CHAPTER 1

The Greenhouse Effect

The global warming forecast is not new, nor has it changed much over the last century. The basic physics of the greenhouse effect was described in 1827 by Jean Baptiste Joseph Fourier. Fourier was a mathematician in Bonaparte's army in Egypt. His name is best known for the Fourier transform, a mathematical technique for separating some complicated signal (such as the history of temperature through time, to choose an apropos example) into the sum of simple waves of different frequencies (such as the day/night cycle and the annual cycle), what we call calculating a spectrum.

Fourier's contribution to Earth science is the idea that gases in the atmosphere that absorb infrared radiation could eventually warm up the surface of the earth. He made the analogy of a greenhouse, but the actual name "greenhouse effect" came later.

The temperature of a planet is set by a natural thermostat, which balances the planet's energy budget. Energy comes in to the Earth as sunlight and leaves as infrared. The greenhouse effect of a gas changes the outgoing part of the budget, the infrared. All objects warmer than absolute zero shine in the infrared. A hot heating element glows red that we can see; the same object at room temperature glows in the infrared.

The rate of energy loss from an object as infrared radiation depends on the temperature of the object. According to the Stefan–Boltzmann relation, the object loses energy at a rate of σT^4, where σ is the Stefan–Boltzmann constant (just a number one can look up in a reference book) and T^4 is the temperature of the Earth in kelvins raised to the fourth power. When the object is hot, it sheds energy much more quickly than when it is cool.

The planet balances its energy budget by warming up or cooling down until the energy loss to space equals the energy gain from the Sun, as in the top panel of Figure 1. The thermostat is a by-product of the need to balance the energy budget. The idea is analogous to water running through a sink, as in Figure 1, bottom panel. The faucet is on, and water is falling into the sink. The drain is open at the bottom of the sink, and the higher the water level is in the sink, the faster it will drain.

When the faucet is initially turned on, water flows in faster than it flows out, and the water level in the sink rises. The sink fills until water is running down the drain as quickly as it is coming out of the faucet. If we start the sink experiment with too much water in the sink, water would drain faster than it filled until it reached that same balancing water level.

If we give our no-atmosphere planet the same energy input from sunlight that the Earth enjoys, it would have an average temperature of about 3°F or −16°C, sub-freezing temperatures around the world. Fourier's greenhouse effect is what's keeping Earth so much warmer than this poor cold naked planet.

Fourier's insight was to add a layer of atmosphere to the planet, which absorbs and emits infrared radiation (Figure 2). The Earth's surface receives energy from the Sun, as before, and it also receives energy from infrared radiation shining down from the atmosphere. The temperature of the Earth's surface rises to about 86°F or 30°C. That's a bit on the high side, but much closer to the real temperature of the Earth.

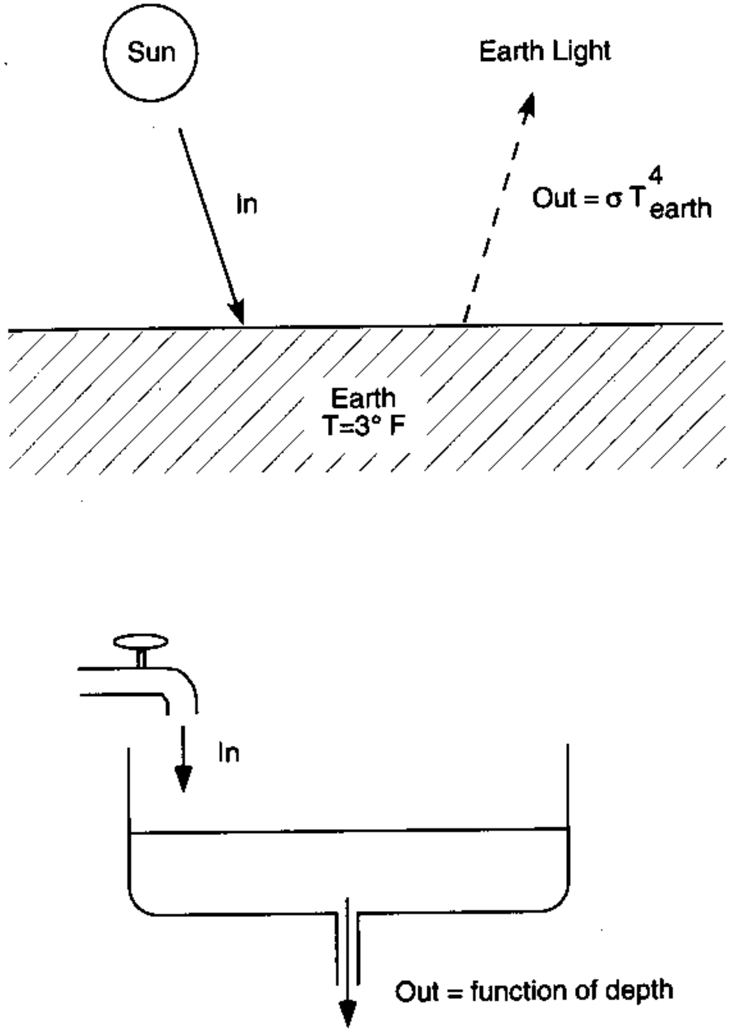

FIGURE 1. Top is a diagram of the energy balance of a planet with no atmosphere. The temperature of the planet finds the value at which energy outflow as infrared balances energy influx from the sun. The bottom is a sink, with water flowing in from a faucet and out down the drain. The rate of flow down the drain depends on the water level in the sink. The water level finds the value at which outflow balances inflow.

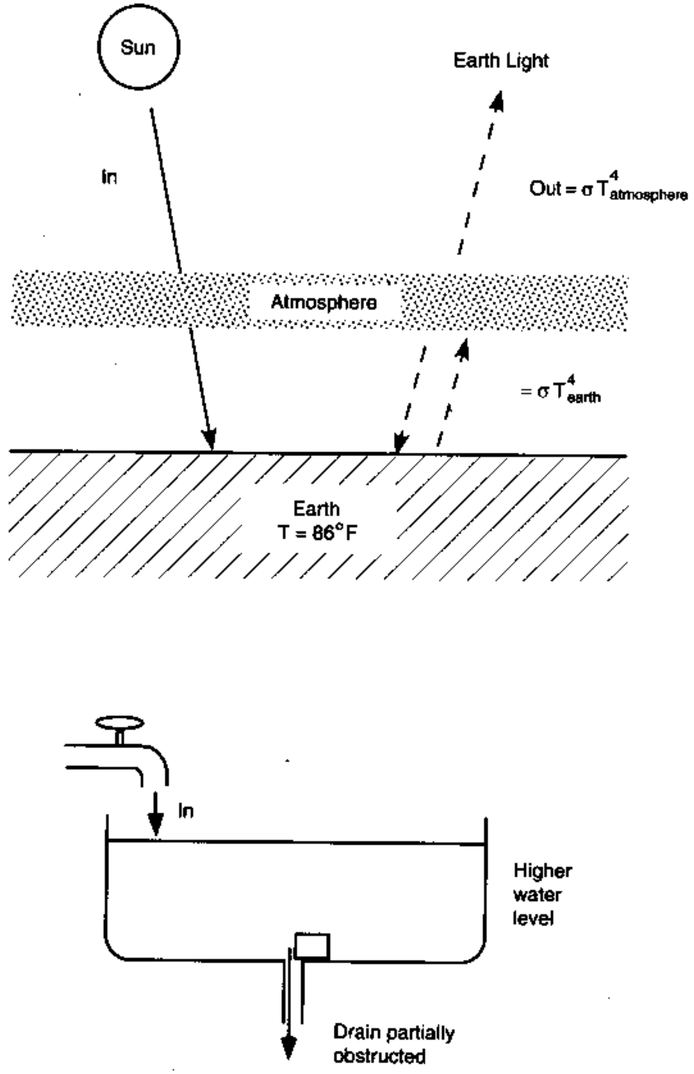

FIGURE 2. Top. A pane of glass analogous to the atmosphere absorbs infrared radiation from the ground, and radiates infrared at its own temperature. The atmosphere is colder than the ground, so the infrared radiation is impeded by the presence of the atmosphere. This is analogous to partially blocking the drain in the sink (bottom), which causes the water level in the sink to rise.

The greenhouse gas in Earth's energy balance is analogous to a partial obstruction of the drain at the bottom of the sink. A grape or piece of cucumber falls into the strainer, slowing down the drainage. The water level in the sink rises until it gets deep enough to force water through the obstructed drain as quickly as it flows in from the faucet. Let's hope the sink reaches a new balanced water budget before it overflows.

Just over a century ago, in 1896, Svante Arrhenius, a Swedish chemist, took the most astonishing leap I have ever read in climate science. Arrhenius used measurements of the brightness of infrared radiation from the moon to predict the temperature change you would get from raising CO_2. Arrhenius estimated a quantity which we now call the climate sensitivity, abbreviated as ΔT_{2x}. This is defined as the amount of warming that the Earth would undergo, on average, from a doubling of the atmospheric CO_2 concentration. The climate sensitivity is probably the first benchmark that two climate scientists in a bar would use to compare two different climate models.

The moonlight infrared data came from Samuel Pierpont Langley, who was trying to determine the temperature of the moon. The hotter the moon is, thought Langley, the brighter it shines in the infrared, the same story as for the Earth. "The dark rays" as they called them, were separated into different bands of wavelengths (different colors, if we could see them) by using a prism made of salt, because salt is one of the few solid substances that doesn't absorb infrared radiation. The intensity of the different invisible beams was measured using something called a bolometer, a device that measures the rate at which the invisible incoming light warms up a thermometer. It all must have seemed rather spooky.

Arrhenius used the data in a way that was not intended by Langley. Arrhenius looked for changes in the intensity of the "dark rays" that vary with humidity, and with the moon angle overhead, which affects the amount of atmosphere the light had to go through. In the moonlight data, more moonlight is absorbed when the light passes through more CO_2 or more water vapor. Arrhenius used this relationship in the data to predict how much the Earth would warm if you doubled CO_2. It was as though, analyzing the water flowing through our sink, Arrhenius calculated precisely how much the flow would slow down if you put a piece of carrot on the drain trap, obstructing the flow of water through the drain, and how much higher the water level in the sink would be.

The surface of the Earth does not all have the same temperature, though, the way that a sink has only one water level. Arrhenius did his calculation on a latitude and longitude grid, just as climate models do today, writing, "I should certainly not have undertaken these tedious calculations if an extraordinary interest had not been connected with them." After two years of pencil-and-paper arithmetic, he concluded that doubling the CO_2 concentration of the atmosphere would lead to 4 to 6°C of warming. Today, with the benefit of a century of innovation, hard work, and exploding computing power, we now estimate that doubling CO_2 would lead to about 2.5 to 4°C of warming. There have been revisions, discoveries, missteps, and wrong directions, as in any science, but on the whole not much has changed in the past century.

So what have climate scientists been doing in the meantime? Climate science has really exploded in the past few decades, as global warming grew from a prediction into an observation in the real world. Globally, about 2 billion dollars per year are being spent on climate change research, 50% of this in the United States. This sounds like a lot of money, and it is, but to put it into perspective, it amounts to only about 5% of the profits from the Exxon Mobil Oil Company. Much of the climate research money is used to pay for satellites that monitor various aspects of the climate of the earth from space. Satellites are expensive. Meanwhile, thousands of scientists worldwide, at universities mostly, are hard at work developing climate models and theory, analyzing meteorological data, and reconstructing climates of the past. This form of research has an entrepreneurial feel to it: individuals or small groups, looking for the new angle that will get them funded and published. The scientific literature about global warming has exploded in the last decades, rising from about a hundred papers per year in the 1980s to a thousand per year today (Figure 3).

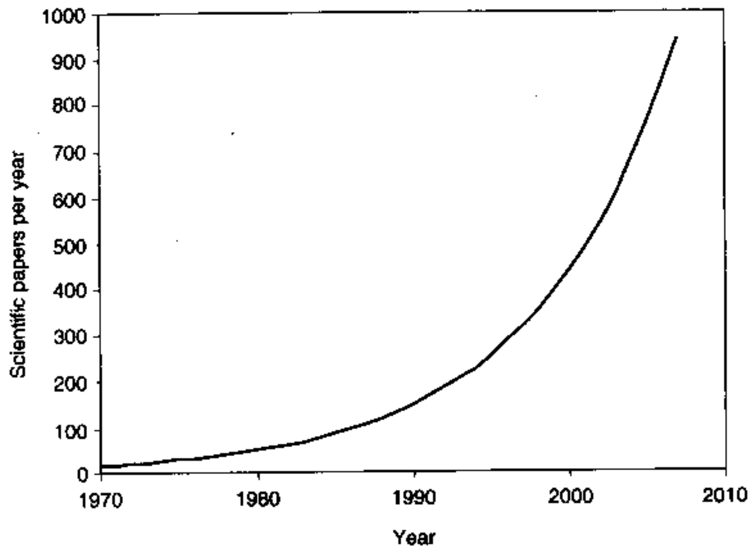

FIGURE 3. The rate of publication of scientific papers about climate in the past 35 years.

Climate science is interdisciplinary enough that it is a challenge to synthesize the bits and pieces. For example to understand climate change in the Arctic requires soil science, forestry, atmospheric and ocean physics, polar bear biology, and other scientific specialties. The state of the warming forecast for the entire globe encompasses so much information that no one human mind could hold it all at one time (not mine, anyway).

In response to warnings of the threat of global warming, the World Health Organization created an organization of scientists charged with the task of summarizing the state of the science, called the Intergovernmental Panel on Climate Change or IPCC. The function of IPCC is not to do new research, but rather to summarize and synthesize all the published scientific papers into coherent reports. The scientists who do the actual work for IPCC are mostly employed by universities and national research laboratories around the world like NASA and NOAA. Working Group I of the IPCC writes the Scientific Assessment reports, while Working Group II reports on Impacts of climate change, and III on Mitigation (reducing CO_2 emissions, mostly). The most recent IPCC reports were released in spring of 2007. The projections and impacts of global warming as presented in the next two chapters are based on information from this report.

Most of the major ingredients in the global warming forecast were there in the results of Arrhenius' tedious calculations. One important example is called the ice albedo feedback. The word albedo describes the reflectivity of a planet to visible sunlight. Clouds reflect sunlight, as does ice and snow. When sunlight is reflected to space, it would be analogous to water from the faucet in the sink analogy that splashes onto the floor. Since that water doesn't have to go down the drain, the water level in the sink decreases. The water level is analogous to Earth's temperature, which falls if more incoming sunlight is reflected back to space instead of being absorbed.

The coupling between ice and light works out to be a loop of cause and effect called a feedback. The air warms for some external reason like rising CO_2, and as a result ice and snow melt on the land or ocean surface. The ice and snow are very reflective, which helped keep the planet cool, but the ground or ocean underneath have a greater tendency to absorb incoming sunlight, so the planet warms more than it would have. This is an example of a positive, amplifying feedback.

The Arctic warms more intensely than the tropics, because ice melts in the Arctic and the bare ground absorbs more sunlight than the ice did. You can see it in Arrhenius' results, you can see it in the Arctic climate records of the past few decades, and you can see it in the global warming forecast for the future. Full disclosure: where you can't see it is in Antarctica. It's a bit of a mystery how cold it's been in Antarctica; it may have something to do with the ozone hole.

Another amplifying feedback to global warming involves water vapor. Water vapor is a greenhouse gas, responsible for capturing more of the outgoing infrared radiation in the atmosphere than CO_2 does. The fact that water vapor is a stronger greenhouse gas than CO_2 does not mean we needn't worry about rising CO_2 concentrations. The concentration of water vapor in the atmosphere is controlled by the fact that if the humidity gets too high, it rains. Warmer air can carry more water vapor than cool air can, so warming from rising CO_2 could lead to more water vapor in the atmosphere. Water vapor warms the Earth still further, because it is a greenhouse gas.

Like the ice albedo feedback, the water vapor feedback is an amplifier of global warming. Unlike ice albedo, which is confined to high latitudes, the water vapor feedback has a rather more uniform effect around the globe, and it about doubles the temperature change we expect from rising CO_2 alone.

There is uncertainty about how strong the water vapor feedback is. The question is whether a warmer world could be drier or wetter than we expect it to be. The average relative humidity of the Earth's surface is about 80%. Arrhenius assumed that the atmosphere would remain 80% relative humidity as it warmed. A relative humidity of 80% represents more actual molecules of water in the warm atmosphere than in the cooler atmosphere, because warm air holds more vapor than cool air. Modern climate models also predict that the relative humidity will not change much with global warming. If the real atmosphere turns out to get wetter with rising CO_2 than models predict, for example, the real water vapor feedback would be stronger than we expect.

Though the answer hasn't changed much, the quality of the answer has certainly improved in the last century. Many pieces that Arrhenius simply had to guess at can now be predicted based on a mechanistic understanding of how things work. Just as important, the models that make the predictions have been tested against reality. In the 1930s, scientists were excited by a theory that sunspots controlled climate by changing the intensity of the Sun. A prediction was made, based on the weather patterns of the recent past, that it ought to get drier in Africa during the sunspot minimum of the 1930s. It turned out that Africa got wetter during the sunspot minimum, so that was it for sunspot theory. The intensity of the Sun is currently thought to have a large impact on century-timescale climate fluctuations such as the Medieval Optimum and Little Ice Age climates, described in Chapter 3. But variations in solar intensity in the last few decades have been weak compared with the change in climate forcing from greenhouse gases.

One problem that might seem like a show-stopper for climate forecasting is the discovery in the 1960s by Edward Lorenz that the weather is fundamentally unpredictable beyond a time horizon of a week or two. One popular name for this phenomenon is "chaos," and another is the "butterfly effect." The idea is that two nearly identical states of the weather, differing only a little bit, will tend to diverge from each other, so that a small initial difference between the two will grow with time. Small imperfections in a model of the weather today will grow, until eventually all that is left in the model is amplified garbage. The weather forecast for tomorrow is pretty good, and my impression is that the forecasts have been getting better every year. The weather forecast for 10 days from now however has always been and continues to be pretty much useless. How can we expect to forecast the weather in 100 years, let alone in 100 millennia, if we can't do 10 days?

The answer is that no one is attempting to forecast the particular weather to expect on a particular day a century from now. Individual fluctuations of weather are chaotic, but the time-averaged weather, called climate, is not. Drawing once again on our sink analogy (Figure 1), waves on the surface of the water could be called weather, while the average water level in the sink would be climate. Climate is constrained by the simple, system-wide energy budget just as the water level in the sink is constrained by water through-flow. Predicting weather would be like predicting the waves in the sink, which requires that you know a lot more about the water in the sink than just the through-flow.

Climate models have their own weather, which is a useful estimate of the statistics of future climate, the frequency of storms and things like that. And long-time averages, say a temperature average over 10 Januaries, can be compared between the model and the real world. Perhaps this is as good a definition for the word "climate" as any; those aspects of the weather that can be predicted far in the future, in spite of the fact that weather is chaotic.

The energy budget of the surface of the Earth varies from place to place, because the temperature varies from place to place. The sink in our analogy only had one water level, but the Earth has a range of surface temperatures. The energy flowing in as sunlight might get transported to another location by the winds or ocean currents before it is lost to space as infrared.

Energy tends to be exported from the tropics, where the sunlight is most intense, to higher latitudes. The high latitudes act like cooling fans of the planet, keeping the tropics cool by carrying heat away and helping to ditch it into space. Dramatically, if the tropics were isolated from the high latitudes, unable to use the poles as cooling fans, the oceans in the tropics could boil in a phenomenon called a runaway greenhouse effect. We are not in danger of experiencing a runaway greenhouse effect on Earth, but it happened on Venus.

Heat is carried around the surface of the Earth by fluid flow, which is tricky to simulate or understand. Simple, slow, gooey flows like molasses are fairly easy to describe, but when the flow becomes turbulent, forget about it. Turbulent flow is one of the great challenges for computation, because there is interesting stuff going on at a huge range of spatial scales. On Earth, circulation patterns range in scale from millimeters up to the size of the Earth.

Some phenomena in nature, such as the ballistic arc of a baseball, can be described pretty well by simple equations. Unfortunately, there are no simple equations that capture most types of fluid flow. Fluid flow can be simulated on a computer by chopping the domain of the problem (the Earth's atmosphere or oceans) into pieces or blocks on a 3-dimensional grid. Each block has a single temperature, one wind velocity, a water vapor content, etc.

We can stretch the sink analogy to correspond to our multitemperatured Earth, although the analogy is starting to get contorted. We would need an array of sinks, each with drains, and separate faucets, and water that would be allowed to flow from sink to sink. Each sink could have a slightly different water level from its neighbors. Some sinks would generally have more water than others, but there would also be a good deal of sloshing around (weather).

Climate science is climbing a brick wall in the fact that increasing the detail of the simulation makes the computer program run much much more slowly. If you want to double the model resolution, it requires twice as many grid points in each of three dimensions, equaling eight times more work to do per time step. Making matters worse, time steps have to get shorter as the grid boxes get smaller, or else the model crashes. Doubling the resolution of a calculation results in a model that will run sixteen times slower.

Clouds are probably the toughest challenge to simulate in a rigorous, mechanistic, first-principles way. The character of the cloudy skies on Earth depends on collisions between droplets on spatial scales of millimeters, on upward and downward gusting winds on spatial scales of meters, the convergence of winds on the storm scale of 100 kilometers, and on the global atmospheric circulation. Doing this right would take a lot of grid points.

The ideal thing would be to put all of this complexity into a computer model that only knows fundamentals of physics and chemistry, and have the model predict what clouds should look like. Talk about tedious calculations; this would be too much even for the fastest computer. For many years (an eternity in computer time), the fastest computer on Earth was a Japanese machine called the Earth Simulator; this machine was not nearly fast enough to resolve all of the physics of clouds and turbulence in the Earth's climate system. Even with the explosive growth in computing power described by Moore's law, computers in the foreseeable future are never going to be up to doing the calculation that climate scientists would be most happy with.

Plan B is to program into the model the large-scale behavior of, say, the cloudiness of the atmosphere, based on observations of cloudiness. Each grid cell box in an atmosphere model keeps track of the temperature and vapor content of the air in it, as well as the number of cloud droplets per cubic meter, the total water content, and maybe something about the size distribution of the droplets. The cloud subroutine makes a guess about how much water evaporates or condenses in each time step, and how the droplets coalesce. The scheme does not rely solely on the underlying fundamental mechanisms for the process, as would be ideal, but rather tries to capture the observed behavior in a more made-up way. The name for this approach is parameterization.

The law of supply and demand, in economic models, could be described as a parameterization. Supply and demand curves describe an emergent behavior of an economic system, a description of a result rather than a fundamental mechanism. The fundamental mechanism in this case would have to do with individual investors, which would have to be simulated in the computer program, gloating and gnashing their teeth and feeling envy and fear and greed and social ambition and reading their horoscopes. Climate may seem computationally intractable, but it is much easier to model climate than it is to model economics.

A climate scientist might come up with a scheme to describe clouds, and see that it captures the variability of the real Earth. The real earth does span a wide range of variation, from the tropics to the poles, deserts to jungles, mountains and plains. If the scheme is able to predict all of the variations in cloudiness on Earth today, then perhaps it will also capture any change in cloudiness as Earth's climate changes. However, because a parameterization is not built from the ground up using only the fundamental building blocks of physics and chemistry, it comes with no guarantees that it will change realistically if the climate upon which it's based changes too much.

There are many different ways to cook up a parameterized cloud, and it is done differently in different models. Some parameterizations are better than others. Often the best reassurances that these parameterizations are not deluded come from formal model intercomparison projects. In the 2007 IPCC Scientific Assessment, there are intercomparisons between nineteen different climate models, each developed by separate, competing groups of scientists. The models are also compared with measurements from the real world, present-day measurements or inferred climate parameters from the past. In practice the "duplicate, compete, and compare" approach seems to function fairly well at rooting out mistakes and bias.

Uncertainty in the science of climate change is often used as an argument not to worry about global warming. That logic makes intuitive sense if one thinks from a reference point of an unchanging climate. The forecast says "warming" but it could be wrong, therefore there might not be warming.

But it is known with certainty that CO_2 affects the climate. If Fourier's greenhouse effect were wrong, Earth's natural climate would be much colder than it is. It is certain also that CO_2 levels in the atmosphere are rising. The response to rising CO_2 is certainly some degree of warming; no effect or a cooling effect can be ruled out.

So the forecast calls for warming, but the warming could be more or less than the forecast calls for. In general, past climate changes, described in Section 2 of this book, are more intense than we would have expected. The future could also be worse than expected. Uncertainty in the climate forecast, when we think about it carefully and honestly, is no argument for complacency.