
Thursday night—Third, perhaps most powerful, storm looms.
Tuesday—Another storm headed toward California is expected to intensify as it draws moisture from humid air off Hawaii.
Today's storm
Hawaii

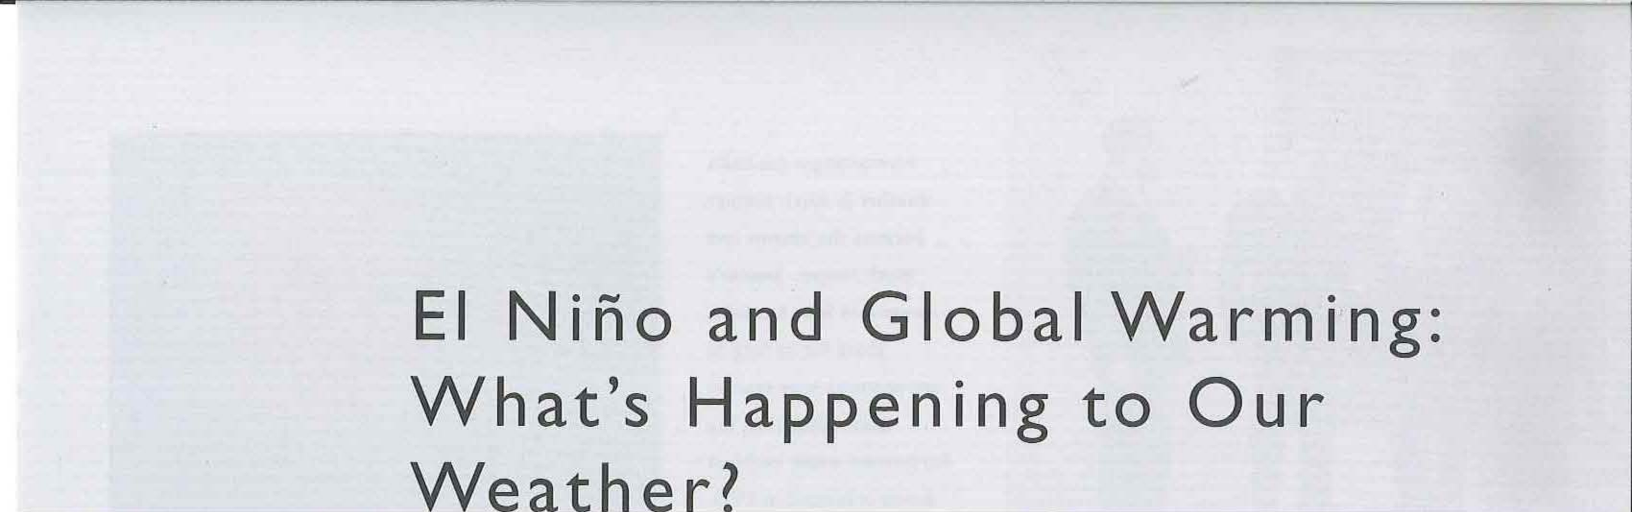

El Niño and Global Warming: What's Happening to Our Weather?

by Andrew P. Ingersoll

Satellites are the forecaster's best friend. This view of the Pacific Ocean, which ran in the *Los Angeles Times* on Saturday, February 14, shows three storms in procession from Japan (left) to L.A. (right) and their predicted arrival times. Satellite photo courtesy of the National Oceanic and Atmospheric Administration (NOAA), forecast by WeatherData, Inc. But satellites can't see unborn storms—on Monday, February 23rd, a fourth one rolled into town and engaged Ed Lewis (PhD '42), Morgan Professor of Biology, Emeritus, and Nobel laureate, in a tug-of-war for his umbrella.

As you all know, the wet weather we've been having was actually predicted half a year ago. For example, on August 20, 1997, the *Los Angeles Times* ran an article headlined "Southland Prepares for Worst Winter in Decades—Up to 300 percent of normal rainfall is expected from El Niño. Agencies scramble to be ready." Now they weren't predicting that we would have a big storm on any particular day; they were just predicting that we'd have a wet winter. They were quite right about the latter, and they wouldn't have dared to do the former. Since the days of Noah, no one has succeeded in predicting the weather, to the day, six months in advance. There are reasons for that, and I'll tell you what they are. (We can predict Jupiter's day-to-day weather six months in advance, and I'll also talk about that, but it doesn't work here on Earth.) But there are certain kinds of long-term weather predictions that we *can* make, such as El Niño and global warming, and that's my primary subject.

Is the Kyoto agreement just another rain dance, irrelevant to what's actually driving our climate? Or do we know enough now to say that this is really the right action?

We are getting better at forecasting the weather a few days ahead. Thirty to forty years ago, you could predict tomorrow's weather, and you could make some kind of wild estimate about the day after. Beyond that, you were guessing—you might as well have read the *Farmer's Almanac*. But now we make reasonably reliable six-day forecasts. Again, for example, the *Los Angeles Times* for Saturday, February 14, predicted drenching rain for that day, to be followed by another storm on Tuesday and a third storm Thursday night. The *Times* being a morning paper, the forecast was actually made on the previous day, Friday. And Thursday night the third storm came in, right as expected. This is the kind of thing that makes meteorologists proud. A 90 percent versus an 80 percent chance of rain isn't really what they live for—it's the long-range forecast that shows off who's good.

It's not that the IQ of weather forecasters has gone up; it's just that they have better tools nowadays. One important tool is a set of satellites that gives global coverage of the planet and fills in the gaps between the ground stations. In the old days, the only midocean data you had were from wherever a ship or an island happened to be. At our latitude, storms basically move from west to east, so if you see one out in the middle of the Pacific today and you know how fast it's moving, you can extrapolate forward and say when it's going to hit. It's like a merry-go-round going from left to right, and the storms are the horses—if you have a little child on the merry-go-round, you can sit and read your book and, as the child comes around, look up and wave. Weather forecasting is tougher because the horses keep vanishing, and new horses appear in different places. Thus the theoretical limit to how far ahead you can forecast the weather is set by the lifetime of the storms. It's probably about two weeks at best—we don't yet know exactly where the limit is, because we haven't got the tools to really test it. And because you don't know when and where storms are going to appear and disappear, you can't just put your data into a computer model—another important new tool—and fast-forward the model to print out six months' worth of weather predictions.

At this point, I promised some friends that I would read from the scriptures. But this is Caltech, so the scriptures are *The Feynman Lectures*

Forecasting a gas ball's weather is much simpler because the storms last much longer. Jupiter's Great Red Spot has been there for as long as astronomers have trained telescopes on it; the Earth-sized white oval just below it formed in 1939.

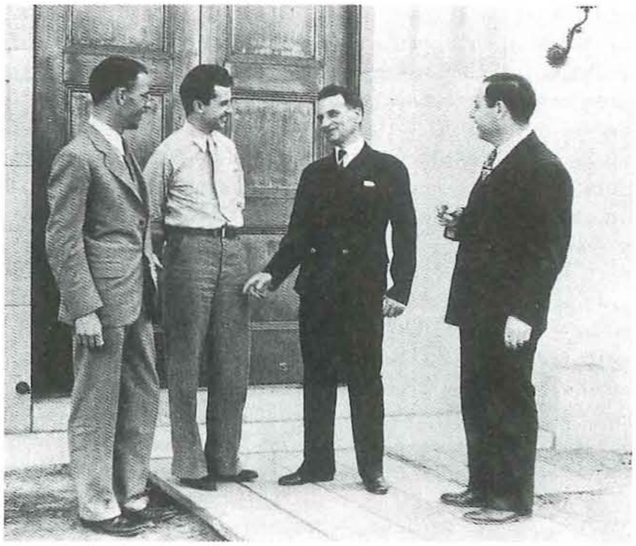

Caltech's first meteorology course, on atmospheric structure, was taught in the geology department by seismologist Beno Gutenberg in 1930. (Gutenberg was interested in acoustic waves in the atmosphere as well as seismic waves.) The meteorology program began in the fall of 1933 under the aegis of the aeronautics department. Besides Gutenberg, the instructors included (from left) Clark Millikan (PhD '28), Irving Krick (MS '33, PhD '34), Theodore von Kármán, and Arthur Klein BS '21, MS '24, PhD '25). The program eventually became a freestanding department with Krick, a grad student of both Gutenberg and von Kármán, as its chair.

on Physics—the bane of Caltech undergrads in the 1960s and '70s. Feynman understood why complicated classical systems, as opposed to quantum-mechanical systems, are basically unpredictable. Let me read from the Book of Feynman, Chapter 38, Page 9: "If we knew the position and the velocity of every particle in the world, or in a box of gas, we could predict exactly what would happen.... Suppose, however, that we have a finite accuracy and do not know *exactly* where just one atom is, say to one part in a billion. Then as it goes along it hits another atom, and because we did not know the position better than to one part in a billion, we find an even larger error in the position after the collision. And that is amplified, of course, in the next collision, so that if we start with only a tiny error it rapidly magnifies to a very great uncertainty." That's it, folks. That's exactly why weather forecasting is so hard. That's why no computer will ever foretell the birth or death of a specific storm. Weather forecasters call this the Butterfly Effect: the flapping of a butterfly's wing in Brazil might eventually cause a blizzard in Helsinki.

Caltech had a meteorology department back in the '30s and '40s, and the faculty bandied about the idea of a theoretical limit to predictability. It was not clear then that there was such a thing. (And there really isn't on Jupiter, as I said.) In fact, the department chair, who maintained that it was possible to predict the weather months in advance by matching observed weather patterns with historical ones, supposedly predicted the weather for D-Day in December 1943. Caltech abolished its meteorology program shortly after the war, partly because President DuBridge, who took office in September 1946, wanted to focus the Institute on basic instead of applied science.

Several decades passed, and Caltech hired a few planetary scientists with some atmospheric-science background, one of whom was me. And a funny thing happened—we started predicting the weather on the giant planets months in advance. I was a member of the Voyager imaging team, and I was in charge of Jupiter's atmosphere. We knew that in the last two days before the spacecraft zoomed past Jupiter, we would get a chance to photograph some of its storms up close. Voyager would be so close that only a small portion of Jupiter would fit in the camera frame, so we had to figure out where the storms were going to be in time to send commands up to the spacecraft saying, at such and such a time, aim the camera at such and such a place, and we promise there will be a storm there. We had to give the engineers the aim points three weeks in advance. That's how long it took the engineers to integrate our aim points into everything else the spacecraft was doing, write up the whole command sequence, test it, and send it to the spacecraft. (Later, for Galileo, we had to provide a rough forecast for Jupiter six months in advance, so that the mission team would know which side of the planet the Great Red Spot would be on.)

Anyway, during Voyager's long approach to Jupiter, the spacecraft was snapping pictures every day, as were telescopes on Earth. And we knew that the storms on Jupiter rode the merry-go-round for a long, long time—the Great Red Spot, for example, is at least 300 years old. The storms

Neptune is too far away for its storms to be easily visible from Earth, but it had ample time to establish a track record during Voyager's leisurely approach. The Great Dark Spot is at center; Scooter is the white cloud halfway between the dark spots.

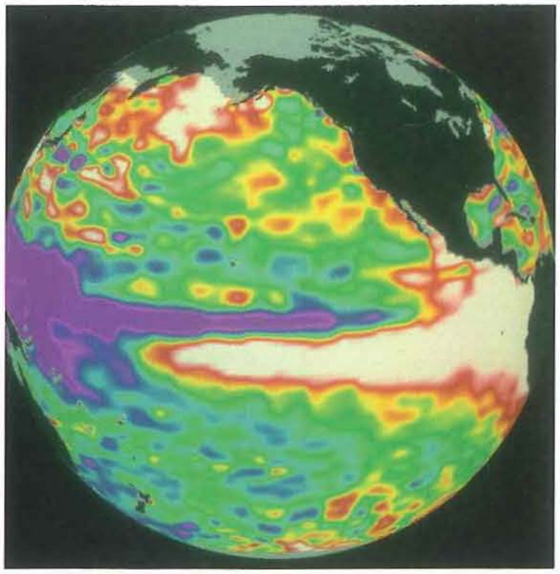

Above: A portrait of El Niño from October 23, 1997. This data is from JPL's TOPEX/Poseidon satellite, which doesn't actually measure ocean temperature, but instead measures the expansion of the ocean—a good proxy for temperature, because warm water bulges the ocean's surface upward. (See *E&S*, Spring 1995.) Yellow-green represents the ocean's normal height. Yellow is five centimeters above normal, red is 10, and white is 15; blue and magenta are below normal, with magenta being -15 centimeters. Below: NOAA seven-day average temperature data from October 26 – November 1, 1997. This data is compiled from buoys, ships, and satellites that measure the infrared radiation from the topmost millimeter of seawater (which, unfortunately, is sensitive to winds, clouds, sunlight, and evaporation).

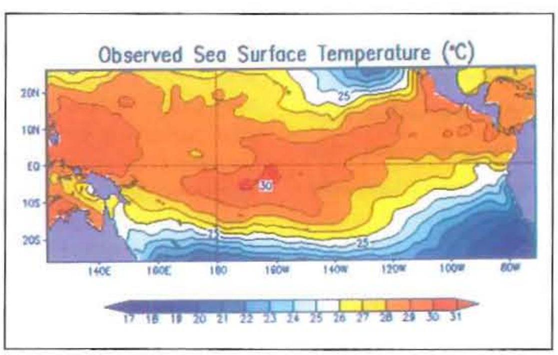

are all moving relative to one another, and of course the planet is rotating, so we took the data from the pictures, plotted the storms' positions as a function of time, laid a ruler on the graph, and extrapolated where the storms were going to be. The storms do change—they churn and boil, they fade and brighten; their appearance changes daily. And smaller storms come and go. But we hit just about every target, and that's not because we were brilliant people. It's just that Jupiter is very different from Earth. Predicting the weather on the giant planets is simpler—probably because they have no solid surfaces, no topography, and no oceans to complicate the circulation patterns.

Ten years later, Voyager 2 was at Neptune. Neptune is a little more complicated because, while Jupiter's storms move at relative velocities of tens of meters per second, Neptune's storms zip past one another at velocities of up to several hundred meters per second. For example, the two dark spots north and south of a storm we nick-named Scooter lap each other every five days. (Neptune's storms may also be shorter-lived—the Great Dark Spot seems to have disappeared from Hubble Space Telescope images taken in 1991.) But we could still make our three-week forecasts with junior-high-school mathematics. We didn't have any supercomputers or fancy stuff, but it worked. We got wonderful photos. By contrast, at the same time, August 1989, Hurricane Hugo was threatening the Carolina coast like a prize-fighter—dancing around, faking left, faking right. The meteorologists on the East Coast were issuing 12-hour forecasts, trying to predict where Hugo would come ashore. But the hurricane kept stopping dead and veering off in another direction, leaving them flat-footed.

Enough about day-to-day forecasting—let's move on to predicting El Niño six months in advance. El Niño is a sloshing of warm water from the western side of the Pacific Ocean eastward toward the American side. There's a lot of mass involved, and the ocean currents move slowly, and it's this ponderousness that makes long-range predictions possible—people know that once the warm water piles up on our side, it's going to stay here for a while. This affects our weather because warm water evaporates faster, and more water vapor in the atmosphere means more rainfall and more storms. Meanwhile, in the western Pacific, the water is colder than normal, which causes droughts and fires—you may remember that Indonesia had terrible problems with both.

The upper figure at left is a picture of this year's El Niño—I'm sure you've seen pictures like this. The angry, highly colored region is five or more degrees centigrade above normal. But this isn't really the way the ocean temperatures look—the actual sea-surface temperatures are shown in the map at left. The warmest water is south of Hawaii in the central Pacific, near the equator, where the

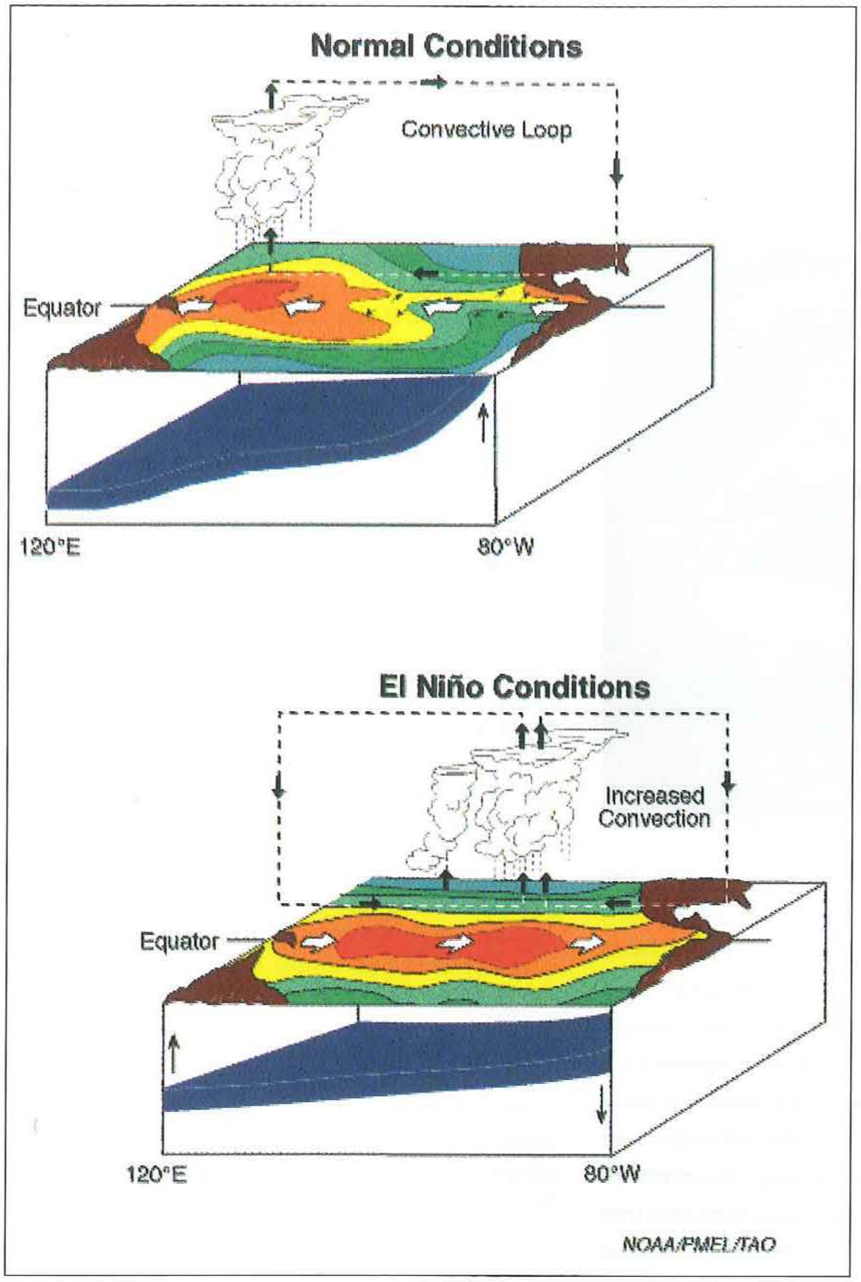

The sloshing thermocline. Under normal conditions (top), the trade winds (white arrows) blow from east to west. The warm surface water (orange and red contours) piles up in the western Pacific, pushing the thermocline (the blue layer) down there, while allowing it to rise in the eastern Pacific. In an El Niño year (bottom), the wind slackens or even reverses direction, and the warm surface water remains evenly distributed across the entire ocean. Then the thermocline becomes almost horizontal.

most sunlight falls. Well, so what—wouldn't you expect that? This is actually abnormal because, during normal years, the warm water is all piled up in the western Pacific. The trade winds, which blow from east to west at the equator, drive the warm water westward. So if we take the abnormal pattern and subtract from it the normal pattern, you get the picture we're used to seeing. The American coast looks warm, because the water there is normally much colder. In fact, the American coast is the warmest anomaly of all—the largest departure from normal.

It's not really the ocean's surface that's sloshing, but something called the thermocline, which lies about a hundred meters deep. The thermocline is the interface between the upper ocean, which is relatively warm (up to 30° C), and the cold water below. Most of the ocean is barely above the freezing point. Normally, the trade winds blow the warm surface water toward the west, depressing the thermocline in the western Pacific. Pushing the warm water westward means the thermocline rises to the surface in the east, so that there's cold water off the coast of Peru. But for some reason, the trade winds periodically slacken or even reverse, blowing the warm water the other way. The thermocline gets shoved down in the east, and there's warm water all the way across.

So the thermocline sloshes back and forth, like water in a kid's bathtub, and the frequency of the sloshing depends on the density difference above and below the thermocline. This difference is not very great, so the frequency is very slow. It's like that parlor toy you may have seen that's supposed to relax you—the long, horizontal container filled with two different-colored fluids of almost the same density. You tip the container, and waves slosh back and forth very slowly.

But there are several mysteries connected with El Niño. The natural period of the bathtub mode is a little less than a year, which is too short to explain the observed frequency of El Niño. El

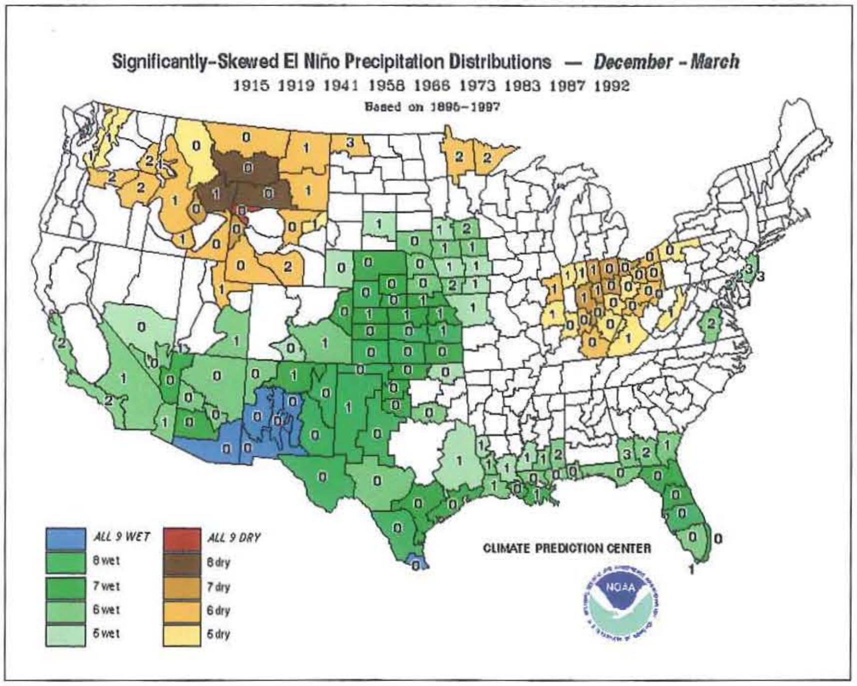

Above: The colored areas in this map of the United States show regions that have been particularly wet or particularly dry in the nine El Niño years of this century. (The white areas got normal precipitation.) The number in each colored area shows how often an El Niño year bucked the trend—a dry year in a region that usually gets extra rainfall during El Niño, for example.

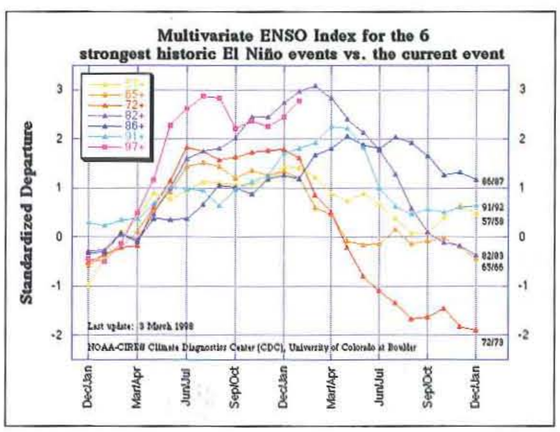

Above: Although El Niño's arrival was predicted successfully, predicting its strength is still dicey. Early indications were that it would be very strong indeed—the fierceness of its onset outstripped the one of the winter of 1982, which did tremendous damage. Fortunately, the current El Niño has not lived up to its advance billing. The "multivariate index" is a composite of such variables as air pressure and temperature, wind speeds, ocean temperature, etc.

Niños come, on average, every four years, but they can be as few as two or as many as seven years apart. Also, the bathtub mode doesn't take into account the trade wind's changing direction, which obviously has something to do with El Niño. This leads to another problem—when water vapor condenses into rain, the vapor gives up heat and warms the air. The warm air rises, causing a convective motion that draws in more air down at the surface. So when the trade winds slacken and the convection centers drift eastward toward Peru, they augment the eastward-blowing winds along the surface. The ocean should get stuck in the El Niño mode, with all the winds blowing east, and never get out. Or it should get stuck in the opposite, normal position, with all the winds blowing west, and the air rising near Australia. So we're at a loss to understand why the trade winds weaken and allow the water to slosh back. We have lots of empirical theories, but no grand understanding behind them.

I started to get a little tired of the media hype last fall, and I decided to see what past El Niños had really done to the weather. So I checked the Web site of the National Oceanic and Atmospheric Administration (NOAA), and I found the map at left. Researchers divided the United States into geographical areas, and for each area they took a hundred years of weather data for December through March, which they divided into thirds. So, by definition, one-third of the years were wet years, one-third were dry, and the rest were medium. Now, what wet means in Arizona is different from what it means in Florida, but still, each geographic area has its definition of wet, dry, and medium. And then the researchers asked, of these nine El Niño years, how many were wet? how many were dry? how many were medium? You can see from the colors that the southern part of the U.S. typically had wet El Niño years, but notice that Southern California only had six wet years out of nine, which is not overwhelming odds. And the figures on the map tell you the number of El Niño years that went the opposite way—in our case, dry years. Southern California had six wet years, two dry, and one in-between. That's hardly a slam-dunk for El Ninõ. So all we can say, based on past experience, is that we've got six out of nine chances that this year will be in the wettest one-third. I went around saying that, and I offered to bet one of my colleagues that this winter would be a dry one, if he would give me 4:1 odds—his $4 to my $1. He didn't take me up on it, which is good, because I clearly would have lost.

Let's move on to global warming. It has been predicted that if we add carbon dioxide, methane, freon, and some other gases to our atmosphere, which we are doing—no question about that!—then Earth will warm up, and in 50 to 100 years we'll have some very costly changes in our climate. These gases are called greenhouse gases, because a

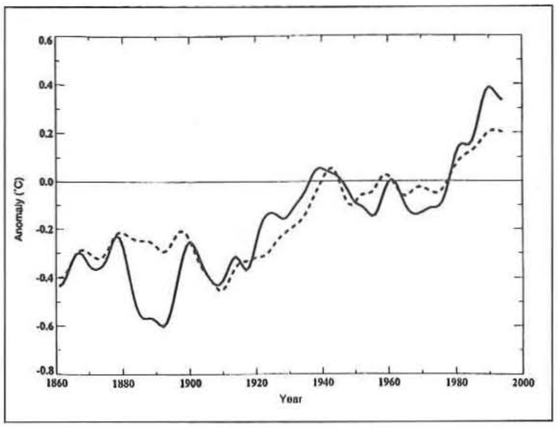

Above: The global mean annual temperature from 1861 to 1994, as compared to an arbitrary "normal" temperature of about 15° C, shown as 0.0 on the graph. Thus, for example, -0.4 is really 0.4° C below normal. The solid line is air temperatures averaged over the land masses, and the dashed line is average sea-surface temperatures. Data after the 1995 Intergovernmental Panel on Climate Change (IPCC) report.

Below: Atmospheric carbon dioxide levels in parts per million, measured at Mauna Loa, Hawaii (triangles) and from air bubbles trapped in the ice near Siple Station, Antarctica (squares). Data from the 1990 IPCC report.

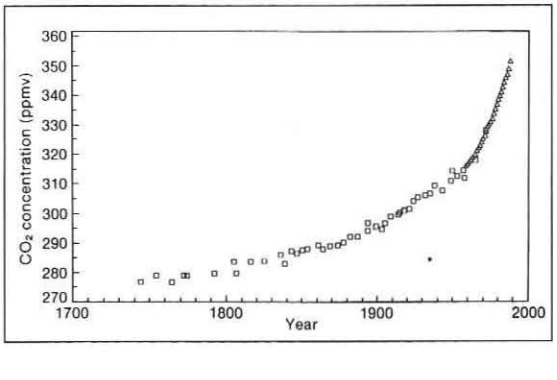

greenhouse stays warm even in the winter; its windows trap the heat inside it. These gases do the same for planet Earth. But cutting down on carbon dioxide production means burning less oil and coal—our civilization's basic energy source. That's going to hurt the world economy, so there are sacrifices involved—we're playing for real stakes here. And last December, delegates from all over the world met in Kyoto, Japan, to hammer out an agreement about who should sacrifice how much. You might ask, have we finally gotten to the point where we're having such an impact on the weather that we have to make these great sacrifices? There have been rain dances for as long as there have been people growing crops. There were, on occasion, serious sacrifices—people were killed; cattle were slaughtered. Is the Kyoto agreement just another rain dance, irrelevant to what's actually driving our climate? Or do we know enough now to say that this is really the right action?

Let's look at the evidence. For one thing, 1997 ranks as the warmest year of the century. And why not? There's got to be a warmest year, so why not 1997? But this is really quite unusual, because five of the century's warmest years have been in this last decade. Clearly, it's getting warmer. Is the buildup of greenhouse gases, mainly carbon dioxide, responsible? The mean annual temperature for the planet, compiled from daily temperature data from several hundred stations around the world, is shown at top left. You can see lots of bumps and wiggles—for example, it went up to a maximum around 1940, and then back down again. But for the last few decades, it's been going up steadily.

The amount of atmospheric carbon dioxide in parts per million, as measured directly at Mauna Loa and from air bubbles trapped in the Greenland and Antarctic ice sheets, is shown at left. The level was pretty constant until nearly our century, when combustion took off—carbon dioxide is now

Left: Europe's glaciers have been retreating for more than 100 years. The engraving (top), circa 1850–1860 by an unknown artist, shows the front of the Argentière glacier lying close to the church in the village of the same name, near the Swiss–Italian border. In the photograph (bottom), taken from the same vantage point in 1966, the glacier has receded to the mountain's shoulder. From *Times of Feast, Times of Famine* by Emmanuel Le Roy Ladurie.

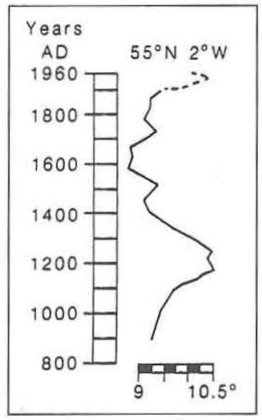

Left: The mean annual temperature, in degrees Centigrade, in the vicinity of Newcastle-Upon-Tyne in northern England. Data after H. H. Lamb.

at 370 parts per million, and rising. So it's tempting to associate the two curves, especially when we know that carbon dioxide traps heat. It's a fairly easy calculation to say how much heat it traps, but what's difficult is calculating all the other elements of the climate system. If you warm Earth up a little bit, you might get more clouds—clouds are bright and reflect sunlight, and might cool Earth back down. Or there might be more thunderstorms—thunderheads condense at relatively high altitudes and would carry heat up into the atmosphere, cooling the surface. Clouds, thunderstorms, and turbulence in general are basically unsolved problems, so associating the carbon dioxide buildup with the temperature rise is a tough business.

Another reason it's a tough business is that the climate varies naturally. There are changes of several degrees going on over hundred- and thousand-year timescales. In the temperature data above, you can see that the period from about A.D. 1400 to 1850 was approximately a degree and a half colder than the periods before or since. That's true in Michigan and England, in Canada and California. This period is known as the Little Ice Age, and it really was a little ice age. There are old pictures of Swiss glaciers reaching way down into the valleys, and if you go to the same spots today the glaciers are gone. They've retreated up into the mountains. There were great midwinter parties in London, where they rolled big logs out onto the ice in the middle of the River Thames and roasted oxen. The Thames never freezes now. But the Little Ice Age had nothing to do with human impact—in fact, no one quite knows what caused it. Maybe the sun dimmed a little; maybe the Gulf Stream stopped carrying warm, equatorial water northward. A lot of things might have happened. And if you look farther back into time, there are even bigger changes—20,000 to 40,000 years ago, there was ice a mile thick covering Chicago, with lots of rapid changes in between.

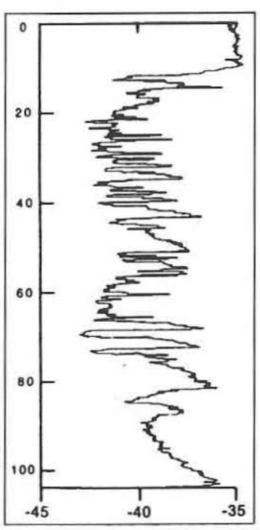

Left: Temperature variations over the last 100,000 years, as deduced from the ratio of oxygen isotopes in a core from the Greenland ice cap. The vertical scale is marked in thousands of years before the present; an increase of five units on the horizontal scale is equivalent to a temperature increase of 6° C. Note the frequent variations of several degrees over time periods of 1000 years or less. As moist air cools on its poleward journey, water molecules containing the heavier ^{18}O tend to fall out faster than those containing the lighter ^{16}O. The colder it gets, the less ^{18}O is left aloft. Comparing the ancient ice's $^{18}O/^{16}O$ ratios to ratios measured around the planet today gives a measure of how cold the ice was when it froze—a technique invented by Caltech's Sam Epstein in the 1950s. Data after Dansgaard et al.

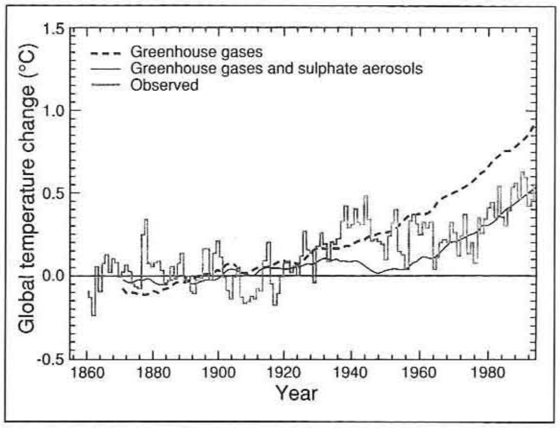

Left: If your computer model only includes greenhouse gases, its predictions (dashed line) don't match the real-world data (gray line). But if you add a soupçon of smog (solid line), the fit is much better. Data from the Hadley Center, UK Meteorological Office's climate model, published in the 1995 IPCC report.

Then, about 12,000 years ago, the ice suddenly melted, and it's been relatively warm for the last 10,000 years. Earth was about 7° C colder during the depths of the last Ice Age. By contrast, the warming in our own century has been about 0.7° C.

So how did we come to predict global warming? We used computer models of Earth's climate. (The three main models in the U.S. are at the National Center for Atmospheric Research, the Goddard Space Flight Center, and the NOAA labs at Princeton. There are other models that don't have as much funding but have some very smart people working on them, including a model developed at UCLA.) These models all divide the globe up into a grid, and put pressure and temperature and moisture content and whatnot into each box in the grid. There are equations for how these parameters interact, and how air moves from one box to the next, and how land and sea and the passing of the seasons affect the air. There are even equations for how plants suck carbon dioxide out of the air as they photosynthesize. We set the model in motion, gradually add carbon dioxide, and watch what happens.

But if we just model the rising levels of greenhouse gases, we overpredict the warming—if we start the model at, say, 100 years ago, it tells us that the planet should be hotter today than it actually is. But if we add in some aerosols—shiny particles, smog basically—that reflect sunlight, we don't get as much warming, and the model tracks the historical data pretty well. So it seems we might be pulling the weather in one direction with heat-trapping greenhouse gases, and pushing it the other way with sunlight-reflecting aerosols. The average of the models' predictions, if carbon dioxide doubles in the next 75 years, is a global mean temperature increase of 2.5° C. That's about a third of the warming that occurred from the end of the last Ice Age to the present. The human race survived that, so we should be able to survive another 2.5 degrees. Some models say 1.5 degrees; others say 4.5 degrees. There's a lot of uncertainty, and just about every element in the models is under debate. They make different assumptions about turbulence, for example, and the effects of clouds. But in the end, we have to use the models—they're all we've got. We just don't trust them to the last decimal point. We always quote an uncertainty.

There's currently a lot of debate about whether we've already seen the signature of global warming. I would say that debate is not a terribly productive one. Global warming may or may not account for the little upturn of the last few decades, but I'm quite confident that we'll see its effect in the next century. The effect is just beginning to rise above the noise of natural planetary variability. If it turns out that the current upturn was because the ocean hiccuped, it doesn't mean that global warming is going to go away.

Global warming may or may not account for the little upturn of the last few decades, but I'm quite confident that we'll see its effect in the next century. The effect is just beginning to rise above the noise of natural planetary variability.

Of course, things can happen. For example, a good-sized volcano such as Mount Pinatubo can fill the stratosphere with shiny, highly reflective particles that could cool Earth and stave off global warming for a time. However, while the aerosols stay up for a few years, the carbon dioxide lasts for centuries.

So now we come to politics. The economists, meteorologists, and everyone in between are all trying to say what the world should do. The report by the Intergovernmental Panel on Climate

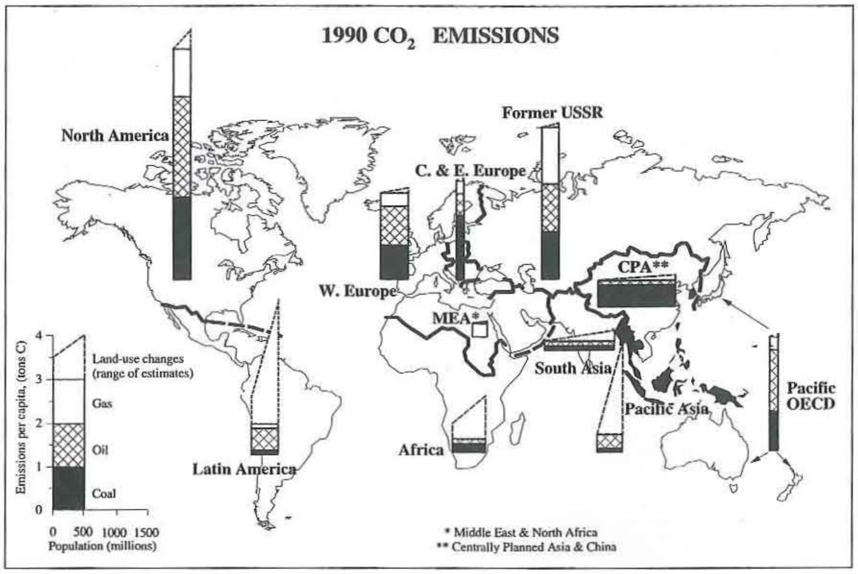

Above: Global output of carbon dioxide from fossil fuels and from deforestation (labeled as "land-use changes" on the legend; note the large range of uncertainties in the tropical estimates). Since the number of tons of carbon emitted per capita is plotted vertically and each region's population is plotted horizontally, the area of each bar gives the total amount of carbon dioxide emitted by each region. Pacific OECD stands for Pacific Organization for Economic Cooperation and Development. Data from the 1995 IPCC report.

Change formed most of the basis for the debate in Kyoto. The meteorologists predicted that if we warm the planet by 2.5° C, the world's desert areas will expand. Louisiana might become a desert, and Montana might become a lush agricultural area. (Of course, the rainfall predictions are just as uncertain as the temperature predictions.) The economists took that data and said, well, how much would those changes cost the world? There'll be losses to agriculture, and the increased use of water for irrigation may drive up its cost for all users. As the deserts expand, trees and other vegetation will die. There'll be some 30 centimeters of sea-level rise, which will affect ports, beachfront property owners, and coastal wildlife. There'll even be the cost of extra air-conditioning. Typical estimates for the U.S. alone were that it might cost us $50–100 billion a year. That's not a lot—it's one or two percent of our economy. But India and China, for example, would be much more vulnerable, because their economies are weaker and they're more dependent on hand-to-mouth agriculture. Their losses could be 10 percent of their economy. Taken overall, the losses will be a few percent of the world economy.

Then you have the question of who should pay. Well, who's doing the damage? North America, Western Europe, the former Soviet Union, and China are the big players, as you can see from the graph above. So you might say we should pay according to how much carbon dioxide we produce. The Chinese say that's baloney—that they've got many times more people than we do, that they have the right to pollute as much as we do on a per capita basis, that they want to build up their economy to be on a par with ours. And if you look at how much carbon dioxide each country produces per capita, China looks very good and we look very bad. In fact, we're a lot worse than Western Europe and the former Soviet Union.

The decision was finally made to reduce the United States' emissions to 7 percent below their 1990 levels over the next 10 years. (The treaty hasn't been ratified by Congress.) If we do ratify it, the cost to the U.S. economy to achieve these reductions will be about 1 or 2 percent—the same as the cost of global warming. The European Union is to reduce their emissions by 8 percent; Canada and Japan by 6 percent. India and China carried their point, and are not required to make any reductions under the treaty.

I don't think that the scientific issues are as uncertain as the economic and political ones. It's quite possible that in 75 years, we will have developed solar energy, clean nuclear fuel, wind power, or who knows what. [See *E&S*, 1997, No. 3] Then the debate will disappear, because we won't burn coal and oil any longer. In which case, why worry now? Let's just wait for that wonderful future. The other possibility is that we'll be so overrun with wars, famines, and plagues, that we'll have much worse problems to worry about. There, too, we don't have to do anything, if we're waiting for the end of the world. It's only as long as we believe in something in between that we have to do something. I'm serious, and I do believe in something in between—my children and grandchildren. But on the other hand, I like to defer my taxes. I especially don't pay taxes today that I won't owe until 10 years from now—that would be foolish.

So I think we should start stimulating our economies to develop those wonderful technologies the optimists think might happen. We have to work on conservation and stimulate the marketplace to prepare for limits on combustion by developing those other power sources. We could stimulate the marketplace by imposing a tax on people who exceed some emissions quota, or allowing people to sell credits to produce carbon dioxide. Let's not clamp down on the economy and send it into a depression—let's push it a little bit instead, so that this wonderful world of cheap, clean energy will actually come to pass. □

Professor of Planetary Science Andrew P. Ingersoll earned his BA in physics from Amherst College in 1960, and his AM and PhD in atmospheric physics from Harvard in 1961 and 1966, respectively. He immediately came to Caltech as an assistant professor, becoming an associate professor in 1971 and a full professor in 1976. At one time or another, Ingersoll has worked on the atmospheres of Venus, Earth, Mars, Jupiter, Saturn, Uranus, and Neptune as part of the Mars Global Surveyor, Cassini, Galileo, VEGA, Voyager, Nimbus, Viking, and Pioneer teams. He has five children and six grandchildren and is optimistic about the future.

This article was adapted from a recent Watson lecture.