Global warming in a nonlinear climate - Can we be sure?

Tim Palmer, European Centre for Medium-Range Weather Forecasts, Reading, UK

Picture this. It's Sunday morning and you are standing on the first tee of your local golf club, ready for that all-important first drive of the day. Settling yourself over the ball, and freeing your mind of extraneous thoughts, you begin your swing. But, half way through, a big drop of rain hits you squarely on the top of your head! It's too late to abort the swing, and with your state of concentration momentarily disturbed, your golf club strikes the ball well off centre and it ends up in a thick patch of rough, 50 yards to the right of the tee. As you walk up to the ball, the rain drops begin falling more frequently. By the time you hit your second shot, which only succeeds in moving the ball into yet deeper rough, the rain becomes persistent. At this point, you think to yourself: wait a minute, they never said anything about rain on the weather forecast last night! And so you continue your round, cursing in equal measure, your inelegant golf swing, the miserable weather and those lousy weather forecasters. Eighteen torrid holes later, you return to the club house, soaked through and in a generally foul mood. Someone turns on the TV to catch the 1 O'Clock news. Top stories are the Prime Minister's commitment to tackle the problem of climate change, and a related report that the Government's Chief Scientific Advisor has proclaimed climate change as the most serious problem facing humanity - more serious even than the terrorist threat. You've never had a view on climate change - until today. You mutter under your breath "Nothing but hot air - those bloody weather forecasters can't even get tomorrow's weather right!"

Another climate sceptic is born!

The golfers among us will no doubt sympathise with the plight of our protagonist, but as physicists, should we sympathise with the logic of his final argument? The conventional scientific view is that we should not; weather forecasting is an initial value problem, while the global warming problem focuses on estimating how long-term weather statistics are affected by some prescribed climate forcing associated, say, with a doubling of atmospheric CO_2.

However, I will conclude that whilst the global warming problem indeed poses a severe threat to humankind - with a quantified risk that the Greenland and Antarctic ice sheets would melt in the coming century - the magnitude of global warming in the coming century is still very uncertain. These uncertainties are in fact linked to those in weather prediction, through an inability to simulate key kilometre-scale phenomena in global weather and climate models. This inability arises from lack of supercomputer capacity rather than scientific capability. I believe that significant reduction in uncertainty in forecasts of climate change will require investment in significantly greater computing resources than currently available. This investment will benefit weather forecast accuracy too, and will help link weather and climate forecasting more quantitatively. Like CERN, this is an investment which may be unaffordable at the national level and which in any case will benefit from international cooperation.

Uncertainty in Weather Prediction

Perhaps the one thing all physicists know about the weather is that it is chaotic - its evolution is sensitive to initial conditions. The prototype (Lorenz, 1963) chaotic model is given by the three equations:

$$\begin{aligned} \dot{X} &= -\sigma X + Y \\ \dot{Y} &= -XZ + rX - Y \qquad (1) \\ \dot{Z} &= XY - bZ \end{aligned}$$

For suitable values of the parameters (eg $r = 28$, $\sigma = 10$, $b = 8/3$) the model has a chaotic attractor (see background in Figure 1). As discussed below, this model doesn't describe weather at all; however, like the weather it is a nonlinear dynamical system.

Because of nonlinearity[1], the growth of uncertainty depends on the initial state. This dependence is illustrated in Figure 1; initial uncertainty is represented by a small circle, and the evolution of this uncertainty is estimated by integrating equation (1) from an ensemble of initial states on this ring. Depending on the starting conditions, the forecasts can be "very predictable", "somewhat predictable", "or very unpredictable".

Figure 1 highlights the problem facing weather forecasters every day. Is the weather evolving through a stable or unstable part of state space? If the former, then they can be confident about the forecast; if the latter, they can't be confident. Let's take a liberty with the Lorenz model and suppose that all states on the left hand regime of the attractor are "dry" states, and all states on the right hand regime are "rainy" states. Then the predictability of the weather forecast that our golfer recalled on the first hole of his doomed round was, perhaps, like that in the top right of Figure 1: the most likely evolution is for the weather to stay dry, there is a significant probability of rain. If the golfer had been aware of the full ensemble distribution, rather than its mode, he would have known there was a quantifiable risk of rain.

Can forecasters use this ensemble technique in reality to determine the predictability of their forecasts ahead of time? In practice this is a computationally challenging problem, as the state space of a contemporary weather forecast model is several millions of times larger than that of the Lorenz model (see below). Nevertheless, in recent years, with the development of supercomputers, ensemble weather forecasts are becoming commonplace (Palmer, 2000); these ensembles start from sets of initial conditions, each of which is consistent with the available weather observations (see Fig 2 for a realistic example of an exceptionally unpredictable weather event).

However, why are the initial conditions of weather forecasts uncertain? Can we reduce this uncertainty? At a superficial level it may appear that initial-condition uncertainty arises because the thermometers, barometers and satellite instruments that remotely sense the atmosphere, do not have perfect accuracy. However, that is only part of the reason. To understand this better, and to see the link with climate prediction, we need to discuss weather and climate forecast models in more detail.

Let us return to equation (1). Lorenz derived these from the equations for a thermally-convecting laboratory fluid (the two regimes in Figure 1 correspond to distinct modes of convective

[1] If F in $\dot{X} = F[X]$ is nonlinear, then dF/dX in $\delta\dot{X} = dF/dX\,\delta X$ must depend on X

Article available at http://www.europhysicsnews.org or http://dx.doi.org/10.1051/epn:2005202

▲ Fig. 1: The background shows the attractor of equation (1). The foreground shows finite-time evolutions of (1) from three different rings of initial conditions. These initial rings represent a measure of uncertainty in the initial state. The figure illustrates the fact that in a nonlinear system, growth of uncertainty is dependent on initial state. In the top left there is essentially no growth in uncertainty - hence a prediction in evolution from the left to right-hand lobe of the attractor can be made with certainty. In the top right, there is significant growth of uncertainty, and only probabilistic predictions of regime change can be made. In the bottom left, the growth of uncertainty is explosive.

overturning). The fundamental fluid equations are nonlinear partial differential equations (PDEs); they describe a potential continuum of interacting circulations.

The climate system shares with the laboratory fluid the property that it too is described by nonlinear PDEs - the atmosphere is also a thermally-forced turbulent system. In this respect, weather and climate prediction models are similar to the Lorenz model; they correspond to a reduction of some underlying PDEs to a finite set of finite difference equations with quadratic nonlinearities. Hence the state vector of the Lorenz model is the three-component $\mathbf{X}_L = (X, Y, Z)$, whilst the state vector of a weather forecast model

$$\mathbf{X}_W = (X_1, X_2, X_3 \ldots X_{10^7}) \qquad (2)$$

has as many as 10^7 components. These components can be thought of as representing the amplitudes of circulations of different spatial scale (eg associated with a spherical harmonic decomposition of the basic fluid variables). Hence X_1 describes a planetary-scale circulation feature in the atmosphere whilst X_{10^7} represents a circulation feature near the truncation scale of the model, some tens of kilometres in the horizontal for a weather forecast model, and some hundreds of kilometres for a climate model.

The accuracy to which a particular component of the atmospheric circulation can be simulated, depends on how close this component is to the truncation scale; at the truncation scale itself, the model's representation of the atmospheric circulation is obviously very poor indeed. This issue is central to the problem of assimilating observations into a weather forecast model in order to create an initial state $\mathbf{X}_W^{(i)}$. Data assimilation is achieved by a sophisticated form of least-squares fitting, which, given the observations, finds, by minimisation, a time-evolving solution of the model equations using *a priori* error estimates of the observational instruments, and the short-range forecasts (eg Courtier et al, 1994). In such schemes there is a clearly a fundamental problem when we try to assimilate observations which are strongly influenced by circulations near or below the model's truncation scale. No matter how accurate these observations are, it is impossible to represent them properly in the model - the model equations simply do not allow the corresponding weather elements to be simulated accurately.

That is, one of the key uncertainties in the specification of the initial state arises, not because the observations themselves are uncertain, but because the model equations do not allow observations of small-scale weather elements to be properly assimilated.

At initial time, this uncertainty is confined to scales near the truncation scale. However, as the forecast progresses, the uncertainty in the initial state propagates upscale non-linearly and starts to infect the forecast accuracy of larger cyclone/anticyclone scales, ultimately affecting the forecast accuracy of the planetary scales. This is, of course, the butterfly effect.

The key to improving the accuracy of weather-forecast initial conditions is, therefore, not just better observations, but more powerful computers, so that models can better resolve the scales which (for example) the satellite sensors are capable of observing, and hence so that these observations can be better assimilated into the models. As discussed in the next section, this inability to assimilate observations of small-scale weather is also leading to major uncertainty in our predictions of climate change.

Uncertainty in Climate Prediction

If we use the Lorenz model by way of illustration, then the weather forecast problem consists of finding finite-time trajectories on an underlying attractor, as shown in Figure 1. The climate change question, on the other hand, can be thought of as the problem of estimating how the attractor as a whole will change as a result of some prescribed perturbation to the equations of motion (Palmer, 1998). However, like the weather forecast problem, there are fundamental uncertainties in predicting the impact of such a perturbation. As in weather forecasting, these uncertainties arise from the limited resolution of our models.

As mentioned, there are similarities between the atmospheric envelope in which we live, and Lorenz's laboratory convecting fluid. However, there are also some key differences. One very important difference is that the atmosphere is not a single phase fluid; water exists in the atmosphere in all three phases. Water vapour is the primary greenhouse gas in the atmosphere; feedbacks between water vapour and CO_2 are of primary importance in determining the magnitude of anthropogenic global warming. Moreover, the release of latent heat associated with the condensation of water vapour into cloud liquid water is one of the principal processes which fuels atmospheric circulations. Such latent heating occurs in frontal regions in extratropical latitudes and in deep-convective thunderstorm clouds in both tropics and extratropics. Anyone who has been in a thunderstorm knows how awesome these "heat engines" can be; anyone who has been near a multi-cell convective storm - the sort that can spawn flash floods (such as the one that caused such devastation to Boscastle in Cornwall, UK last summer), or form intense downbursts from which aircraft cannot escape, or from tornadoes which can throw vehicles and houses into the air, will not be surprised that the latent-heat energy released in one of these convective systems is comparable to that of an fission bomb. A schematic of such a multi-cell convective system is shown in Figure 3.

These deep convective systems have multiple roles in the climate system. They cool the Earth's surface, leading to significantly lower surface temperatures than would be expected from estimates of pure radiative balance. They transport moisture from the surface to the free atmosphere, and therefore are intimately associated in the determination of the feedback between CO_2 and atmospheric water. Also, because the atmosphere is a nonlinear system, the kinetic energy in these convective systems can trigger yet larger-scale circulations. For example, it is believed that organised deep convective storms over the tropical West Pacific can trigger eastward-propagating wave-like disturbances (known as Madden-Julian Oscillations) with scales of thousands of kilometres. In turn these wavelike disturbances can trigger the occurrence of El Nino events where much of the central and east Pacific Ocean warms. In turn, these El Nino events lead to an elevation of global mean surface temperature. This is an example of upscale energy transfer from small to planetary scales that one would expect to occur generically in a turbulent system like climate.

Because of the computational constraints discussed above, global climate models cannot resolve and hence simulate the sort of deep convective systems illustrated in Figure 3. As has been discussed, the inability to resolve kilometre-scale weather elements is a source of initial condition uncertainty for weather forecasts. What impact does this have for predicting global warming?

The problem faced by anyone trying to model weather and climate is that we cannot totally ignore the unresolved scales of motion. Following work of Osborne Reynolds in the 19th century and Lewis Fry Richardson in the early 20th century, we try to represent the unresolved scales in climate models by imagining an ensemble of sub-grid processes in approximate secular equilibrium with the resolved flow. The ensemble-mean (or "bulk") effect of these sub-grid processes is then given by a set of relatively simple (eg diffusive-like) deterministic formulae. We call such formulae "parametrisations" of the sub-grid processes.

In practice these parametrisations have a number of free parameters that influence the distribution of water vapour and cloud liquid water in the grid box, as well as the amount of water that precipitates out. Because of the feedbacks between CO_2-induced global warming and the distribution and amount of water in the atmosphere, these parameters can strongly influence the magnitude of simulated global warming arising from a doubling of CO_2, a quantity known as "climate sensitivity". Since the values of these parameters are uncertain, there is a corresponding uncertainty in climate sensitivity. The blue curve in Figure 4 (from Murphy et al) gives a probability distribution of simulated climate sensitivity based on a 53-member ensemble of integrations of the Hadley Centre climate model, made by varying some of the key parameters in the sub-grid parametrisations, within prescribed confidence intervals. The probability distribution of global warming is rather broad. Between the extremes of the distribution, global warming could vary from less than 2K to 8K.

A more recent study by Stainforth et al (2005) using a multi-thousand ensemble of climate change forecasts (making use of the distributed climate*prediction*.net experiment) has estimated uncertainties in forecasts of global warming which are even larger than in the Murphy et al (2004) study, with climate sensitivities ranging from less than 2K to more than 11K. An 11K warming would be catastrophic for humankind; the implications of the sea-level rise implied by the melting of the Greenland and Antarctic ice sheets are enough to focus the mind.

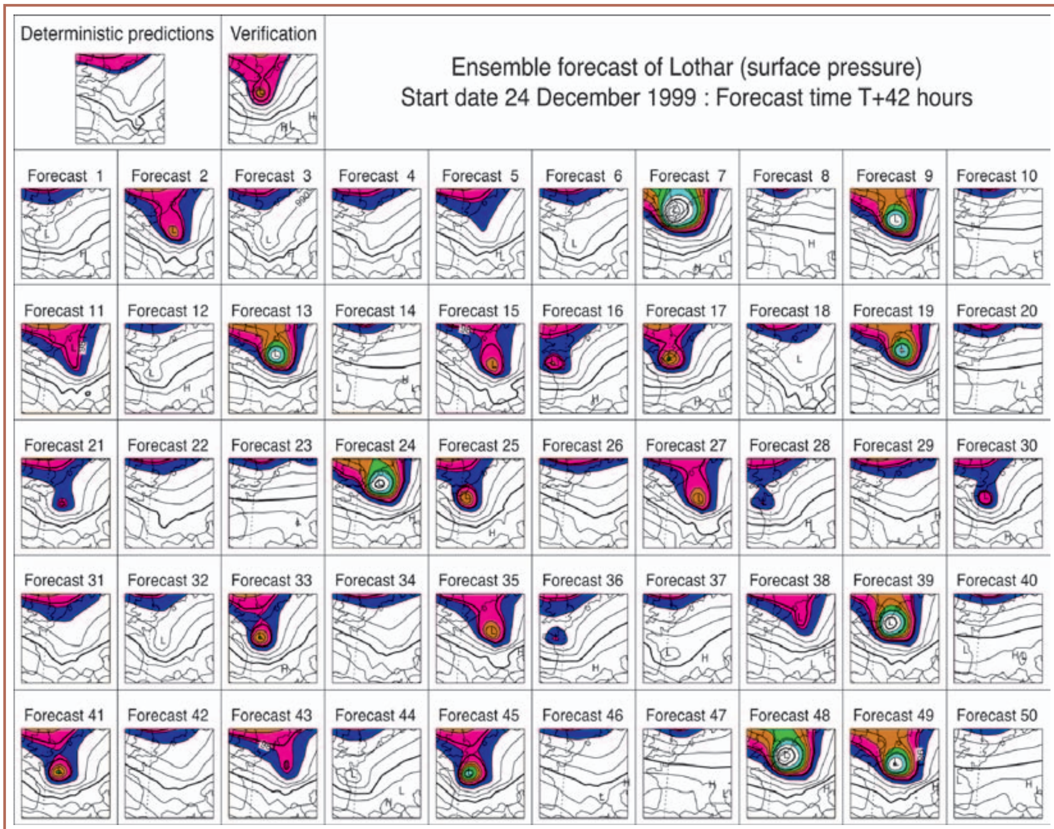

◄ **Fig. 2:** Isobars (lines of constant surface pressure) from a 51-member ECMWF ensemble forecast of the exceptionally severe storm "Lothar", based on initial conditions 42 hours before the storm crossed northern France in December 1999.Part of the coastline of Europe is shown for reference. The top left shows the forecast made from the best estimate of these initial conditions. This did not indicate the presence of a severe storm. The 50 forecasts which comprise the ensemble indicate exceptional unpredictability and a significant risk of an intense vortex.

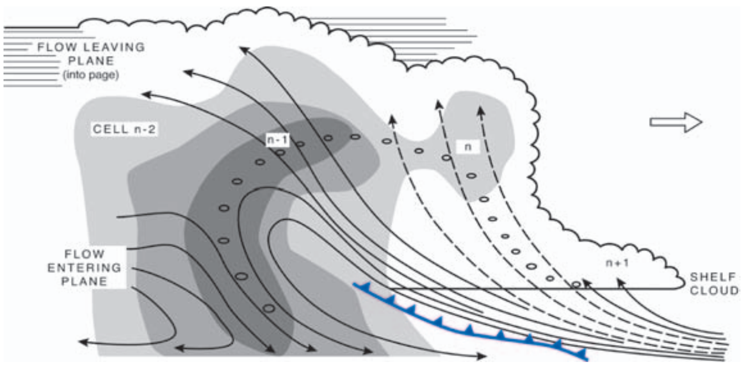

◄ **Fig. 3:** Schematic cross section of a mature multi-cell convective system. A typical horizontal scale for such a system is a few tens of kilometres. The gray shades denote the strength of radar reflectivity. Such systems are fuelled by temperature gradients in the lower few kilometres in the atmosphere, arising from the heating of the land (and ocean) surface by solar radiation. From Browning et al (1976)

These two studies, Murphy et al (2004) and Stainforth (2005), demonstrate conclusively that global warming is a potential threat that must be taken very seriously indeed. On the other hand, the range of uncertainty is also disconcerting. In view of the unprecedented seriousness of this problem for humankind, there is an urgent need for scientists to try to make these forecast probability distributions sharper. How could we do this?

It would seem reasonable to suppose that the key to solving the problem is to reduce the uncertainties in the climate model parameters. What about more and better observations? Improving the observational base for climate is important for many reasons, however, many of the key model parameters are not directly observable. The fundamental reason for this is that the assumptions that underlie parametrisation theory are not always satisfied in any quantitative sense (Palmer et al, 2005). For example, rarely can one find, within the region defined by a model grid box, an ensemble of deep convective systems in secular equilibrium with the large-scale flow; indeed when one does, these are not the most energetically important types of convective systems.

Another way to try to determine the values of the free parameters is through a procedure that one could perhaps call "tuning". The idea is to vary sets of parameters until the climate-model simulation of large-scale well-observed features agrees with the observed global climate of, say, the last 100 years. Unfortunately, this procedure does not discriminate between different sets of parameter values as well as one might like. For example, the red curve in Figure 4 is based on a weighting of the ensemble members of the Hadley Centre ensemble, based on how well a model with particular set of parameters, fits the observations of the large-scale climate. Doing this has not decreased the uncertainty in climate sensitivity - all that has happened is that low climate sensitivity has become less likely and high climate sensitivity has become more likely.

Towards Petaflop Supercomputing

The recent ensemble-based results of Murphy et al (2004) and Stainforth et al (2005) show that uncertainties in processes which are unresolved in the current generation of climate models, lead to substantial uncertainty in forecasts of global warming (even more so, uncertainties in regional climate change). As discussed above, we cannot use observations to reduce substantially the uncertainty in these parameters, because the assumption behind parametrisation theory, that unresolved small-scale climate processes can be treated by statistical mechanical methodologies, does not stand up to close scrutiny.

My view is that to reduce significantly the uncertainty in forecasts of climate change, global climate models should be able to resolve more of the atmospheric processes that directly determine the cloud cover and water vapour distribution in the atmosphere. In order to resolve deep convection requires grid lengths on the order of 1 km. How far are we from being able to do this? Currently the dedicated Japanese Earth Simulator (the leading climate computer at the end of 2004) runs global climate models with best sustained speeds of around 20 Teraflops (10^{12} floating point operations per second). Their standard climate change models have atmospheric grid spacing between one and two hundred kilometres. Experimental cloud-resolving models are being developed on the Earth Simulator. However, such models take a significant fraction of one day's wall-clock time to run one day of simulated climate - it would clearly be impossible to integrate such cloud-resolving climate models on the Earth Simulator for long enough to make meaningful forecasts of climate change.

To quantify climate change with cloud-resolving climate models will require computers with substantially higher performance - we must start looking towards machines with sustained speeds in the Petaflop range (10^{15} floating point operations per second).

In fact improved resolution is only one of four reasons why higher computer capacity is urgently needed for climate prediction research. The second reason is the continued need to run large ensembles of climate integrations. Even with models that can resolve deep convective processes, the effects of uncertainties in cloud microphysical parameters will have to be quantified using ensemble integration techniques. The third reason is associated with the growing complexity of climate models. Increasingly it has become recognised that climate models must include much more sophisticated representations of bio-geochemistry, aerosols, the cryosphere and so on. Such comprehensive models are now referred to as "Earth-System" models. Each new variable increases the complexity of such Earth-System models, requiring more processor power, and greater memory requirements for efficient integration. The final reason for enhanced computing is associated with the fact that one of the most severe tests of an Earth-System model will be its ability to simulate the paleo-climatic record, including the abrupt changes found in ice-core data (Alley, 2002). This will require model integrations of thousands of years.

It can be noted that, apart from the ensemble methodology, Grid technology alone does not offer a viable way forward to attack these issues. When a climate model is run over multiple processors, the sustained performance depends critically on how rapidly and efficiently the different processors can communicate with each other. Acceptable performance will be impossible to achieve with remotely distributed current-generation processors.

Can we be sure that investment in Petaflop computers for climate will lead to a reduction in uncertainty in global warming? I would say "yes" for two reasons. Firstly, one would reduce

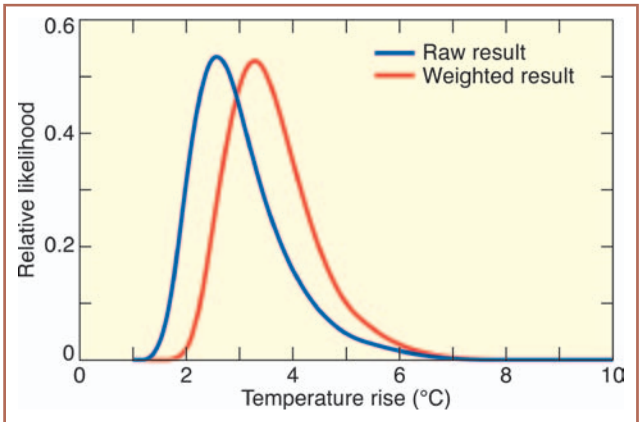

▲ **Fig. 4:** Probability distributions of global average annual warming associated with a 53-member ensemble for a doubling of carbon dioxide concentration. Ensemble members differ by values of key parameters in the bulk formulae used to represent unresolved processes, in a version of the Hadley Centre climate model. Blue curve: based on "raw model output". Red curve: with the probability distribution weighted according to the ability of different model versions to simulate observed present day climate. From Murphy et al (2004).

uncertainty in simulating one of the key climatic processes: deep convection in the atmosphere. Secondly, if global climate models can be formulated at resolutions comparable with the best numerical weather prediction models, the sorts of techniques that are currently used to validate short-range numerical weather forecasts (eg Klinker and Sardeshmukh, 1992; see also www-pcmdi.llnl.gov/projects/capt/index.php) could be brought to bear on climate models. For example, recall that many of the key uncertainties in climate models are associated with poor representations of cloud processes. These processes (as opposed to the feedbacks between cloud and CO_2) have intrinsic timescales of hours. Hence, by initialising the climate models with the sorts of sophisticated data assimilation procedures used in weather prediction, as described above, and by running the models for just a few hours from many different initial states, it may be possible to reject either the high and low climate sensitivity models through their relative fit to detailed observations of cloud, humidity and temperature. At present, it is likely that climate models have too coarse resolution for these tests to be sufficiently discriminating. From this latter perspective, investment in computing for climate could be seen as important requirement for getting value for money from the corresponding investment in space-based observations of climate.

Petaflop computing is not science fiction - the main high-performance computing manufacturers are actively working towards this goal and are expected to reach it in the coming few years. In many countries, resources for such climate computing will be beyond national funding levels, eg within Europe. In any case, what would be the point of duplicating such a climate computing resource in many different national countries - like CERN, this is exactly the sort of project that should be resourced at the international level.

Postscript

Imagine this. It's Sunday morning, some years in the future. Your thirst for that perfect golf swing is undiminished, but your ability to play in the rain has deteriorated. The alarm rings; it's time to get up if you want to play golf. You glance at your laptop, set up on the bedside table. Based on the latest cloud-resolved weather-forecast ensembles, the predicted probability of mesoscale convective storms over your golf course in the coming five hours has increased to 40%; 4% higher than the predicted figure last night when you went to bed. Most likely it won't rain, but you don't want to take the risk; you get so bad tempered when it rains. With little further thought, a text message you prepared last night is sent to your golfing partners' mobile phones: "Not playing today…feeling terrible…stayed out much too late last night!"

You switch on the bedside TV. The Prime Minister has announced agreement on a new climate treaty, endorsed by all governments. As the PM explains, the change in heart has come from the latest ensemble forecasts produced at the new international petascale climate computing centre; uncertainties have been reduced to such an extent that humankind can now agree on a course of action.

It's raining outside. You smile… and go back to sleep.

About the author

Dr Tim Palmer leads a research division at ECMWF based in the UK. He was lead author of the 2001 assessment report of the Intergovernmental Panel on Climate Change, and has coordinated two European Union climate projects. He is a Fellow of the Royal Society and will give a public lecture based on this article at the Royal Society in London on the evening of April 26th.

References

Alley, R.B. 2002. The two-mile time machine: ice cores, abrupt climate change, and our future. Princeton University Press. pp229.

Browning, K.A., Fankhauser, J.C., Chalon J.-P., Eccles P.J., Strauch R.G., Merrem F.H., Musil D.J., May E.L. and Sand W.R., 1976. Structure of an evolving hailstorm. Part V: Synthesis and implications for hail growth and hail suppression. Mon.Wea.Rev., 104, 603-610.

Courtier, P., J.-N. Thépaut and A. Hollingsworth, 1994: A strategy for operational implementation of 4D-var using an incremental approach. Quart. J. Roy.Meteorol.Soc., 127, 1469-1488

Lorenz, E.N., 1963. Deterministic non-periodic flow. J.Atmos.Sci., 42, 433-471.

Klinker, E. and P.D. Sardeshmukh, 1992: The diagnosis of mechanical dissipation in the atmosphere from large-scale balance requirements. J.Atmos.Sci., 49, 608-627.

Murphy, J.M., Sexton, D.M.H., Barnett, D.N., Jones G.S., Webb, M.J. et al, 2004: Quantifying Uncertainties in climate change using a large ensemble of global climate model predictions. Nature, 430, 768-72.

Palmer, T.N., 1998: Climate change from a nonlinear dynamical perspective. J.Clim., 12, 575-591.

Palmer, T.N., 2000: Predicting Uncertainty in Forecasts of Weather and Climate. Rep.Prog.Phys., 63, 71-116.

Palmer, T.N., G.J.Shutts, R. Hagedorn, F.J. Doblas-Reyes, T.Jung and M.Leutbecher, 2005. Representing Model Uncertainty in Weather and Climate Prediction. Annual Reviews in Earth and Planetary Science. To appear.

Stainforth, D.A., T. Aina, C.Christensen, M.Collins, D.J.Frame, J.A.Kettleborough, S.Knight, A.Martin, J.Murphy, C.Piani, D.Sexton, L.A.Smith, R.A.Spicer, A.J.Thorpe, M.R.Allen, 2005:Uncertainty in Predictions of the Climate Response to Rising Levels of Greenhouse Gases. Nature, 433, 403-406.