in Figure 7.

To test our understanding of the relevant forcings and parameters we apply the optimal fingerprinting technique described in Allen and Stott (2003). This method relates the observations, \vec{y} , to m model scenario response patterns, \vec{x}_i where $i = 1, \dots, m$, through a linear regression equation:

$$\vec{y} = \sum_{i=1}^m (\vec{x}_i - \vec{v}_i) \beta_i + \vec{v}_0. \quad (7)$$

The observations are contaminated with sampling noise \vec{v}_0 , so this term must be included in the regression. Similarly, the model scenario response patterns are uncertain due to the limited sampling size used in their estimation, and this uncertainty is included in the regression through the \vec{v}_i noise terms. The goal is to estimate the regression coefficients β_i and their PDFs from this information. This is performed using the total least squares (TLS) algorithm described in Allen and Stott (2003).

In this example we only consider the one plausible scenario outlined above. The TLS algorithm produces a best estimate of the regression coefficient β of 1.03, with a 95% confidence interval of 0.83–1.23. The fact that this interval excludes $\beta = 0$ indicates that we have, at the 5% level, detected an effect of the forcing change on our observed history. Moreover, since $\beta = 1$ is within this range, our understanding of the forcing changes and their effects is consistent with the change in the observed variables.

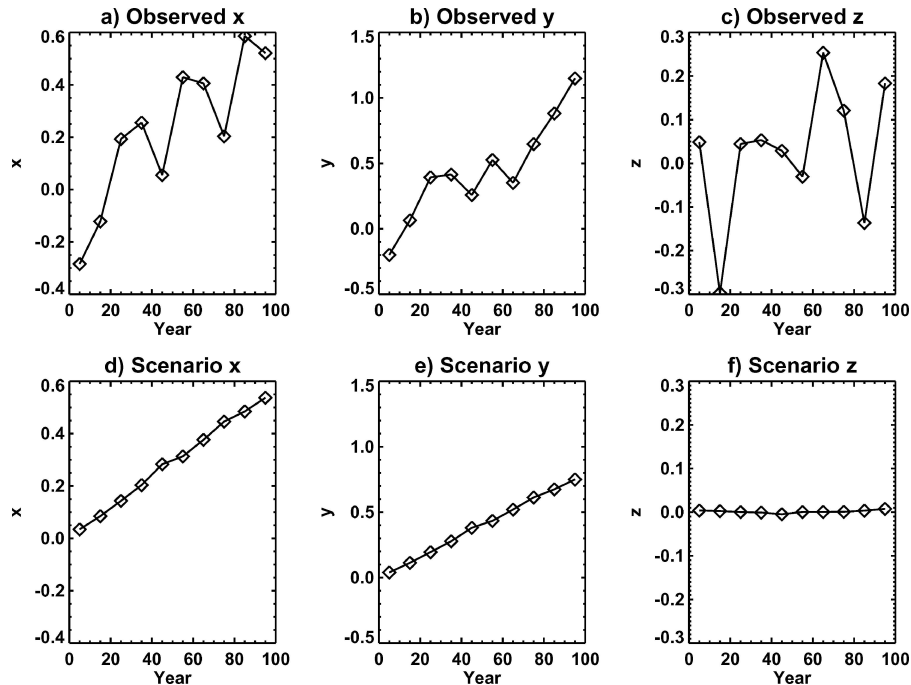


Figure 7. (a–c) The evolution of decadal averages of the ‘observed’ variables, including measurement uncertainty. Values are anomalies from the means of the control simulation. (d–f) The average evolution of the 100 model simulations in response to a presumed scenario of linearly increasing forcing.

Through this test we have established that the imposed forcing has indeed caused a change in the observed system (which of course is the case here by design). While this result only verifies our understanding of the relevant forcings and parameters in the context of the means of the variables we are examining, sampling restrictions force us to assume that this result holds for the extremes of the variables as well. Thus, to attribute the occurrence of an extreme event to this forcing we would now proceed with the experiment detailed in Section 3.4 and Figure 5, which is not repeated here. It should be noted though that there are two options for this second step. The first would be to use the models by themselves, as done in Section 3.4. The second would be to use the PDF of β arising from the observational constraints to scale the model output (Stott et al., 2004), the advantage being that our uncertainties would depend on the observational constraints and thus would be less dependent on the specific model formulation.

5. Discussion

We have transferred a simple technique for the attribution of sporadic events to external factors from the epidemiological field to the climate change problem and

demonstrated it with a simple nonlinear dynamical system. Implicit in this transfer was a recognition that the estimation of this attribution cannot be inferred from the single sample of the observational record, and therefore must be derived from a large ensemble of climate and impact model simulations representing the range of our uncertainties. The uncertainties in the model formulation are evidently difficult to quantify since the model itself is an approximation of the climate system, yet they must be estimated if we are to produce a useful result.

The heat wave in Europe in summer 2003 was given as motivation for this approach. For this problem the first step would be to verify that our chosen climate model, when forced with all of the external forcings as we understand them, can reproduce weather that is consistent with what has been observed. This step was demonstrated in Section 4 on a different problem. For the heat wave, an obvious measure would be summer mean temperature over Europe, since comparisons of measures of very rare extreme events would be unreliable. This step is necessary in order to determine whether the model is appropriate for the problem at hand. Using model simulations, we would then estimate the probabilities of the heat wave in climates with and without the anthropogenic forcing, and estimate the FAR from these. This full analysis of the 2003 heat wave is done in Stott et al. (2004).

A concern with this methodology is the selection effect. The motivation for this approach assumes a recent extreme event, but extremes of some form happen somewhere all the time. This lends itself to a bias toward mistaken attribution since cases which cannot be attributed will tend to be forgotten. While it is not clear how this bias can be solved, it can be minimised. For instance, Stott et al. (2004) examine the probability that a record heat wave, rather than the 2003 heat wave, occurred. Furthermore, they choose an established measure in mean temperature and a domain defined partly independently of the heat wave. Of course, ultimately a specific harm should be considered and its relationship to climate should be explicitly modelled as well.

In this probabilistic framework, all statements concerning the change in attributable risk must be qualified by the usual statistical interpretation. Thus in any statement attributing an event we would have to be able to state, for instance, that a given estimate of attributable risk was statistically significant from zero at some level α . This is equivalent to stating that a $1 - \alpha$ fraction of the PDF of the FAR must be above zero. From a practical point of view, however, there lies the question as to the value at which the FAR becomes substantial. A well constrained and statistically significant value of $\text{FAR} = 0.01$ will only be of academic interest. A natural threshold is $\text{FAR} = 0.5$, in other words when over half of the risk of a given event is attributable to the cause. Indeed, Grossman (2003) states that US courts tend to interpret a doubling in the risk as acceptable in civil tort cases. Thus a more practical evaluation would be the statistical significance over an attributable risk of 50%. In the end, of course, the critical value for the null hypothesis will depend on the problem at hand.

Finally, the question must be asked as to whether the detailed probabilistic treatment outlined in Section 2.2 is necessary in real world situations. Dessai and Hulme (2003) argue that the use of probabilistic information is dependent on the context. However, they only consider the prediction of future climate and not the attribution of current weather. An examination of Figure 5 indicates that in the case of attributable risk this probabilistic information is vital in assessing the robustness of the result. A PDF of the FAR as in Figure 5 contains useful information that can only be revealed in a probabilistic approach. From it we can be confident that some of the risk is attributable to the external factor, but we are rather uncertain as to the exact degree of this attribution. Naturally it is unlikely that we will ever be able to fully quantify the effect of all sources of uncertainty, but the unacceptable alternative is to assume absolute certainty in a scalar result. Of course, in the long run we would hope to attain more precise results, resembling Figure 2, in which case the uncertainty has little impact on the derived conclusions. In the meantime, however, statements attributing the occurrence of extreme weather events to external forcing must include appropriate assessments of their robustness.

Acknowledgements

The authors wish to thank David Anderson, Suraje Dessai, David Grossman, Bruce Hewitson, Jamie Kettleborough, Hugo Lambert, Pardeep Pall, Tim Palmer, Edmund Phillips, Prashant Sardeshmukh, Peter Stott, and Alberto Troccoli for stimulating discussions, and also Gary Yohe and two anonymous reviewers for helpful comments. Financial support for this research from the Department for Environment, Food, and Rural Affairs and from the Wellcome Trust is gratefully acknowledged.

References

- Allen, M.: 2003, 'Liability for climate change', *Nature* **421**, 891–892.
- Allen, M. R. and Stainforth, D. A.: 2002, 'Towards objective probabilistic climate forecasting', *Nature* **419**, 228.
- Allen, M. R. and Stott, P. A.: 2003, 'Estimating signal amplitudes in optical fingerprinting. Part I. Theory', *Clim. Dyn.* **21**, 477–491.
- Barsugli, J. J., Whitaker, J. S., Lough, A. F., Sardeshmukh, P. D., and Toth, Z.: 1999, 'The effect of the 1997/1998 El Niño on individual large-scale events', *Bull. Am. Meteorol. Soc.* **80**, 1399–1411.
- Dessai, S. and Hulme, M.: 2003, *Does Climate Policy Need Probabilities?*, Technical Report, Tyndall Centre Working Paper No. 34.
- Greenland, S. and Rothman, K. J.: 1998, 'Chapter 4: Measures of effect and measures of association', in Rothman, K. J. and Greenland, S. (eds.), *Modern Epidemiology*, 2nd edn., Lippincott-Raven Publishers, Philadelphia, USA, 737 pp.
- Grossman, D. A.: 2003, 'Warming up to a not-so-radical idea: Tort-based climate change litigation', *Colum. J. Envtl. L.* **28**, 1–61.

- Institut de veille sanitaire: 2003, 'Impact sanitaire de la vague de chaleur d'août 2003 en France. Bilan et perspectives', 25 November, http://www.invs.sante.fr/publications/2003/bilan_chaleur_1103/.
- Lorenz, E. N.: 1963, 'Deterministic nonperiodic flow', *J. Atmos. Sci.* **20**, 130–141.
- Mitchell, J. F. B., Karoly, D. J., Hegerl, G. C., Zwiers, F. W., Allen, M. R. and Marengo, J.: 2001, 'Chapter 12: Detection of climate change and attribution of causes', in Houghton, J. T., Ding, Y., Griggs, D. J., Noguer, M., van der Linden, P. J., Dai, X., Maskell, K. and Johnson, C. A. (eds.), *Climate Change 2001: The Scientific Basis, Contribution of Working Group I to the Third Assessment Report of the Intergovernmental Panel on Climate Change*, Cambridge University Press, Cambridge, UK, 881 pp.
- Palmer, T. N.: 1999, 'A nonlinear dynamical perspective on climate prediction', *J. Clim.* **12**, 575–591.
- Palmer, T. N. and Räisänen, J.: 2002, 'Quantifying the risk of extreme seasonal precipitation events in a changing climate', *Nature* **415**, 512–514.
- Stott, P. A., Stone, D. A. and Allen, M. R.: 2004, 'Human contribution to the European heatwave of 2003', *Nature* **432**, 610–614.

(Received 21 May 2004; in revised form 11 August 2004)