

Global Warming: Understanding the Forecast

David Archer
d-archer@uchicago.edu

University of Chicago

6/13/05

Table of Contents

Part I The Greenhouse Effect

- Chapter 1 Blackbody Radiation
- Chapter 2 The Layer Model
- Chapter 3 Greenhouse Gases
- Chapter 4 Convection
- Chapter 5 Heat, Winds, and Currents
- Chapter 6 Feedbacks

Part II The Carbon Cycle

- Chapter 7 Carbon on Earth
- Chapter 8 Fossil Fuels
- Chapter 9 The Perturbed Carbon Cycle

Part III The Forecast

- Chapter 10 Is It Reliable?
- Chapter 11 The Forecast
- Chapter 12 Decisions, Decisions

Part I. The Greenhouse Effect

Chapter 1. Blackbody Radiation

Summary

Energy travels through a vacuum between the sun and the earth by means of light. Objects can absorb energy from light, and they can also emit light energy, if the vibrations of their chemical bonds generate oscillations of the electrical field. An object that can emit all wavelengths of light is called a blackbody. The spectrum of light from a blackbody sparked the development of quantum mechanics, a revolution in physics and natural philosophy.

Heat and Light

Heat is simply the bouncing-around energy of atoms. You've heard that before, but have you ever wondered how we can feel such a thing? Atoms are tiny things. We can't really feel individual atoms. But it doesn't take any state-of-the-art laboratory technology to tell how fast the iron atoms in your stove burner are bouncing around. All you have to do is touch it, if you dare. Actually, another method may be occurring to you, that you could look at it; if it's hot, it will glow. It glows with blackbody radiation, which we'll get back to later.

You can feel the hot stove because the energetic bounciness of the stove atoms gets transferred to the atoms in your nerves in your fingers. The fast-moving atoms of the burner bounce off the atoms in your finger, and the fast ones slow down a bit and the slower ones bounce back with a little more energy than they came in with. Biological systems have evolved to pay attention to this, which is why you can feel it, because too much bounciness in our atoms is a dangerous thing. Chemical bonds break when they are heated too much, and we could cook ourselves. Burning your finger by touching a hot electric stove burner is an example of heat **conduction**, the easiest form of heat flux to imagine.

Energy Through a Vacuum

A thermos bottle is designed to slow the flow of heat through its walls. You can put warm stuff in there to keep it warm, or cool stuff to stay cool. Thermos bottles have two walls; an inside and an outside. In between the two walls is an insulator. The best insulator you can get is a vacuum. A vacuum ideally has no molecules or atoms of gas in it. If there are no atoms, there is nothing to conduct heat between the inner and outer walls of the thermos. Of course there will still be heat conduction along the walls, but a thermos bottle is not really a perfect thought experiment. Let's think about a planet. There are no walls, however thin, connecting a planet with anything. The space between the earth and the sun is not a perfect vacuum, but the amount of heat you can carry in the solar wind is negligible. Peanuts. We know how warm it is in the sunshine, so we know

that heat flows from the sun to the earth. And yet separating the sun from the earth is 150 million kilometers of vacuum. How can heat be carried through a vacuum?

Light carries heat through a vacuum. Electrons and protons have a property called **electric charge**. What an electric charge is, fundamentally, no one can tell you, but it interacts with a property of the vacuum called the **electric field**. A positive electric field attracts a negatively charged electron. This is how a TV tube hurls electrons toward the screen in a picture-tube television, in the old days, before electronic flatscreens, by releasing an electron and then pushing it around with the electric field. The electric field interacts with another property of the vacuum called the **magnetic field**. If the strength of the electric field at some location, measured in volts, changes with time, if the voltage is oscillating up and down for example, that will cause a change in the magnetic field, such as a compass would point to. This is how an electro-magnet works. If the magnetic field changes, it can produce an electrical field. This is how a generator works.

It turns out that the electric and magnetic fields in the vacuum fit together to form a closed cycle like the ringing of a bell. We will see other such self-sustaining, “ringing” oscillations, so let’s assign them a name: **self-sustaining oscillations**. An up-and-down oscillation in the electric field will cause a complementary oscillation in the magnetic field, which reinforce the electric field in turn. The two fields “ring” together. Such a little bundle of electric and magnetic waves can in principal hurl through the vacuum forever, carrying energy with it.

The ringing of the electromagnetic field in light differs from the ringing of a piano string, in that light can come in any frequency. **Frequencies**, of oscillators or of light waves, have units of cycles/second and are denoted by the greek letter ν , pronounced “new”. It turns out that all different frequencies of light travel at the same speed in a vacuum. Within some non-vacuum medium such as air, water, or glass, different frequencies of light might vary in their speeds a little bit, which is how a prism separates white light into its component colors. But in a vacuum, all light travels at the same speed. The **speed of light in a vacuum, c**, is another fundamental constant of nature. The constancy of the speed of light in a vacuum is a gift to us from the universe, because it allows us to relate the frequency of the light to its **wavelength**, the distance between the crests of a wave. We can figure out what the relationship is between frequency and wavelength by thinking geometrically, imagining the wavy line in [Figure 1-1](#) to be moving past us at some speed c . If the crests are 1 cm apart and they were moving at 10 cm/second, then 10 crests would move past us every second. Alternatively, we can make use of units. Pay attention to units, and they will lead you to virtue, or at least to the right answer. Problem: assemble the two things we know, ν and c , such that the units combine to be the same as the units of λ . Solution:

$$\lambda \left[\frac{cm}{cycle} \right] = \frac{c \left[\frac{cm}{sec} \right]}{\nu \left[\frac{cycle}{sec} \right]}$$

Don't take my word for it, check and make sure the units are the same on both sides of the equation.

Scientists who discuss infrared light often use a third way of describing different colors, called the **wave number**, defined as the number of cycles per cm of length. It's simply a matter of convenience; when infrared light is described by its wave number, it will have a nice, easy-to-remember number to carry around in our memories. We will use the letter **n** to designate the wave number. How can we construct the units of n from the building blocks of λ , ν , and c ? All we need is λ ,

$$n \left[\frac{\text{cycles}}{\text{cm}} \right] = \frac{1}{\lambda \left[\frac{\text{cm}}{\text{cycle}} \right]}$$

Different frequencies of light all have the same fundamental nature; they are all waves of the same essential physics. **Figure 1-2** shows the names we have assigned to different types of light, based on their frequencies. Of course, if we know the frequency of light, we know its wavelength and wavenumbers also, so we have added "mile markers" of in these units too.

Our eyes are sensitive to light in what we sensibly call the visible range. Light of higher frequencies than this we call ultraviolet, or **UV**, then x-rays, then gamma rays. Different frequencies of light are all waves of the same fundamental type, but higher frequencies of light carry more energy than do lower frequencies of light. This is why ultraviolet rays cause sunburn, and x-rays are even more dangerous. Objects that emit the highest frequency light are considered radioactive. Extremely short wavelength light is called gamma radiation, and we might encounter it coming from processes that happen in the nuclei of atoms, or processes that happen in space such as the explosions of stars, coming to us as cosmic rays. (Other common forms of "radiation" are made of flying electrons (beta radiation) or bundles of two each protons and neutrons, called alpha radiation.)

At longer wavelengths than visible we find the infrared range, or **IR**. Later we will find that objects at about room temperature glow with IR light. Heat lamps at the skating rink warm us by shining invisible IR light on us.

All that energy whizzing around space in the form of coupled electric and magnetic field waves would be of little interest to the energy balance of a planet, if they did not give up or carry off energy. Neutrinos carry energy, but don't interact with matter. They can zing right through the earth and never even slow down. Neutrinos do not contribute to global warming. Light, in contrast, does exchange energy with matter. Light is the way that energy travels through the vacuum inside a thermos bottle, or surrounding a planet.

There are a number of different mechanisms by which light may interact with matter, but infrared light interacts mostly with vibrations of the chemical bonds of a molecule.

Light interacts with matter by means of the electric field that they share (Figure 1-3). Let's imagine matter as constructed with charged oscillators all over its surface; little charged weights sticking out from the surface of the matter on little springs that stretch and contract. This little oscillator will have a frequency with which it will "ring", a self-sustaining oscillation such as we defined above. (The other example was the light wave). Imagine one of those African thumb harps, with little keys of metal which are plucked with the thumbs. There are a series of keys, of different lengths or weights to ring at different frequencies (different notes). Imagine matter as looking something like an African thumb harp, with electric charges on the plucking ends of the harp keys.

Incoming energy in the form of light brings with it an electric field oscillating back and forth: voltage goes plus, minus, plus, minus. If the frequency of the cycles of voltage in the light is the same as the frequency of a key on the thumb harp, the light can be absorbed. Its energy is transferred into **vibrational energy** of the matter.

Important fact to remember: This mechanism of energy transfer is a two-way street. If energy can flow from the light to the oscillator, it will also be able to go the other way, from the oscillator to light. The vibrational energy of the oscillator is what we have been calling its temperature. Any matter that has a temperature above absolute zero (zero degrees on the Kelvin scale) will have energy in its oscillators that it may be able to use to create light. The two-way street character of this process is important enough that it is given a name of **Kirchoff's law**.

You can think of it as analogous to the vibrations of the keys in the thumb harp or the strings of a piano interacting with the pressure field of the atmosphere to generate sound waves. You may have tried the experiment of singing into a piano with dampers off to hear the strings echo back the note you sang into it. This wave to oscillator energy transfer is a two-way street, as well.

Blackbody Radiation

So where can we see electrical energy traveling the other way, from matter into light? One example: a red hot electric burner shines light you can see. The light derives its energy from the vibrations or thermal energy of the matter. We normally don't think of it, but it turns out that your electric burner continues to shine even when the stove is at room temperature. The difference is that the room temperature stove emits light in colors that we can't see, down in the infrared.

A real African thumb harp will not interact with all possible frequencies of sound. There are only a few notes the instrument can make, and the vast majority of possible notes it cannot make. If we imagine our conceptual chunk of matter as having all of the notes, it would be able to absorb or emit all different frequencies of sound. We have a word for that: we call that object a **blackbody**. The light that is emitted by a blackbody is called **blackbody radiation**. In this book we are going to talk a lot about matter at room temperature and colder generating light. Objects in this temperature range, like the turned-off burner, emit light in the infrared. So for our purposes, when we call an object a blackbody we will mean that the object radiates energy at all frequencies of the

infrared. And of course, if an object radiates at some frequency, it must also absorb at that frequency, so a blackbody is a perfect infrared absorber as well.

Blackbody radiation is made up of a characteristic distribution of frequencies (colors) of infrared light. **Figure 1-4** shows a plot with axes of the intensity of light in the y-direction and frequency in the x-direction. The units of intensity look a bit complicated; they are $\frac{W}{m^2 \text{ wavenumber}}$. The first part, W/m², you recognize. Dividing by wavenumbers allows us to add up the area under the curve to obtain the total energy intensity in W/m². The intensity of light at each frequency is called a spectrum. The infrared light emission spectrum of a blackbody depends only upon the temperature of the blackbody. There are two things you should notice about the shapes of the curves in **Figure 1-4**.

First, as the temperature goes up, the curves are getting higher, meaning that light of each frequency is getting more intense (brighter). When we start talking about planets we will need to know how much energy is being radiated from a blackbody in total, over all wavelengths. The units in **Figure 1-4** were chosen specifically so that the total energy being carried by all frequencies of light is equal to the area under the curve of the spectrum. As the temperature of the object goes up, the total energy emitted by the object goes up, which you can see by the fact that the curve in **Figure 1-4** is bigger. There is an equation which tells us how quickly energy is radiated from an object. It is called the **Stefan-Boltzmann equation**, and we are going to make extensive use of it. Get to know it now! The equation is

$$I = \epsilon \sigma T^4 \tag{1-1}$$

The **intensity** of the light, **I**, is defined as the number of Watts of energy carried per square meter of surface area that the light is shining on. The greek letter epsilon (ϵ) is the **emissivity**, a number between zero and one describing how good a blackbody the object is. For a perfect blackbody, $\epsilon = 1$. Sigma (σ) is a fundamental constant of physics which never changes, a number you can look up in reference books, called the **Stefan-Boltzmann constant**. T is the temperature in Kelvins, and the superscript 4 is an exponent, indicating that we have to raise the temperature to the fourth power.

One of the many joys of thinking scientifically is the importance of units. Approach this with the same willingness to be intrigued as you might approach a crossword puzzle. Try. Let us examine equation 1 again, with units of the various terms specified in the brackets.

$$I \left[\frac{W}{m^2} \right] = \epsilon [unitless] \sigma \left[\frac{W}{m^2 K^4} \right] T [K]^4$$

The unit of energy is Watts [W], the same kind of Watts that describe your hairdryers and audio amplifiers. The meters squared on the bottom of that fraction is the surface area of the object that's radiating. The area of the earth, for example, is $5.14 \cdot 10^{14} \text{ m}^2$.

Temperature is in Kelvins, and ϵ has no units; it is just a number. Here's what I wanted to point out: The units on each side of this equation must be the same. On the right-hand side, K^4 cancels leaving only W/m^2 to balance the left-hand side. In general, if are unsure how to relate one number to another, the first thing to do is to listen to the units. They will guide you to virtue. We will see many more examples of units in our discussion, and you may rest assured I will not miss an opportunity to point them out.

Infra-red sensitive cameras allow us to see what the world looks like in infra-red light. The cheeks of the guy in [Figure 1-5](#) are warmer than the surface of his glasses, which are presumably at room temperature. We can estimate how much more light is shining from the cheeks than from the glasses surface using equation 1, to be

$$\frac{I_{cheek}}{I_{glasses}} = \frac{\epsilon_{cheek} \sigma T_{cheek}^4}{\epsilon_{glasses} \sigma T_{glasses}^4}$$

The Stefan-Boltzmann constant σ is the same on both top and bottom; σ never changes. The emissivity ϵ might be different between the cheek and the glasses, but let's assume they are the same. This leaves us with the ratio of the brightnesses of the skin and glasses equal to the ratio of temperatures to the fourth power, maybe $(285 \text{ K} / 278 \text{ K})^4$, which is about 1.1. The cheeks shine 10% more brightly than the surface of the coat, and that is what the IR camera shows us.

The second thing to notice about the effect of temperature on the blackbody spectra in [Figure 1-4](#) is that the peaks shift to the right as the temperature increases. This is the direction of higher frequency light. As we discussed above, higher frequency light must have a shorter wavelength. Higher frequency light is also higher energy light. This makes some kind of intuitive sense to me, anyway; I imagine it would take a lot more energy to oscillate a field quickly than to oscillate it slowly. Then it makes sense that a hotter object, which has more bouncing-around energy contained in its various ways of bouncing around, would generate light of higher energy also.

You already knew that a hotter object generates higher energy, shorter wavelength light, because you know about red hot, white hot. Which is hotter? White hot, of course; any kid on the playground knows that. A room temperature object (say 273 K) glows in the infrared, where we can't see it. A stove at stovetop temperatures (400-500 K) glows in shorter wavelength light, which begins to creep up into the visible part of the spectrum. The lowest energy part of the visible spectrum is red light. Get the object hotter, say the temperature of the surface of the sun (5000 K), and it will fill all wavelengths of the visible part of the spectrum with light. [Figure 1-6](#) compares the spectra of the earth and the sun. You can see that sunlight is visible while "earth light" (also referred to as **terrestrial radiation**) is infrared. Of course, the total energy flux from the sun is much higher than it is from earth. Repeating the calculation we used for the infrared photo, we can calculate that the ratio of the fluxes is $(5000 \text{ K} / 273 \text{ K})^4$, or about 10^5 . The two spectra in [Figure 1-6](#) have been scaled by dividing each curve by the maximum value that the curve reaches, so that the top of each peak is at a value of one. If we hadn't done that,

the area under the earth spectrum would 100,000 smaller than the area under the sun spectrum, and you would need a microscope to see the earth spectrum on the figure.

It is not a coincidence that the sun shines in what we refer to as visible light. Our eyes evolved to be sensitive to visible light. The infrared light field is a much more complicated thing for an organism to measure and understand. For one thing, the eyeball, or whatever light sensor the organism has, will be shining IR light of its own. The organism measures light intensity by measuring how intensely the incoming light deposits energy into oscillators coupled to its nervous system. It must complicate matters if the oscillators are losing energy by radiating light of their own. Infrared telescopes must be cooled in order to make accurate IR intensity measurements. Snakes are able to sense IR light. Perhaps this is useful because their body temperatures are colder than those of their intended prey.

The Ultraviolet Catastrophe

We are about to take a diversion, for the sake of general science education. We are passing by a topic that sparked one of the most profound revolutions in the history of science, and I can't see passing by without stopping for a moment to tell you this story. The shape of the blackbody radiation spectrum was a huge mystery, until it was solved by the development of quantum mechanics. There are two points to telling this story. One is the way that theories and observations are combined to make progress in science. The other is that quantum mechanics is just plain weird.

Scientists observe and measure things, often times seemingly obscure things like the amount of light coming from an object at room temperature as a function of its wavelength. Scientists also make theories which attempt to explain these observations. If the theory fits the data, that seems great; you can hand in your lab report or publish your paper, no problem. However, paradoxically, the only time that a scientist really has any hope of learning anything new is when the theory fails to reproduce the observations. Then you've learned, without a doubt, that something is going on in nature that isn't going on in your theory. If the theory reproduces the data beautifully, you haven't learned anything for sure, because the theory could be getting the right answer for the wrong reasons.

Physicists in the nineteenth century had theories that did a pretty good job of explaining most observations. These theories are collectively called "classical" physics, and they included things like Newton's laws of motion, and rules governing energy and how it can be transferred from one form to another. One set of observations which classical physics was unable to explain was the blackbody spectrum. We can describe the situation without going through the mathematics. Imagine a blackbody as an object with charged oscillators on its surface. By the definition of a blackbody, the object has oscillators which vibrate at all different frequencies, and so in principle the object has the potential to radiate light at all different frequencies.

One property of classical mechanics is called **equipartitioning of energy**. This principal holds that, in a system that is all at the same temperature, each different

oscillator will hold, on average, the same amount of energy. This quantity of energy is a simple function of temperature

$$\Delta E = 1/2 k T$$

where k is the **Boltzmann constant**, a number that never changes that you can look up in reference books.

Here comes the part where we're going to skip over doing the physics and mathematics properly, and I'm just going to tell you the result. If we know the amount of energy in an oscillator, we can calculate the intensity of light that will be generated by that oscillator. So, we can assign a temperature to our blackbody and calculate the intensity of light that should be coming from it at all different frequencies. When this calculation is done and compared with the measurements of IR light from a real blackbody, the theory fails, magnificently, spectacularly (**Figure 1-7**). The classical theory predicts that the intensity of light coming from the blackbody ought to increase as the frequency of light increases, forever. The observations show that there is a cutoff; a room temperature object does not shine light in the visible range and beyond. The classical theory says that our room temperature stove burner should be glowing brightly, not only in visible but in ultraviolet, x-rays, and gamma rays. The classical theory says that the object should be radiating away energy at an infinite rate! There is just absolutely no mistaking it; the real world is doing something that the classical theory is not doing.

The missing ingredient in the classical theory is called **quantum mechanics**. Quantum mechanics was developed in the early 20th century to explain the blackbody spectrum and a few other little mysteries, such as the way that light can be converted to electricity by a photoelectric cell. The basic premise of quantum mechanics is that energy comes in undividable chunks rather than being a continuous quantity. The idea is analogous to a glass of water. Water looks continuous, as though you could divide the water into two droplets of whatever sizes you wanted to. At the molecular level, smaller than we can see with our eyes, water comes in chunks, molecules, and viewed in this way, we can see that the division of our glass of water into two cannot be in any proportion we choose, because the water molecules are not dividable. We can remove ten water molecules or eleven, but not ten and a half.

Light comes in waves, but the energy of the light in waves comes in chunks called **photons**. You can have ten photons of light or eleven, but not ten and a half. It turns out that the energy of a photon of light increases as the frequency of the light increases, according to the equation

$$\Delta E = h \nu \tag{1-2}$$

ΔE is the energy of the chunksize, in Joules, h is Planck's constant, in $J \text{ sec cycle}^{-1}$, and ν is the light frequency in cycles/sec.

Vibrational energy in an oscillator comes in chunks, also. It is as if you were swinging on a swing, and you were unable to swing at a height of four feet above the resting position, but could either swing at three feet or five feet. In real life, just like the molecular chunks of water, the chunks are too small to perceive, but at the level of blackbody oscillators and beams of light, they are there and have important effects. The chunks of vibrational energy in an oscillator are given by the equation

$$\Delta E = h \nu \quad (1-3)$$

Equation 1-2 looks just like equation 1-3, doesn't it? When we first introduced the idea of light radiating from a charged oscillator, we stated that the frequency of the light must equal the frequency of the oscillator

$$\nu_{\text{light}} = \nu_{\text{oscillator}}$$

This makes good physical sense; to push a child on a swing you must push at the swing's natural frequency. Quantum mechanics must also explain that observation, which it does by supposing that the energies of the chunks are the same function of frequency for both light waves and for vibrational energy.

$$E_{\text{light}} = E_{\text{oscillator}}$$

The problem with classical mechanics was its assumption that each oscillator had its $1/2 k T$ equipartition allotment of energy (Figure 1-8). Under the new rules of quantum mechanics, it may not be allowed for each oscillator to have exactly $1/2 k T$ worth of energy; energy comes in chunks of size $h \nu$. For the slow oscillators, the frequency ν is small, and so the chunk size $h \nu$ is also small, much smaller than the equipartition energy $1/2 k T$. Therefore the oscillator is free to load up with chunks of vibrational energy until it is pretty close to $1/2 k T$. If the chunks are small enough, they will be like our water molecules and have no impact at all. As we go higher in frequency, however, the interesting thing happens. The chunks get larger and larger as the frequency ν increases, until at some point, the chunks are bigger than the thermodynamic energy $1/2 k T$. Imagine some children's party game which rewards their skill with cash. The game is set up so that the average child will win 25 cents. However, the smallest denomination of money you possess is \$20 bills, and unless a child is some kind of prodigy and does well enough to win \$20, performing completely off-scale above average, he or she gets nothing. The situation with the high-frequency oscillators is analogous. Each would aspire to, on average, $1/2 k T$ amount of energy, but the smallest available chunk size is much larger, $h \nu$. So they go without. The blackbody spectrum follows classical theory for low frequency oscillators and light, where the chunks are small enough not to matter, but at high frequency, the existence of the energy chunks prevents the oscillators from oscillating above its ground state at all. As the temperature of the blackbody increases, so does the thermal energy, and so a few oscillators that were out of reach before become available to radiate light. This explains the red hot, white hot shift toward higher frequency light with increasing temperature of the blackbody.

Box: Quantum Weirdness and Spooky Action at a Distance

The story of blackbody radiation and quantum mechanics is interesting for the way it demonstrates the relation of theory and observation in science, but it is also worthy of a short diversion simply because quantum mechanics is so weird. This is a well-recognized mainstream scientific expression: Quantum Weirdness. The best description of quantum weirdness I have ever heard comes from an experiment called the two-slit experiment. The weirdness arises from the hybrid wave / particle nature of light and also of things that we think of as particles by default, like electrons.

Imagine a light source from which light emerges which is all the same wavelength, and all the peaks and troughs are lined up. The light encounters a board which has two slits cut into it. Light emerges from the two slits to shine onto a screen behind the board with the slits (Figure 1-9).

*The beams of light from the two slits interact with each other. If two beams come together on the viewing screen such that their peaks and troughs coincide with each other, they combine into a brighter wave. If the peak of one wave coincides with the trough of another, they cancel each other out. These are called **constructive** and **destructive interference**. The geometry of the two-slit apparatus is such that bands of light and dark will appear on the viewing screen where the light beams constructively and destructively interfere. This is proof of the wave nature of light. We have not encountered anything terribly weird yet.*

Now let's replace the light source with an electron gun, which shoots electrons instead of light. Picture tube televisions and computer screens have electron guns in them. Again the electrons hit the board with the slits in it, and a few make it through the slits to hit the screen behind it. Let's replace the screen with some sort of electron detector which shows us where the electrons hit, such as the front screen of a TV in which phosphors glow when hit by an electron. Quantum mechanics claims that light comes in chunks, which we could also call particles. Quantum mechanics also claims that things like electrons, which we would most naturally consider particles, are also waves. The wavelength of a particle depends on its mass, so the larger the particle, the smaller its wavelength. For this reason, the wave character of things we like to think of as particles is not obvious to a person walking down the street, just as the particle character of stuff like energy or water is too small to be perceived. In the two-slit experiment, however, the wave character of the electrons becomes important, because the electrons from the two slits interfere with each other. There will be bright bands on the phosphors, and dark bands. Do two electrons whose waves happen to be out of synch just annihilate each other? This is getting a bit stranger to our everyday experience.

Now let's dial the electron gun down so that it is only shooting one electron at a time. Most of the electrons hit the board, but a few make it through the slits to the other side. Let's record where the electrons land on the screen behind by locking

the phosphors after once they are triggered by an electron. What we see is rather startling; the same pattern of light and dark bands emerges on the phosphor screen. Electrons from the two slits are interfering with each other, even though electrons are only being fired one at a time. The only conclusion is that each electron is somehow going through both slits at once. That's starting to get mighty strange.

You refuse to believe it, so you rig a pair of sensors on each slit, to see for yourself if each electron is going through two slits simultaneously. What you see confirms your common sense, that an electron goes through one slit or the other, but never both. But then, another startling thing emerges on the phosphor screen. The interference bands have disappeared. The simple act of watching to see if electrons go through one slit or another has changed the pattern of where the electrons arrive at the phosphor screen.

*Physicists describe this behavior by saying that the act of observing the electrons at the slits **collapses the wave function** of the electron. The unobserved electron exists in this ghostly wave of probability, which can be precisely calculated and verified by experiment. The probability wave tells you where you are most likely to find that electron. When an observation is made, however, the universe has to make a decision; where is the electron going to be? The probability wave collapses to a single point, since you now know where the electron is, the probability of it being where you found it is one and the probability of it being somewhere else is zero. It's no longer a smeared-out wave of probabilities. Here's the weird part. When the electron existed as the wave of probabilities, it interfered with itself. The electron did go through two slits at once, but once we tried to really catch it in the act by observing it, it got its act together and began to act sensibly. The expression "collapsing the wave function" does not imply that anyone knows why things should behave this way. The expression is part of the **Copenhagen Interpretation** of quantum mechanics, which just assures us that this is how things work, and we should try not to worry about how and why too much. Get over it. Another way of explaining the same information is called the **Many Worlds** interpretation, which supposes that the entire universe splits into two parallel universes, just like an old Star Trek rerun, every time an observation forces a probability wave to make a decision about where the electron is: that is every time a wave function collapses. The many worlds interpretation sounds insane but it is just as consistent with the weird observations as any other theory.*

Take-Home Points

Light carries heat through space.

If an object can absorb light, it can also emit light.

An object that can emit all frequencies of light (a blackbody) emits light energy at a rate equal to σT^4 .

Exercises

1. A joule is an amount of energy, and a watt is a rate of using energy, defined as $1 \text{ W} = 1 \text{ J} / \text{s}$. How many joules of energy are required to run a 100 W light bulb for one day? Burning coal yields about $30 \cdot 10^6 \text{ J}$ of energy per kg of coal burned. Assuming that the coal power plant is 30% efficient, how much coal has to be burned to light that light bulb for one day?

2. This is one of those job-interview questions to see how creative you are, analogous to one I heard, "How many airplanes are over Chicago at any given time". You need to make stuff up to get an estimate, demonstrate your management potential. The question is: What is the efficiency of energy production from growing corn?

Assume that sunlight deposits 250 W/m^2 of energy on a corn field, averaging over the day/night cycle. There are 4.186 joules per calorie. How many calories of energy are deposited on a square meter of field over the growing season? Now guess how many ears of corn grow per square meter, and guess what is the number of calories you get for eating an ear of corn. The word "calorie", when you see it on a food label, actually means "kilocalories", thousands of calories, so if you guess 100 food-label-calories, you are actually guessing 100,000 true calories or 100 kcal. Compare the sunlight energy with the corn energy to get the efficiency.

3. Hoover Dam produces $2 \cdot 10^9$ Watts of electricity. It is composed of $7 \cdot 10^9$ kg of concrete. Concrete requires 1 MJ of energy to produce per kg. How much energy did it take to produce the dam? How long is the "energy payback time" for the dam?

The area of Lake Mead, formed by Hoover Dam, is 247 mi^2 . Assuming 250 W/m^2 of sunlight falls on Lake Mead, how much energy could you produce if instead of the lake you installed solar cells that were 12% efficient?

4. It takes approximately $2 \cdot 10^9 \text{ J}$ of energy to manufacture 1 square meter of crystalline-silicon photovoltaic cell. (Actually the number quoted was 600 kWhr. Can you figure out how to convert kilo-watt hours into Joules?) Assume that the solar cell is 12% efficient, and calculate how long it would take, given 250 W/m^2 of sunlight, for the solar cell to repay the energy it cost for its manufacture.

5. Infrared light has a wavelength of about 10 microns. What is its wavenumber in cm^{-1} ?

Visible light has a wavelength of about 0.5 microns. What is its frequency in Hz (cycles per second)?

FM radio operates at a frequency of about 100 MHz. What is its wavelength?

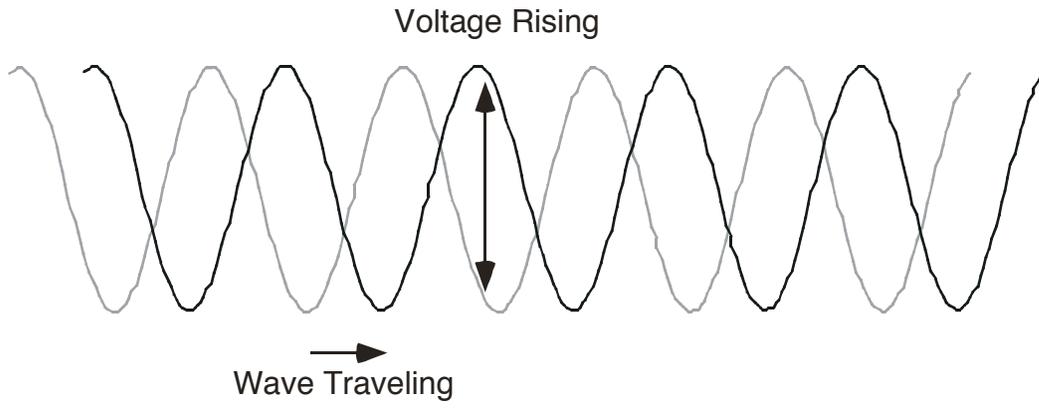


Figure 1-1

Frequency, Hz

Wavelength, nm

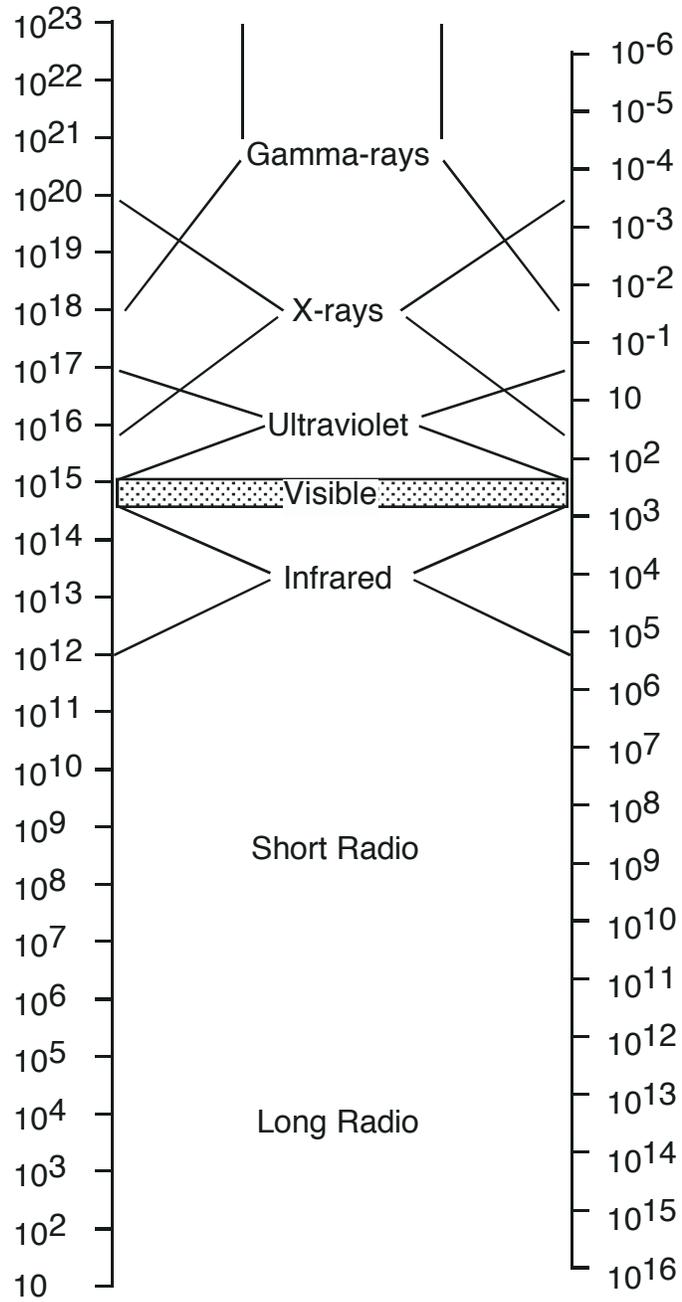


Figure 1-2

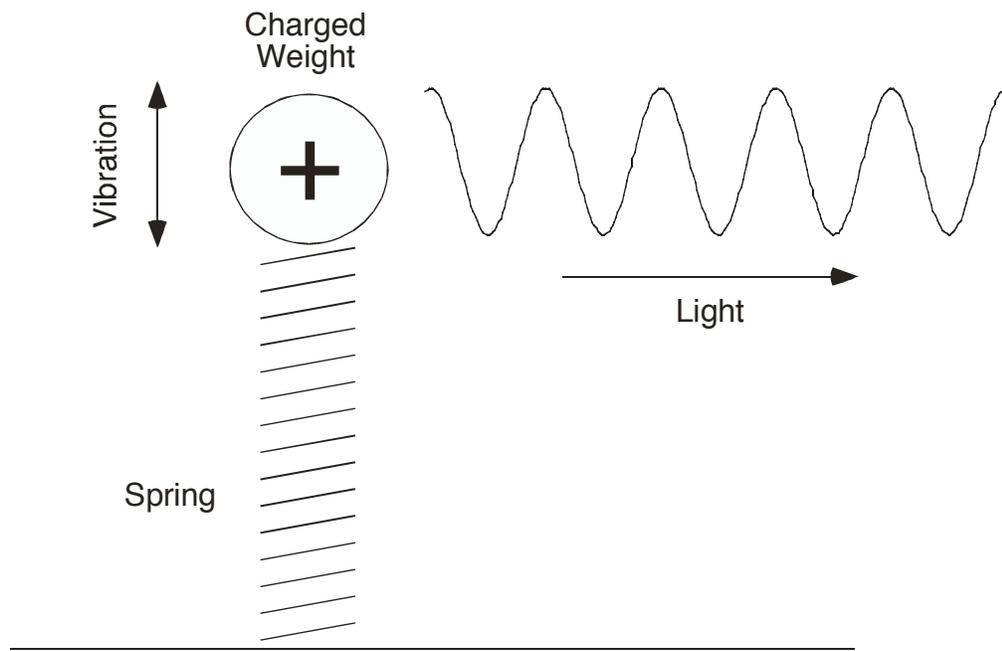


Figure 1-3

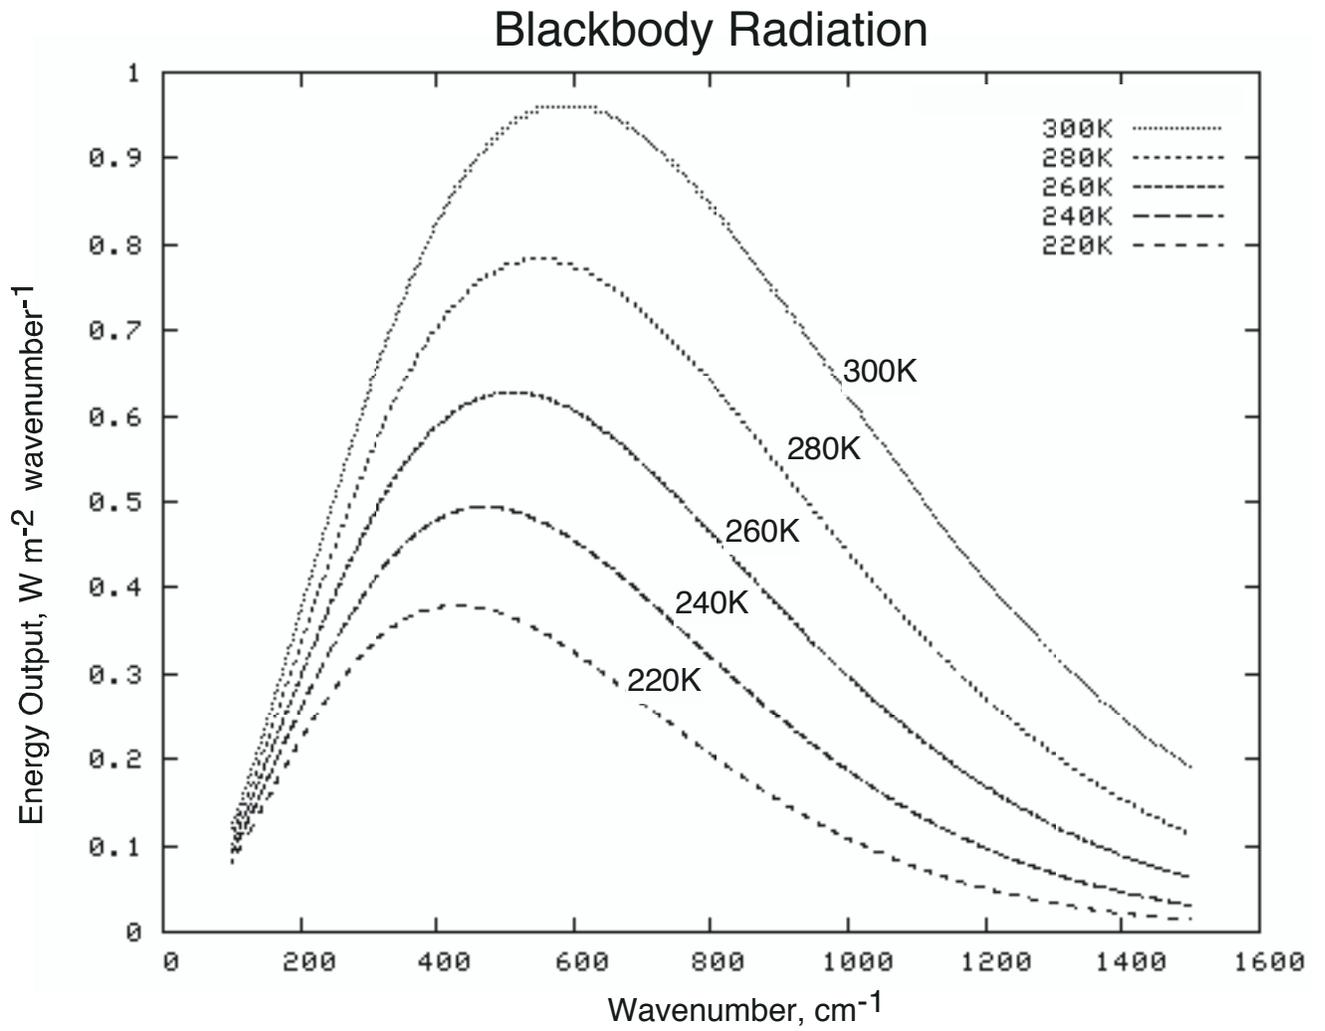


Figure 1-4



Figure 1-5

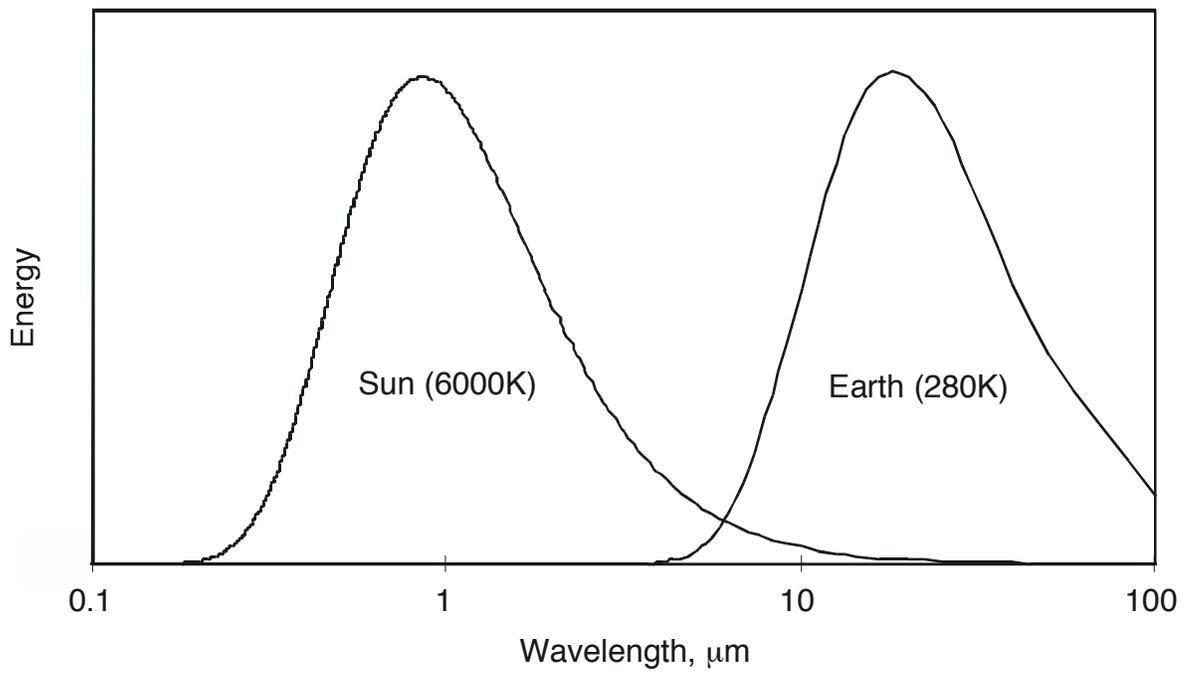


Figure 1-9

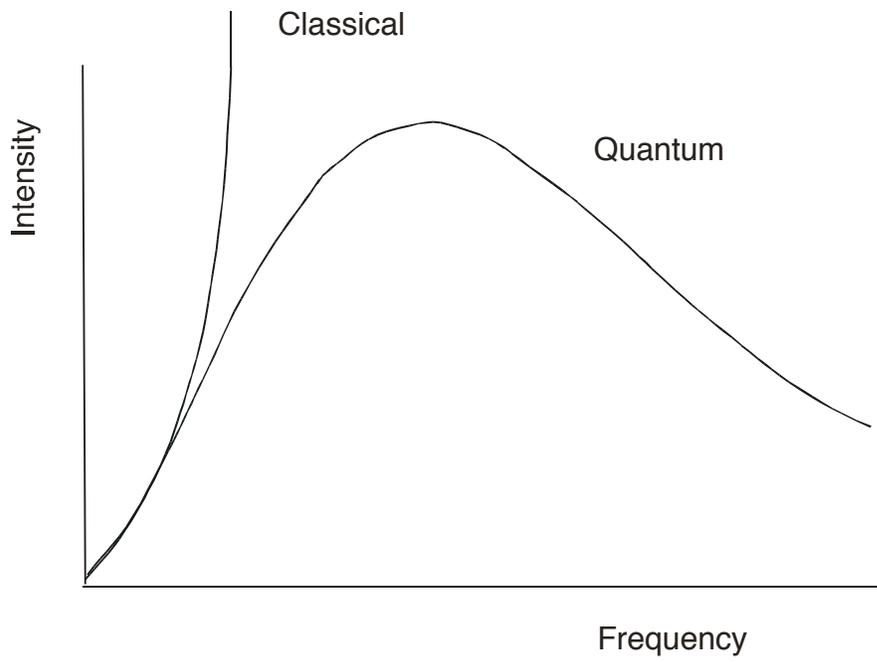


Figure 1-7

redrawn from <http://hyperphysics.phy-astr.gsu.edu/hbase/mod6.html>

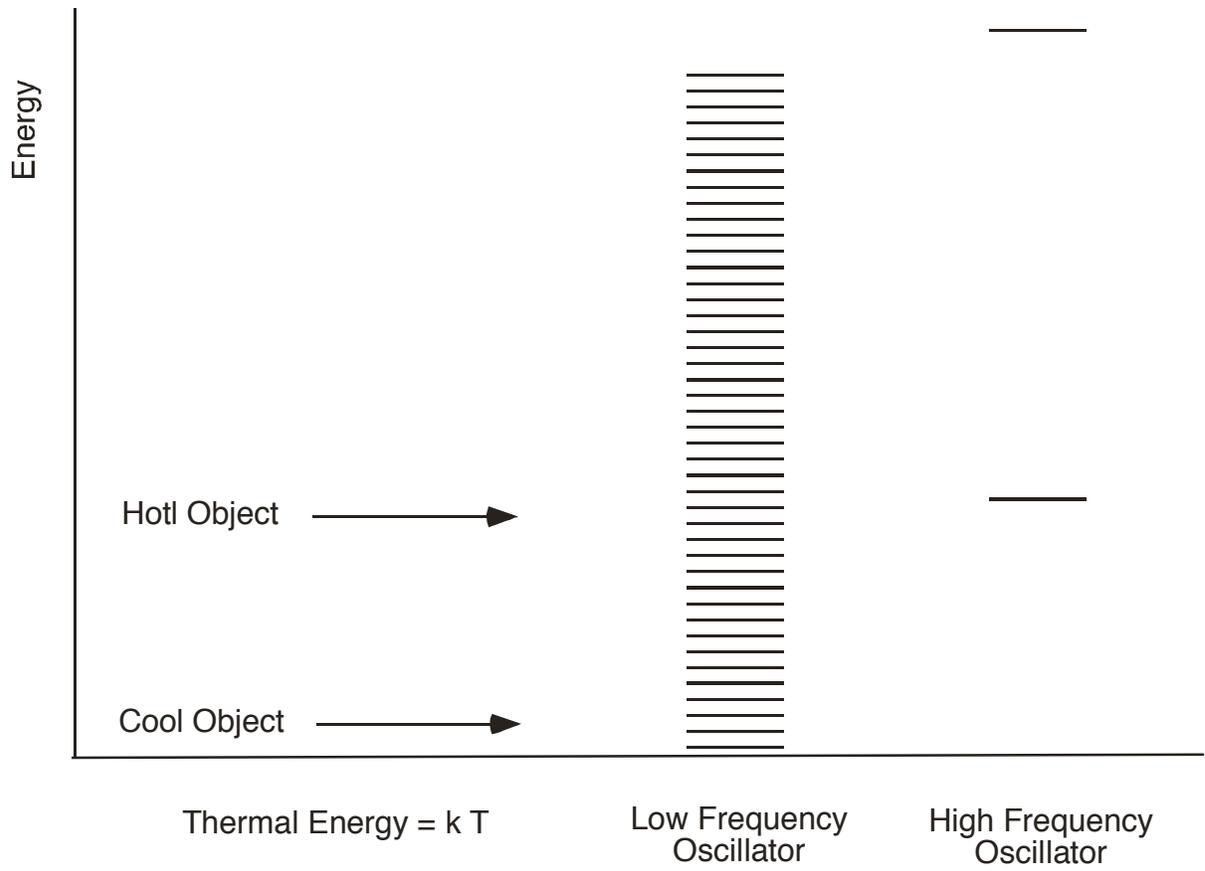


Figure 1-8

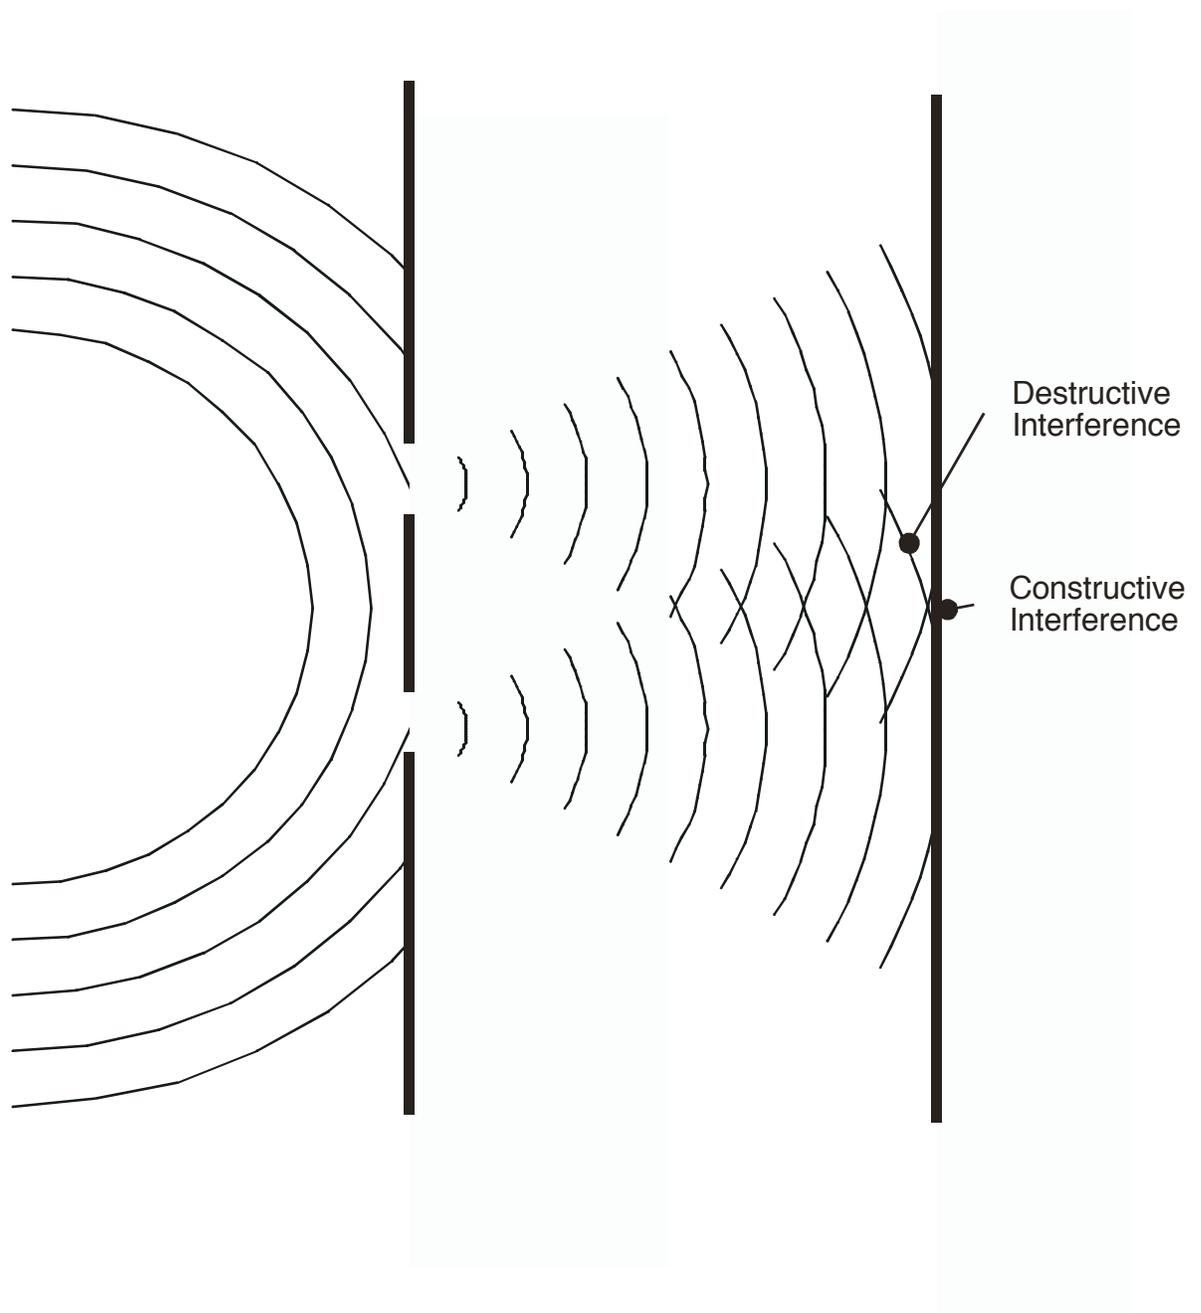


Figure 1-9

Chapter 2. The Layer Model

Summary

This is an algebraic calculation of the effect of an infrared absorber, a pane of glass essentially, on the mean temperature of the surface of the earth. By solving the energy budgets of the earth's surface and the pane of glass, the reader can see how the pane of glass traps outgoing IR light, leading to a warming of the surface. The layer model is not an accurate, detailed model suitable for a global warming forecast, but the principal of the greenhouse effect cannot be understood without understanding this model.

The Bare-Rock Layer Model

The temperature of the surface of the earth is controlled by the ways that energy comes in from the sun and shines back out to space as infrared. The sun shines a lot of light, because the temperature at the visible surface of the sun is high and therefore the energy flux $I = \epsilon \sigma T^4$ is a large number. Sunlight strikes the earth and deposits some of its energy into the form of vibrations and other bouncings-around of the molecules of the earth. Neither the earth nor the sun are perfect blackbodies, but they are both almost blackbodies, as are most solids and liquids. (Gases are terrible blackbodies, as we will learn in the next chapter). The earth radiates heat to space in the form of infrared light. Earth light is much lower frequency and lower energy than sun light.

We are going to construct a simple model of the temperature of the earth. The word **model** is used quite a bit in scientific discussion, to mean a fairly wide variety of things. Sometimes the word is synonymous with "theory" or "idea", such as the Standard Model of Particle Physics. For doctors, a "model system" might be a mouse that has some disease that resembles a disease that human patients get. They can experiment on the mouse rather than experimenting on people. In climate science, models are used in two different ways. One way is to make forecasts. For this purpose, a model should be as realistic as possible, and should capture or include all of the processes that might be relevant in nature. This is typically a mathematical model implemented on a computer, although there's a nifty physical model of San Francisco Bay you should check out if you're ever in the neighborhood. Once such a model has been constructed, a climate scientist can perform what-if experiments on it that could never be done to the real world, to determine how sensitive the climate would be to changes in the brightness of the sun or properties of the atmosphere, for example.

The simple model that we are going to construct here is not intended for making predictions, but is rather intended to be a toy system that we can learn from. The model will demonstrate how the greenhouse effect works by stripping away lots of other aspects of the real world that would certainly be important for predicting climate change in the next century or the weather next week, but make the climate system more complicated

and therefore more difficult to understand. The model we are going to explore is called the **layer model**. Understanding the layer model will not equip us to make detailed forecasts of future climate, but one cannot understand the workings of the real climate system without first understanding the layer model.

The layer model makes a few assumptions. One is that the amount of energy coming into the planet from sunlight is equal to the amount of energy leaving the earth as infrared. The real world may be out of energy balance for a little while or over some region, but the layer model is always exactly in balance. We want to balance the energy budget by requiring that the **outgoing energy flux F_{out}** must equal the **incoming energy flux**,

$$F_{in} = F_{out}$$

Let's begin with incoming sunlight. The **intensity of incoming sunlight I_m** at the average distance from the sun to the earth is about $1350 \text{ W} / \text{m}^2$. We'll consider the Watts part of this quantity first, then next the m^2 part. If you've ever seen Venus glowing in the twilight sky you know that some of the incoming visible light shines back out again as visible light. Venus' brightness is not blackbody radiation; it is hot on Venus but not hot enough to shine white-hot. This is **reflected** light. When light is reflected, its energy is not converted to vibrational energy of molecules in the earth and then re-radiated according to the blackbody emission spectrum of the planet. It just bounces back out to space. For the purposes of the layer model, it is as if the energy had never arrived on earth at all. The fraction of a planet's incoming visible light that is reflected back to space is called the planet's **albedo** and is given the symbol α (greek letter alpha). Snow, ice, and clouds are very reflective, and tend to increase a planet's albedo. The albedo of bright Venus is high, 70%, because of a thick layer of sulfuric-acid clouds in the Venusian atmosphere, and low, 0.15, for Mars because of a lack of clouds on that planet. Earth's albedo of about 0.3 depends on cloudiness and sea ice cover, which might change with changing climate.

Incoming solar energy that is not reflected is assumed to be absorbed into vibrational energy of molecules of the earth. Using a present-day Earthly albedo of 0.3, we can calculate that the intensity of sunlight that is absorbed by the earth is $1350 \text{ W} / \text{m}^2 (1 - \alpha) = 1000 \text{ W} / \text{m}^2$.

What about the area, the m^2 on the bottom of that fraction? If we want to get the total incoming flux for the whole planet, in units of W instead of W/m^2 , we need to multiply by a factor of area,

$$F_{in} [W] = I \left[\frac{W}{m^2} \right] \cdot A [m^2]$$

What area shall we use? Sun shines on half of the surface of the earth at any one time, but the light is weak and wan on some parts of the earth, near dawn or dusk or in high latitudes, but much more intense near the equator at noon. The difference in

intensity is caused by the angle of the incoming sunlight, not because the sunlight, measured head-on at the top of the atmosphere, is much different between low and high latitudes (Figure 2-1). How then to add up all the weak fluxes and the strong fluxes on the earth to find the total amount of energy that the earth is absorbing?

There's a nifty trick. Measure the size of the shadow. The area we are looking for is that of a circle, not a sphere. The area is

$$A[m^2] = \pi r_{earth}^2$$

Putting these together, the total incoming flux of energy to a planet by solar radiation is

$$F_{in} = \pi r_{earth}^2 (1 - \alpha) I_{in}$$

Our first construction of the layer model will have no atmosphere, only a bare rock sphere in space. A real bare rock in space, such as the moon or Mercury, is incredibly hot on the bright side and cold in the dark. The differences are much more extreme than they are on Earth or Venus where heat is carried by fluid atmospheres. Nevertheless, we are trying to find a single value for the temperature of the earth, to go along with a single value for each of the heat fluxes F_{in} and F_{out} . The real world is not all the same temperature, but we're going to ignore that in the layer model. The heat fluxes F_{in} and F_{out} may not balance each other in the real world, either, but they do in the layer model.

The rate at which the earth radiates energy to space is given by the Stefan-Boltzmann equation,

$$F_{out} = A \epsilon \sigma T_{earth}^4$$

As we did for solar energy, we are here converting intensity I to total energy flux F by multiplying by an area A . What area is appropriate this time? Incoming sunlight is different from outgoing earthlight in that sunlight is all traveling in one direction, whereas earthlight leaves earth in all directions (Figure 2-2). Therefore the area over which the earth radiates energy to space is simply the area of the sphere, which is given by

$$A_{sphere} = 4 \pi r_{earth}^2$$

Therefore the total outgoing energy flux from a planet by blackbody radiation is

$$F_{out} = 4 \pi r_{earth}^2 \epsilon \sigma T_{earth}^4$$

The layer model assumes that the energy fluxes in and out balance each other

$$F_{out} = F_{in}$$

which means that we can construct an equation from the “pieces” of F_{out} and F_{in} which looks like this:

$$4 \pi r_{\text{earth}}^2 \varepsilon \sigma T_{\text{earth}}^4 = \pi r_{\text{earth}}^2 (1 - \alpha) I_{\text{in}}$$

This equation tells us how the temperature of the earth ought to respond to factors such as the intensity of sunlight and the albedo of the earth.

The first thing to notice is that a few terms appear in common on both sides of the equation, which means that we can cancel them by dividing both sides of the equation by those factors. These are π and r_{earth}^2 . Before we toss these factors aside, however, we need to stop for another general science education tangent break and talk about π .

Box: π and transcendental numbers

The number π is one of a very special type of numbers called transcendental numbers. The simplest numbers are counting numbers; 1,2,3, etc. If we add digits after the decimal place, such as 2.7 and 3.24, we get rational numbers. A rational number can run on forever, such as the decimal expression for the fraction 1/9, which is 0.11111... with an infinite number of ones following the decimal point. They go on forever, but there is some repetition in the sequence of digits, likes “ones forever” in this example. Then there are the irrational numbers, such as the square root of two ($\sqrt{2}$), which can be approximated as 1.41421356.... These numbers go on forever like the rational numbers may do, but they never repeat. Pythagoras thought that numbers like this were an abomination, and their discovery was kept a closely guarded secret. There is no precise way to write the exact decimal value of $\sqrt{2}$, but we can get a handle on the number using geometry. In this case (and how Pythagoras came across $\sqrt{2}$), we can use the Pythagorean theorem to construct a line segment which is exactly $\sqrt{2}$ in length (Figure 2-3).

Finally, the most elusive and exalted of them all are the transcendental numbers like π . The decimal expression of π , 3.1415927... goes on forever as many rational and all irrational numbers do. The digits after the decimal never repeat, like the irrational numbers. Transcendental numbers differ from irrational numbers is that there is no way using a compass and straightedge to construct a straight line whose length is π . Of course, it's easy to construct a curved line of length π by drawing a half of a circle using a compass, but that doesn't count; we are looking for a straight line. This endeavor has been named “squaring the circle” and has been the focus of much devotion among amateur mathematicians and kooks over the centuries, along with perpetual motion machines and the fountain of youth. A resident of my home state of Indiana managed to convince the Indiana State Legislature in 1897 that he had squared the circle to discover that the true value for π . He persuaded them to recognize the truth of his proof in

state law, in exchange for allowing them to use his value of π without paying royalties or license fees. Hoosier generosity!

Before we move on, reflect for a moment that the number π is a transcendental number, one of the most elusive of all numbers. With modern computers we can approximate π as closely as we wish, but we can never write the value of π exactly. The exact area of the shadow of the earth, πr_{earth}^2 , will be forever out of our reach, as will the exact area of the surface of the sphere, $4 \pi r_{earth}^2$. Yet we can write the ratio of these two unreachable quantities, with the simple number 4. Perhaps this is an indication that I don't get out enough, but I find this beautifully fascinating.

OK, end of tangent. Time to do our first calculation of the temperature of the earth (Figure 2-4). Eliminating the factors of πr_{earth}^2 , and dividing by 4 on both sides so that we are left with units of Watts per area of the earth's surface, we get

$$\varepsilon \sigma T_{earth}^4 = \frac{(1 - \alpha) I_{in}}{4}$$

We know everything here except the T_{earth} . If we rearrange the equation to put what we know on the right hand side, what we don't on the left, we get

$$T_{earth} = \sqrt[4]{\frac{(1 - \alpha) I_{in}}{4 \varepsilon \sigma}} \quad (4)$$

What we have constructed is a relationship between a number of crucial climate quantities. Changes in solar intensity such as the sunspot cycle or the Maunder Minimum (Chapter 10) may affect I_{in} . We shall see in the next chapter that greenhouse gases are extremely selective about the wavelengths of light that they absorb and emit; in other words they have complexities relating to their emissivity ε values. The albedo of the planet is very sensitive to ice and cloud cover, both of which might change with changing climate. All kinds of interesting possibilities.

If we calculate the temperature of the earth, we get a value of 255 K or about -15°C . This is too cold; the temperature range of earth's climate is -50° to about $+35^\circ\text{C}$, but the average temperature, what we're calculating using the layer model, is closer to $+15^\circ\text{C}$ than -15°C . Table 2-1 gives the values we need to do the same calculation for Venus and Mars, along with the results of the calculation and the observed average temperatures. In all three cases, our calculation has erred on the side of too cold.

Table 2-1. The temperatures and albedos of the terrestrial planets.

	I_{solar} , W/m ²	α	T_{bare} K	T_{observed} K	$T_{\text{1 layer}}$ K
Venus	2600	71%	240	700	285
Earth	1350	33%	253	295	303
Mars	600	17%	216	240	259

The Layer Model With Greenhouse Effect

Our simple model is too cold because it lacks the greenhouse effect. We had no atmosphere on our planet; what we calculated was the temperature of a bare rock in space, like the moon. Of course the surface of the moon has a very different temperature on the sunlit side than it does in the shade, but if the incoming sunlight were spread out uniformly over the moon's surface, or if somehow the heat from one side of the planet conducted quickly to the other side, or if the planet rotated real fast, our calculation would be pretty good. But to get the temperature of the earth, Venus, and Mars, we need a greenhouse effect.

In keeping with the philosophy of the layer model, the atmosphere in the layer model is simple to the point of absurdity. The atmosphere in the layer model resembles a pane of glass suspended by magic above the ground [Figure 2-5](#). Like glass, our atmosphere is transparent to visible light, so the incoming energy from the sun passes right through the atmosphere and is deposited on the planet surface, as before. The planet radiates energy as IR light according to $\epsilon \sigma T_{\text{ground}}^4$, as before. In the IR range of light, we will assume that the atmosphere, like a pane of glass, acts as a blackbody, able to absorb and emit all frequencies of IR light. Therefore the energy flowing upward from the ground, in units of W/m² of the earth's surface, which we will call $I_{\text{up, ground}}$, is entirely absorbed by the atmospheric pane of glass. The atmosphere in turn radiates energy according to $\epsilon \sigma T_{\text{atmosphere}}^4$. Because the pane of glass has two sides, a top side and a bottom side, it radiates energy both upward and downward, $I_{\text{up, atmosphere}}$ and $I_{\text{down, atmosphere}}$.

The layer model assumes that the energy budget is in steady state; energy in = energy out. This is true for any piece of the model, such as the atmosphere, just as it is for the planet as a whole. Therefore we can write an **energy budget for the atmosphere**, in units of Watts per area of the earth's surface, as

$$I_{\text{up, atmosphere}} + I_{\text{down, atmosphere}} = I_{\text{up, ground}}$$

or

$$2 \epsilon \sigma T_{atmosphere}^4 = \epsilon \sigma T_{ground}^4$$

The budget for the ground is different from before, because we now have heat flowing down from the atmosphere. The basic balance is

$$I_{out} = I_{in}$$

We can break these down into component fluxes

$$I_{up, ground} = I_{in, solar} + I_{down, atmosphere}$$

and then further dissect them into

$$\epsilon \sigma T_{ground}^4 = \frac{(1-\alpha)}{4} I_{solar} + \epsilon \sigma T_{atmosphere}^4$$

Finally, we can also write a **budget for the earth overall** by drawing a boundary above the atmosphere and figuring that if energy gets across this line in, it must also be flowing across the line out at the same rate.

$$I_{up, atmosphere} = I_{in, solar}$$

The intensities are comprised of individual fluxes from the sun and from the atmosphere

$$\epsilon \sigma T_{atmosphere}^4 = \frac{(1-\alpha)}{4} I_{solar}$$

There is a solution to the layer model, for which all the budgets balance. We are looking for a pair of temperatures T_{ground} and $T_{atmosphere}$. Solving for T_{ground} and $T_{atmosphere}$ is a somewhat more complex problem algebraically than the bare-rock model with no atmosphere we solved above, but we can do it. We have two unknowns, but we appear to have three equations, the budgets for the atmosphere, for the ground, and for the earth overall. According to the rules of linear algebra, we can solve a problem if we have the same number of independent constraints (equations) as we have unknowns. If we have more constraints than unknowns, there might not be a solution. Can't get there from here. In our case, any two of the equations will suffice to pin down the two unknowns, and any third turns out to be just a combination of the other two. The budget equation for the earth overall, for example, is just the sum of the budget equations for the ground and the atmosphere (verify this for yourself). Therefore the third equation contains no new information that wasn't already contained in the first two, and we have a perfectly well-posed algebraic system, where the number of unknowns is just balanced by the number of independent constraints or equations.

So we are free to use any two of the three budget equations to solve for the unknowns T_{ground} and $T_{atmosphere}$. There are laborious ways to approach this problem, and

there is an easy way. Shall we choose the easy way? OK. The easy way is to begin with the energy budget for the earth overall. This equation contains only one unknown, $T_{\text{atmosphere}}$. Come to think of it, this equation looks a lot like equation 4, describing the surface temperature of the bare planet model above. If we solve for $T_{\text{atmosphere}}$ here, we get the same answer as when we solved for $T_{\text{bare earth}}$. This is an important point, more than just a curiosity or an algebraic convenience. It tells us that the place in the earth system where the temperature is the most directly controlled by the rate of incoming solar energy is the temperature at the location that radiates to space. We will call this temperature the **skin temperature** of the earth.

What about temperatures below the skin, in this case T_{ground} ? Now that we know that the outermost temperature, $T_{\text{atmosphere}}$, is equal to the skin temperature, we can plug that into the budget equation for the atmosphere to see that

$$2 \epsilon \sigma T_{\text{atmosphere}}^4 = \epsilon \sigma T_{\text{ground}}^4$$

or

$$T_{\text{ground}} = \sqrt[4]{2} T_{\text{atmosphere}}$$

The temperature of the ground must be warmer than the skin temperature, by a factor of the fourth root of two, an irrational but not transcendental number that equals about 1.189. The ground is warmer than the atmosphere by about 19%. When we do the calculation T_{ground} for Venus, Earth, and Mars in [Table 1](#), we see that we are getting Earth about right, Mars too warm, and Venus not yet warm enough.

The blackbody atmospheric layer is not a source of energy, like some humungous heat lamp in the sky. How then does it change the temperature of the ground? I am going to share with you what is perhaps my favorite earth-sciences analogy, that of the equilibrium water level in a steadily filled and continuously draining sink. Water flowing into the sink, residing in the sink for a while, and draining away is analogous to energy flowing into and out of the planet. Water drains faster as the level in the sink rises, as the pressure from the column of water pushes water down the drain. This is analogous to energy flowing away faster as the temperature of the planet increases, according to $\epsilon \sigma T^4$. Eventually the water in the sink reaches a level where the outflow of water balances the inflow. That's the equilibrium value, and is analogous to the equilibrium temperature we calculated for the layer model. We constrict the drain somewhat but putting a penny down on the filter. For a while, the water drains out more slowly, and the water level in the sink rises because the water budget imbalance. The water level rises until the higher water level pushes water down the drain fast enough to balance the faucet again. A greenhouse gas, like the penny in the drain filter, makes it more difficult for the heat to escape the earth. The temperature of the earth rises until the fluxes balance again.

Take-Home Points

The outflow of IR energy from a planet must balance heating from the sun.

The planet accomplishes this act of energetic housekeeping by adjusting its temperature.

Absorption of outgoing IR light by the atmosphere warms the surface of the planet, as the planet strives to balance its energy budget.

Lab

1. The moon with no heat transport. The layer model assumes that the temperature of the body in space is all the same. This isn't really very accurate, as you know that it's colder at the poles than it is at the equator. For a bare rock with no atmosphere or ocean, like the moon, the situation is even worse, because fluids like air and water are how heat is carried around on the planet. So let's make the other extreme assumption, that there is no heat transport on a bare rock like the moon. What is the equilibrium temperature of the surface of the moon, on the equator, at local noon, when the sun is directly overhead? What is the equilibrium temperature on the dark side of the moon?

2. A two-layer model. Insert another atmospheric layer into the model, just like the first one. The layer is transparent to visible light but a blackbody for infrared.

a) Write the energy budgets for both atmospheric layers, for the ground, and for the earth as a whole, just like we did for the one-layer model.

b) Manipulate the budget for the earth as a whole to obtain the temperature T_2 of the top atmospheric layer, labeled Atmospheric Layer 2 in [Figure 2-6](#). Does this part of the exercise seem familiar in any way? Does the term skin temperature ring any bells?

c) Insert the value you found for T_2 into the energy budget for layer 2, and solve for the temperature of layer 1 in terms of layer 2. How much bigger is T_1 than T_2 ?

d) Now insert the value you found for T_1 into the budget for atmospheric layer 1, to obtain the temperature of the ground, T_{ground} . Is the greenhouse effect stronger or weaker because of the second layer?

3. Nuclear Winter. Let's go back to the 1-layer model, but let's change it so that the atmospheric layer absorbs visible light rather than allowing to pass through (Figure 2-7). This could happen if the upper atmosphere were filled with dust. For simplicity, let's assume that the albedo of the earth remains the same, even though in the real world it might change with a dusty atmosphere. What is the temperature of the ground in this case?

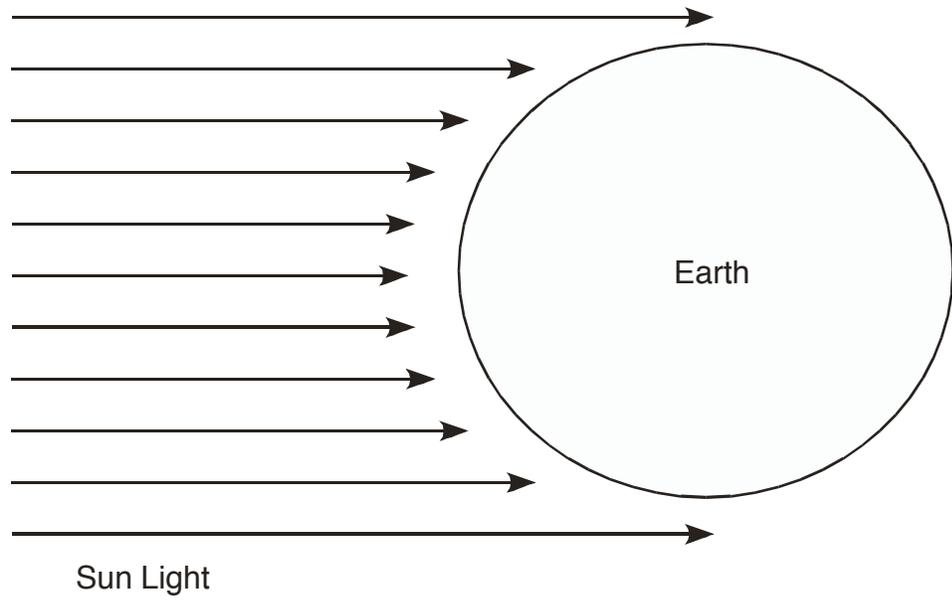


Figure 2-1

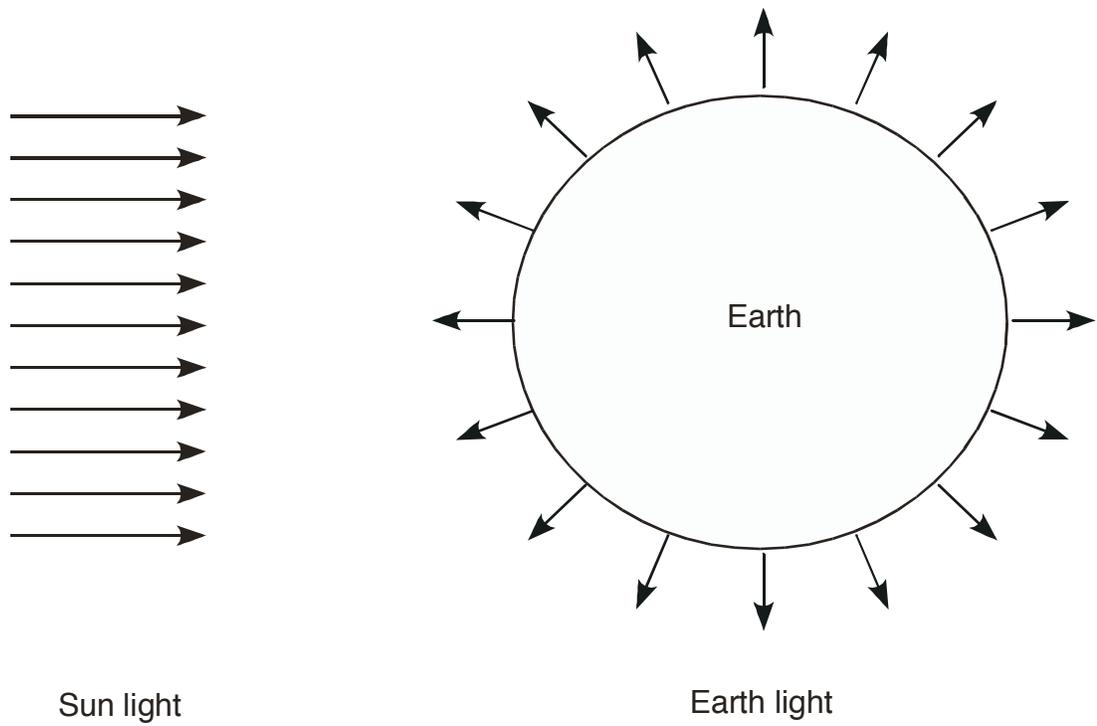


Figure 2-2

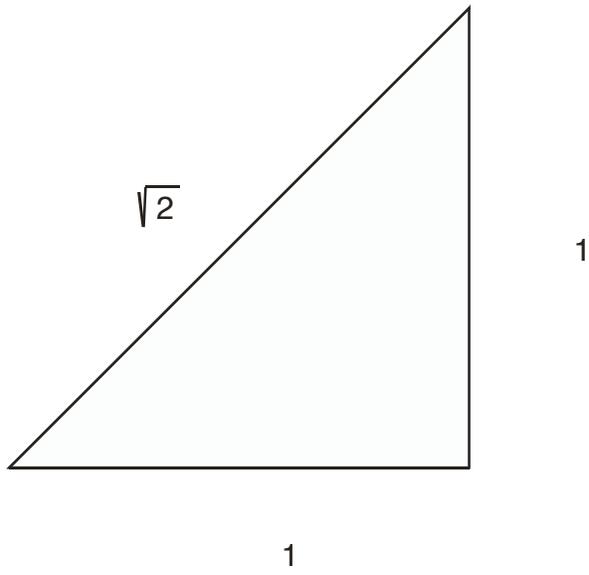


Figure 2-3

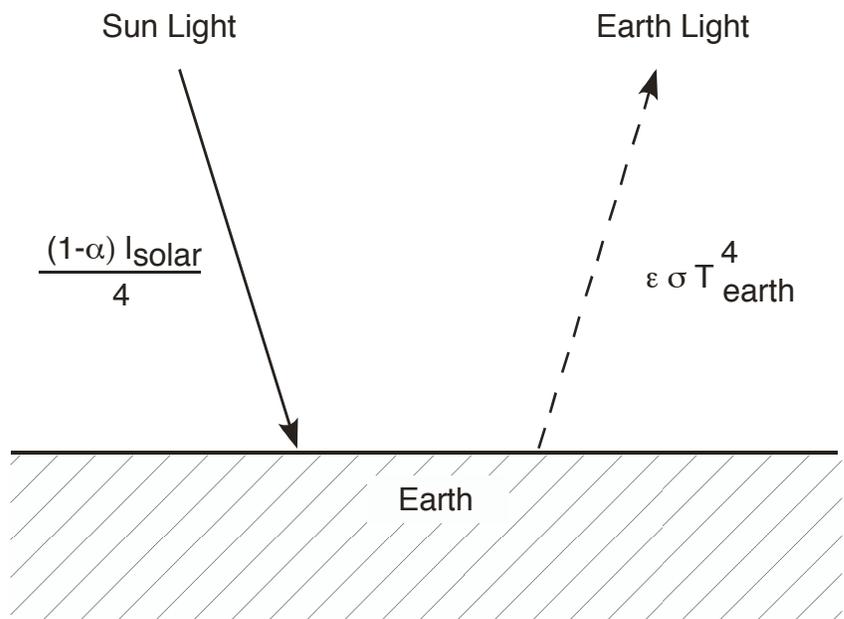


Figure 2-4

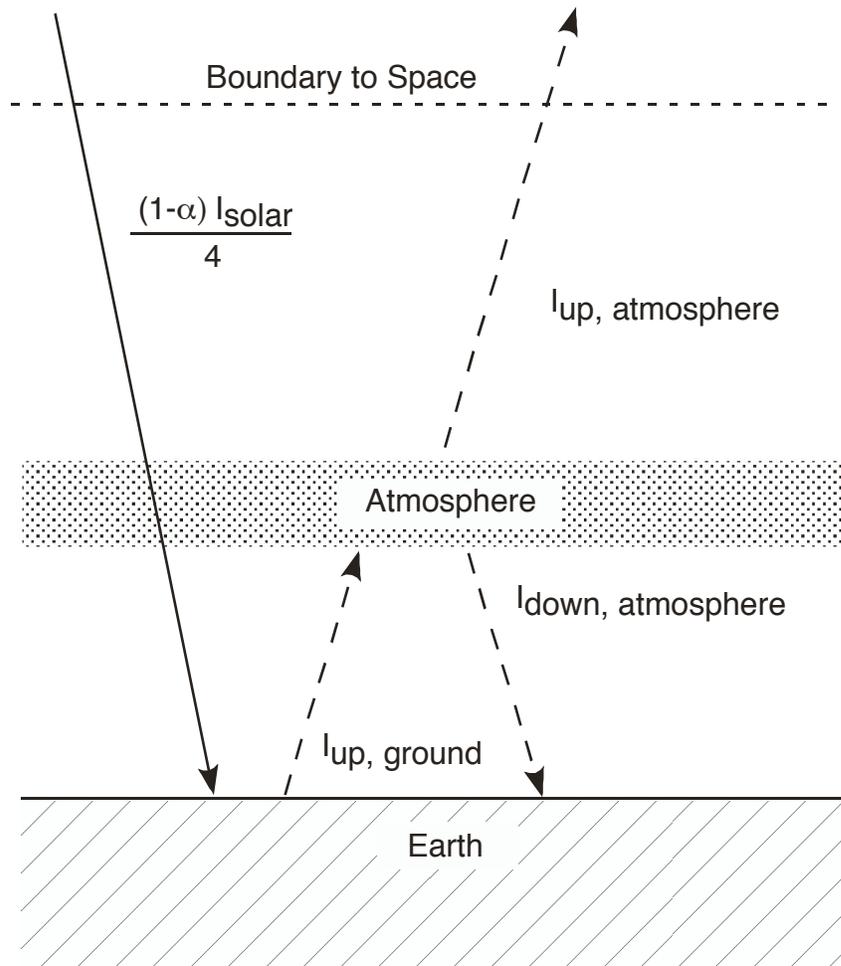


Figure 2-5

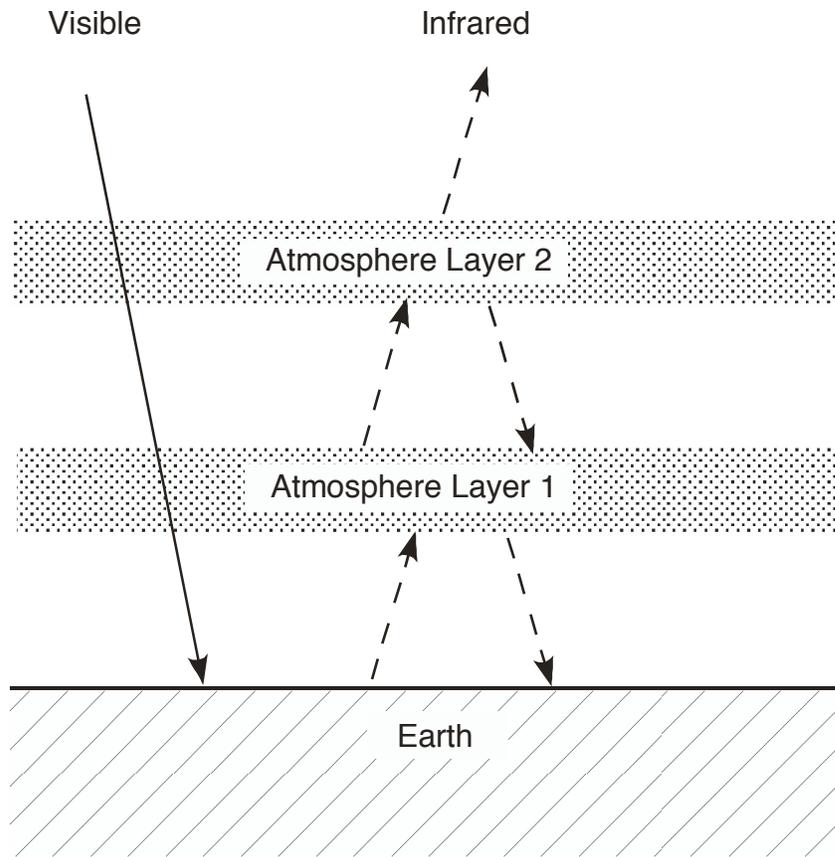


Figure 2-6

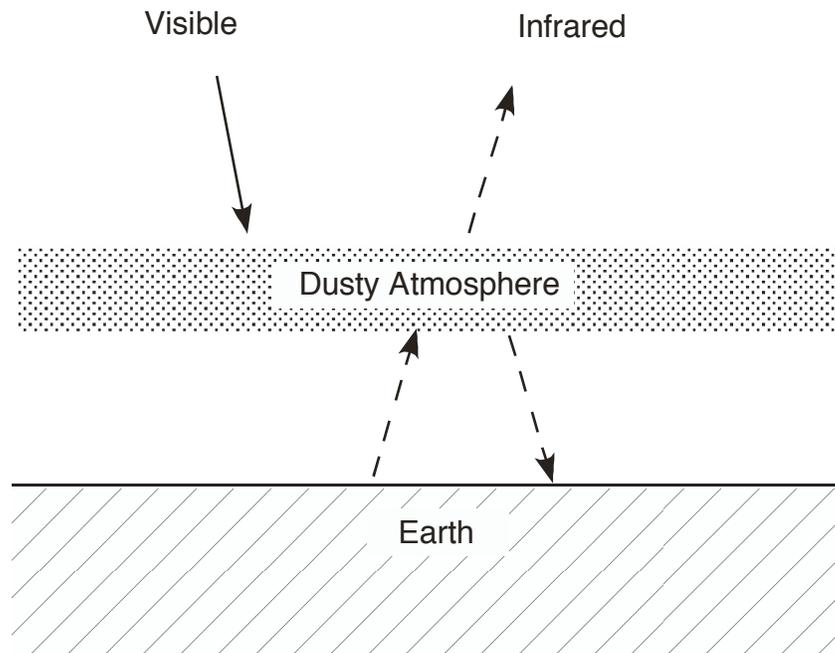


Figure 2-7

Chapter 3. Greenhouse Gases

Summary

The layer model assumes that the atmosphere acts as a blackbody in the infrared, absorbing and emitting all frequencies of IR light. In reality, gases absorb IR light selectively, and most of the gas in the atmosphere doesn't interact with IR light at all. The difference can be understood in terms of the effect of molecular vibration on the electromagnetic field. Because gases absorb IR selectively, there are some radiation bands that are completely absorbed (the gases are saturated), and others such as the atmospheric window, where no gases absorb. This leads to much higher greenhouse forcing per molecule from some trace gases, such as freons, SF₆, or to a lesser extent methane, than to major gases such as CO₂. Some absorption bands fall in the middle of the IR emission spectrum of the earth's surface, while other bands fall outside this spectrum and are therefore irrelevant to the heat budget.

About Gases

The layer model is what we call an idealization of the real world. Now that we understand the core mechanism of the greenhouse effect, by understanding the layer model, we can add things one at a time from the real world, and see how they affect the way that earth's temperature is controlled. The first modification we have to make to the layer model is to think more about real gases in the atmosphere.

Let's begin by defining different ways of describing the amounts of gases in the atmosphere. The word **concentration** means number of molecules within some volume. The difficulty this raises for gases in the atmosphere is that there are fewer molecules per volume overall as you go up in the atmosphere. The major gases in the atmosphere are pretty well mixed, so that the concentrations of these gases go down proportionally with altitude. This is why it's hard to breath oxygen quickly enough on Mount Everest.

It is often more convenient to talk about proportions of gases, like oxygen is about 20% of the molecules of gas in the atmosphere, and nitrogen almost 80%. The proportion of CO₂ is 0.037%, or we express that in a different way by saying 370 parts per million or **ppm**. This number is called a **mixing ratio**.

The weight of gases above you in the atmosphere determine the **pressure** that the gases exert on objects like the walls of our thermos bottle with a vacuum inside. If the walls of the thermos were too thin the atmosphere could just crush it. We will talk about pressure in units of atmospheres or **atm**, where one atm is the average pressure at sea level. Here's what you might not have expected: the pressure exerted by each molecule of a gas is the same, even if they have different weights. It is as if you were shaking a big box full of bowling balls and ping pong balls and the ping pong balls were pushing out on

the cardboard walls of the box as hard, per ball, as the bowling balls. The way this works is that all different types of gas molecules have the same average energy, if they are at the same temperature. That amount of energy is the same was stored in vibrational modes, $1/2 k T$ for each of three dimensions, or $3/2 k T$. A lighter molecule like a ping-pong ball can go faster with this amount of energy than a heavier one like a bowling ball. So when the lighter molecule hits the wall, it has less mass but a higher velocity, and these two factors just cancel to make the lighter and the heavier molecule molecules each exert the same pressure on the wall. What this means is that we can easily divide the total pressure into **partial pressures**. If 20% of the atoms of the 1 atm pressure of gas is oxygen, then the partial pressure of oxygen is 0.2 atm. 370 atoms out a million are CO_2 , and so the partial pressure of CO_2 is 370 μatm , where the greek letter mu at the beginning means 10^{-6} . This unit is pronounced micro-atmospheres. The partial pressure of CO_2 is abbreviated **pCO₂**. The pCO₂ of the atmosphere in μatm is the same number as the mixing ratio in ppm, 370, rising by about 1.5 ppm or μatm per year.

Gases, Vibrations, and Light

Most of the mass of an atom is in its nucleus, which resembles the massive sun at the center of the solar system. Electrons float in ghostly quantum mechanical probability clouds, called **orbitals**, around the nucleus. Two nuclei of two different atoms always repel each other, because of their positive charges. The orbitals for the electrons fit together better, however, with certain numbers of orbitals than with others. Electrons from two different atoms may be able to combine their orbitals in such a way that they are lower energy, as if happier, when they share, a **chemical bond**. A chemical bond is like a spring in that the two nuclei on either end of the bond have some freedom to move closer or farther apart. There is an optimum distance for the nuclei to be from each other. Closer, and the positive nuclei will start to repel each other. Farther, and you get less energy gain from sharing the electrons. A bond vibrates when the distance between the nuclei oscillates between the nuclei being too close together, then too far apart.

Gases are the simplest type of molecule, and they only vibrate in very particular ways. Vibrations in a gas molecule are like vibrations of a piano string in that they are fussy about frequency. This is because, like a piano string, a gas molecule will only vibrate at its “ringing” frequency. The ringing frequency of an oscillator made of weights and springs depends on two things: the amount of weight on the ends and the strength of the spring holding them together. Heavy weights will have enough inertia to keep a bond growing in the wrong direction for longer than will a pair of light weights, so the frequency of the vibration will be slower. If the spring is very strong, it will reverse the velocity of a vibration more quickly, and the frequency of the oscillation will be higher. Vibrations in chemical bonds depend on the mass of the nuclei and on the energy penalty for having the nuclei too close or too far apart: the springiness of the chemical bond.

However, the vibrations of many gas molecules, such as the major gases in the atmosphere oxygen and nitrogen, are invisible to the electromagnetic field. They don't shine light or absorb infrared light; we say they are not **infrared active**. Oxygen and

nitrogen are not greenhouse gases, because they are transparent to infrared light. These molecules are invisible because when you stretch one, it doesn't change the electric field. These are symmetric molecules, made of two identical atoms. Neither atom can hold the electrons any more tightly than the other. In general, symmetrical molecules with only two atoms are not greenhouse gases.

We can break the symmetry, making a molecule of NO for example. This is a very reactive molecule, an ingredient for producing urban smog, but that's another story. NO has one atom of each element, and as a result has a slight imbalance in its distribution of electrons. One side of the molecule will have a slight positive charge, and the other will be slightly negative. We could oscillate the electric field simply by rotating an NO molecule. Also, if we vibrate an NO molecule, the steepness of the transition from slightly positive to slightly negative will oscillate with time. By these mechanisms, NO could be a greenhouse gas, but it turns out not to be a very important one because there is not very much of it.

Molecules with more than two atoms have more than one chemical bond. All of their bonds ring together rather than each bond ringing with its own characteristic frequency. Water, H₂O, is a molecule that is bent in its lowest energy state (Figure 3-1). This is because several of the electron orbitals stick off in the direction that appears in my diagram to be empty space. Hydrogen atoms hold their electrons more loosely than oxygen atoms, and so each hydrogen has a slightly positive charge (marked in Figure 3-1 using the lowercase greek letter delta, as $\delta+$). The oxygen end of the molecule has a slight negative charge. Just as for the NO molecule, rotating an H₂O molecule would oscillate the electric field and generate light. Because the arrangement of the nuclei in H₂O are more complex than for NO, there are several modes of vibration of the water molecule, including a symmetric stretch and a bend. These modes are also infrared active.

The CO₂ molecule is shaped in a straight line with carbon in the middle (Figure 3-2). It is a symmetric molecule; the oxygen atom on one end pulls the electrons just as tightly as the other oxygen on the other end. Therefore rotating the molecule at rest has no effect on the electric field. Nor does a symmetric stretch. However, there are two modes of vibration which do generate an asymmetry in the electric field. One is an asymmetric stretch, and the other is a bend. The bend is the most climatically important one, as we shall see next.

How a Greenhouse Gas Interacts with Earth-Light

We have seen that gases are terrible blackbodies, because they are very choosy about which frequencies they absorb and emit. What we will now see is that some frequency bands are more important to the climate of the earth than others. There are two factors to consider. One is the concentration of the gas, which we will discuss in a bit. The other is the frequency of the absorption band relative to the blackbody spectrum for the earth.

Figure 3-3 shows blackbody spectra again for temperatures ranging from 300 K, a hot summer day, down to 220 K, which is about the coldest it gets in the atmosphere, up near the troposphere at about 10 km altitude. There is also a jagged-looking curve. This is the

intensity of light that an infrared spectrometer would see if it were in orbit over the earth, looking down (Figure 3-4). Figure 3-3 is not data, but rather a model simulation from one of our on-line models. You can point a web browser at <http://geosci.uchicago.edu/~archer/cgimodels/radiation.html> to run this model yourself. We will do so in the exercises.

The spectrum of the light leaving the earth going into space ranges between two different blackbody spectra, a warmer one of about 270 K, and a colder one from about 220 K.

The parts of the spectra that seem to follow the colder blackbody curve come from greenhouse gases in the upper atmosphere. They follow the colder blackbody curve because it is cold in the upper atmosphere. The most pronounced of these absorption bands, centered on a wave number of about 700 cycles/cm, comes from the bending vibration of CO₂. Light of this intensity that shines from the surface of the earth is absorbed by the CO₂ in the atmosphere (Figure 3-4). The CO₂ in the atmosphere then radiates its own light at this frequency. Remember from Chapter 1 that light emission and absorption is a two-way street.

Other parts of the spectrum, most notably the broad smooth part around 1000 cycles/cm, follow a warmer blackbody spectrum. These come directly from the ground. The atmosphere is transparent to infrared light in these frequencies. This band is called the **atmospheric window**.

The situation is analogous to standing on a pier and looking down into a pond of water. If the water were very clear, you could see light coming from the bottom; you would see rocks or old tires or whatever in the reflected light. If the water were murky, the light you would see would be scattered light coming from perhaps just a few inches down into the water. The old tires would be invisible, alas.

Remember we said that the total energy flux from one of these spectra can be “eyeballed” as the total area under the curve. The areas of the pure blackbody curves are going up proportionally to the temperature raised to the fourth power, because of the Stefan-Boltzmann equation (our equation 1-1 in Chapter 1). The area trick works with our new jagged spectrum as well. The effect of an atmospheric absorption band is to take a bite out of the blackbody spectrum from the earth’s surface, decreasing the area and therefore decreasing the outgoing energy flux.

Compare the CO₂ absorption band at 700 cycles/cm with the absorption band of methane at around 1300 cycles/cm. The CO₂ band has a lot more room to change the outgoing infrared energy flux than does the methane band, simply because the earth and the atmosphere radiate a lot more energy near 700 cycle/cm than near 1300 cycles/cm. Both blackbody spectra are pretty low intensity in the methane band.

Band Saturation

The core of the CO₂ bend absorption band, between 600 and 800 cycles/cm, looks smooth rather than jagged and it follows a blackbody spectrum from about 220 K. This is about as cold as the atmosphere gets, and if we change the amount of CO₂ in the atmosphere, the intensity of light in this range does not get any lower (Figure 3-5). We call this phenomenon **band saturation**. You can see it in a series of model runs in which the CO₂ concentration of the atmosphere goes up from zero to 1000 ppm. The current concentration of CO₂ in the atmosphere is about 370 ppm, as we will learn more in Section II. If there were no CO₂ in the atmosphere, the atmosphere would be transparent to light of around 700 cycles/cm, as it is in the atmospheric window. Adding the first 10 ppm of CO₂ has a fairly noticeable impact on the shape of the outgoing light spectrum, but increasing CO₂ from say 100 to 1000 has a somewhat subtler effect.

I have plotted the total energy intensity I_{out} in W/m² as a function of the concentration of CO₂ in the atmosphere in Figure 3-6. Changes in CO₂ concentration have the greatest effect if we were starting out from no CO₂ and adding just a bit. The first 10 ppm of added CO₂ changes I_{out} by as much as going from 10 to 100, or 100 to 1000 ppm. We can understand why by analogy to our murky pond or by looking back at Figure 3-4. As we increase the murkiness of the water, we decrease the distance that a photon of light can travel before it is absorbed. It doesn't take much murk in the water to obscure the old tire on the bottom, shifting the depth to which we can see from the bottom at say 3 meters to maybe only one meter. If we make the pond a lot murkier we will only be able to see a few cm down into the water. Making it murkier still will limit our view to only on cm. The change in depth is getting less sensitive to the murkiness of the pond. In the same way, the changes in the temperature at which the atmosphere radiates to space get smaller as the CO₂ concentration of the air gets higher. You just see the coldest light that you can get.

The band saturation for CO₂ makes CO₂ a less potent greenhouse gas than it would be if we had no CO₂ in the air to start with. Let's revisit our comparison of the CO₂ and methane as greenhouse gases. Methane had a disadvantage because its absorption band sort of fell in the suburbs of the earth-light spectrum whereas CO₂ fell right downtown. Now we see the advantage shifting the other way. Methane has a much lower concentration in the atmosphere. You can see from the jagged edges of the methane peak in Figure 3-3 that the methane absorption band is not saturated. For this reason, in spite of the suburban location of the methane band, a **molecule of methane** added to the atmosphere is **20 times more powerful than is a molecule of CO₂**.

If the edges of the absorption bands were completely abrupt, as if CO₂ absorbed 600 cycles/cm light completely and 599 cycles/cm light not at all, then once an absorption band from a gas was saturated, that would be it. Further increases in the concentration of the gas would have no impact on the radiation energy budget for the earth. CO₂, the most saturated of the greenhouse gases, would stop changing climate after it exceeded some concentration. It turns out that this is not how it works. Even though the core of the CO₂ band is saturated, the edges of the band are not saturated. When we increase the CO₂

concentration, the bite that CO₂ takes out of the spectrum doesn't get deeper, but it gets a bit broader.

The bottom line is that the energy intensity I_{out} in units of W/m² goes up proportionally to the log of the CO₂ concentration, rather than proportionally to the CO₂ concentration itself (we say linear in CO₂ concentration). The logarithmic dependence means that you get the same I_{out} change in W/m² from any doubling of the CO₂ concentration. The radiative effect of going from 10 to 20 μatm pCO₂ is the same as going from 100 to 200 μatm, or 1000 to 2000 μatm.

The sensitivities of climate models are often compared as the average equilibrium temperature change from doubling CO₂, a diagnostic number that is called ΔT_{2x}. Most models have ΔT_{2x} between 2 and 5 K, which is the same as 2 to 5°C. You can use ΔT_{2x} to estimate a temperature change resulting from some change in CO₂. Note that this is the ultimate temperature change, after hundreds or even thousands of years have passed (see [Chapters 6 and 11](#)). The equation is

$$\Delta T = \Delta T_{2x} \times \frac{\ln\left(\frac{\text{new } p\text{CO}_2}{\text{orig. } p\text{CO}_2}\right)}{\ln(2)} \quad (3.1)$$

where ln is the natural log, the reverse operation of the exponential function e^x. The symbol e denotes a number which has no name other than simply e. We will meet e again in Chapter 4. The exponential function is to raise e to a power of x. If

$$e^x = y$$

then

$$y = \ln(x)$$

Equilibrium temperature changes from changes in CO₂, assuming various ΔT_{2x} values, are shown in [Figure 3-7](#).

Greenhouse Gases and the Planetary Energy Budget

What happens to the energy balance of the earth if we add a greenhouse gas to its atmosphere? If the energy budget was in equilibrium before, it isn't any more, because the greenhouse gas has decreased the amount of energy leaving the earth to space. We can see this visually as the big bite out of the spectrum going from the top to the middle diagram in [Figure 3-8](#). The decrease in energy flux is proportional to the area of that bite, the difference between the top and middle figures. Remember back to chapter 2, the premise of the layer model is that the energy coming into and going out of the planet must balance, and the planet accomplishes this feat by adjusting its temperature. If we want to re-balance the energy flux after kicking it by adding CO₂, we do that by increasing the temperature of the ground. Using the on-line model, we find that a temperature change of 8.5 K or 8.5° C brings us back to the same energy output I_{out} as we

had before. Looking at the bottom result in Figure 3-8, we see that the new, warmer output spectrum has risen everywhere compared to the middle figure. Visually, we have cut some area out of the CO₂ absorption band, and added it in the atmospheric window and other parts of the spectrum, until the overall area under the curve is the same as it was initially. Adding the CO₂ caused the planet to warm.

Take-Home Points

Gases absorb / emit infrared light if they vibrate at the frequency of the light, and if its vibration has a dipole moment that affects the electric field. O₂ and N₂ are not greenhouse gases. All molecules of three or more atoms are infrared active.

A greenhouse gas has a stronger impact on the radiative balance of the earth if it interacts with light in the middle of the earth-light spectrum.

Band saturation: A greenhouse gas at high concentration like CO₂ will be less effective, molecule per molecule, than a dilute gas like methane.

Lab

Answer these questions using the on-line model at <http://geosci.uchicago.edu/~archer/cgimodels/radiation.html> . The model takes CO₂ concentration and other environmental variables as input, and calculates the outgoing IR light spectrum to space, similarly to Figures 3-3, 3-5, and 3-7. The total energy flux from all IR light is listed as part of the model output, and was used to construct Figure 3-6.

1. Methane. Methane has a current concentration of 1.7 ppm in the atmosphere, and it's doubling at a faster rate than is CO₂.

a) Is ten additional ppm of methane in the atmosphere more or less important than ten additional ppm of CO₂ in the atmosphere at current concentrations?

b) Where in the spectrum does methane absorb? What concentration would it take to begin to saturate the absorption in this band? (How do you identify saturation of a band, on a spectrum plot?)

c) Would a doubling of methane have as great an impact on the heat balance as a doubling of CO₂?

d) What is the "equivalent CO₂" of doubling atmospheric methane? That is to say, how many ppm of CO₂ would lead to the same change in outgoing IR radiation energy flux as doubling methane? What is the ratio of ppm CO₂ change to ppm methane change?

2. CO₂.

a) Is the direct effect of increasing CO₂ on the energy output at the top of the atmosphere larger in high latitudes or in the tropics?

b) Set $p\text{CO}_2$ to an absurdly high value of 10,000 ppm. You will see a spike in the CO_2 absorption band. What temperature is this light coming from? Where in the atmosphere do you think this comes from?

3. Earth Temperature. The greenhouse effect theory presumes that an increase in the temperature at ground level will lead to an increase in the outgoing IR energy flux at the top of the atmosphere.

a) How much extra outgoing IR would you get by raising the temperature of the ground by one degree? What effect does the ground temperature have on the shape of the outgoing IR spectrum and why?

Move or kill b) More water can evaporate into warm air than cool air. By setting the model to hold the water vapor at constant relative humidity rather than constant vapor pressure (the default) calculate again the change in outgoing IR energy flux that accompanies a 1 degree temperature increase. Is it higher or lower? Does this make the earth more sensitive to CO_2 increases or less sensitive?

c) Now see this effect in another way. Starting from a base case, record the total outgoing IR flux. Now increase $p\text{CO}_2$ by some significant amount, say 30 ppm. The IR flux goes down. Now, using the constant vapor pressure of water option, increase the Temperature Offset until you get the original IR flux back again. What is the change in T required? Now repeat the calculation but at constant relative humidity. Does the increase in CO_2 drive a bigger or smaller temperature change? This is the water vapor feedback.

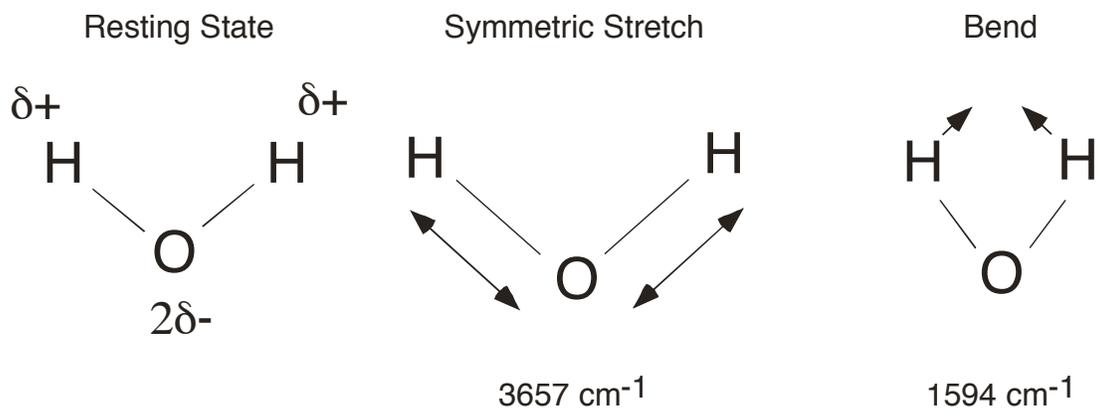


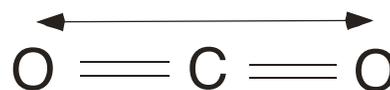
Figure 3-1

Resting State



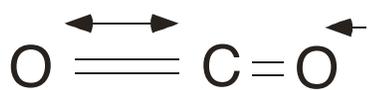
No Resting Dipole

Symmetric Stretch



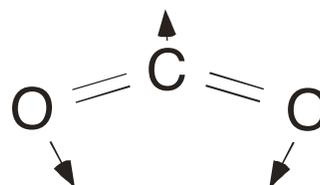
IR Inactive

Asymmetric Stretch



3760 cm^{-1}

Bend



1595 cm^{-1}

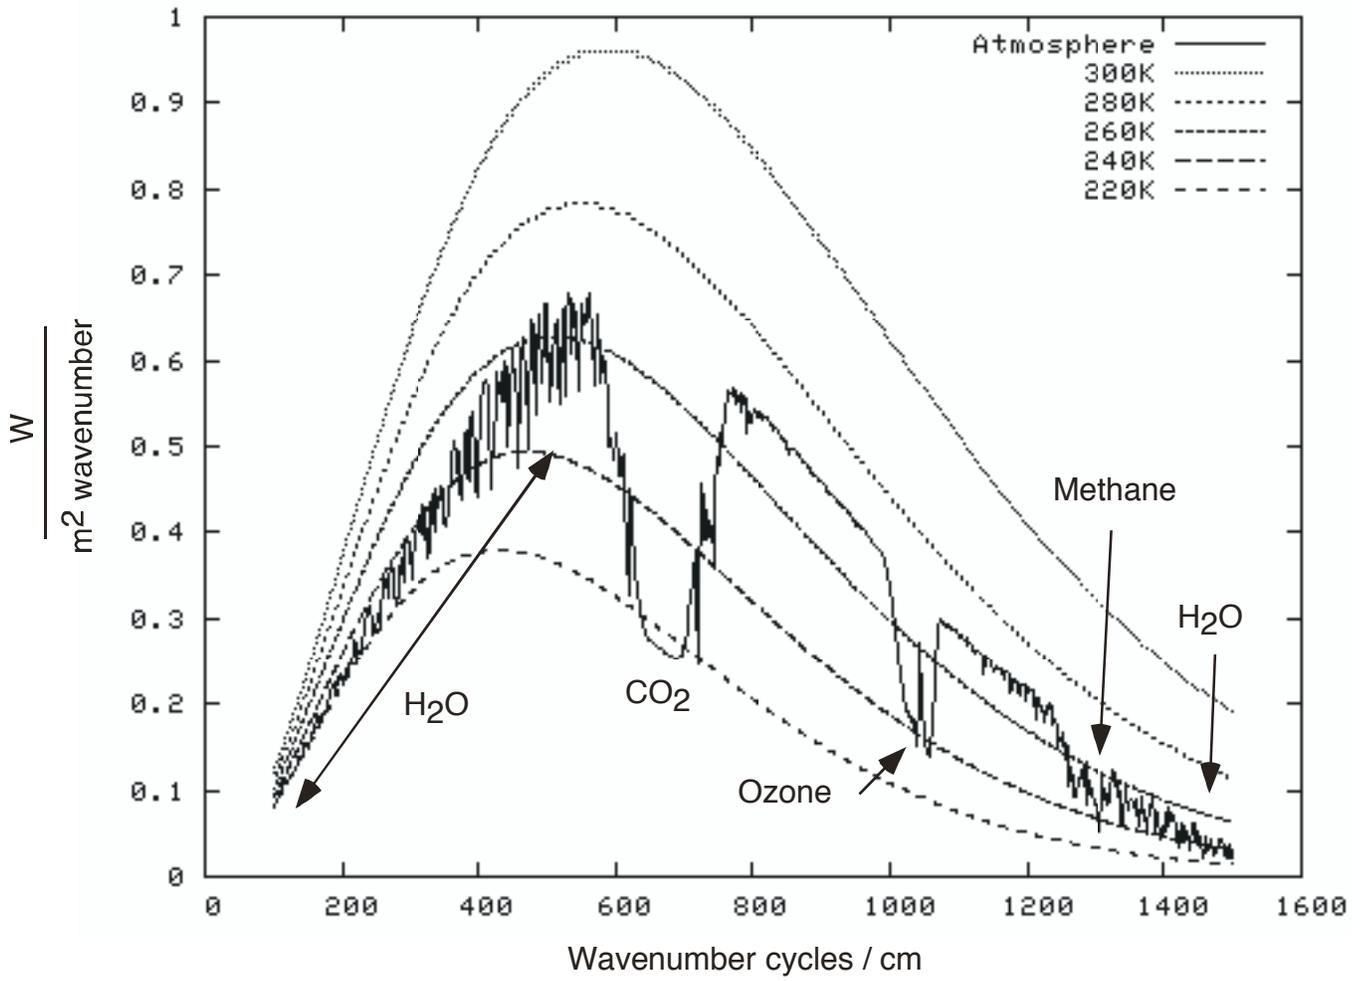


Figure 3-3

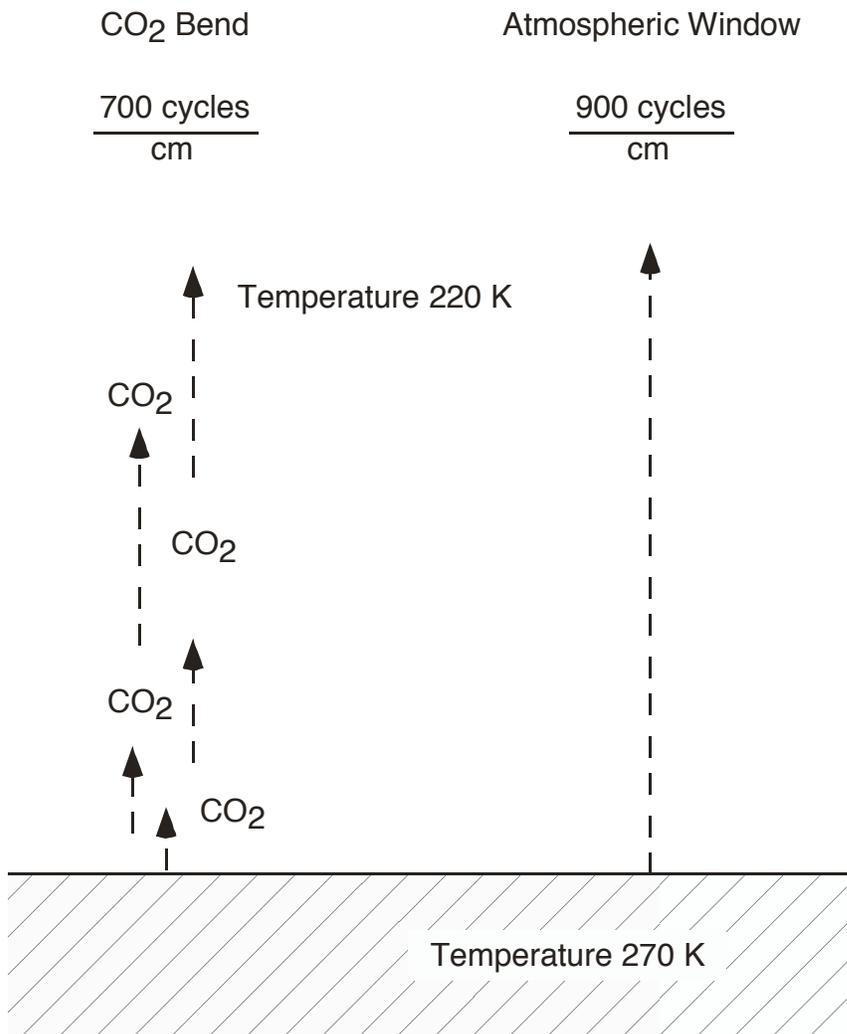
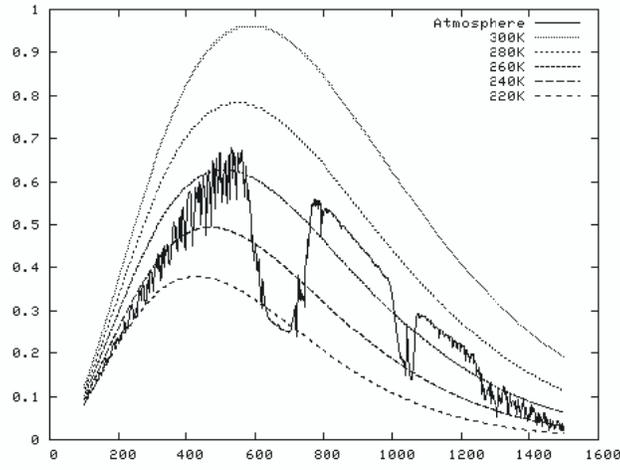
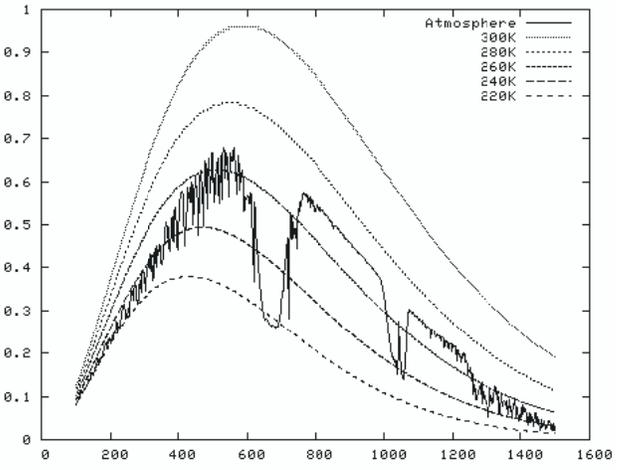
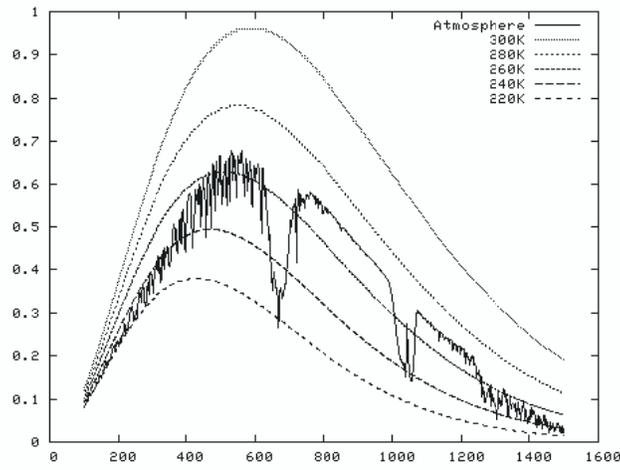
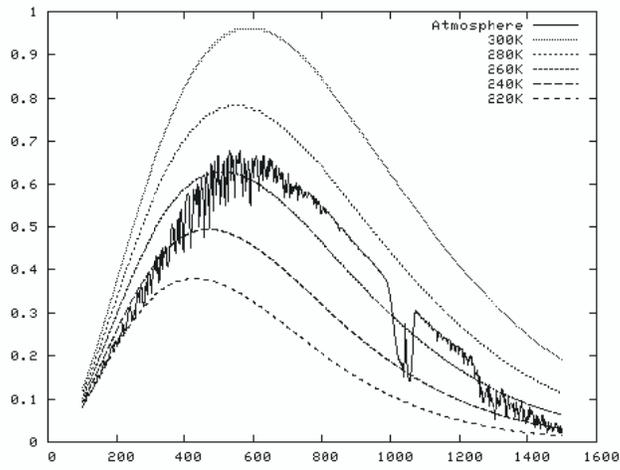


Figure 3-4

W
m² wavenumber



Wavenumber cycles / cm

Figure 3-5

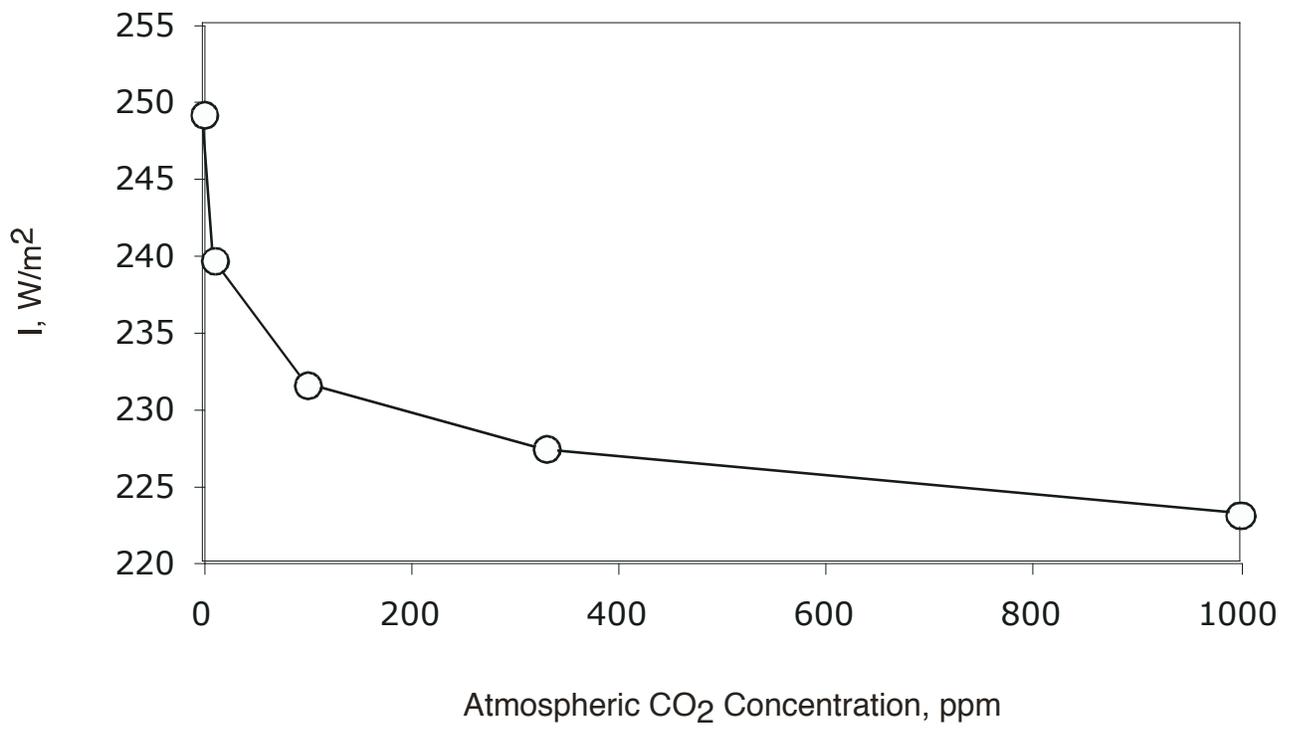


Figure 3-6

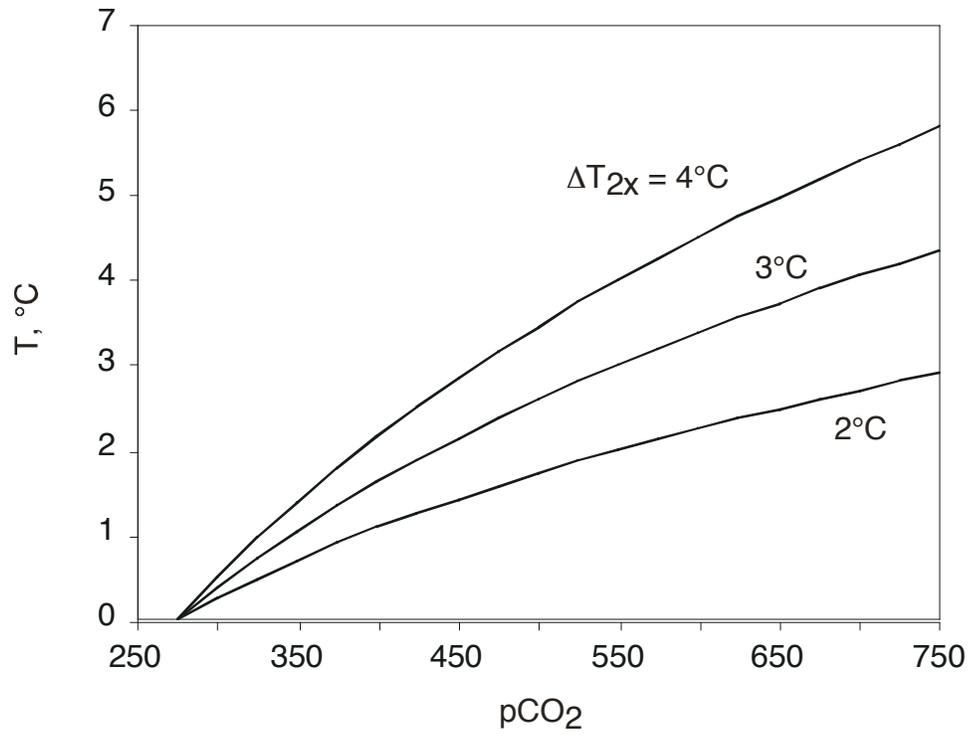


Figure 3-7

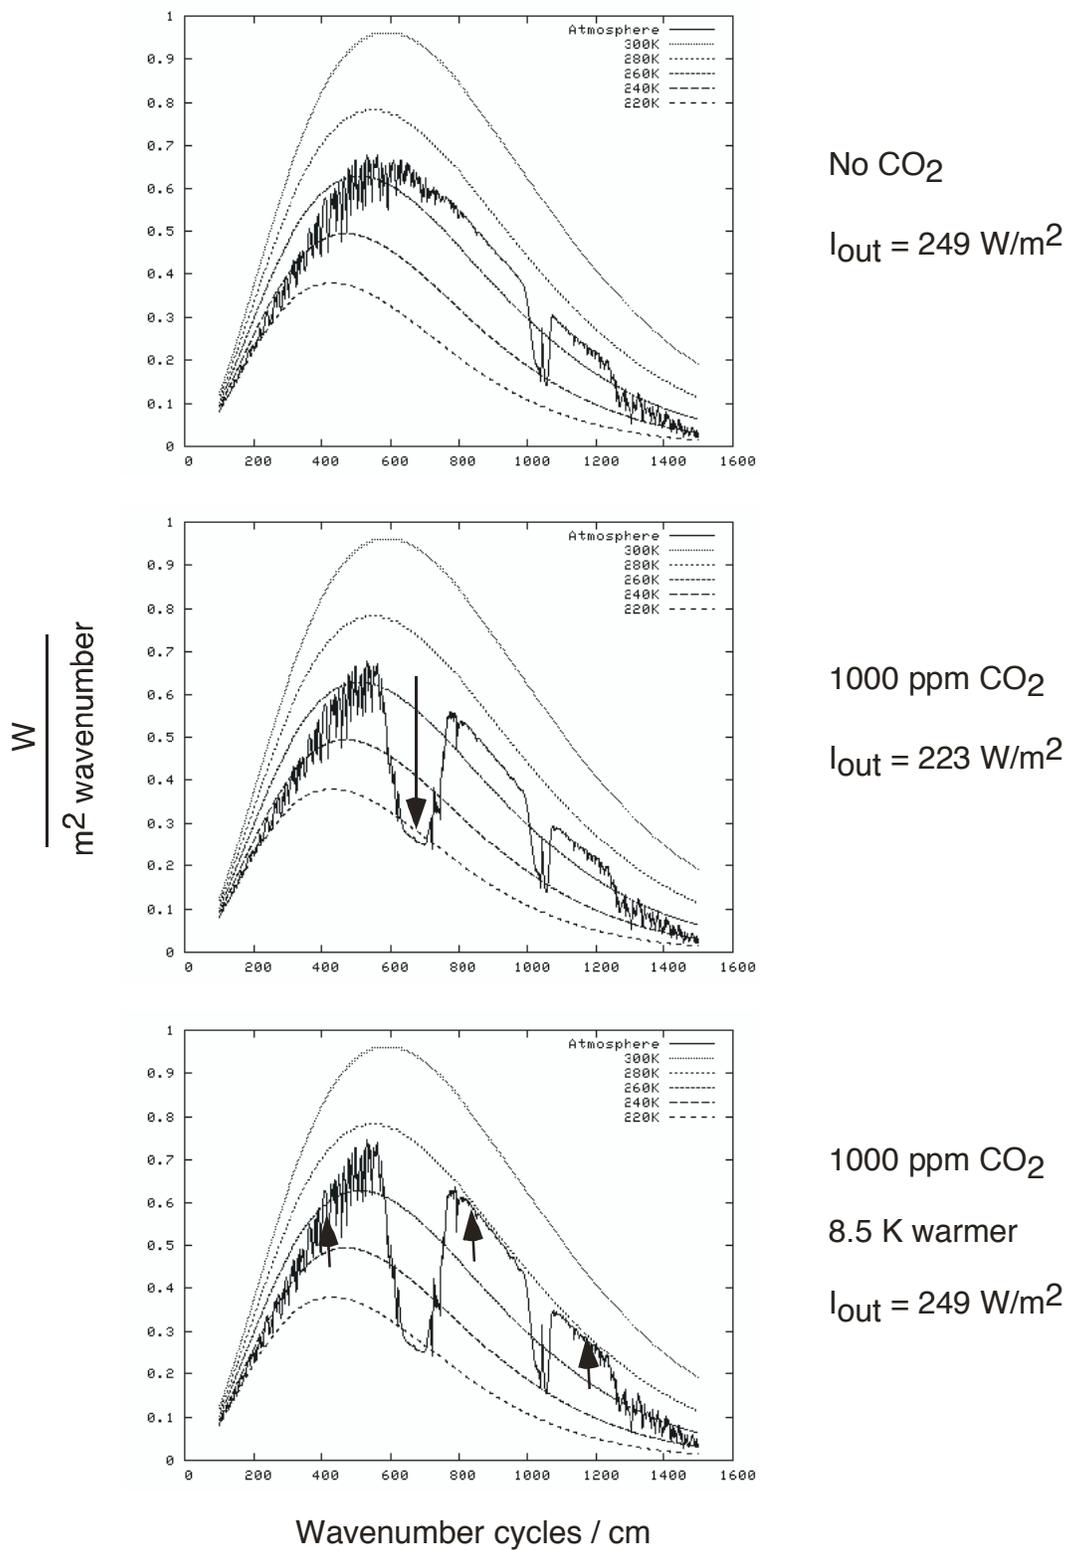


Figure 3-8

Chapter 4. Convection

Summary

The Layer Model assumes that heat flows between the earth's surface and the atmosphere by radiation. In reality, the temperature of the atmosphere is coupled to the temperature of the ground by convection. Warm air rises and carries heat. As it expands, it cools, leading to a decrease in temperature with altitude. If the atmosphere were incompressible, like water, there would be very little change in temperature with altitude, such as occurs in a pan of water heating up on a stove. If this were the case, then there would be no greenhouse effect, because the outgoing IR intensity would be the same regardless of whether the light on average comes from the ground or from high in the atmosphere.

Meet the Atmosphere

If you have ever climbed a mountain, you know that the temperature decreases as you go up (Figure 4-1). This part of the atmosphere is called the **troposphere**, and it contains 90% of the air in the atmosphere and all of its weather. If you were Superman and could climb high enough, the temperature would reach its coldest point at an altitude called the **tropopause**, and would begin to rise above that in the **stratosphere**. Above the stratosphere are wispy layers called the mesosphere and the exosphere which affect the way that ham radio signals propagate around the world, but don't affect our climate story very much.

Explaining the temperature structure of the atmosphere requires weaving together a number of what may seem to be disparate threads. Stick with me and everything will all tie together in the end.

Pressure as a Function of Altitude

Gases and liquids exert pressure on surfaces of solids that are immersed in them, simply the force of the atoms bouncing off of the solid surface. The pressure gets lower as you climb higher in the atmosphere. As we ascend, we decrease the amount of fluid that is above us, decreasing the pressure that we feel.

Scuba divers know that diving 10 m deep increases the pressure by about 1 atm. Each 10 m of depth is the same 1 atm pressure increase: descending from 30 to 40 m would increase the pressure by the same 1 atm as descending from 0 to 10 m. We say that pressure is **linear** with depth (Figure 4-2). The pressure can be calculated as

$$P = 1 \text{ atm} + \frac{-1 \text{ atm}}{10 \text{ m}} \cdot z[\text{m}]$$

where we are using the variable z to denote the vertical position, as before, with positive numbers upward, as before. So a depth is a negative height. At the water surface, $z=0$, and we have 1 atm pressure from the weight of the atmosphere. The increase in pressure with depth, from the weight of the water, is linear with depth below the surface, because the factor $-1 \text{ atm} / 10 \text{ m}$ of height is constant.

Box: the story of e

*The pressure in the atmosphere is **nonlinear** with altitude, in that a climb of one meter at sea level changes the pressure much more than one meter of climb up at the tropopause. The equation to describe pressure as a function of height in the atmosphere is based on the **exponential function**, which is a number called e raised to a power. The value of e is approximately 2.71828.... e is another transcendental number, like π . Let's take a break from pressure for a moment and talk about e and the exponential function. The exponential function was invented by bankers to calculate compound interest for bank accounts. Let's say your account pays you 10% interest per year, and you start out the year with \$100 in your account. One way to calculate how much money you have after a year is*

$$\text{Balance}(1) = \text{Balance}(0) + \text{Balance}(0) \cdot \frac{10\%}{\text{year}} \cdot 1 \text{ year}$$

for a total of \$110. The problem is that we get a different answer if we update the account twice a year instead of once. After six months we get

$$\text{Balance}(0.5) = \text{Balance}(0) + \text{Balance}(0) \cdot \frac{10\%}{\text{year}} \cdot 0.5 \text{ year}$$

which would be \$105, and the second six months we get

$$\text{Balance}(1.0) = \text{Balance}(0.5) + \text{Balance}(0.5) \cdot \frac{10\%}{\text{year}} \cdot 0.5 \text{ year}$$

*which yields \$110.25. Bankers call this updating calculation **compounding the interest** in the account. Compounding the interest twice results in a higher balance at the end of the year because the interest we make in the first half of the year itself grows during the second half of the year. If we compound the interest 10 times, every 36.5 days, we end up with \$110.46. Compounding 100 times gives \$110.51. However, there is a limit to how much money we can squeeze from the bank by this trick. If we compounded the interest an infinite number of times, the amount of money we would collect at the end of the year could be calculated using the exponential function as*

$$\text{Balance}(1.0) = \text{Balance}(0.0) \cdot e^{\left(\frac{0.10}{\text{yr}} \cdot 1 \text{ yr}\right)} \tag{4-1}$$

where the yielding \$110.52.

The exponential function comes up time and time again in the natural sciences. Population growth and radioactive decay are two examples. In each case, the rate of change of the variable depends linearly on the value of the variable itself. Population growth is driven by the number of babies born, which depends on the number of potential parents to beget them. The rate of radioactive decay that you would measure with a Geiger counter depends on the number of radioactive atoms present. The growth rate of your bank account depends on its size.

The multiplier in the exponent describes the relative rate of change of the variable, in this case, the interest rate. If the exponent is positive, as for population, the growth accelerates as it progresses (Figure 4-3). One could reasonably call this type of growth an explosion. For decay, the exponent is negative, and the quantity of radioactive atoms gets ever closer to zero with time, but mathematically never gets there.

The atmospheric pressure varies as a function of altitude according to an exponential decay type of equation in height.

$$P(z) = 1 \text{ atm} \cdot e^{\frac{-z[\text{km}]}{8 \text{ km}}}$$

The height z is zero at sea level, leaving us with e^0 which equals 1, so the pressure at sea level is 1 atm. At an altitude of 8 km, pressure is lower than at sea level by a factor of e^{-1} or $1/e$, about 37%. We call that altitude the **e-folding height**. Most of the mass of the atmosphere is contained in the e-folding height. In fact, if the entire atmosphere were at 1 atm pressure, instead of smeared out in a decaying exponential function, it would fit into one e-folding height exactly. If we were tracking the decay of a radioactive chemical with time, the scaling factor in the exponential would be called an **e-folding time**. This quantity is similar to a half-life for radioactive decay but rather than decaying to half the initial quantity, we're waiting until we're 37% lower than the original. It takes 44% longer to decay to $1/e$ than it does to $1/2$ of the original number of atoms.

From the appearance of the exponential function in the equation for pressure, you could probably guess that the rate of change of pressure with altitude must depend on the pressure itself in some way. This would be astute. The rate of change of pressure depends on pressure because at high pressure, gas is compressed, and so a climb of one meter through gas at high pressure would rise above more molecules of gas than would a climb through a meter of gas at low pressure. Imagine a wall made of compressible bricks (Figure 4-4). A row of bricks is thinner at the bottom of the wall, because they are compressed. Batman and Robin climbing up the wall would pass more bricks per step at the bottom than at the top. For incompressible bricks (the normal kind), the rows are all the same height and the mass of wall above you is a linear function of height.

Adiabatic Expansion

Here begins the second apparently unrelated thread of our story. If we compress a gas its temperature goes up. This occurs even if we don't allow any heat to enter the gas

or leave it, say if we had gas inside an insulated piston which we compress or expand. The condition that we are describing, a closed system with no heat coming in or out, is called **adiabatic**. If gas is compressed adiabatically, it warms up. If you ever let the air out of a bicycle tire by holding the little valve open with your thumb nail, you may have noticed how cold your thumb nail got. The gas expanded as it flowed from the high-pressure tire into the lower-pressure atmosphere.

Figure 4-5 shows the temperature change that a parcel of dry surface air would experience if it were carried aloft adiabatically (an **adiabatic trajectory** or **adiabat**). We'll worry about water vapor in a moment, but for now we are looking at the line labeled Dry Adabat (foreshadowing).

Why should the temperature go up in a gas as you compress it? It takes work to compress a gas. You have to squeeze our hypothetical insulated piston in order to compress it. Your muscles push against molecules bouncing against the walls of the cylinder. The work you put in is transformed into bouncing-around energy of the molecules of gas, its temperature. The situation is a little harder to envision for expansion, because we tend to ignore the atmosphere around us and think of an empty room as empty space, but when the piston expands, it has to push back the atmosphere, in other words it must do work. The energy to do that work comes from the thermal energy of the molecules of the gas, so the gas cools as it expands.

Water Vapor and Latent Heat

Here comes seemingly unrelated thread number three. Water molecules can exist together in any of three **phases**: gas, liquid, or solid. A transformation from liquid or solid to gas requires an input of a considerable amount of energy. One could write a chemical reaction for water as



If you have ever burned your skin with steam from a teapot, this reaction will have meaning for you. Steam from a teapot is probably at the boiling temperature of water, 373 K or 100°C. This is the temperature you set for your oven to warm bread; you can easily stick your hand in to get the bread without burning yourself. Steam burns, not because it is so hot, but because it deposits its heat when it condenses on your skin. The heat that the steam is carrying is called **latent heat**. You charge up an air parcel with latent heat when you evaporate water into it, and you get the heat back again when the water condenses. A thermometer doesn't measure the latent heat content of an air parcel unless it is released by condensation.

If you set a tray of water or ice in a closed box with some extra space around it, some of the water molecules are going to evaporate, becoming a gas. Here's an interesting tidbit; it makes almost no difference whether there are other gases in that space or just vacuum. The same number of water molecules will make the jump into the gas phase regardless; that is to say, the partial pressure of water vapor doesn't care how much oxygen or nitrogen is there. The partial pressure of water vapor depends primarily on the

temperature. If the temperature is high, many molecules will have the energy required to jump into the vapor phase, and the vapor pressure of water will be high. If it is cold, the vapor pressure will be lower (Figure 4-6).

At any given temperature, the water vapor pressure will drift toward some **equilibrium** pressure. If the amount of water vapor is lower than equilibrium, that is to say **undersaturated**, water will tend to evaporate. If water vapor is **supersaturated**, meaning the pressure is higher than equilibrium, then water will tend to condense, perhaps into raindrops or snowflakes. Equilibrium is the state of lowest energy; a dead battery is in equilibrium. However a chemical system can be out of equilibrium for a long time if the reaction rates are very slow. In the clean remote marine atmosphere, forming droplets can be very slow, and water vapor 25% supersaturated has been observed. The **relative humidity** is the water vapor pressure divided by the saturation value; from the last example 25% supersaturated would be 125% relative humidity.

Convection

Now we are ready to weave our seeming disparate threads of story together into a picture of what controls the temperature as a function of altitude in the atmosphere. The pieces are assembled into a process called **convection**. Convection takes its place among conduction and radiation which we've already discussed as means of carrying heat in the environment. Convection occurs when a fluid medium is heated from the bottom or cooled from the top, and the heavy water atop light water causes the fluid to circulate. Imagine a lava lamp, in which the fluid is heated from below by the light bulb (a ridiculous energy-wasting incandescent bulb; you could never make a lava lamp work with a nice efficient compact fluorescent bulb, because fluorescent bulbs do not generate so much waste heat). The fluid at the bottom becomes warmer than the fluid overlying it. Warmer molecules bounce around more energetically, pushing all of the molecules somewhat farther apart from each other. For this reason the fluid **expands** as its temperature increases. As it expands its density (mass per volume) decreases. If we stack some fluids of different density together in a column, the stable configuration is to have the densest ones on the bottom. Think of oil and water; the oil always floats on the water because it is less dense. So our expanding parcel of fluid at the bottom of the lamp begins to rise. It floats to the top of the lamp until it cools and sinks back to the bottom again.

Forget the lava lamp now and think of a pan of water on a stove, which is a simpler case because there is only one fluid instead of two. Figure 4-7 plots the temperature of the water as a function of the height in the pan, which we call a temperature **profile**. If we thoroughly mix the water, the temperature will be the same throughout the pan of water. That sounds obvious but we'll find in a moment that when we mix a column of air, the temperature is not uniform throughout. A well-mixed, uniform-temperature water column is called **statically stable**, because if the fluid is left alone, it won't feel any need to circulate. Any parcel of water has the same density as its neighbors.

Next we'll turn on the burner, warming and expanding the water at the bottom of the pan. Buoyant water at the bottom tends to rise; this situation is called **convectively**

unstable. The buoyant water from the bottom could rise to the top like the lava in the lava lamp, in which case we would end up with denser water underlying lighter water, which we call **stratified**. Pubs in Ireland serve a concoction called a black and tan, with warm dark Guinness Porter overlying cold Bass Ale. The two types of beer remain unmixed because they are stratified by temperature. Alternatively (back to the pan analogy, alas) the rising warm water could mix with the rest of the water, as generally occurs to a large extent in the atmosphere and ocean. In this case we end up with a second statically stable profile at a higher temperature than the first.

Convection in a compressible gas is analogous to convection in an incompressible fluid like our pan of water (Figure 4-8). We begin from the temperature profile of static stability. As for the water column, we can construct a statically stable temperature profile by mixing the gas thoroughly. No gas parcel will be more or less dense than its neighbors because it's all the same stuff. The pressure decreases as you ascend the gas column, and so a gas parcel raised from the bottom of the column expands, and its temperature drops. After it does so, our parcel finds that it is still exactly the same temperature as the gas it finds itself surrounded by. The parcel was neutrally buoyant at the bottom of the column, and it is neutrally buoyant aloft.

Convection is driven by heating at the bottom of the column, such as by sunlight hitting the ground. The warmed air parcel from the ground finds itself warmer than the air immediately above it, even when it has expanded to the pressure of the air above it. The rising parcel follows its own adiabat, which is higher than the temperature profile of the gas column, and so the rising parcel has the ability to rise to the top of the gas column if it does not mix with other gas on the way up. If the gas does mix, the temperature profile of the entire column will rise to a new adiabat, all completely analogous to the incompressible case of the pan of water.

Moist Convection

Convection in a column of a compressible gas is more complicated than it was in our nice simple pan of incompressible water, but I hope it helped to think about the incompressible water case first. To get to convection in the real atmosphere, though, we need to add one more ingredient, and that is latent heat from water vapor. Imagine a source of water vapor at the earth's surface, so that the air above the ground has a relative humidity of 100%. If we raise this air parcel aloft, it will decompress and cool. The saturation water vapor pressure decreases with temperature (Figure 4-6), and so the air parcel becomes supersaturated, carrying more water than it wants for equilibrium. The excess water condenses into cloud droplets or perhaps rain. As the water condenses, it releases its latent heat, just as the steam from the teapot burns your skin by releasing latent heat. The effect of the latent heat is to warm the air parcel up a bit in the face of its decompression cooling (Figure 4-9). We distinguish between the dry and the wet cases by referring to a **dry adiabat** or a **moist adiabat**. One more definition: the decrease in temperature with altitude in the atmosphere is called the **lapse rate**. The lapse rate of a dry convective atmosphere would be about 10 K per km of altitude, whereas the lapse rate of a wet atmosphere is only about 6 K per km.

Convection in the Layer Model

The layer model that we constructed in Chapter 2 did not have convection. Think about a layer model with multiple atmospheric layers, such as you constructed in the lab. The temperatures of the layers decrease with altitude, just like the real atmosphere does, but the only way that heat is carried between layers is by blackbody radiation. One could construct a model with a continuous atmosphere, with temperature varying smoothly with altitude like the real atmosphere, but where radiation is the only way of moving heat. The temperature profile you would get from a model like this is known as **radiative equilibrium**. Radiative equilibrium controls the temperature structure inside some stars, and would control the temperature profile in Earth's atmosphere, except that convection kicks in first. A radiative equilibrium lapse rate in our atmosphere would be about 16 K per km of altitude, so steep that it would be convectively unstable.

If we were to try to add convection to the layer model, we would have to add another set of heat flow arrows to the model, representing the heat carried by air as it ascends, and by water vapor as it condenses releasing latent heat. We are not going to create such a model, but [Figure 4-10](#) gives an impression of how it might look. Convection might insist that the temperature of layers aloft must follow a moist adiabat, roughly 6 K per km of altitude.

Lapse Rate and the Greenhouse Effect

The steeper the lapse rate, the stronger the greenhouse effect. If the atmosphere were incompressible like water, and convection maintained a uniform temperature with altitude, as it did in our pan of water on the stove, there would be no greenhouse effect. To understand why this is so, imagine increasing the concentration of CO₂ in an atmosphere. This has the effect of raising the altitude in the atmosphere where light on average escapes to space ([Figure 4-11](#)). You remember from Chapter 3 that reality is a bit complicated in this regard; some IR light goes directly to space from the ground, in the frequency range of the atmospheric window, while at other frequencies, like in the CO₂ band, the light appears to originate from the coldest part of the atmosphere, at the tropopause. But bunching all of that light together in our minds, we can imagine that, on average, the light comes from a higher altitude as the CO₂ concentration goes up. Let's call that altitude the **skin altitude**.

Now think back to Chapter 2 and the layer model. Remember that the outermost part of the atmosphere, the part that radiates directly to space, always had the same temperature in all of those different model configurations. We called that the skin temperature, and it was always 253 K for the albedo and sunlight intensity of the earth, whatever the model configuration. In solving the layer model for the temperatures of all the atmospheric layers and the ground, it was convenient to start from the outer skin layer and work down toward the ground. Let's take that approach again, assuming that the temperature change with altitude (the lapse rate) has been decided by the physics of convection. If the skin altitude were 5 km high and the lapse rate were 6 K per km, then the temperature of the ground would be

$$T_{ground} = T_{skin} + \frac{6 \text{ K}}{\text{km}} \cdot 5 \text{ km}$$

We can visualize this as a line drawn downward from the skin altitude, following the slope of a moist adiabat until it intersects the ground (Figure 4-11). If we increase the CO₂ content of the atmosphere, we raise the skin altitude, and the same moist adiabat slope intersects the ground at a higher temperature. Algebraically, the change in temperature from raising the skin altitude z_{skin} can be calculated as

$$\Delta T = \Delta z_{skin} [\text{km}] \cdot \frac{6 \text{ K}}{\text{km}}$$

Here's the point. If the lapse rate were different than 6 K / km, then a change in CO₂, driving a change in the skin altitude, would have a different effect on the temperature of the ground. If the atmosphere were incompressible, for example, and convection insisted that the temperature should be uniform with altitude, then it would make no difference how high the skin altitude was.

$$\Delta T = \Delta z_{skin} [\text{km}] \cdot \frac{0 \text{ K}}{\text{km}} = 0$$

The temperature of the ground would be the same as the skin temperature no matter what.

The lapse rate in the atmosphere (how quickly temperature decreases with altitude) is determined primarily by convection and the "hidden" heat carried aloft by water vapor. The lapse rate determines the sensitivity of the temperature of the ground to changes in the IR opacity of the atmosphere. If we want to forecast the effect of the rising atmospheric CO₂ concentration on the temperature of the ground, we will have to get the lapse rate right, and any changes in the lapse rate that result from future climate change.

Take-Home Points

Pressure decreases with altitude.

Temperature decreases as a gas expands.

Moisture in the rising gas releases its latent heat as it condenses.

The lapse rate is controlled by the moist adiabat.

The strength of the greenhouse effect depends on the lapse rate.

Lab

1. Lapse Rate. Use the on-line full-spectrum radiation model at http://geosci.uchicago.edu/~rca/toa_balance.html . Adjust the lapse rate in the model and document its impact on the equilibrium temperature of the ground.

2. Skin Altitude. Answer this question using the on-line IR radiation model.

- a. Run the model in some configuration without clouds and with present-day $p\text{CO}_2$. Compute σT^4 using the ground temperature, to estimate the heat flux that you would get if there were no atmosphere. The value of σ is $5.67 \cdot 10^8 \text{ W}/(\text{m}^2 \text{ K}^4)$. Is the model heat flux at the top of the atmosphere higher or lower than the heat flux you calculated at the ground?
- b. Now calculate the “apparent” temperature at the top of the atmosphere by taking the heat flux from the model and computing a temperature from it using σT^4 . What is that temperature, and how does it compare with the temperatures at the ground and at the tropopause? Assuming a lapse rate of $6 \text{ K} / \text{km}$, and using the ground temperature from the model, what altitude would this be?
- c. Double CO_2 and repeat the calculation. How much higher is the skin altitude with doubled CO_2 ?
- d. Put CO_2 back at today’s value, and add cirrus clouds. Repeat the calculation again. Does the cloud or the CO_2 have the greatest effect on the “skin altitude”?

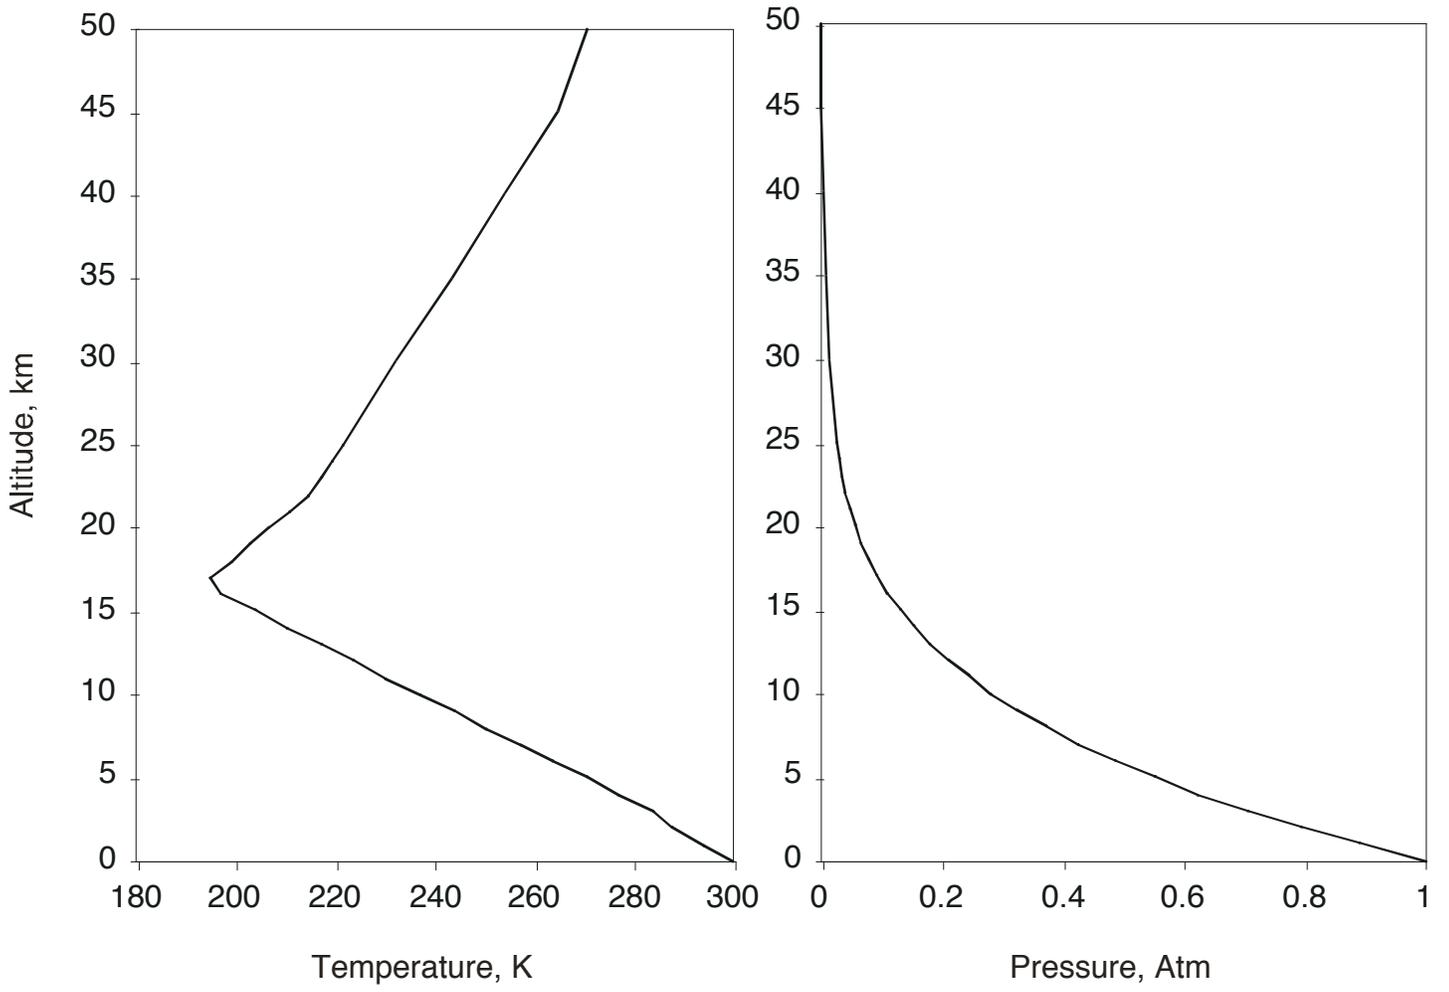


Figure 4-1

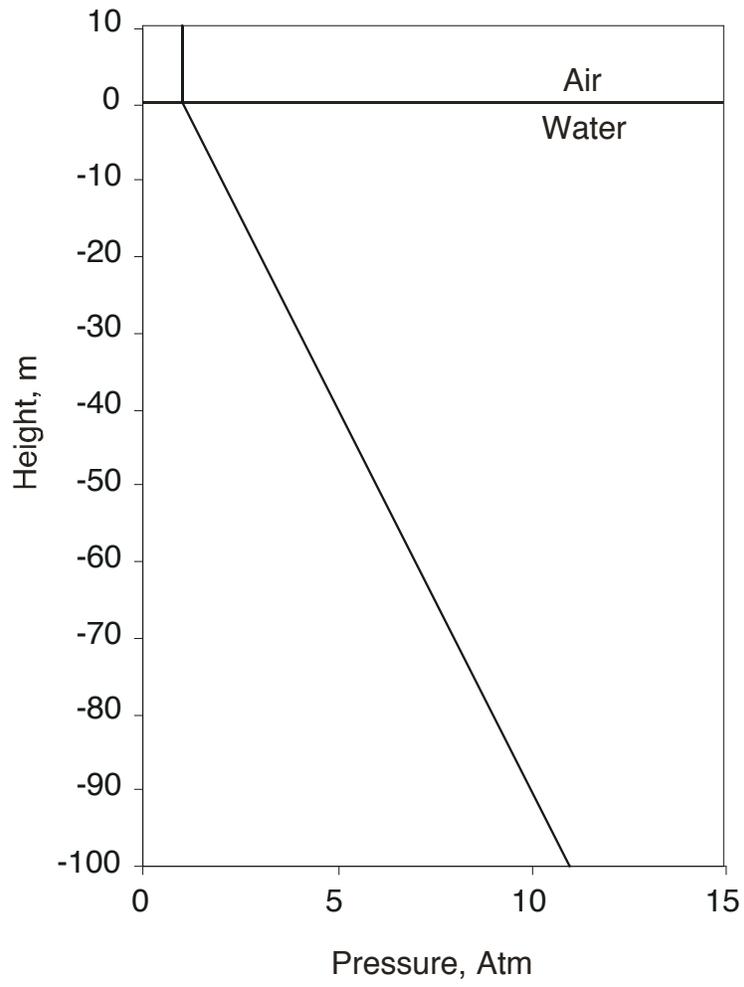


Figure 4-2

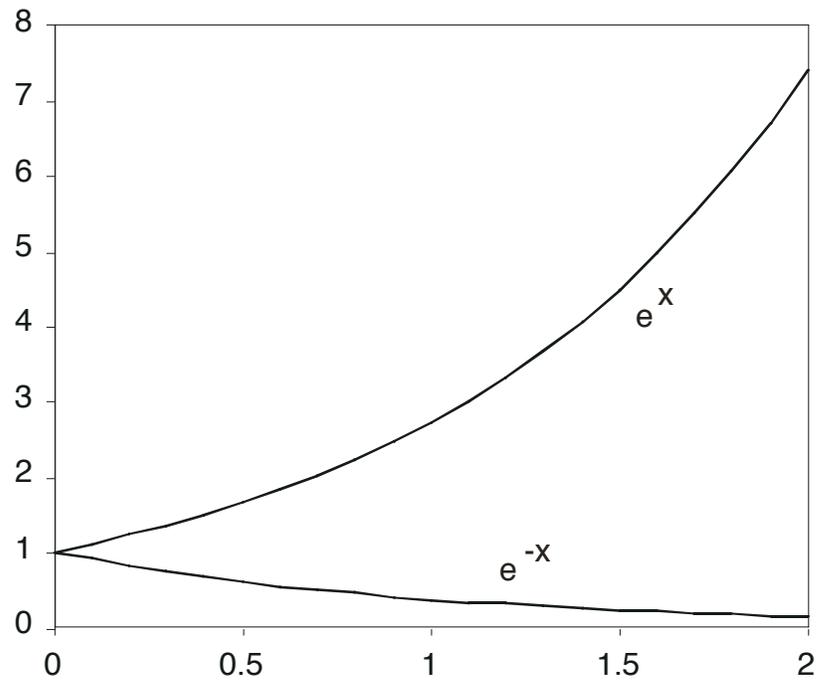


Figure 4-3



Compressible Bricks



Incompressible Bricks

Figure 4-4

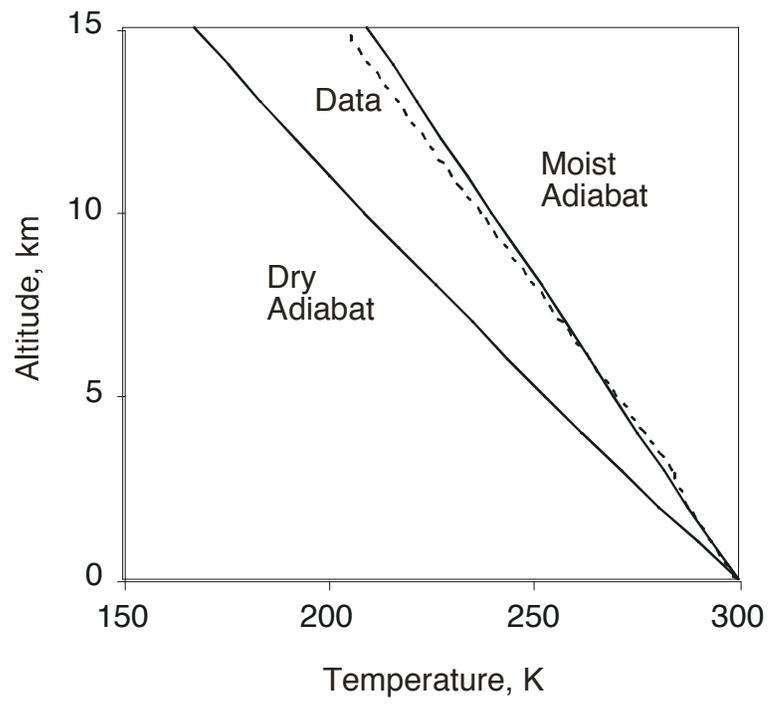


Figure 4-5

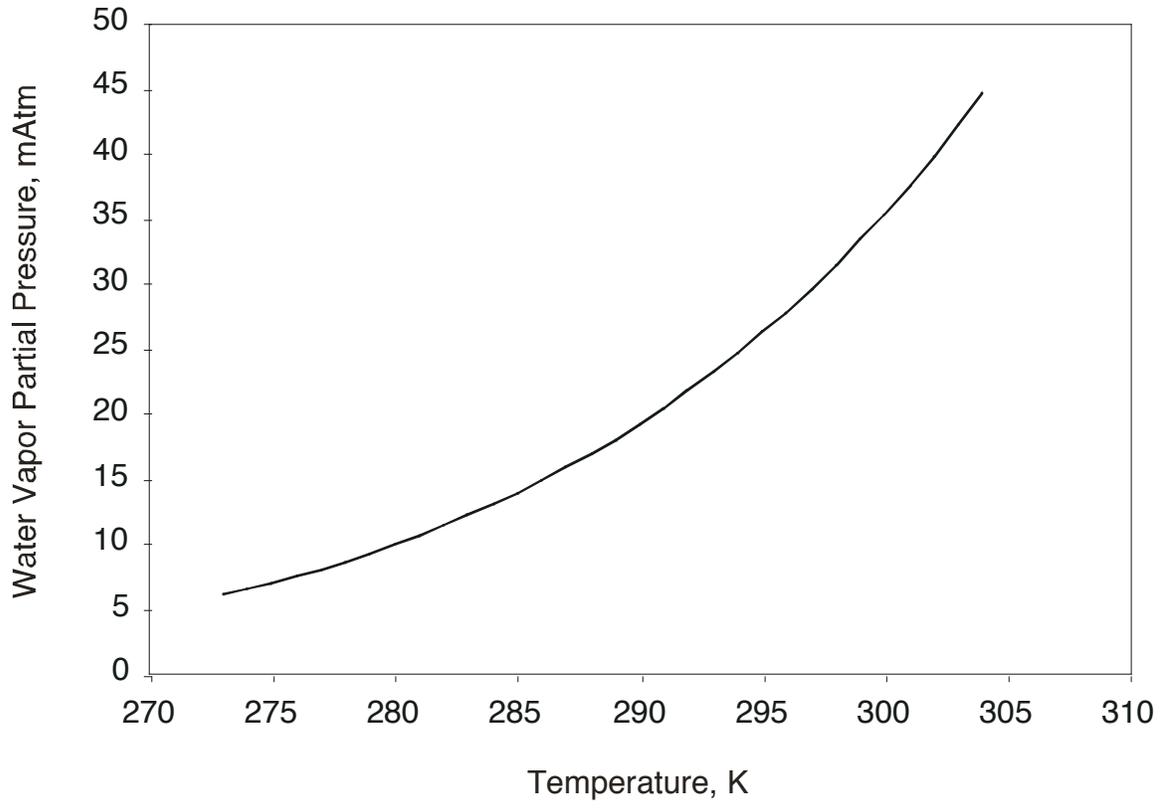


Figure 4-6

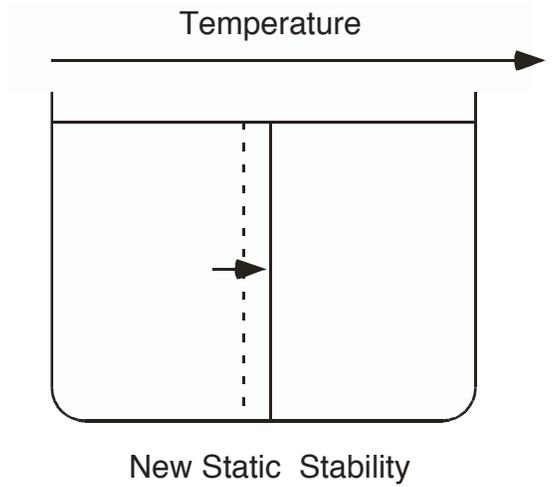
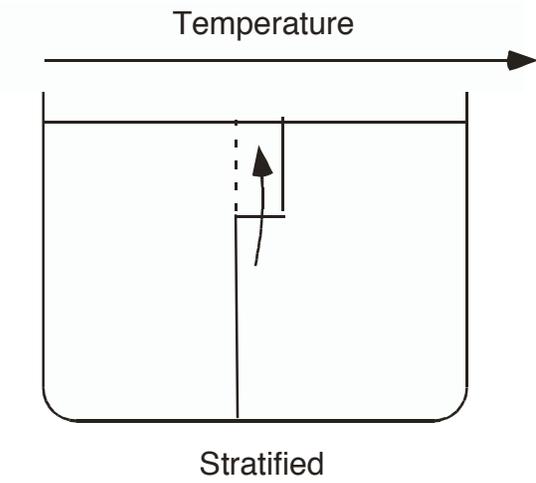
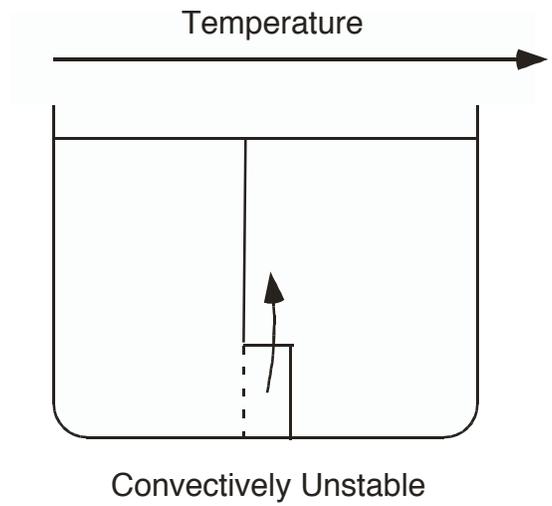
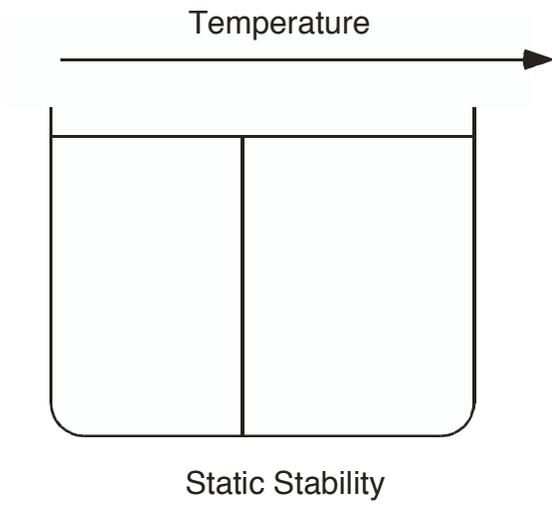


Figure 4-7

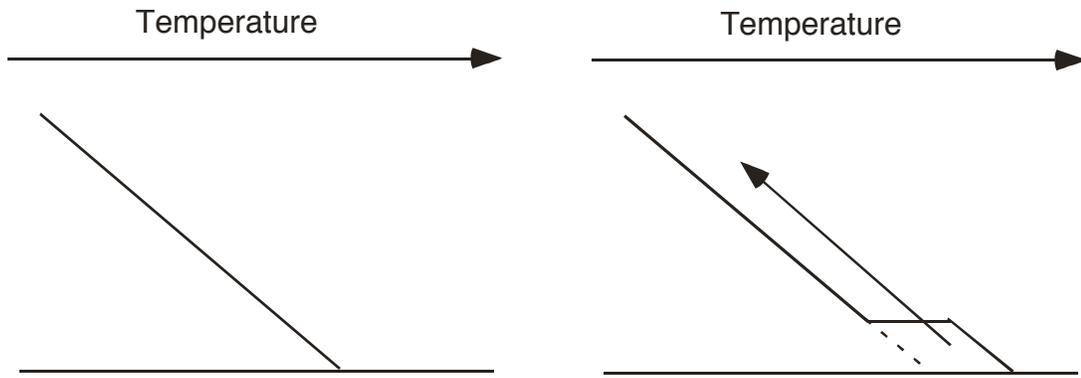


Figure 4-8

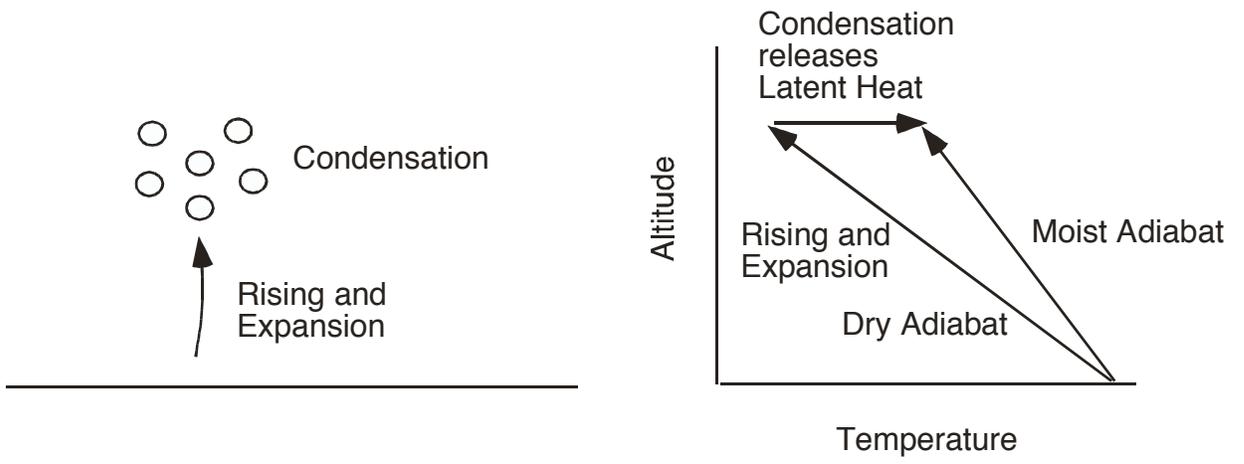


Figure 4-9

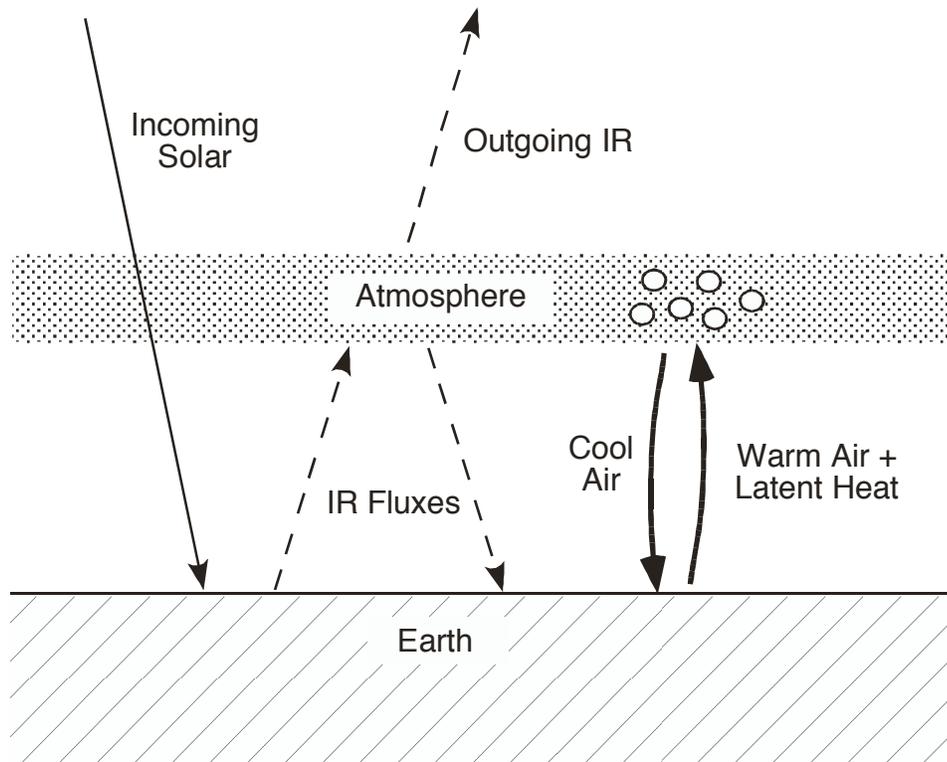


Figure 4-10

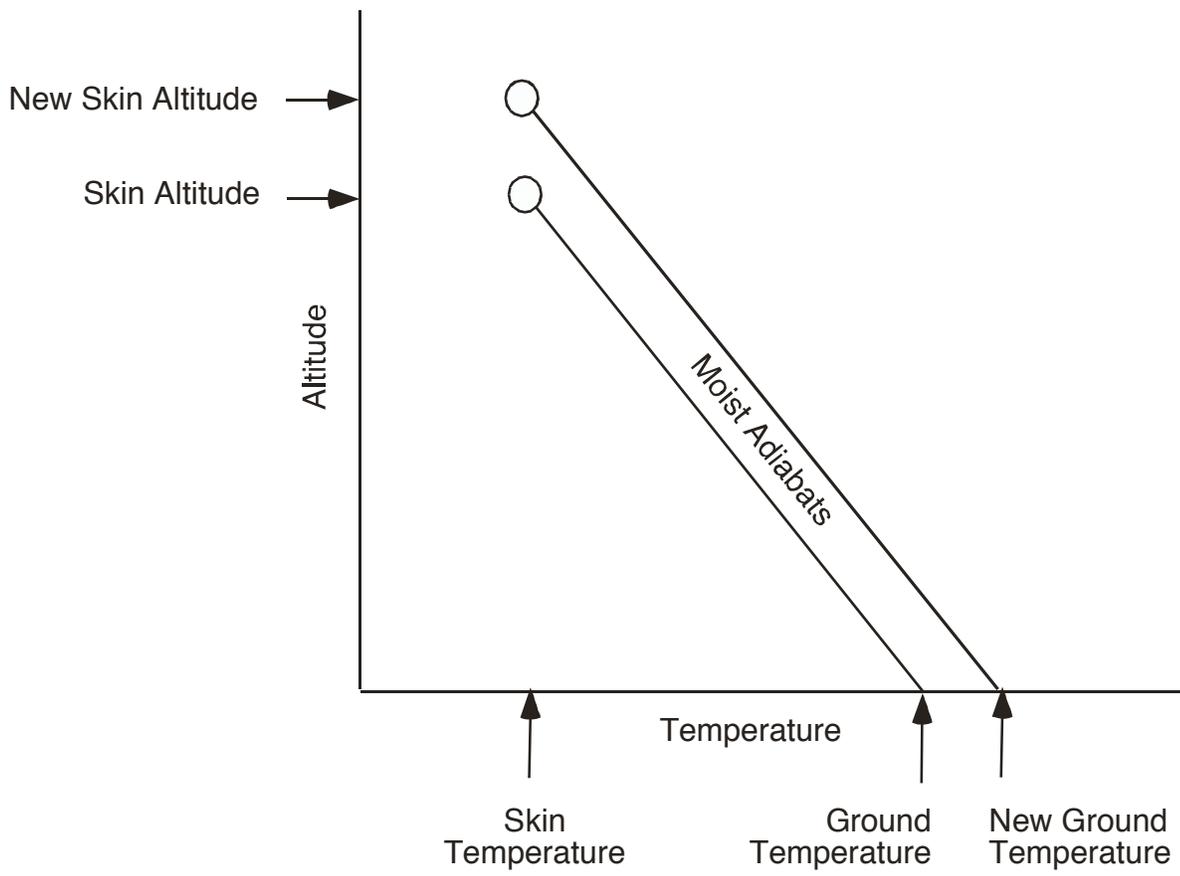


Figure 4-11

Chapter 5. Heat, Winds, and Currents

Abstract

Unlike the layer model, the real world is not all the same temperature. The seasonal cycle drives surface temperatures to values above and below the average value. Also unlike the layer model, energy budgets such as in the layer model do not balance locally even if you average out the time-varying part. There is a net heat input in low latitudes, and transport to high latitudes. This process of heat transport is complex, driven by the turbulent flow of air and water on the rotating earth, requiring complex computer climate models.

Averaging and Aliasing

The layer model assumes that the heat fluxes into and out of the earth must be in balance exactly. What the layer model is after is an average temperature over the entire globe and over time. This is a reasonable assumption to make if we are willing to wait long enough for our answer to be right. On long enough time scales there is simply nowhere else for the heat energy to go; what comes in must go out. As we look in more detail, however, there are all kinds of wild imbalances. It takes heat every spring to warm lake and ocean water and melt the snow. There is also an imbalance between the heat fluxes on the long time average reflecting heat transport from equatorial regions and the higher latitudes to the north and south.

So is an eternal, unchanging, averaging model a reasonable one for a world that is bouncing around like jello? Can you construct the average of the whole system by using averages of the pieces of the system? Or will averaging change the answer? The term for this possibility is **aliasing**. Aliasing issues come up a lot in the natural sciences. In principal, there could be a problem with averaging IR energy fluxes, because they are a nonlinear function of temperature. One way to see the nonlinearity is to look at the equation and see that pressure is proportional to temperature to the fourth power, and not to the first power, which would be linear.

$$I \left[\frac{W}{m^2} \right] = \varepsilon [unitless] \sigma \left[\frac{W}{m^2 K^4} \right] T[K]^4$$

Another way would be to notice that a plot of the function is not a straight line (Figure 5-1). Let's say we wanted to estimate the average energy flux of a planet that had two sides, one at 300 K (rather like Earth) and the other at 100 K (much colder than anywhere on Earth). The outgoing energy flux from the cold side would be about 6 W/m², and from the warm side 459 W/m². The average energy flux would be 232 W/m². Let's now average the temperatures, run the average T through σT^4 , and try to get the same answer. The average temperature is 200 K, and the predicted energy flux is 91 W/m². We'd be off by more than a factor of two. You can see the effect graphically as the difference between the straight line and the curve in Figure 5-1.

For the terrestrially normal range of temperatures (Figure 5-1), we are zooming in on the nonlinear function enough that it looks much straighter than when we considered an absurdly wide temperature range. Still, it could be an important effect, and we certainly wouldn't want to neglect it in the global warming forecast. It doesn't appear that the aliasing is so bad as to undermine the principals demonstrated by the layer model, however, because in this case the function turned out to be fairly linear.

Aliasing could arise when calculating a time-average from a time series of measurements, as if ships of opportunity measured the sea surface temperature more in the summer or the winter, for example. The problem can be corrected for, as long as it is recognized. There are many other **nonlinearities** in the climate system, pairs of variables that are related to each other in nonlinear ways. For example, many effects of the wind, such as exchanging heat with the ocean, depend nonlinearly on wind speed.

Weather versus Climate

If you don't like the springtime weather in Chicago, just wait a few days. It's a grey 5°C out there now, but the forecast for the weekend puts it up to 15°C, which is a little more springlike. The 10-day forecast says showers the weekend after that, but no one believes the end of a 10-day forecast anyway. They're better than they used to be, but 10 days is still something of a crapshoot. And here I am sitting down to write about forecasting the climate 100 years from now. I suppose some would feel that an explanation might be in order.

It is indeed tricky to forecast the **weather** too far in advance, because weather is **chaotic**. To a scientist, the word chaotic brings to mind an extreme sensitivity to initial conditions, so that small differences between two states tend to amplify, and the states diverge from each other. This behavior is called the **butterfly effect**, the reference being to a puff of air from a butterfly's wing eventually resulting in a giant storm somewhere that wouldn't have happened if the butterfly had never existed. The butterfly effect was first observed in a weather simulation model. The model computer run stopped, and the researcher Edward Lorenz restarted it by typing in the values of model variables like temperatures and wind speeds, but with a small seemingly insignificant typo. It didn't take long for the restarted simulation to diverge from the results of the initial simulation. The weather is forecast by constructing an initial condition from meteorological data and past model results, and running this initial condition forward into time using a model. The initial condition will never be perfect, and these imperfections tend to blow up, so that by about 10 days the prediction becomes meaningless. One way to cope with this situation is to run the model lots of times with tiny variations in initial conditions; an **ensemble** of model runs. Then you can average the results of the ensemble as a sort of best-guess forecast. It doesn't fix the problem exactly but it does result in a more reliable forecast.

The **climate** is defined as the time-average of the weather. One could speak of a **climatological** January, which would be the average of many Januaries. How many Januaries? The weather varies on all sorts of time scales, up to hundreds of thousands and millions of years. The definition of climate is a semantic issue; in practice, one could

probably get away with talking about the different “climate” of the Little Ice Age, or even perhaps a “climate shift” from one decade to the next. Let’s say a 10-year average of the weather. Forecasting climate is not as difficult as forecasting weather, because it doesn’t matter if the rain you predict comes on Tuesday or Thursday. The butterfly can’t foil our climate prediction with her storm, if the model gets the large-scale tendency for storminess right. If the model doesn’t follow a storm trajectory this time, it will next time; the butterflies average out. You would also like your climate model to get the frequency of the most extreme events, etc., what climatologists would call getting the **statistics**, right. Even the decade and longer timescale “climate” output from models is sensitive to initial conditions, however, and so the ensemble technique is used for climate forecasting, predictions to 2100 for example, just as they do for weather forecasting.

The heat input forcing to the real world does not sit unchanging at the eternal average value, but rather it bounces around pretty wildly. On the day-night cycle for example, the night-time side of the earth receives no solar input at all, while infrared energy loss to space continues around the clock (Figure 5-2).

The heat input varies over the seasonal cycle as well. The **seasons** are caused by the tilt of the earth’s axis of rotation (Figure 5-3). The issue is not the intensity of the incoming sunlight at the top of the atmosphere, I_{solar} , because the size of the earth is negligible compared to the distance from the sun. Rather, the issue is that the summertime hemisphere, the south in Figure 5-3, gets more than its share of the solar energy, and the winter hemisphere gets less. Figure 5-4 shows a sort of map of the intensity of solar heating as a function of latitude and day of the year. The colors and contours show us the intensity of sunlight per m^2 on the ground averaged over 24 hours. Because the earth completes one rotation in 24 hours, any location on some line of latitude, say 42°N which goes through Chicago will get the same solar influx as any other at that same latitude, like Barcelona or Vladivostok. A map of solar influx in regular latitude and longitude would look like a bunch of horizontal stripes. Therefore we are showing two “dimensions” over which the solar influx does vary; latitude and day of the year.

The beginning of the year is southern hemisphere summer, and we see red colors in the south and blues in the north. Northern hemisphere summer is the middle of the year, centered around day 180. Now we see red, representing high energy flux, in the north and blue in the south. The color pattern arises by two mechanisms. First, the intensity of sunlight per m^2 of ground area is greater at noon in the summer than in the winter, because the ground is at less of an angle to the sun in the summer (recall Figure 5-3). Secondly, days are longer in summer, and this increases the 24-hour average energy influx.

It is interesting to note that the highest daily-average energy fluxes you find anywhere are at the north or south pole during north or south summer. The sun never sets in midsummer at the pole, it just whirls around in a circle over the horizon. In winter the pole gets no sunlight for months on end. The poles do not turn into tropical garden spots in summer because it takes time for temperatures to respond to changes in heat forcing.

This **thermal inertia** tends to damp out the temperature swings of the day/night cycle, the seasonal cycle, and any global warming temperature trends as well. Liquid water, it turns out, has much greater temperature stabilizing power than land or ice does, because water mixes, bringing a much greater mass of water into thermal exchange with the atmosphere, than land can bring. Cooling of surface waters drives convection, the same way the warming of surface air brings convection in the atmosphere. The seasonal cycle of temperatures in soils reach only a few feet down, while the temperatures farther down, in caves for example, reflect the yearly average, about 10 C in my part of the world. The seasonal cycle of temperature changes at the ocean reaches a hundred meters down and more. For this reason, the seasonal cycle is much more intense in the middle of large continents than it is in “maritime” areas impacted by the temperature stabilizing effects of liquid water.

Even if we average out the seasonal cycle, the energy budget of a local region of the earth may be far from in balance, because heat energy is redistributed around the earth’s surface by wind and water currents (Figure 5-5). There is a net influx of heat in the tropics; sunlight brings in energy faster than it is radiated back to space. Heat is carried poleward by flow in the atmosphere and the ocean. In high latitudes, the earth vents the excess tropical heat, as excess radiative heat loss to space over direct solar heating.

So We Have to Simulate the Weather

The upshot of this chapter is that in order to forecast global warming, we will have to simulate the time and space variations and imbalances in the energy budget, and the way that the earth’s climate responds to this forcing, by storing or transporting heat around. The layer model will not do. Unfortunately, the physics which govern the flow of air and water are complex and difficult to simulate by their very nature. The difficulty is that fluid flows such as in the atmosphere and the ocean take place on a wide range of size scales. Large-scale currents or winds tend to spawn smaller-scale vortices and eddies. If we are thinking about air flow around an automobile or water flow in a cup of coffee, the tiniest flows eventually get absorbed by the friction of the fluid. Some fluids like molasses have more friction than others like water; they are said to have higher **viscosity**. This process of large-scale flow breaking up into smaller and smaller eddies until they get so small that they are absorbed by the fluid is called the **turbulent cascade**. The physicist Lewis Richardson summed it up in verse

*Big whorls have little whorls,
Which feed on their velocity,
And little whorls have lesser whorls,
And so on to viscosity.*

At the heart of it, the physics governing fluid flow and turbulence is fairly simple. Newton’s laws of motion describe how an object of some mass will respond to a force pushing on it in some direction. Fluids are basically urged to flow downhill. Fill up a bathtub with water and then by magic pull up on the water surface on one side, making a slope (Figure 5-6), and then let it go. The water will flow to flatten out the surface, of

course. Another way to see what just happened is to think about the pressure within the fluid relative to the gravity field of the earth. When we say “downhill” we mean downhill relative to a surface that is flat, which we call a **geopotential surface**. The surface of a pool table had better be a good geopotential surface, or the balls won’t roll straight. Let’s say that the bottom of the bathtub is a geopotential surface. The pressure is higher on the deep side because there is more water overhead. The **pressure gradient** in the bathtub could be calculated as

$$\text{Pressure Gradient} = \frac{\Delta P}{\Delta x} = \frac{P_{\text{shallow}} - P_{\text{deep}}}{1.5 \text{ m}}$$

The pressure gradient is high in the bathtub. Our intuition is that the fluid should flatten out. When that happens, the pressure gradient is zero on the geopotential surface. The bathtub is behaving the way Newton’s laws say it should; a pressure force pushes fluid, and the fluid begins to flow in the direction the force is pushing.

Flows in the atmosphere or ocean differ from those in the bathtub, however, in that those in the atmosphere and ocean persist for long enough that they are steered by the rotation of the earth. Let’s ponder this topic sitting on a merry-go-round at a playground ([Figure 5-7](#)). Two kids sit opposite each other on a merry-go-round tossing a ball back and forth while spinning. The thrower (T) tosses the ball, and it goes straight as Newton would like. Before the ball reaches the catcher (C) on the other side, the merry-go-round has rotated. The catcher misses the ball. What the kids see is that the ball curved, relative to their rotating world on the merry-go-round.

The same situation applies to a weatherman attempting to describe the motion of the winds on the rotating earth. The way to interpret a weather map is to think of the earth as if it were fixed relative to the stars. Then add one other magic force, a fudge force called the **coriolis acceleration**, which gives you all the same effects as the rotation of the earth that we’re ignoring. The coriolis acceleration is a fake force which we imagine is applied to any object in motion. Let’s return to the ball on the merry-go-round. Our fake force should change the direction of the ball without changing its overall velocity. To do this, the fake force must be at exactly 90° to the direction that the object is moving. In [Figure 5-7](#), the coriolis fake force would have to be directed to the right of the direction of motion, so that it would turn the trajectory of the ball to the right. On Earth, the coriolis force goes to the right in the northern hemisphere, and to the left in the southern hemisphere.

Now we have to jump from the flat merry-go-round to the gravitating sphere. When we rotate the sphere, what does that do to the people trying to read their weather maps? How much rotation do they feel? The best way to envision this is to imagine a Foucault’s pendulum. Foucault pendulums are usually set up in an atrium or stairway so that the weight can hang on a wire 10 or 20 meters long. The wire is so long that the pendulum swings with an enchantingly long, slow period of maybe 10 or 20 seconds. The weight is rather heavy at 100 kg or so. The wire is mounted on a swivel to allow the pendulum to rotate freely if it chooses to. Once the pendulum has started swinging it goes all day. At

the base of the pendulum, some museum employee sets up a circle of little dominos which the pendulum knocks over as it swings. Over the course of the day, the swing direction of the pendulum changes, knocking over a new little block every hour, on the hour or some other crowd-pleasing rhythm. Leon Foucault installed the first of these into the Pantheon in Paris for the 1850 Paris Exposition.

Now we can use the pendulum to think about the question of rotation on the surface of the earth. Let's set up a Foucault pendulum on the North Pole (Figure 5-8). The earth spins relative to the fixed stars. The pendulum maintains its swing orientation to be stationary relative to the fixed stars. The pendulum knocks over dominos on both sides of its swing, so they will all be knocked over after the planet has rotated 180° , which will take 12 hours.

Next let's move the pendulum to the equator. The trajectories of the pendulums cannot remain completely motionless with respect to the fixed stars, because the direction "down" keeps changing as the planet rotates. If we start the pendulum swinging in a North-South direction, the pendulum will keep its alignment with the stars by continuing to swing North-South. The swing of the pendulum does not appear to rotate to an observer on earth. The dominos are safe.

What we are looking for here is the alignment of the axis of rotation with the local direction of gravity (down). At the equator, the two are at right angles to each other, and there is no rotation. In the middle latitudes, the rate of rotation is in between these extremes of the pole and the equator. Actually, the rate of rotation scales with the sin of the latitude, a trigonometric function. [Figure? Explain sine?](#)

On the merry-go-round, the strength of the apparent coriolis acceleration depends on how fast the merry-go-round was rotating. The same applies on the rotating sphere. The rate of rotation about a vertical pole is faster in high latitudes than in low. Therefore the apparent coriolis force is stronger in high latitudes than in low.

If we release the sloping surface of the water in the bathtub in Figure 5-6, the water will tend to flow in the direction that the pressure is pushing it, to flow downhill. The driving force for an ocean flow could also be that the wind is blowing (Figure 5-9). As the fluid begins to flow, the coriolis acceleration begins to try to deflect the flow to the right (in the northern hemisphere). After a while the fluid is flowing somewhat to the right of the direction that we're pushing it. Eventually, if we wait long enough, the flow will reach a condition called a **steady state**, a condition in which the flow can persist indefinitely. In the steady state, the coriolis force balances the driving force, and so there is no net force acting on the fluid. The astonishing implication is that in a rotating world the fluid will eventually end up flowing completely cross-ways to the direction of the forcing! This condition is called **geostrophic flow**. It is as if the sloping surface in the bathtub would drive a flow across the bathtub. Which direction? Let's stop and figure it out. The pressure from the sloping surface would drive the flow from left to right. 90° to the right of that would be flowing straight out at the reader, from the page. The water in the bathtub does no such thing, of course, because (1) it doesn't have time to adjust to the steady state condition in Figure 5-10, and (2) the bathtub is not infinitely wide, and water

cannot flow through the walls. Great ocean currents which persist for longer than a few days, though, do flow sideways to their driving forces. Sea level on the east side of the Gulf Stream in the North Atlantic is about one meter higher than it is on the west side. Did I get that right? Figure out the direction that the pressure would be pushing the water, and verify that the flow is 90° to the right.

You can see geostrophic flow on a weather map as cells of high and low pressure with flow going around them (Figure 5-10). A low pressure cell in the northern hemisphere has a pressure force pointing inward all around its circumference. 90° to the right of that, the winds flow counter-clockwise (in the northern hemisphere). Meteorologists call this direction of flow **cyclonic**. Flow around a high pressure cell is **anticyclonic**. The fluid flows around the cells rather than into or out of the cells, preserving rather than smoothing out the pressure field that drives it. The ball rolls sideways around the hill rather than directly down the hill as we'd expect. Pressure bumps like hurricanes tend to persist in the atmosphere like balls that never run downhill.

Some sets of mathematical equations governing some systems, like the layer model, can be solved directly using algebra to get the exact right answer. The equations of motion for a turbulent fluid are not that easy to solve. We are unable to calculate an exact solution to the equations, but must rather approximate the solution using a computer. A solution we are after consists of temperatures, pressures, flow velocities, and other quantities related to the flow or the drivers of the flow. These quantities vary from place to place, and as a function of time. To represent them in a **computer model**, the atmosphere and the ocean are diced up into a three-dimensional **grid** of numbers. There is a temperature, a pressure, velocities etc. associated with each grid point. The computer model calculates the pressure force at each grid point by comparing the pressures at adjacent grid cells. The model steps forward in time to predict the response of the fluid velocities to the pressure forces. The flows carry around fluid, altering the pressure field, and so the simulation snakes it way forward through time. A typical time step for a climate model may be a few minutes.

Weather simulation has always been a problem ideally suited for computers. Even before computers, researchers like Lewis Richardson dreamed of using human computers to perform the immense numbers of calculations that would be required (Figure 5-11). ENIAC (Figure 5-12), one of the first electronic computers constructed in 1945, was used for a range of scientific computation projects including weather forecasting. Weather and climate simulation can be done fairly efficiently on clusters of computers. Each computer in the cluster is called a **node**. Each node is responsible for some area of the globe, which is divided up like a checkerboard. Each node keeps track of the values of velocity, etc. at the grid points of the nodes adjacent to it on the globe. Whenever a node alters the values in its little piece of the world, it essentially sends email to its neighbor nodes, a technique called **message passing**.

There is an unavoidable tragic trade-off in numerical climate simulation that will prevent us from doing the simulations we would really like to do until computers get much faster than they have now. The problem is that the mechanisms that govern fluid

flow often operate at fairly small spatial scale. If we wanted a model to include everything that governs how cloud drops form, we would have to have grid points every few meters in clouds for example. Fronts in the ocean can be only a few meters wide. To get the answer right, we'd like to have lots of grid points. The trouble is that if we increase the number of grid points by, say, a factor of 10, in each dimension, the total number of grid points goes up by a factor of $10 \cdot 10 \cdot 10$ or 1000. That's 1000 times more mathematical operations the computer has to perform in order to do a time step. To make matters worse, as the grid gets more closely spaced, the time step has to get shorter. A grid 10 times finer would require a time step about 10 times shorter. So perhaps it would take 10,000 times longer to do a model year of time. State-of-the-art computer models of climate are run at higher resolution all the time, and in general they look more realistic as resolution increases. But they are still far from the resolution they would like, and will be so for the conceivable future.

Some of the physics of the real atmosphere cannot be explicitly resolved, so they must be faked in some way. The real formation of cloud droplets depends on small scale turbulent velocities and particle size spectra and other stuff like that which the model doesn't have. So the model is programmed to use what information the model does have, ideally in some intelligent way that might be able to capture some of the behavior of the real world. Cloud formation may be assumed to be some simple function of the humidity of the air, for example, even though we know reality is not so simple. The code word for this approach is **parameterization**. The humidity of the air is treated as a parameter that controls cloudiness. Other important and potentially weak-link parameterizations include the effects of turbulent mixing, air / sea processes such as heat transfer; the list goes on. We will discuss these in more detail in Chapter 10.

Take-Home Points

The energy budget of the earth fluctuates with time on daily and seasonal timescales (in contrast to the layer model).

The annual average energy budget may not balance locally, either, because excess heat from the tropics is carried to high latitudes by winds and ocean currents.

The global warming forecast requires simulating the weather, which is a real computational challenge.

Lab

1. The Orbit and Seasons. Answer this question using an on-line model of the intensity of sunlight as a function of latitude and season at <http://geosci.uchicago.edu/~archer/orbit.html>. The model calculates the distribution of solar heating with latitude and season depending on the orbital parameters of the earth. Enter a year A.D. and push calculate. The eccentricity is the extent to which the orbit is out-of-round; an eccentricity of 0 would be a fully circular orbit. Obliquity is the tilt of the earth's axis of rotation relative to the plane of the earth's orbit. The third number, the longitude of the vernal equinox, determines the location on the orbit (the date of the year)

where Northern hemisphere is tilted closest to the sun. Using the present-day orbital configuration, reproduce Figure 5-4. Now straighten the tilt of the earth's axis of rotation by setting obliquity to 0° . What happens to the seasons?

2. Heat Transport. Answer these questions using an on-line full-spectrum radiation model at http://geosci.uchicago.edu/~archer/cgimodels/full_spectrum.html .

a. The incoming solar radiation at the equator, averaged over the daily cycle, is about 420 W/m^2 . What would the temperature be at the equator if there were no heat transport on earth? The default albedo in the web interface is 30%, but the real albedo in the tropics may be closer to 10%. What happens then? How much heat transport is required to get a reasonable temperature for the equator? What fraction of the solar influx is this?

b. Repeat the same calculation for high latitudes. Estimate the annual average heat influx at 60 degrees latitude by running the orbital model from the first problem. Just eyeball the fluxes through the year to guess at what the average would be. Now plug this into the full-spectrum light model to see how cold it would be up there if there were no heat transport. If there were no transport and also no storage, how cold would it be in winter?

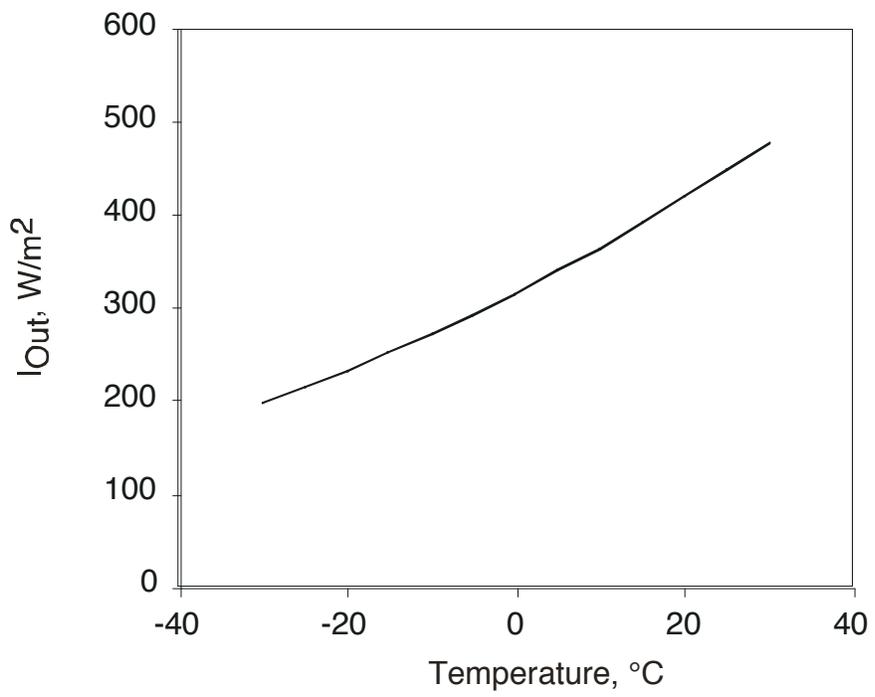
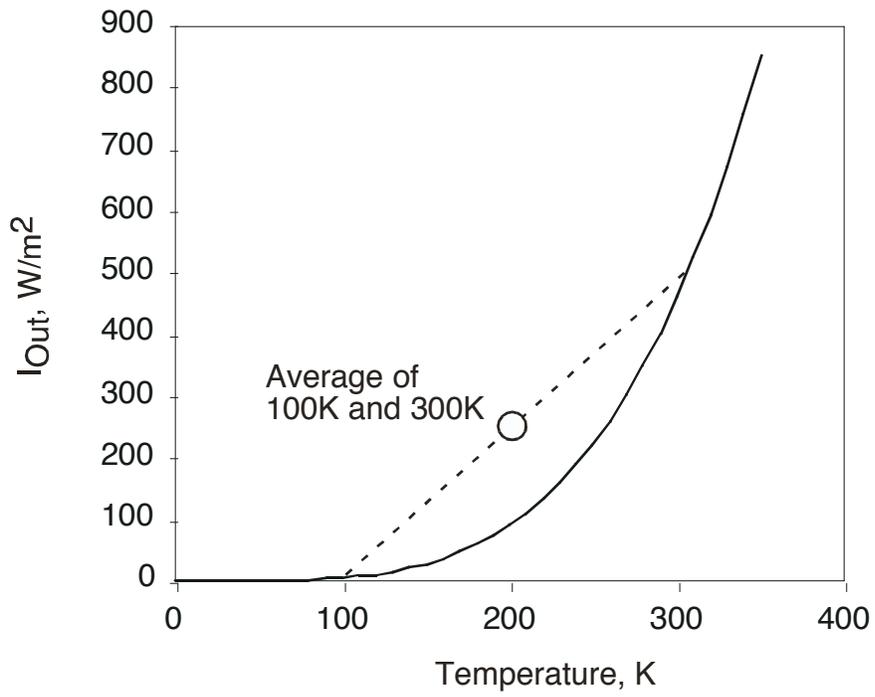


Figure 5-1

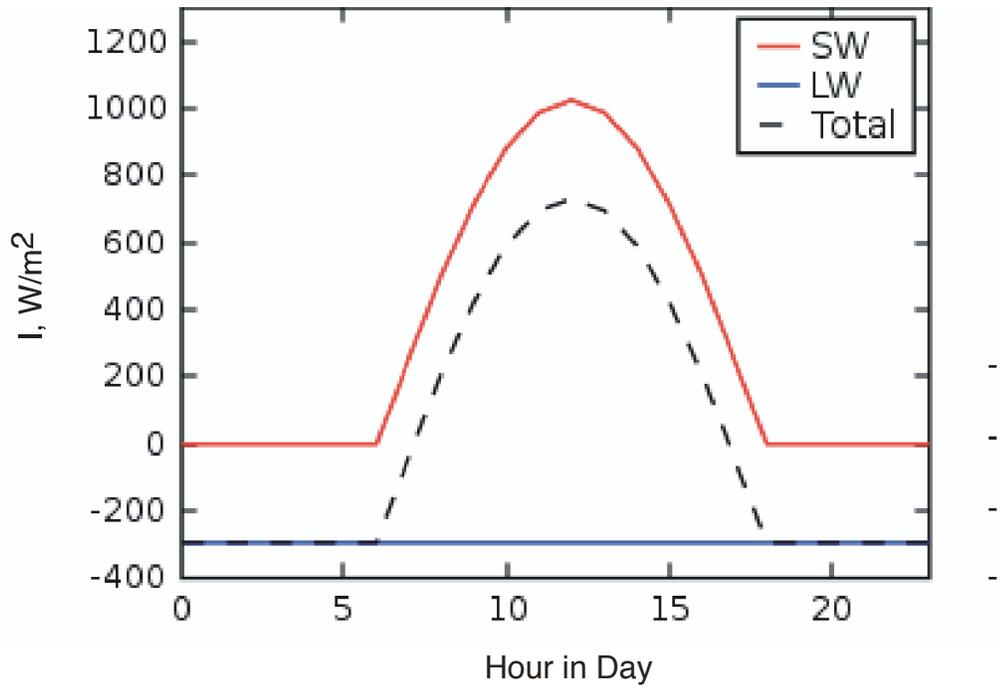


Figure 5-2

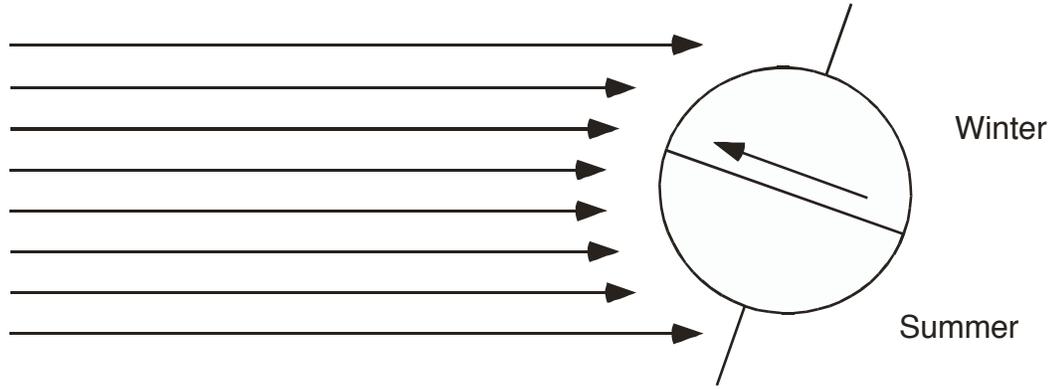


Figure 5-3

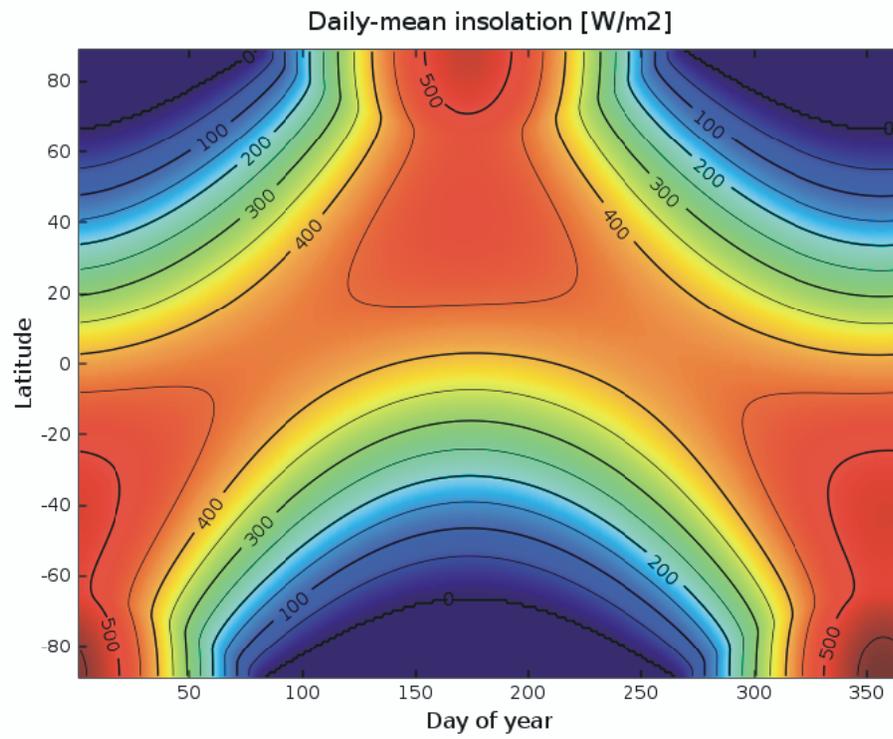


Figure 5-4

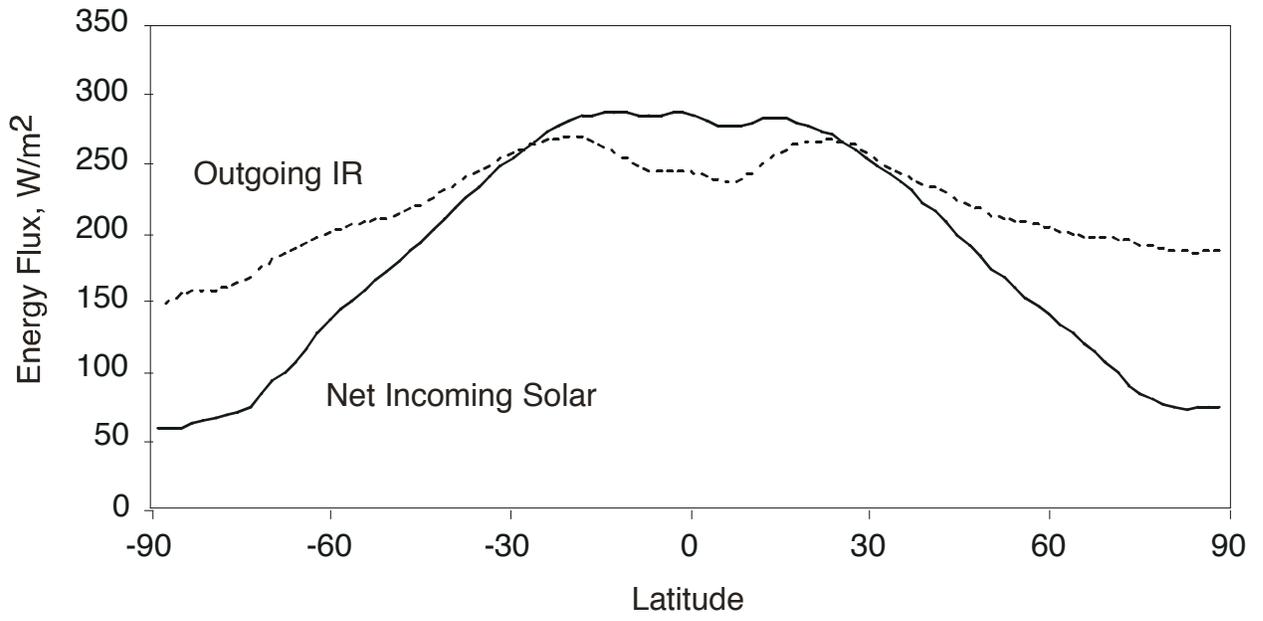


Figure 5-5

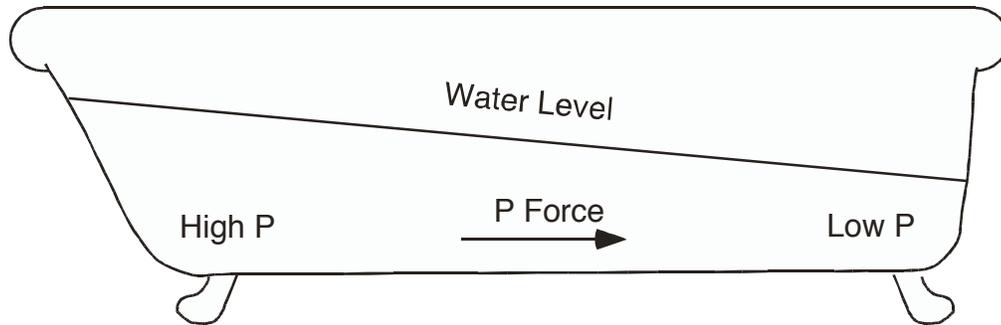
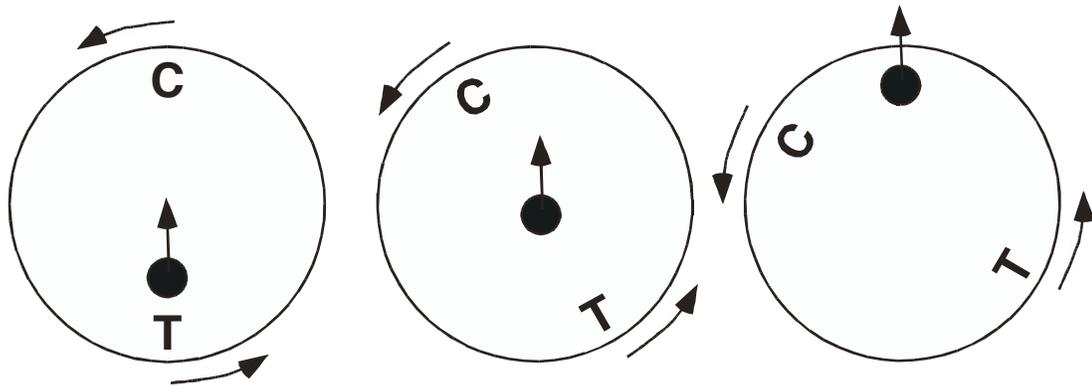
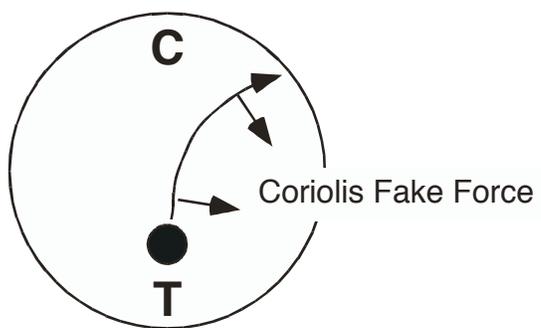


Figure 5-6



Stationary Observer

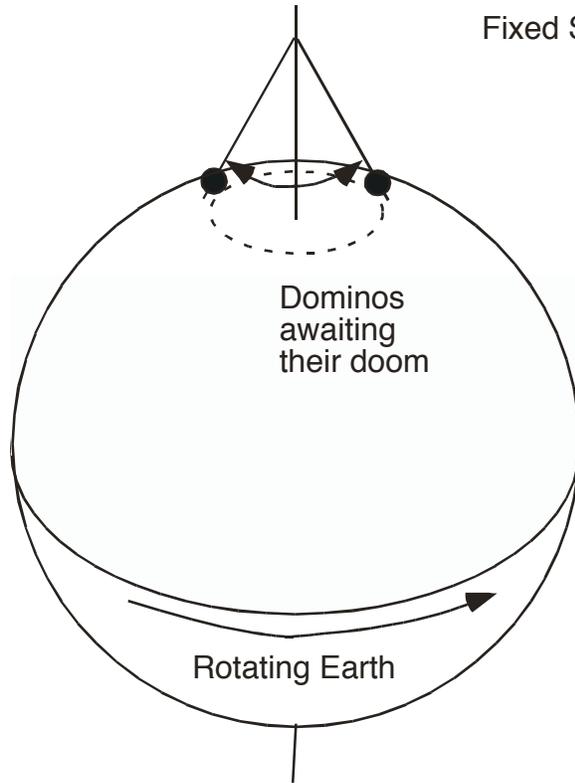


Rotating Observer

Figure 5-7

Fixed Stars

Fixed Stars



Fixed Stars

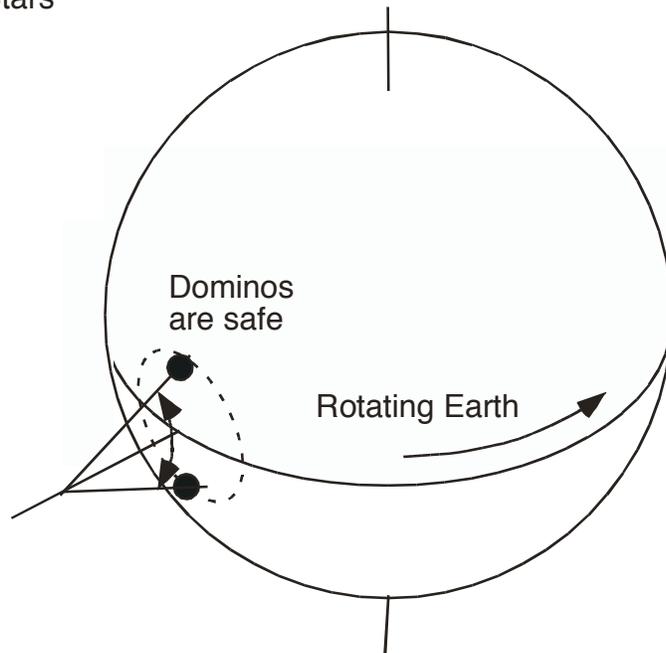
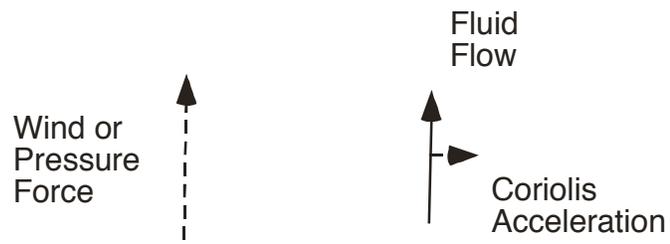
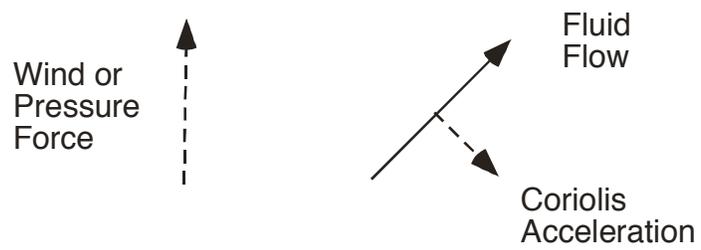


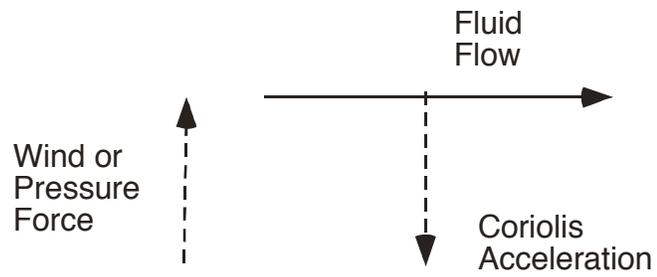
Figure 5-8



When You First Turn on the Wind



After a Few Hours



The Eventual Steady State

Figure 5-9

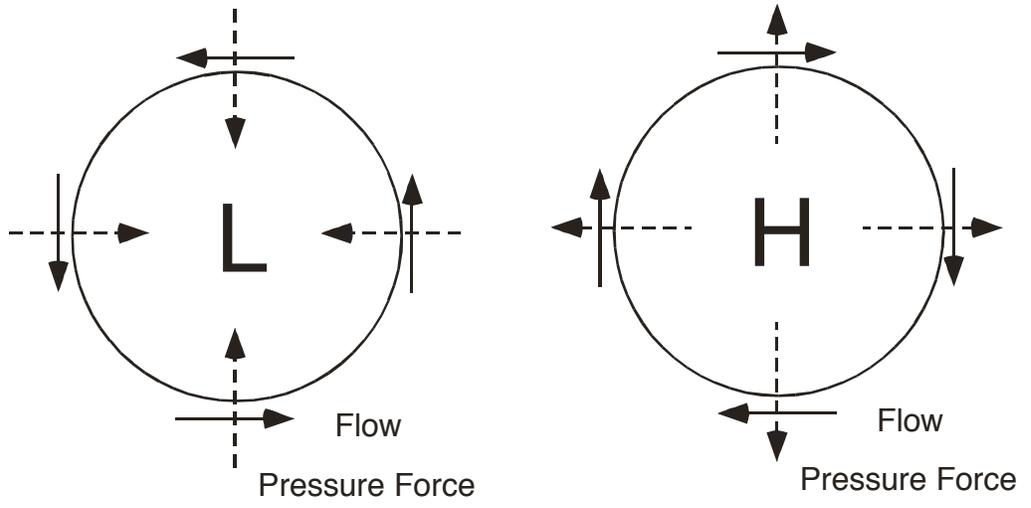


Figure 5-10



Figure 5-11

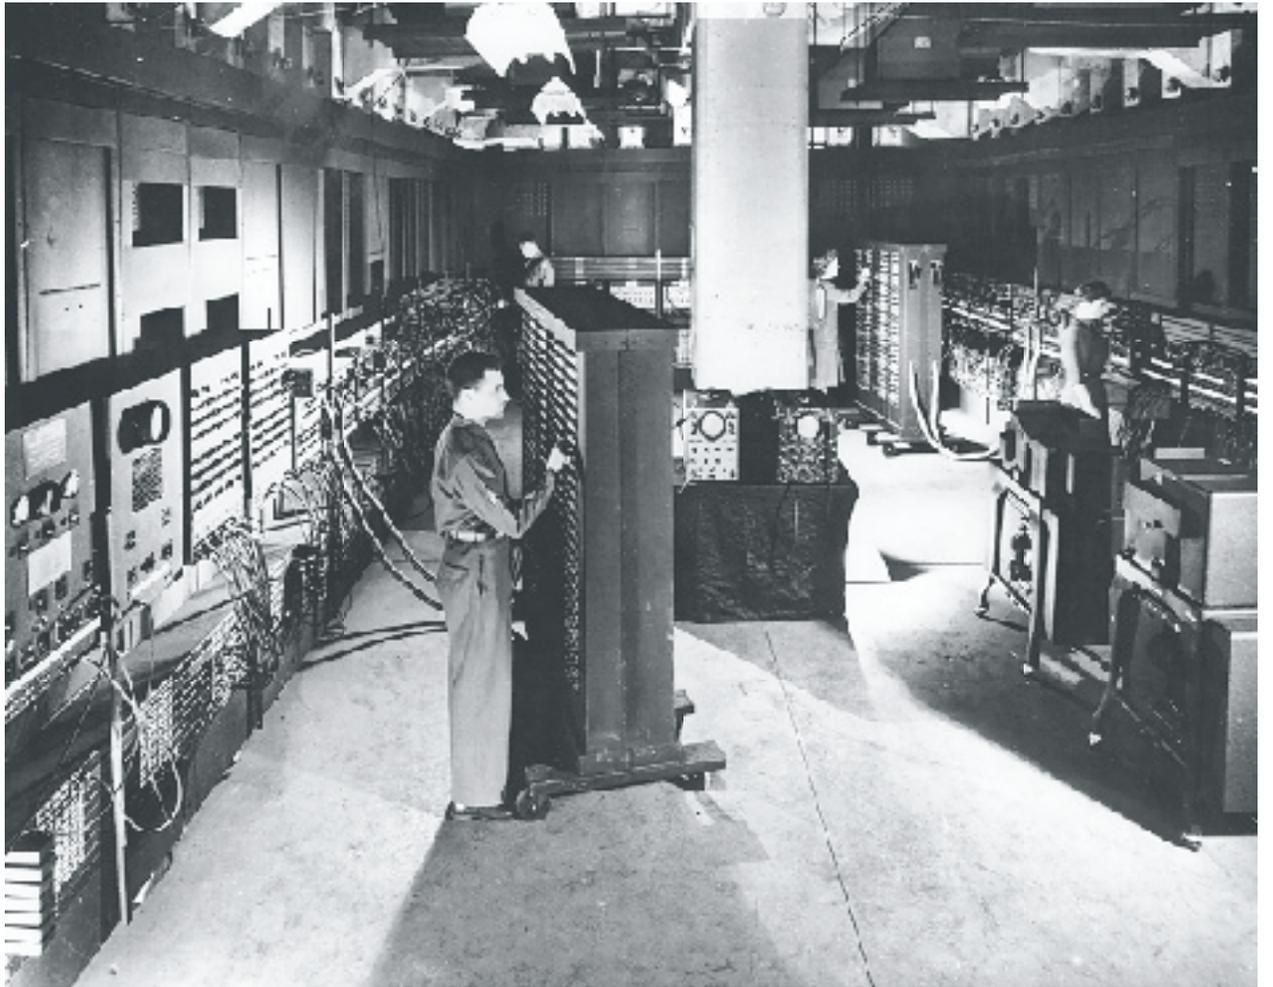


Figure 5-12

Chapter 6. Feedbacks

Abstract

One tricky part about modeling climate is the way that different parts of the climate system interact with each other. Positive feedbacks tend to amplify the variability of climate, while negative feedbacks provide stability. Water evaporates into warmer air, acting as a greenhouse gas to amplify the initial warming. Melting ice changes the albedo of the earth, acting as a positive feedback to temperature variations. Clouds have a huge influence on climate by interacting with both visible and IR light. We can point to present-day human influence on cloudiness, but it is impossible to predict how clouds will feed back to long-term climate change. Ocean circulation has feedback interactions with climate, in the equatorial Pacific and in the North Atlantic. In general, the real climate appears to be tippier than our model climates, because it is difficult to capture all of the feedbacks in models.

Positive and Negative Feedbacks

The challenging thing about modeling climate is that there are **feedbacks**. A feedback is a loop of cause-and-effect (**Figure 6-1**). It will be easiest to wade through this particular thicket if we make a few abstruse-sounding definitions here. The temperature is the quantity that gets fed-back-to in many examples, although feedbacks also control populations and chemistry and water vapor lots of other things. So we can be general, let's call the quantity which is being fed-back-to the **state variable**. The loop of causality is called the **feedback loop**. The arrow coming in from stage left is an **input perturbation**.

One example of a feedback loop is the **ice albedo feedback**, operating on the state variable of temperature. An input perturbation, such as a rise in greenhouse gases, drives temperature up a bit. Ice melts, reducing the albedo, which drives temperature up a bit more. This is called a **positive feedback** because the directions of the input perturbation and the feedback loop agree with each other. A positive feedback can work in the opposite direction as well, during glacial time for example: colder, more ice, colder still. The point is that the feedback loop, in this case, reinforces the direction of change in the input. A positive feedback loop is an amplifier.

Negative feedbacks tend to stabilize things. One example from Chapter 7 will be the rate of igneous rock weathering as a feedback to the state variable of the CO₂ concentration in the atmosphere. CO₂ gets too high, weathering will pull it back down. The water vapor concentration in the atmosphere at a particular temperature, which we will run into in the next section, is another example.

Emergent behavior & physical vs. social science comparison?

Water Vapor Feedback

Water vapor is responsible for more greenhouse heat trapping on Earth than CO₂ is, and yet your impression is that global warming is primarily about CO₂, not water vapor. No one screams about global warming when you run your lawn sprinkler. What's up with that? That's because, at some given temperature, the amount of water vapor is controlled by negative feedbacks with the rate of evaporation and with rain. Our lawn sprinkler may be depleting our fresh water supply, but it is not going to lead to global warming, because any water we drive into the atmosphere will just rain out next week. A negative feedback stabilizes the water vapor concentration, at any given temperature.

The positive feedback arises when we start to think about water vapor affecting a state variable of temperature. An initial increase in temperature will allow more water to evaporate, where its presence as a greenhouse gas feeds back to increase temperature still more. Where does this end, you may wonder to yourself. Does that feedback ever stop, or does it explode? It can; it is called a **runaway greenhouse effect**. It happened on Venus, but it won't happen on Earth for a few billion years yet. The easiest way to visualize the water vapor feedback stopping or not stopping is to imagine trajectories on a phase diagram (Figure 7-2). You may not have guessed, at this very moment, that "trajectories on a phase diagram" could possibly be an easy way to visualize anything, but let's give it a try. The **phases** of water are solid, liquid, and gas. The phase diagram for water which we'll be looking at describes which phase or phases of water you will find, if you subject water to some pressure and temperature. If it's hot and low pressure, you expect to find vapor. Vapor likes it hot. As you decrease the temperature and increase the pressure, you get liquid, and when it gets cold enough you get ice. On the boundaries between the regions you get two phases, like vapor and water, or ice and water. At the triple point, a certain pressure and temperature with no latitude for variation in either, you get all three phases coexisting together, like a glass of boiling ice water. Of course, you can throw ice cubes in a pot of boiling water on your stove any time you like, but that's not at equilibrium.

To this phase diagram we add some planetary water vapor and temperature trajectories. We'll start our thought experiment from an initial condition of no water vapor, then we'll watch as water evaporates. The pressure axis of the plot is actually the partial pressure of water vapor. As water vapor evaporates, its partial pressure in the atmosphere goes up. Water vapor is a greenhouse gas, so the temperature warms as the vapor pressure increases. This you can see as the leftward tilt of the trajectory, toward higher temperatures as water vapor increases. Eventually, the trajectory runs into the water stability field. Guess what happens then, it rains. And the water vapor feedback has done all it is going to do.

The trajectory for Mars is similar to that from Earth except that it starts out colder, and it intersects the ice field rather than the liquid. Venus is the fun part of this story. Venus is close enough to the sun that its water vapor feedback trajectory escapes; it never hits equilibrium with either of the condensed phases ice or liquid. This is a runaway greenhouse effect. The oceans boil dry. As a ton of dirt on the coffin, if water vapor

reaches the upper atmosphere, it will get zapped apart by ultraviolet light into oxygen and hydrogen atoms. The hydrogen atoms are light enough that they are able to escape to space, and the water is lost for good. This has never happened on Earth because the upper troposphere is cold, like a giant cold trap preventing our water from leaking to space. Marvelous!

A runaway greenhouse is not likely to happen on Earth, not for a long time anyway. If the equator were isolated from the poles, blocked from any heat transport and forced to balance its energy fluxes using outgoing IR only, the tropics would be a runaway greenhouse. Lucky thing we have the high latitudes as cooling fins.

The role that the water vapor feedback plays in the global warming forecast can be visualized (see, that was easy, wasn't it?) on this same diagram, although global warming is tiny compared with Venus, so I've drawn it horribly exaggerated so we can see it. An initial small temperature perturbation moves the earth a smidgeon toward Venus, in the initial condition with no water vapor pressure at the base of the plot. Trace the trajectory up as water evaporates, and it hits the liquid saturation line at a bit hotter temperature than the first Earth did. The temperature difference between the two raining Earths is greater than the temperature difference between the dry Earths. The water vapor has amplified the temperature perturbation that you would get from CO₂ alone.

On the real earth, mercifully, the relative humidity is not everywhere 100%. One mechanism which controls a lot of the atmospheric water vapor franchise is called **Hadley circulation** (Figure 6-3). Warm air at the equator rises convectively. Water condenses as the air rises and cools. This column of air has a lot of water vapor in it, in general. The air spreads out at high altitude, then begins to **subside** in the subtropics, about 30° latitude North and South. That air has been through the wringer, the cold tropopause, and there's not much water vapor left in it. The great deserts of the world are strategically located under these dry air blowers. The air flows back equatorward along the surface, picking up water vapor as it goes. Globally, the average humidity of surface air is about 80%. One could imagine changing the pattern of the winds, the Hadley circulation for example, and changing the humidity of the surface air. Turn the dry air blowers up a little higher, and the surface air will dry out a bit. Could a negative feedback mechanism like that (Figure 6-1) save us from global warming? This is a difficult question to answer for sure, because it hinges on issues of turbulence and fluid flow, impossible to model perfectly (Chapter 5). Models tend to agree on this issue, however. The relative humidity doesn't change with global warming, is their general consensus.

Ice Albedo Feedback

The ice albedo feedback is a positive feedback that operates mostly in the high latitudes, because that is where the ice is. According to our current climate models, global temperature changes will be amplified in the high latitudes, by a factor of about 3 or 4. The strength of the ice albedo feedback also gets stronger in colder climate states such as the glacial climate. In addition to changing the albedo, the presence or absence of ice on the ocean has a huge effect on the local air temperature. Air over water can't get

much colder than freezing, 0°C, but ice won't prevent surface air from getting much colder than that, any more than land would. As we learned in Chapter 5, water absorbs and emits heat, buffering the extremes of temperature, whereas land and ice does not so much. Sea ice changes may have a lot to do with fast, extreme changes in climate reflected in Greenland ice cores during glacial climate.

The ice albedo has the potential to get stuck in either of two stable states under **the same forcing**. Imagine removing greenhouse gas from the atmosphere gradually until suddenly the temperature drops below freezing and ice appears, cooling the temperature still further (**Figure 6-4**). Now raise the CO₂ slowly, and for a range of CO₂ concentrations, the ice will persist, where it was ice-free on the way down. The ice decreases the albedo and makes the local climate cold enough for the ice to stay frozen. This is an example of a **hysteresis** loop.

Clouds

Clouds are the sleeping giant of the climate system. Most of the variation in sensitivities of different climate models comes from clouds. Clouds interact and interfere with both visible and infrared light. It turns out that according to visible light, clouds should cool the planet, while the infrared thinks that clouds warm the planet. It gets even more complicated; which of these two effects wins depends on the altitude of the clouds. High altitude clouds don't reflect much visible light, but they are pure hell in the infrared. Lower clouds are the reverse; less of an infrared effect because they are warm, but a large change in albedo because they are thick and opaque. We will keep a table of comparisons between high and low clouds, to keep all this straight.

You have already encountered the infrared side of the story, in Chapter 3. Clouds are tolerably good blackbodies, emitting a nice clean blackbody spectrum of light at whatever temperature it finds itself. Light escapes to space from the tops of the clouds, so the crucial thing to know is how high the tops of the clouds are, and therefore what the **cloud top temperature** is. High clouds emit infrared light at cold temperatures, blocking the bright infrared light from the warm ground. Clouds with low tops have a smaller effect on the outgoing IR light.

Clouds also reflect incoming solar radiation, increasing the albedo of the earth. The electric field of incoming visible light stimulates the electric field in the water droplet to oscillate, and this oscillation emits light of its own, at the same frequency as the incoming light. The name for this process is **scattering**. This interaction of light with matter is not the same as light absorption, because the energy never gets converted to thermal energy of the water. Absorbed light would be re-radiated in the infrared, but scattered light in effect bounces off of the water droplets on photons at the same frequency they rode in on. Some of the light energy is absorbed also, and the fraction of scattered versus absorbed light depends on many things. As a result, the albedo of clouds varies widely between different types of clouds. Scientists are struggling to figure out how real clouds work and how they might be changing already. Modeling them perfectly is currently out of our reach. The best we can do is very crude approximation.

One important factor in determining the albedo of a cloud is the amount of water it contains. Low clouds contain more water than high clouds, in general. Clouds that are about to rain have more water in them than non-raining clouds. Another factor is the size of the cloud drops. Scattering is more efficient in smaller drops if the drop size is smaller. A third factor is the directionality of the scattered light. Most of the light scattered by spherical liquid drops continues in a more-or-less forward direction, which if the light is coming down vertically, would mean to continue downward toward the earth. Ice crystals are better at actually reversing the direction of light, sending it back up into space, increasing the albedo.

As air rises it cools, and the amount of water vapor the air can hold decreases, which is to say, the **saturation vapor pressure** decreases. We heard about this process in Chapter 4, but let's look at the cloud formation process in a little more detail. As soon as the relative humidity reaches 100%, there would be an energy payback if the water were to condense. But it might not happen right away. The water can get "stuck" supersaturated, that is humidity can exceed 100%, by quite a bit, if the air is very clean with no particles in it for the droplets to form around. Liquid water has the property that it abhors the gas/liquid surface. Abhors may be a strong word, but the water molecules have a definite preference to be immersed, with other liquid water molecules on all sides, rather than sitting on the surface of the drop exposed to the air. Small particles have a greater proportion of these unhappy, unfulfilled surface molecules. Therefore the water in small droplets has a slightly higher tendency to say, to Hell with this, I'm heading out into the gas phase. When we report relative humidity, we are reporting it relative to water that has a normal, flat surface, or maybe large droplets. The saturation vapor pressure in equilibrium with small droplets is higher than it is for contented water molecules with a flat surface at the same temperature. We will require a relative humidity higher than 100% if we wish to form small cloud droplets.

Most air in the troposphere has particles in it before any cloud drops start to form, around which the water droplets may begin to grow. These particles are known as **cloud condensation nuclei**. Clouds begin to form at lower relative humidity if condensation nuclei are present. The condensation nuclei are of order 0.1 micron in size. Once the droplets form, the large droplets in a cloud tend to grow at the expense of the small ones, because smaller ones have the greater tendency to evaporate. Urbanization of cloud droplets. The droplets in clouds tend to grow fairly quickly to sizes of around 5 microns, larger in clouds that are about to rain. At high altitude, ice has a tendency to form, either by freezing liquid or by forming directly from the vapor. Vapor deposition releases latent heat on a growing ice crystal as it grows, resulting in the amazing symmetric patterns that snowflakes are formed in. No two alike, they say. That would be because no two snowflakes have exactly the same temperature and water vapor supersaturation histories.

There are three main types of clouds (**Figure 6-5**). Cirrus clouds are located up at 8-12 km altitude. These are the only types of clouds up that high. Between cirrus clouds and low-altitude clouds there are several kilometers of air space where you typically do not find clouds. Cirrus clouds contain 10 or 100 times less water per volume than lower altitude clouds typically hold. Because they are so thin, cirrus clouds do not block

incoming solar radiation as effectively as lower clouds do. You can see blue sky right through cirrus clouds. The albedo impact of cirrus clouds is therefore weaker than it is for lower, thicker clouds. The IR effect of cirrus clouds is strong, because the cloud is effective enough as a blackbody to block intense IR from the warm ground, substituting for it weaker IR from the cold upper atmosphere.

The two main types of clouds at low altitude are called **cumulus** and **stratus**. Cumulus clouds are towers, the result of a focused blast of convection. Thunderstorms come from cumulus clouds. Stratus clouds are layered, formed by broad diffuse upward motion spread out over a large geographical area. Both types of fluid motion are difficult to simulate in climate models; the focused motions because they occur on spatial scales that are smaller than the grid sizes of the models, and the broad diffusive upward velocities because they are so slow, they are difficult to predict accurately. Stratus clouds are particularly important to get right for the climate forecast, because they affect the albedo of such a large area of the earth's surface. Two-thirds of the albedo of the earth derives from clouds, stratus clouds in large measure.

There is plenty of scope for human activity to impact the climate by altering clouds. Anthropogenic condensation nuclei are mostly combustion products. The abundance of cirrus clouds on earth may be augmented by the **contrails** left by jet airplanes (**Figure 6-6**). In clean air, in the absence of cloud condensation nuclei, the air can be supersaturated indefinitely without forming clouds. An airplane passes through and creates a trail of air with higher humidity, from combusting the hydrogen in the jet fuel, and with exhaust particles that can serve as condensation nuclei. Droplets form very quickly behind the airplane, and then continue to grow and to spread out. Most of the moisture in the contrail comes from the ambient air, not from the airplane. The cloud particles spread out and eventually become indistinguishable from natural cirrus particles. For this reason it is difficult to know how much impact aircraft have on the cloudiness of the upper atmosphere. Contrails tend to warm the planet, by perhaps a few percent of the total human-induced warming. A difference between contrail warming and CO₂-warming is of course that CO₂ accumulates in the atmosphere while contrails would dissipate in a few days, if we stopped creating them.

Low altitude clouds can be affected by smoke particles from coal power plants, in particular the sulfur in coal. Sulfur is emitted in flue gas as SO₂, and it oxidizes in a week or so in the atmosphere to sulfuric acid, H₂SO₄. The sulfuric acid condenses into small droplets less than 1 micron in size called **sulfate aerosols**. Eventually, the sulfuric acid rains out and is called **acid rain**. Sulfate aerosols are extremely good cloud condensation nuclei, because the strong acid tends to pull water out of the air like those silica gel packets that ship with electronics equipment. Natural condensation nuclei include sea salt, dust, pollen, smoke, and sulfur compounds emitted by phytoplankton. Most parts of the world have enough natural condensation nuclei so the issue is not whether to form a droplet or not, the way it is for cirrus clouds and contrails. Rather, the issue is the size of the droplets. If there are plenty of condensation nuclei around, more droplets will form, and the average drop will be smaller. This is a big deal because smaller drops scatter more efficiently and therefore have a higher albedo. Get this: a change in droplet size

from 10 microns down to 8 microns has the same effect on the radiation balance of the earth as doubling the CO₂ concentration. That's a big deal from such tiny little droplets! The potential for sulfate aerosols to change the albedo of the earth by changing the average size of cloud droplets is called the **sulfate aerosol indirect effect**. IPCC only presents uncertainty in this number, without even an estimate of how big it might be (Figure 9-2 in Chapter 9). All error bar and no estimate. One indication that this might be a real effect comes from ship tracks in the ocean, which are followed by clouds of lower-atmosphere droplets just like contrails (Figure 6-7). Another is an observation that the intensity of sunlight at the ground has been decreasing since the 1950's, a phenomenon called **global dimming**. The change in surface solar intensity is huge, tens of W/m² over 50 years. This is too large to be purely the result of albedo, most of that energy much be absorbed in the atmosphere rather than reflected to space, or we'd in the deep freeze by now. It is at the very least an indication that the optical properties of the particles and droplets in the atmosphere have been changing.

Ocean Currents

The oceans interact with climate, in many ways. One example is el Niño, which every taxi driver in California can tell you about. El Niño is a flip-flop between two self reinforcing states, el Niño and la Niña (Figure 6-8). During el Niño, the thermocline along the equator in the Pacific Ocean slopes upward toward the east, bring cold waters to the sea surface. The temperature contrast between East and West drives winds which travel from East to West, maintaining the slope of the thermocline. The other state, la Niña, flattens the slope in the thermocline and collapses the wind. The coupled atmosphere / ocean system flip back and forth between these two climate states one cycle every 4-7 years. As our taxi driver can tell you, the state of el Niño affects climate to varying extents all around the world. There is a possibility that with global warming, the Pacific may tend to favor the el Niño state.

Another potential feedback from the ocean to climate is in the North Atlantic Ocean. Warm water is carried to the North Atlantic in the Gulf Stream discovered by Benjamin Franklin. As the water cools, it sinks to the deep Atlantic, making room for more warm water to carry more heat to the high northern latitudes. Our climatic and oceanographic reconstructions of the North Atlantic region from the last ice age point to an instability in this system. The climate record in the ice in Greenland shows huge temperature swings called Dansgaard-Oeschger events, temperature changes of 10°C within a few years. Oceanographic reconstructions from sediment cores show that the overturning circulation in the North Atlantic starts and stops in synchrony with these temperature swings. One large climate shift, called the Younger Dryas cold reversal, has been correlated with the catastrophic release of a Great Lake's worth of fresh water into the North Atlantic, as an ice dam holding the lake in finally burst.

The simplest explanation of the data would be that the ocean current carries heat to the North Atlantic, keeping Europe warmer than the equivalent latitudes in western Canada. When the current collapses, Greenland and Europe get cold. This has been shown to be a bit simplistic. Europe would be warmer than western Canada, because of

the way the Rocky Mountains steer the jet stream, even without the Atlantic circulation dumping excess heat in the North Atlantic. The current picture in vogue involves sea ice, which has a profound effect on air temperature as mentioned above. Perhaps the current carries enough heat to melt sea ice, or perhaps the salt carried by the current allows cooling water to sink rather than freeze in the wintertime.

The potential for this abrupt climate change mechanism from the past impacting the anthropogenic future comes from the link with fresh water. Warming, amplified in the high latitudes by the ice albedo feedback, tends to increase rain and snow fall, because warmer air holds more water vapor than cold air. Another potential source of new fresh water to the North Atlantic is by melting of the Greenland ice sheet.

Terrestrial Biosphere Feedbacks

The terrestrial biosphere has the potential to feed back to climate, by altering the albedo of the land, and by altering evaporation. Trees have a much lower albedo than does the bare ground. Cooling might lead to the demise of forests, increasing albedo and driving temperatures down further. Trees also mine ground water and evaporate it from their leaves. The Amazon rain forest is said to perpetuate its own existence, by supplying the water vapor needed to keep the rains coming.

Carbon Cycle Feedbacks

There are many feedbacks within the carbon cycle, the subject of the next three chapters. Warming the atmosphere increases the temperature of the surface ocean, which tends to drive CO₂ to degas from the ocean to the atmosphere. The ice ages are thought to be ultimately triggered by variations in the earth's orbit, but the climate shifts would have been much smaller if the ocean hadn't contrived to pull CO₂ out of the atmosphere as the ice sheets started to grow. The biosphere on land, trees, soils, and so on, may be stimulated to take up or release CO₂ as climate changes. Frozen deposits of methane gas in permafrost could melt, creating a strong positive feedback for future climate change.

The Paleoclimate Record Urges Us to Macabre Creativity in Forecasting the Future

In general, naïve climate models seem to have an inherent tendency to underpredict the extremes of climate variation in the real climate, as inferred from climate records such as ice cores. I believe this is due to the difficulty of capturing feedbacks in climate models. Feedbacks often couple together very different parts of the climate system, requiring a cloud physicist to work together with an oceanographer, for example. It takes creativity to think of some of these relationships! The message is that the future may surprise us, and test our creativity in thinking of couplings, interactions, and feedbacks, that may not be obvious at first glance.

Svante Arrhenius

Many of these feedbacks were discovered using pencil, paper, and moonlight by Svante Arrhenius. A scan of his 1896 paper “On the influence of carbonic acid in the air upon the temperature of the ground” is available at <http://>.

Svante Arrhenius lived in Stockholm from X to Y. He is most famous for the “Arrhenius equation”, taught in freshman chemistry, relating the speed of a chemical reaction to temperature. He was also a believer in panspermia, the idea that life was transplanted to earth from space, from whence we know not.

The title refers to carbon dioxide as “carbonic acid”. As we will learn in Chapter 9, CO₂ reacts with water to form an acid that we currently call carbonic acid.

The greenhouse effect had been proposed by Fourier, who “maintains that the atmosphere acts like the glass of a hot-house, because it lets through the light rays of the sun but retains the dark rays from the ground.” Writing like that is a lost art.

It was known at the time that CO₂ and “coal gas” (methane) absorbed infrared, while nitrogen and oxygen do not. However, it wasn’t known very well how much light each gas absorbed, or at what frequencies. Arrhenius did the beautiful thing of trying to use measurements of the IR intensity of moonlight made by Samuel P. Langley. Langley was trying to figure out how hot the moon is by measuring its blackbody spectrum. He measured the intensity using a device he invented called a bolometer, which measures the heat carried by infrared light as it shines on a temperature sensor. The IR light was separated into different frequencies using a prism made of salt. Most solids are opaque to IR light (they are blackbodies) but a crystal of dry table salt is an exception. The different frequencies of IR light are described in Arrhenius’ paper by the angle that they are refracted to by the prism: “rays of angle 40°” and so on.

Langley measured the intensity of moonlight multiple times. He recorded the humidity and the angle of the moon in the sky. If the moon is directly overhead, it travels through a shorter path of CO₂ and H₂O than it does when the moon is low in the sky. The amount of H₂O the light goes through varies independently from the CO₂, when the humidity changes. Arrhenius developed a simple equation relating the light intensity for each frequency to the amount of CO₂ and H₂O the light travels through. The equation was linear, that is to say,

$$I = x \cdot \text{CO}_2 + y \cdot \text{H}_2\text{O}$$

where x and y are absorption coefficients derived from fitting all the data. The interesting issue is that this equation seems to ignore the effect of IR emission by the gas itself (Chapter 2) a. In Arrhenius’ defense, if the temperature of the moon were much higher than the temperature of the atmosphere, you could perhaps ignore the IR light that comes from the atmosphere. The equation also seems inconsistent with saturation of the IR absorption bands by the gas (Chapter 3), although he is clearly aware of the possibility. Arrhenius has linearized the problem, because he didn’t have enough data to

figure out how it would really go. His equation is an approximate one, which probably has a narrower range where it is applicable than a log function like [Figure 3-6](#).

Arrhenius used his absorption equation within a model that is very much like our layer model ([Chapter 2](#)), but much fancier. He ran the model for a range of latitudes and seasons, applying to each location and time the effects of heat transport ([Chapter 5](#)) and cloudiness ([Chapter 6](#)). Then he doubled the CO₂ concentration in the atmosphere. He realized that increasing the temperature of the earth would increase the water content: the water vapor feedback ([Chapter 6](#)). He included the effect of ice on albedo, and predicted that high latitudes would warm more than low latitudes because of the ice-albedo feedback ([Chapter 6](#)). Overall, his predicted temperature change for doubling CO₂ (ΔT_{2x} , [Chapter 3](#)) was 4-6° C, pretty close to our current estimate of 2.5-5° C ([Chapter 11](#)).

Arrhenius mentions the possibility of human-induced warming, but he was not too concerned. The industrial revolution was “evaporating our coal mines into the air” at a rate of 0.5 Gton C / year, compared with our present-day fossil fuel combustion rate of 5 Gton C / year ([Chapter 8](#)). The 1896 rate, which Arrhenius got right, was indeed too small to have much impact on climate except that it grew exponentially since then ([Chapter 9](#)). Arrhenius was mostly interested in past climates, and predicted that pCO₂ might have been 150-170 ppm during the glacial age. Subsequent ice core data tells us it was in fact 180-200 ppm, compared to 280 ppm a few hundred years ago and 370 ppm today ([Chapter 9](#)). He also predicted pCO₂ 2 to 3 times higher than today during the Tertiary, which scientists still believe, more or less.

No one worried about humankind changing the CO₂ concentration of the atmosphere. This was because the ocean covers such a large area of the earth’s surface and contains so much more CO₂ than the atmosphere does ([Chapters 7 and 9](#)). Roger Revelle raised the issue again in 1957, showing that the chemistry of sea water is important to predicting how much CO₂ the ocean actually will take up.

Lab

Answer these questions using the full-spectrum radiation model at <http://geosci.uchicago.edu/~archer/PS134/LabManual/lab.ncar.html>.

1. Compare Two Codes. You will find that the two radiation codes give very different answers for the temperature sensitivity to CO₂ and water vapor. What is the ΔT_{2x} for each model? Is it the same for doubling from 100 to 200 ppm as it is for 350 to 700 ppm? The model includes the water vapor feedback automatically, but we can turn this off by zeroing the relative humidity. What is the ΔT_{2x} without the water vapor feedback? How comparable are the CCM3 and Chou models?

2. Clouds and Upwelling IR Light. Use the on-line IR radiation model <http://geosci.uchicago.edu/~archer/cgimodels/radiation.html>. What is the effect of clouds on the outgoing IR energy flux of the atmosphere?

a) Are higher clouds or lower clouds the most significant to the outgoing IR energy balance? (Note that this calculation, in fact this entire lab, neglects incoming energy, which clouds can affect by their albedo. We'll deal with that issue in the next lab.)

b) Can you see the effect of the clouds in the outgoing spectra? How is it that clouds change the outgoing IR flux?

3. Clouds and Downwelling IR. Set the Sensor Altitude to 0 km, and choose the Looking Up option, at the bottom of the model web page. Do this with no clouds, and then again with clouds. Explain what you see. Why, at night, is it warmer when there are clouds?

4. Clouds and Full-spectrum Light. Let's look at the effects of clouds. For each radiation code, document the effect of, say, a 100% cloud cover, for high and low clouds (run each separately). Which type of cloud has the largest effect?

b. What is the effect of changing the drop size from 10 to 8 microns in the low clouds? How do these radiative effects compare with the effect of doubling CO₂?

c. Run a set of scenarios with high clouds, using a range of water contents, from 100 down to 2 g/m². Make a plot of the equilibrium temperature depend on the water content. You should get a fairly surprising-looking plot. Can you explain what you see?

Take-Home Points

Positive feedbacks act as amplifiers of variability, while negative feedbacks act as stabilizers.

The water vapor feedback doubles or triples the expected warming due to rising CO₂ concentrations.

The ice albedo feedback amplifies the warming in high latitudes by a factor of three or four.

Clouds have a huge impact on climate. Human activity is already changing the nature of clouds on Earth. Future changes are difficult to predict.

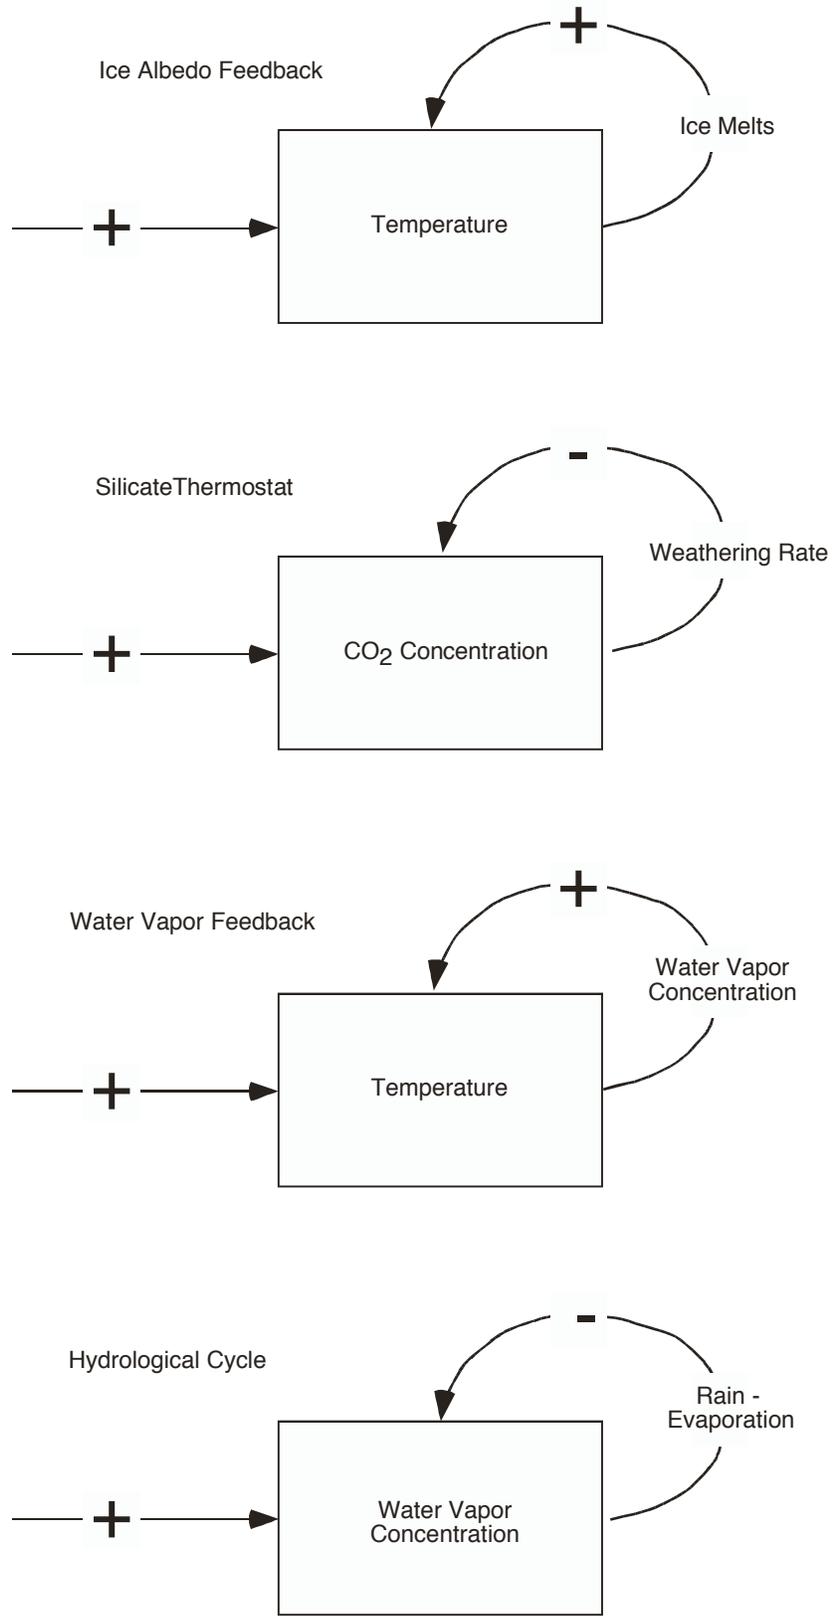


Figure 6-1

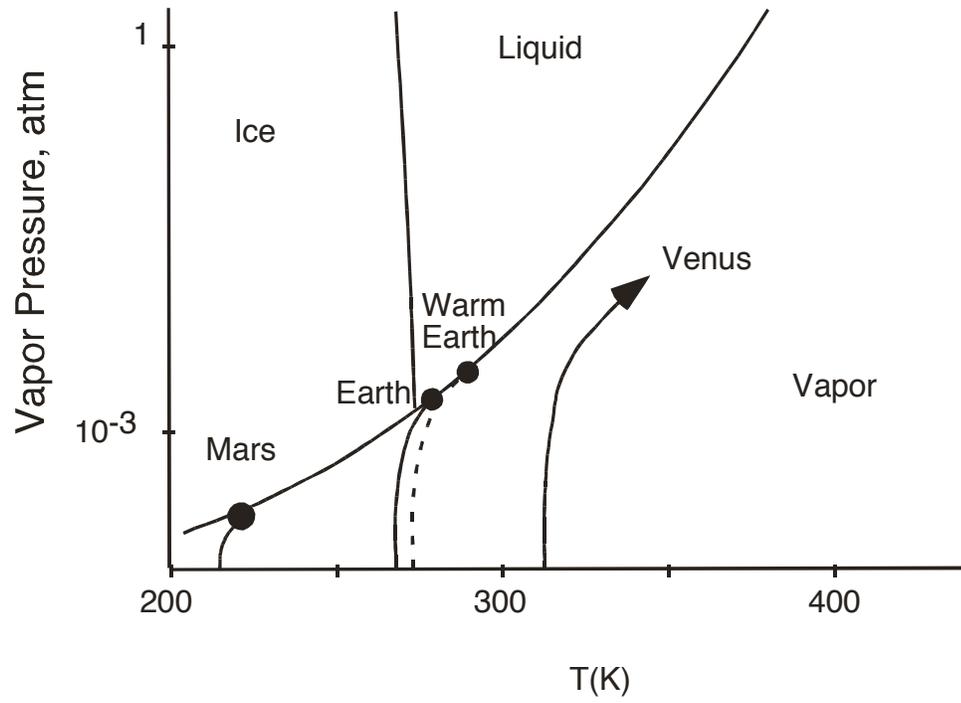


Figure 6-2

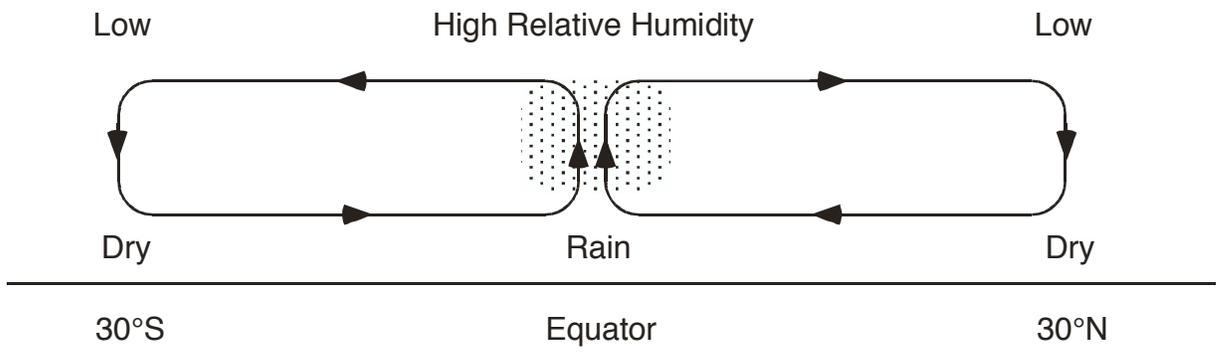


Figure 6-3

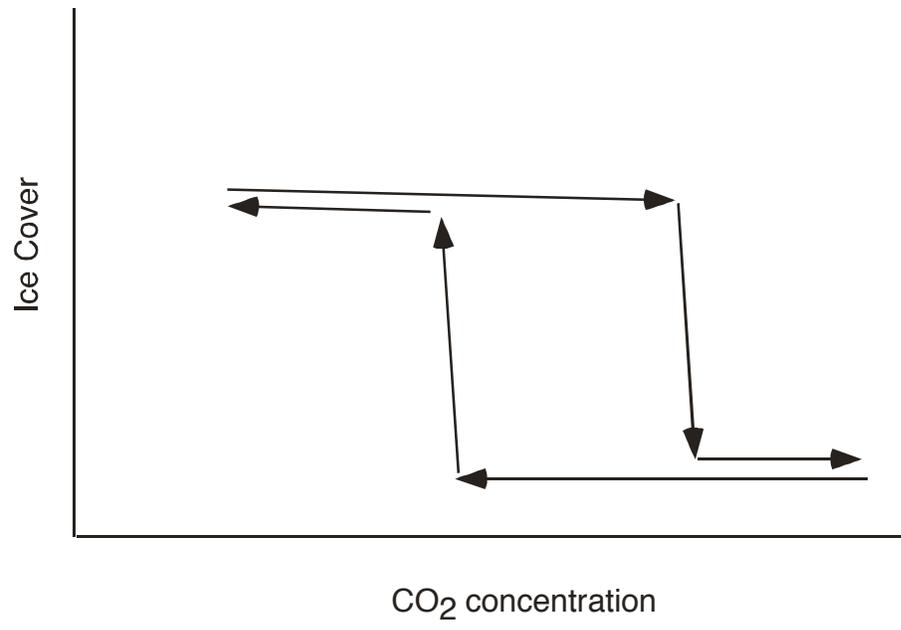


Figure 6-4

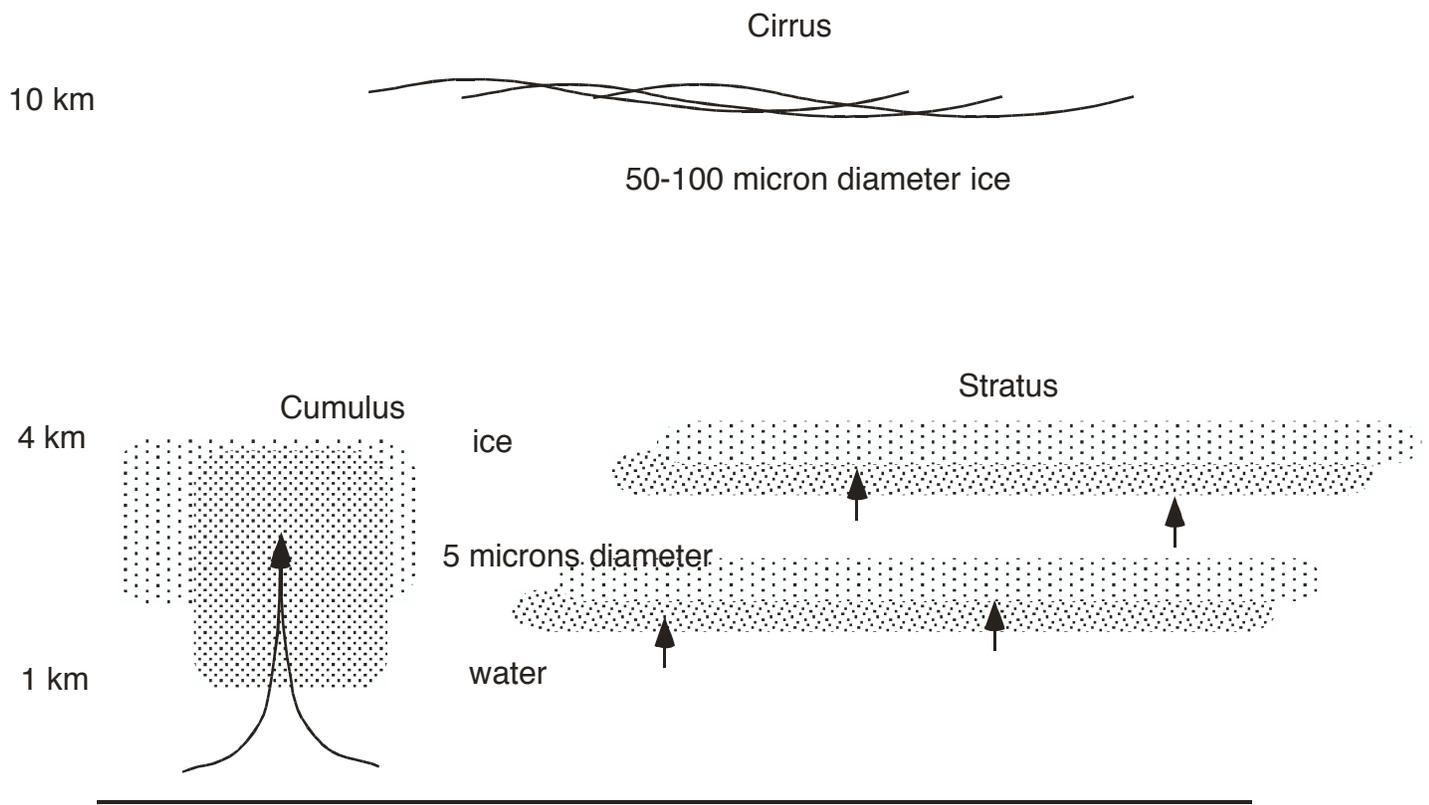


Figure 6-5



Figure 6-6

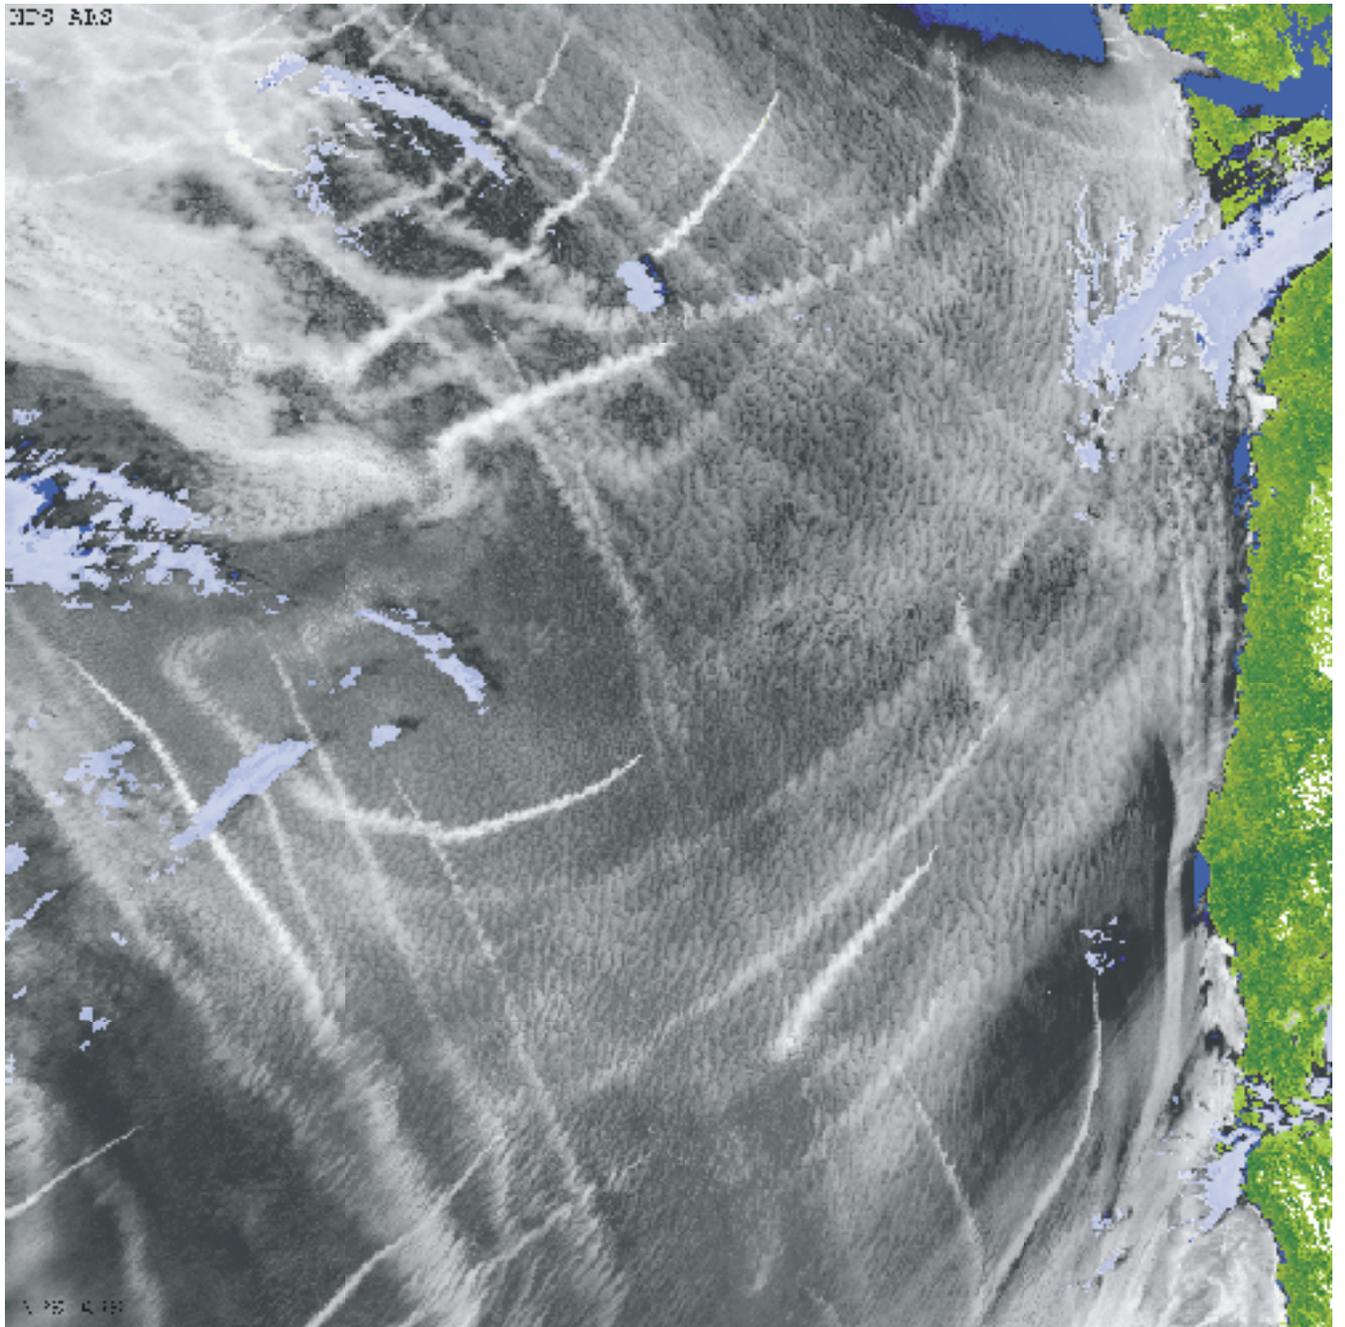


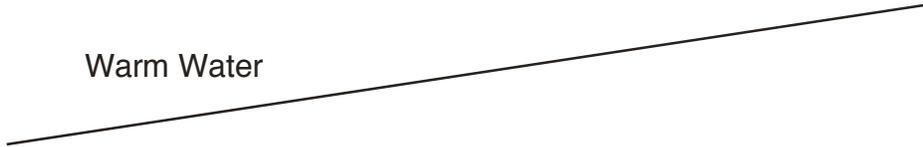
Figure 6-7

el Niño

Strong Wind



Warm Water



Cold Water

la Niña

Weak Wind



Warm Water



Cold Water

Figure 6-8

Part II. The Carbon Cycle

Chapter 7. Carbon on Earth

Abstract

Carbon exists in a range of oxidation states and chemical forms. Carbon is most stable in the oxidized form, on earth mostly as CO_2 and CaCO_3 . Photosynthesis converts carbon to higher-energy reduced form, both to store energy from the sun and because the versatility of reduced carbon chemistry makes a natural building block for life.

The atmosphere contains only a tiny fraction of the carbon on earth. The terrestrial biosphere has several times more carbon, if soil carbon is included in the reckoning. The seasonal breathing of the terrestrial biosphere is measurable in a seasonal cycle of atmospheric CO_2 . The ocean contains 50 times as much carbon as the atmosphere, and apparently is responsible for large changes in atmospheric CO_2 over 100,000-year glacial cycles, although no one is sure quite how. On timescales of millions of years, the weathering of igneous rocks consumes CO_2 , stabilizing atmospheric CO_2 and the climate of the earth in the process.

The Chemistry of Carbon

Carbon has rich, spectacular, wonderous chemistry. There are more scientists studying the chemistry of carbon than any other element, I am sure. Carbon chemistry provides the complexity required for constructing the machinery of life on earth. In addition to serving as the framework and scaffolding for that machinery, carbon chemistry also provides the means of storing energy.

Elements of similar electron configurations are grouped into families which appear as columns in the periodic table. The nearest relative of carbon is silicon directly underneath carbon. Both elements require four additional electrons to complete their filled outer shells of electrons. Silicon is one of the more abundant elements on earth, and the chemistry of the rocks in the mantle is controlled to a large extent by the whims of silicon. Elements that get along well with silicon tend to be retained in the earth's mantle, while elements that do not are the ones that get preferentially squeezed out of the mantle as lava, ultimately generating the earth's crust. Silicon tossed out into the environment is versatile enough to find its way into hundreds of different crystal forms, including the minerals in rocks and clays.

Silicon chemistry is regimented, flat, one could almost say lifeless, compared with the chemistry of carbon. Carbon chemistry is kept highly organized within living things, but after life is finished with it, carbon left in soils, sediments, and rocks forms itself into an

indescribable goo called **humic acids** or **kerogen**. Silicon forms hundreds of crystals, but the number of possible configurations of carbon is essentially infinite. I scoff at the cliché from Star Trek, “carbon-based life forms”. Perhaps I am close-minded, a chauvinist, but I have difficulty imagining life based primarily on any element other than carbon.

There is a fundamental property of elements in molecules called **oxidation state**. You’ve heard of this before: iron oxidizes when it rusts. Three simple carbon molecules demonstrate the spectrum of reduction / oxidation (abbreviated **redox**) chemistry for carbon.

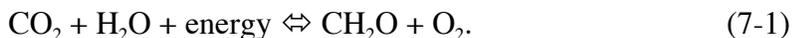
	Oxidized	⇔	Reduced
Simplest Example	CO ₂		CH ₄
Carbon Oxidation State	+4		-4
General Category	Inorganic Carbon		Hydrocarbons
		Carbohydrates	

The first is carbon dioxide, which we learned in Chapter 3 is a gas, a greenhouse gas at that. The last is methane, also a greenhouse gas. CO₂ is the **oxidized** form of carbon in this spectrum, and methane is the **reduced** form. The word oxidation state implies that we must be talking about oxygen in some way. In more general terms, it is a measure of the surplus of deficit of electrons around the carbon. Oxygen is a very greedy element for the two electrons that it requires to find its most stable electronic configuration. Usually oxygen in a molecule is credited with stealing two electrons from its bond partner (an exception is if its bond partner is another oxygen, in which case they must share equally). The carbon in CO₂ has two bonds with oxygen, so it has given up four electrons, two to each oxygen. Since each electron has a charge of -1, a deficit of four electrons leaves the carbon in CO₂ with an oxidation state of +4. Hydrogen, in contrast, is relatively generous with its single electron. The four hydrogens in methane each donate one electron to carbon, so carbon ends up with an oxidation state of -4. The oxidation state of CH₂O, the middle compound in the spectrum, is zero, because carbon gains two electrons from the hydrogens, but donates two to the oxygen.

There are many other carbon compounds that have these oxidation states, and it makes sense to group them together into families in this way. Most of the carbon on Earth is oxidized carbon, also called **inorganic carbon**, in CaCO₃ (limestone) rocks or dissolved in the oceans. Inorganic carbon has a fascinating chemistry (spoken as one who has devoted his life to understanding it), but it is nowhere near complex enough to create the machinery of life. At the reduced end of the spectrum, **hydrocarbons** include oil as well as natural gas. Some biomolecules are hydrocarbons, such as fats. Again, amazing stuff, but life could not be constructed from fats alone. Life is comprised of carbon largely in the intermediate oxidation state, called **carbohydrates**. Sugars for example have formulas that are multiples of that of humble formaldehyde, like glucose for example is (CH₂O)₆. Hydrocarbons and carbohydrates are together referred to as

organic carbon. Here is an case where scientific and common usages of a word diverge strikingly. To a chemist, compounds made of reduced carbon, such as DDT, dioxin, and PCB's, are considered organic compounds. To the common person, "organic" produce at the grocery store is supposed to be free of toxic man-made compounds such as these.

The stablest form of carbon in the chemical environment of the earth is the oxidized form. The carbon pools on Venus and Mars are both nearly entirely oxidized. Life on earth is based on the nifty trick of harvesting energy from sunlight and storing it by creating reduced carbon from oxidized carbon



The forward direction of reaction 7-1 is called **photosynthesis**, and is done by plants. The energy comes from sunlight; photosynthesis cannot proceed without energy from light. Photosynthesis cannot happen in the deep sea for example. The backward direction of the chemical reaction is called **respiration**; we as consumers do that. We consume the products of photosynthesis (indirectly perhaps if we are not vegetarians), and breathe oxygen, to harvest the energy originally from sunlight, exhaling CO₂ and water vapor.

Photosynthesis serves two purposes in the biosphere. First, it creates carbon in the oxidation state at which its chemistry is versatile enough to build the machinery of life. Second, it stores energy from sunlight in the form of reduced carbon. Some biomolecules have a clear use in the living machine, such as proteins which make enzymes which control chemical reactions, DNA which stores patterns for making proteins, or the fat molecules in cell membranes which hold chemicals inside the cells. Others are more obviously for storage of energy, such as glucose in "blood sugar" or the fat deposits underneath our skin.

Nearly all of the organic carbon produced by photosynthesis is respired sooner or later. Peat deposits on land may hold organic carbon for thousands of years, but on geologic time, the only hope an organic carbon molecule has of escaping degradation back to CO₂ is to hide in the sediments of the deep sea. Life on earth has built up a sizable pool of carbon in the reduced form in ocean sediments and **sedimentary rocks** (former ocean sediments that are currently on land). When oxygen-producing photosynthesis first began on Earth, the oxygen released reacted with other reduced species such as dissolved iron in seawater. After the biosphere "rusted" as much as it was going to, oxygen started to build up in the atmosphere. The amount of buried reduced carbon exceeds the amount of O₂ in the atmosphere by about a factor of 10. Oxygen in the atmosphere has built up as a sort of by-product of photosynthesis.

Visionary earth scientist James Lovelock imagines the biota "charging up" the biosphere like a giant battery, by building up a surplus of reactive oxygen and reduced carbon. Certainly we as oxygen-breathing animals benefit from this situation. Bacteria are able to derive energy by reacting organic carbon with other chemicals, for example sulfate respiration producing hydrogen sulfide, rotten-egg smelling stuff, to us one of the stinkiest chemicals in nature. Bacteria prefer O₂ for respiration over sulfate or any other

chemical, because more energy can be gained by using O_2 as one pole of your battery (the food as the other). As for multicellular life, we are unable to respire using anything but O_2 .

The puzzling thing to earth scientists is that the oxygen concentration of the atmosphere seems to have been remarkably stable over the past 600 million years, since the **Cambrian Explosion** when suddenly lots of fossils of multicellular life appear in the fossil record. If atmospheric O_2 had ever dropped to zero, all those multicellular beasts would have perished, leaving multicellular life to start over again. Alternatively, if O_2 had risen to ten times its present abundance, a single spark from lightning would have exploded with the organic carbon on entire continents.

The amount of oxygen in the atmosphere was seemingly determined by the difference between two much larger quantities, the amount of reduced carbon that gets buried, and amount of oxidizable iron and sulfur. This would seem like it would lead to great instability of the oxygen concentration with time if those two quantities do not balance exactly. One is the amount of reduced organic carbon buried in sediments, and the other is the amount of reactive reduced chemicals like reduced iron and sulfur in the ocean. The atmosphere got the leftovers, just about one tenth of the feast. An imbalance in burial of organic carbon could impact atmospheric O_2 in about 2 million years, yet O_2 has remained stable for 600 million years. This eerie stability of the biosphere led Lovelock to pose the **Gaia hypothesis**, that life on earth functions collectively as a single super-organism, able to regulate the earth's environment to its own advantage. Evolutionary biologists have difficulty with the Gaia hypothesis because evolution is essentially a process of trial and error selection of selfish organisms. If the entire biosphere is essentially a single organism, it is difficult to see any room for error. Also, Gaia requires altruism. If Gaia were to call for some bacteria to spend energy producing methane, for example, because Gaia required it in order to balance some geochemical reaction, it is difficult to see why that organism would comply, spending its own metabolic energy for the common good. Lovelock does have a point, though, that the biosphere is uncannily stable over geologic time.

The Atmosphere and Other Carbon Reservoirs

The CO_2 in the atmosphere is only the tiniest fraction of the carbon on earth (**Figure 7-1**). The stable form of carbon in contact with hot rocks, such as in interior of the earth, is as CO_2 , and most of our carbon emerged from the Earth in this form. Most of the carbon degassed long ago, but a trickle still comes out of volcanoes and hot springs at the spreading centers in the deep ocean where ocean plate is formed. At the cooler temperatures of the surface of the earth, CO_2 reacts with the CaO component of igneous rocks to form $CaCO_3$, the stuff that coral reefs are made of. Most of our carbon is now tied up in $CaCO_3$ deposits from ancient sediments.

The largest reduced carbon reservoir is called **kerogen**, which makes up less than 1% by weight of sedimentary rocks. Living carbon on land is called the **terrestrial biosphere**. There is about as much carbon in the terrestrial biosphere as there is in the atmosphere. The **soil carbon pool** in soils is larger by a factor of about three. The soil

carbon pool is largely dead carbon. In the ocean, most of the carbon is not only dead, it is oxidized, called **dissolved inorganic carbon**. This pool is larger than the atmospheric pool by a factor of about 50.

All of these giant reservoirs of carbon interacting with our tiny amount of carbon in the atmosphere mean that there is plenty of room to alter the CO₂ concentration of the atmosphere. CO₂ in the atmosphere is pulled every which way, naturally, on all sorts of time scales, from yearly to time scales of millions of years.

The terrestrial biosphere draws carbon from the atmosphere during the growing season, and returns it during winter. One can see the “breathing” of the terrestrial biosphere by measuring the CO₂ concentration of the atmosphere over the seasonal cycle (Figure 7-2). The seasonal cycle of atmospheric CO₂ in the Southern hemisphere is out of phase with that in the North, because the seasons are reversed in the South. There is much less land in the South, so the seasonal cycle is much smaller than it is in the North.

We will explore the chemistry of seawater in greater detail in Chapter 9 when we think about fossil fuel invading the oceans. For the time being, suffice it to say that CO₂ is released from the ocean in some places, and absorbed by the ocean in other places. These carbon fluxes are huge compared with our release of fossil fuel CO₂, fast enough to replace all of the CO₂ in the atmosphere every five years or so. The CO₂ concentration of the atmosphere drifts toward the concentration at which these exchange fluxes balance each other. Computer model experiments show that it takes hundreds of years for the atmospheric CO₂ concentration to approach the equilibrium value where all the exchange fluxes with the ocean balance each other. The reason it takes so long is that the ocean takes about this long to circulate, for all the waters of the ocean to come to the sea surface somewhere to exchange carbon with the atmosphere.

The ocean tugs on the CO₂ concentration of the atmosphere as part of the glacial / interglacial climate cycles that have been deduced from climate records such as ice cores and deep sea sediments. On time scales of about 100,000 years, ice sheets grow and dwindle, and the earth cools and warms (Figure 7-3). We are currently in an interglacial climate interval. Bubbles of ancient atmosphere trapped in the ice sheet of Antarctica reveal large changes in the CO₂ concentration of the atmosphere that correlate with the ice sheet advances and retreats. The CO₂ concentration during glacial intervals is 180-200 ppm, rising to 260-280 ppm during interglacial intervals before the industrial era. The CO₂ concentration today is about 375 ppm and rising at about 1.5 ppm per year.

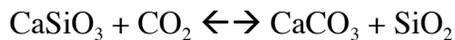
Ice ages are correlated with variations in the earth’s orbit, in particular with the amount of sunlight received in the Northern hemisphere at about 65° N in mid-summer. The idea is that 65° N is where the ice sheets are, in the North because there is land to put ice sheets on in the North, and summer is when the fate of winter snow is decided, to melt back or to remain and grow into an ice sheet. No one is sure how the atmospheric CO₂ manages this trick of going up and down in synchrony with the ice sheets, but the ocean must be the major player. The land carbon reservoirs are not large enough, and are thought to have changed in the wrong direction, putting carbon into the atmosphere as the ice sheets grow rather than taking it out as would be required to bring atmospheric CO₂

down. The carbon in rocks, sedimentary CaCO_3 or kerogen for example, cycles too slowly in the environment to explain the large fast changes. The ocean is the only pool that is large enough and potentially reactive enough to explain the data. Regardless of how the ocean accomplished this trick, the observed changes in CO_2 concentration, and the model-forecast sensitivity of temperature to CO_2 , are required if we are to explain the colder temperatures we infer for the glacial climate.

The Silicate Thermostat

On time scales of millions of years and longer, the release of CO_2 to the environment from the earth, called **volcanic degassing**, begins to be important. Degassing is a tiny flux compared with the back-and-forth exchange fluxes between the ocean and the atmosphere or the breathing of the terrestrial biosphere. The difference is that the carbon coming from the earth, in volcano gases or in the waters of the hot springs at the bottom of the ocean, is new carbon to the environment, rather than internally cycled carbon through the environment.

This new carbon release is balanced, on long enough time scales, by dissolution of rocks on land, chemical reactions that consume CO_2 , called **weathering reactions**. A chemical reaction for weathering that a geologist might recognize would be



where CaSiO_3 is something like the formula for a rock formed at high temperature, called an **igneous** rock. Igneous rocks are formed by cooling and freezing melted rock. Melted rock is called **lava** if it is found at Earth's surface and **magma** if it is subsurface. Real igneous rocks have other elements in them and a wide range of chemical formulas, but simple CaSiO_3 works for illustration. CaCO_3 and SiO_2 , the solid phases on the right-hand side of the reaction are typical sedimentary rocks, from at cold temperatures from elements that were dissolved in water. When an igneous rock weathers, it ends up as dissolved ions in river water, ultimately in sea water. Organisms like corals and shell-forming plankton extract the dissolved ions from seawater and construct solid CaCO_3 and SiO_2 from them. When sedimentary rocks are heated, they release CO_2 . Another way to say that is that high temperatures favor the left-hand side of the reaction. When igneous rocks are left to sit at cold temperature, they in effect absorb CO_2 to become sedimentary rocks. Cold favors the right-hand side of the reaction.

The involvement of CO_2 in this reaction is the issue that concerns us here. When igneous rock weathers, it drags CO_2 ultimately from the atmosphere down into deep sea sediments.

If we wait long enough, say a million years, the rate of CO_2 uptake by weathering has to balance the outgassing of CO_2 from the earth. What comes in eventually has to find its way back out.

The tie to that links this chemical reaction to climate is that the rate of weathering depends on the availability of fresh water to dissolve rocks into. The rate of fresh water

delivery to wash over the rocks depends, in turn, upon the climate of the earth. If the earth gets too cold, water freezes and weathering reactions stop. In this case, the rate of CO₂ uptake by weathering will be slower than the rate of CO₂ degassing from the earth by volcanoes. CO₂ will accumulate in the atmosphere and in the ocean. Eventually enough CO₂ will be in the atmosphere to melt the ice, and weathering will start up again. If the climate were too warm, then weathering might consume CO₂ faster than volcanoes supply it, in which case CO₂ will drop and climate will cool. This climate-stabilizing mechanism is called the **silicate weathering thermostat**, stabilizes the climate of the earth, if we are willing to wait millions of years. This theory helps explain the apparently Gaian stability of the climate of the Earth over geologic time.

Other factors besides CO₂ may affect the rate of silicate rock weathering, including uplift of the rocks by mountain building versus covering the rocks with thick soils, or invading the land by plants which accelerate the rate of CO₂ uptake by pumping CO₂ into soils. If some process encourages the weathering reaction to go faster, say land plants pumping CO₂ into soils, then the CO₂ concentration in the atmosphere falls, and as it does, weathering slows because less water pours over the rocks. After a million years or so, CO₂ in the atmosphere finds the concentration at which the CO₂ fluxes balance. The climate of the earth 100 million years ago was thought to be warmer than today, because CO₂ was driven to higher atmospheric concentration, by mechanisms like this. CO₂ levels may have been ten times what they are at present, and the Earth was tropical up to the poles. As for the glacial cycles, paleoclimatologists require changes in atmospheric CO₂ levels, and the sensitivity of climate to CO₂, if we are to explain the hothouse climates of the age of dinosaurs.

The terrestrial planets Venus, Earth, and Mars provide a perfect Goldilocks parable to end this section. The sin of Venus (the planet) was to be too close to the sun. The water that must have originally been packed with Venus evaporated and its hydrogen lost to space forever. With no water, silicate weathering reactions do not go. With weathering out of the picture, degassing wins. Venus' carbon allotment ended up as CO₂ in the atmosphere, 70 atmospheres of CO₂. At that pressure, CO₂ is no longer strictly speaking a gas; the atmosphere of Venus is actually more like an ocean. And of course Venus is very hot. Mars messed up by being small, about half the diameter of the earth, so that its interior cooled faster than Earth's. Mars boasts the largest volcano in the solar system, Olympus Mons, but today Mars is geologically dead. With degassing out of the picture, weathering wins, and the carbon allocated to Mars has ended up as CaCO₃ rocks. And Mars is cold.

Take Home Points

The most stable form on carbon on earth is oxidized, as CO₂ or CaCO₃.

Photosynthesis stores energy from the sun by producing organic carbon, which also serves as a building scaffolding for the machinery of life.

There is less carbon in the atmosphere than there is in other carbon reservoirs on earth, such as the terrestrial biosphere and the oceans. These other reservoirs tug

on atmospheric CO₂, seasonally for the land, and on glacial interglacial 100,000-year time scales from the ocean.

The weathering of igneous rocks on land controls the pCO₂ of the atmosphere on million-year time scales. This "silicate thermostat" stabilizes climate. The thermostat is broken on Venus, because there is no water left, and on Mars, because there is no active volcanism left.

Lab

Answer these questions using an on-line geologic carbon cycle model. The model predicts the influence of geologic processes on the CO₂ concentration of the atmosphere. To begin with, the atmosphere and the ocean exchange CO₂. It takes about 300 years for CO₂ to find equilibrium between the ocean and the atmosphere, if the chemistry of the ocean were held constant. Chemical reactions with rocks on land and on the sea floor further alter the carbon chemistry of the atmosphere and ocean on much longer time scales. One set of chemical reactions involves CaCO₃, which must be removed from the ocean by burial ("Burial of Carbonates") as quickly as it is supplied to the ocean by dissolution of rocks on land ("Total Weathering"). It takes roughly 5-10 kyr for the ocean to balance its CaCO₃ budget, to come to Carbonate Equilibrium. The Carbonate equilibrium adjustment changes the pH of the ocean and the pCO₂ of the atmosphere. The other geochemical cycle affecting CO₂ involved igneous rocks. The Silicate Thermostat acts to balance the weathering of silicate rocks against degassing of CO₂ from the earth, via volcanoes and hot springs.

The model does an initial spinup phase, to make sure that the initial condition is an equilibrium where all the various budgets balance. You will only see that last few years of this phase in the model output. Then there is a transition to a new set of conditions, say an increase in CO₂ degassing flux from volcanoes. We can watch the transition from one equilibrium state to another.

At the moment of the transition from spinup phase to transient phase, we can also inject a slug of new CO₂ into the system, say by fossil carbon combustion and release to the atmosphere. The high CO₂ concentration at the top of the spike depends a lot on how quickly the CO₂ is released, which is unrealistically quickly in this model, so don't take the spike magnitudes to the bank. Pay more attention to the what happens to CO₂ in the long term after it is released.

1. Weathering as a Function of CO₂. The rate of weathering must balance the rate of CO₂ degassing. Run a simulation where the CO₂ degassing rate increases or decreases at the transition time. Turn off the CO₂ spike (set it to zero), to make things simpler. An increase in CO₂ degassing drives atmospheric CO₂ up or down? Repeat this run with a range of final degassing rates, and make a table of the CO₂ concentration as a function of the CO₂ degassing rate. The CO₂ degassing rate is supposed to balance the CO₂ consumption rate by weathering -- verify that the model achieves this. If so, make a plot of weathering as a function of atmospheric pCO₂.

2. Effect of Solar Intensity. The rate of weathering is a function of CO_2 and sunlight, a positive function of both variables. By this I mean that an increase in CO_2 will drive an increase in weathering, as will an increase in sunlight. The sun used to be less intense than it is now. Turn back the clock 100 million years or 500 million years, to dial down the sun. Weathering has to balance CO_2 degassing, so if sunlight goes down, CO_2 must go up, to balance. Try it. What do you get for the initial steady-state CO_2 , relative to what you get for today's equilibrium value?

3. Plants. Plants pump CO_2 down into the soil gas, possibly accelerating weathering. They also stabilize soils, perhaps decreasing weathering. Run a simulation with a transition from no plants to a world with plants, with no carbon spike on the transition. Figure out if plants in the model, overall, increase or decrease weathering.

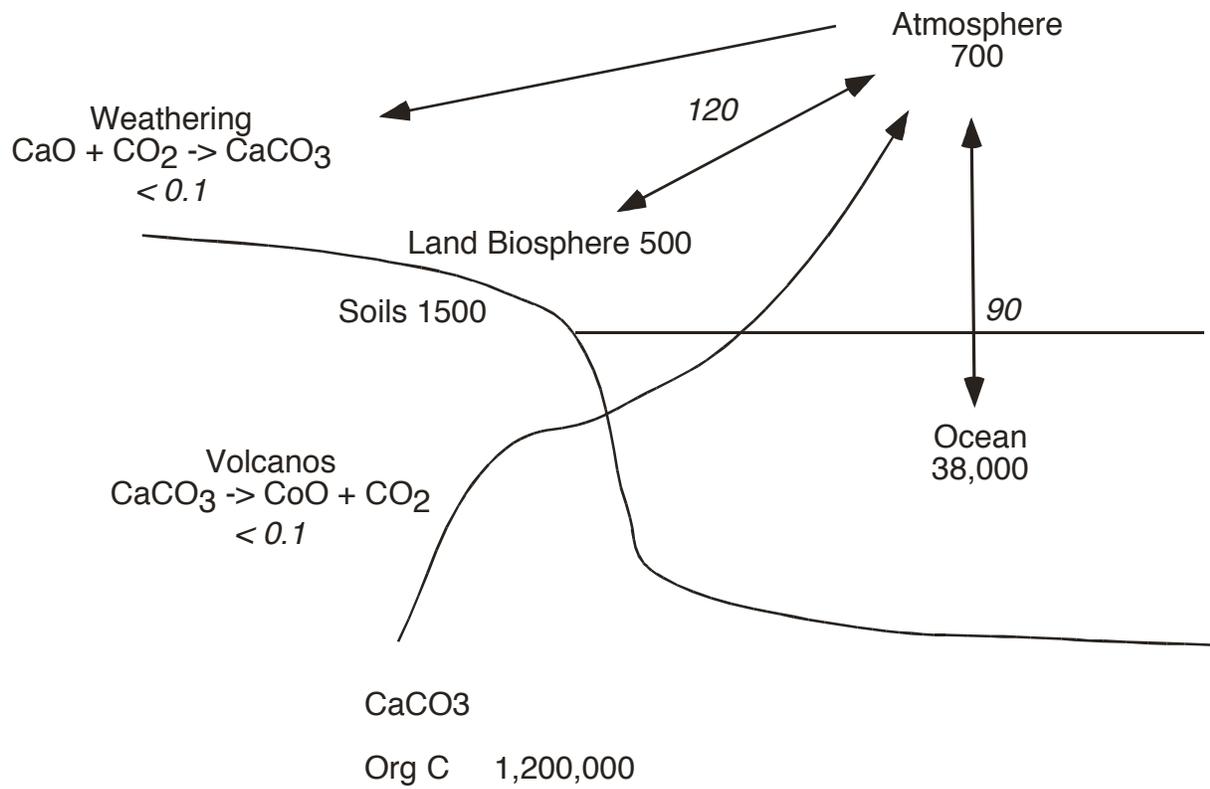


Figure 7-1

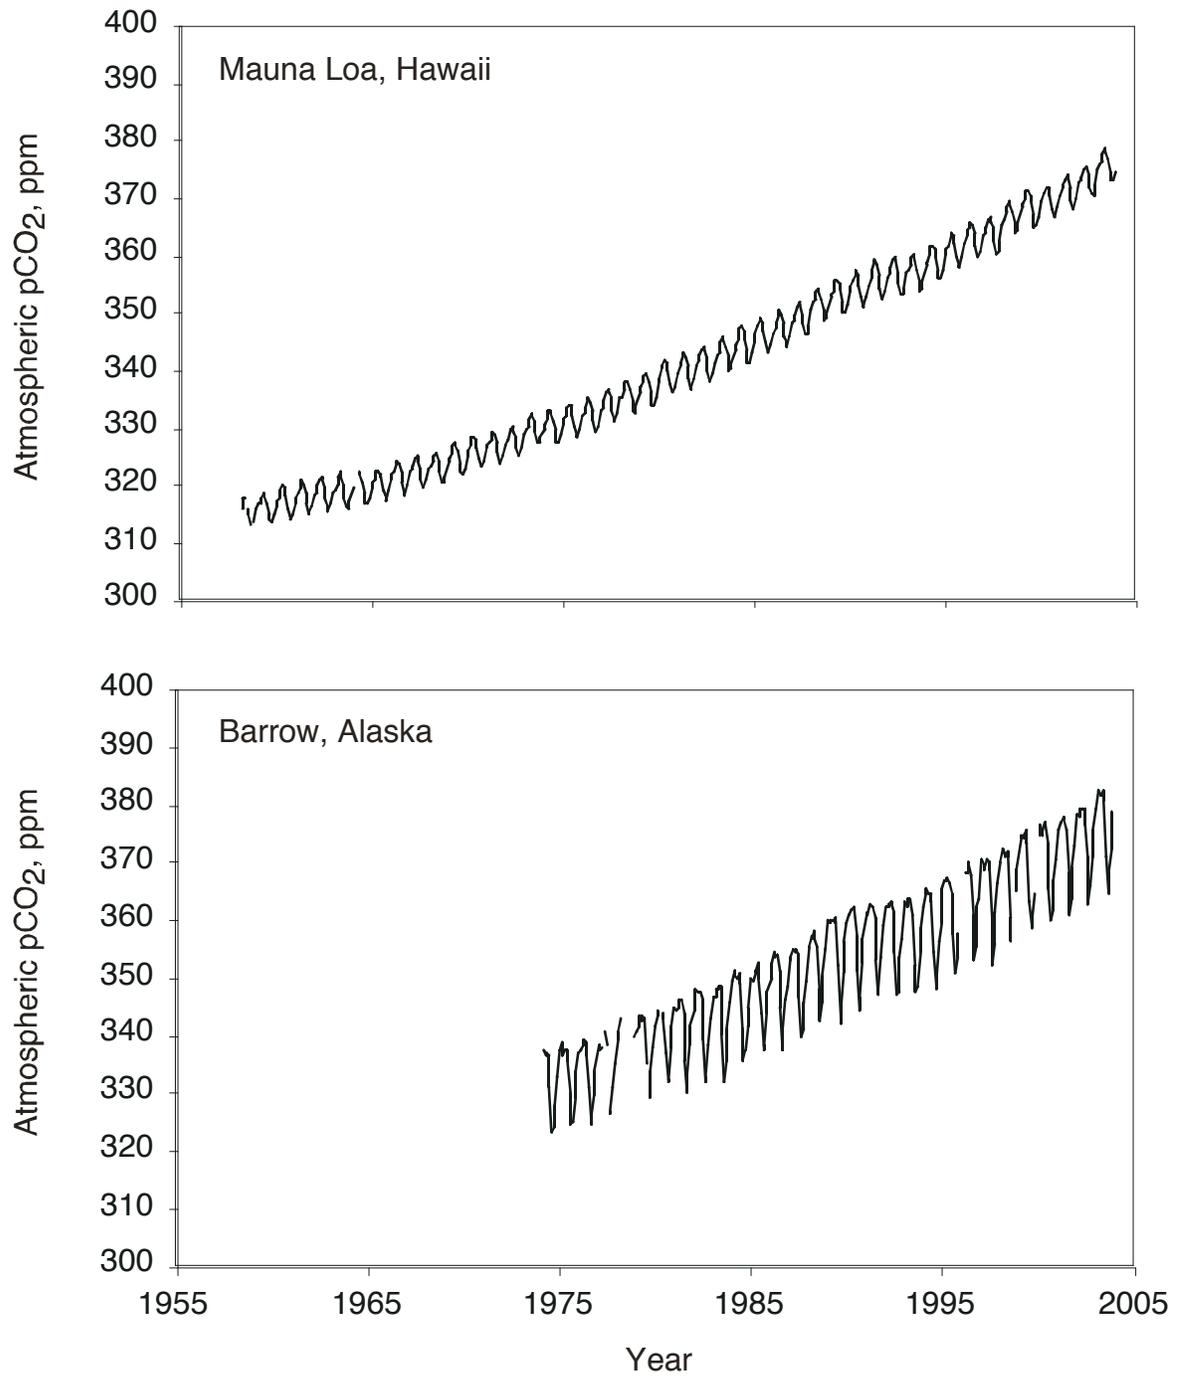


Figure 7-2

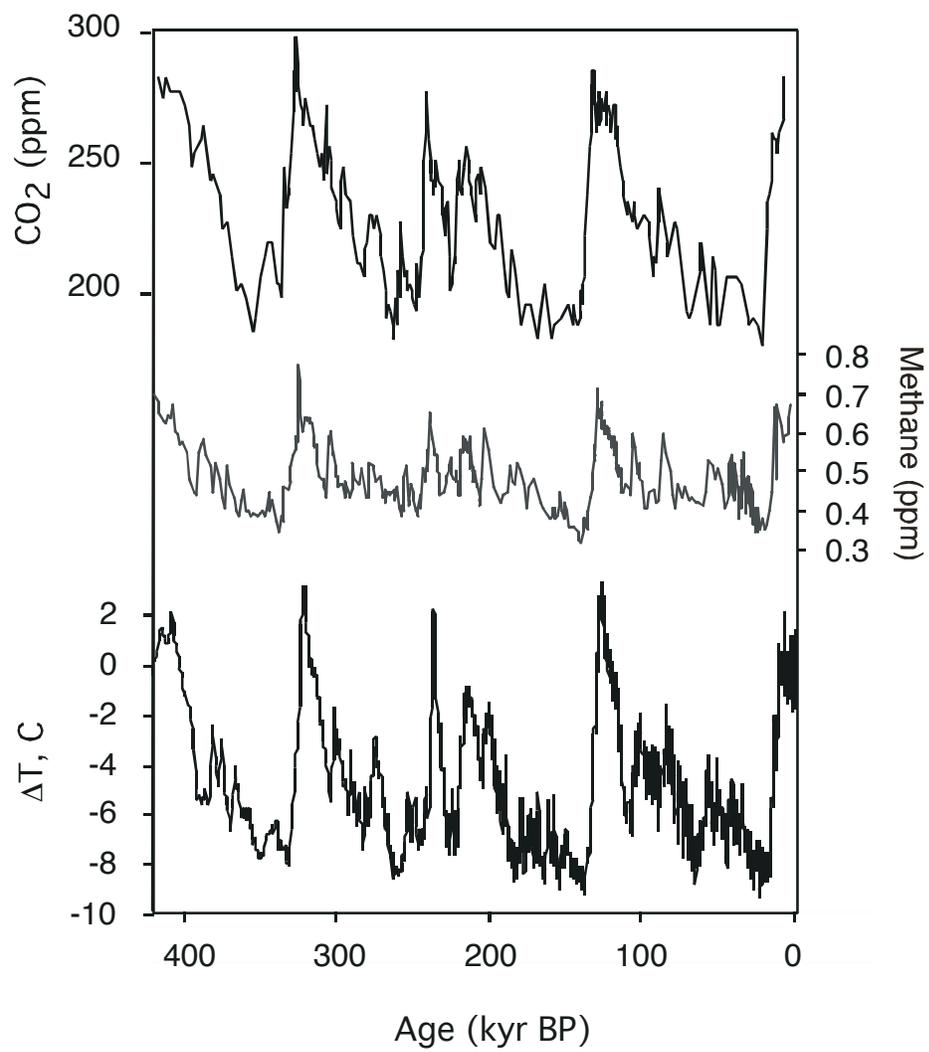


Figure 7-3

Chapter 8. Fossil Fuels and Energy

Abstract

The sun bathes the earth in an immense amount of energy, but most of humankind's energy sources derive from stored energy sources such as fossil fuels or radioactive elements energy rather than instantaneous solar energy. Of the fossil fuels, coal is the most abundant, while oil is more limited and may be depleted in the coming decades. The amount of natural gas as traditionally extracted is comparable to that of oil, but there is an immense amount of gas frozen in ocean sediments. If we wish to stabilize the CO₂ concentration of the atmosphere in the coming century, we need some major new carbon-free source of energy, of a size comparable to our total energy production today or larger.

Energy Sources

Our civilization is based on energy. The sun bathes the Earth in energy, continuously at a rate of about 173,000 terawatts. Photosynthesis captures about 100 TW per year. Mankind is consuming about 13 TW of energy per year. It has been calculated that mankind has captured a sizable fraction, perhaps half, of the photosynthesis, either by agriculture or by preventing photosynthesis by building roads and rooftops and so on. Windmills, hydroelectric dams, and solar cells all derive their energy relatively directly from sunlight. But this is not enough; we are also taking advantage of other forms of stored energy, from fossil fuels and radioactive elements. Fossil fuels derive their energy from past sunlight, photosynthesis that took place millions of years ago.

Over Earth's history, photosynthesis has preserved a sizable amount of reduced carbon, which could in principal be reacted with O₂ to regenerate CO₂ and energy. However, the vast majority of this reduced carbon is so diluted by rocks and clay minerals that it is unusable for fossil fuels. Only in special conditions will the organic carbon be concentrated into a form that we can use for energy. The main forms of fossil carbon suitable for converting to energy are coal, oil, and natural gas.

We also harvest nuclear energy from radioactive elements that were created in the explosions of stars before the formation of the earth and the sun. Geothermal energy, taking energy from the heat inside the earth, also derives much of its horsepower ultimately from radioactive decay.

Globally, most of the energy we use comes from fossil fuels petroleum, gas, and coal (Figure 8-1). These all emit carbon. Methane is the most reduced form of carbon, so as an atom of carbon from methane undergoes the transition to oxidized CO₂, it gives off more energy than would an atom of carbon in coal, which has an oxidation state at the level of carbohydrates. This means that a Watt of energy from methane releases less carbon than a Watt from coal or wood (Figure 8-2).

Non-carbon emitting sources, mostly nuclear and hydroelectric, make up about an eighth of our energy supply. We use renewable sources of energy, defined by the U.S. Department of Energy to mean geothermal, solar, wind, and wood and waste electric power, for about 0.8% of our energy.

The largest traditional fossil fuel reservoir is **Coal**. Coal is solid, almost pure carbon for mature coal. Coal originates from land plants depositing in swamps. The word for freshly deposited plant material of this type is **peat**; peat deposits can be thousands of years old. The British Petroleum [BritishPetroleum, 2004 #2645] estimates of coal reserves are about 1000 Gton C, and [Rogner, 1997 #2636] estimates that the total amount ultimately available might be 10 times that. No matter what assumptions are made about future energy needs, there is enough coal to supply our needs for several centuries, if we choose to use it. In the United States, coal is mostly used for generating electricity (**Figure 8-3**).

Oil and natural gas come from ancient photosynthesis that took place in the ocean. Plants in water are generically called **algae**. There are microscopic algae in the ocean called **phytoplankton**. The phytoplankton in the ocean produce about as much organic carbon by photosynthesis as the terrestrial biosphere on land does. A tiny fraction of the dead phytoplankton end up in sediments, while the rest get eaten by somebody, animal or bacterial, in the water column or in the top layer of sediments. The sediments covering most of the sea floor do not contain enough organic carbon to ever make oil. You find organic-rich sediments nearby continents and under waters which have no oxygen dissolved in them, which are as we say **anaerobic** like the Black Sea today.

If organic-rich sediments are buried to a depth in the earth between 7 and 15 kilometers, they will be heated to temperatures of 500-700 K. This has the effect of converting some of the dead plankton into **oil**, consisting of chains of carbons with hydrogens attached (**Figure 8-4**). The oil may just sit there in little droplets, unreachable by oil companies, or it may flow through the rock if it is porous enough, perhaps to collect in some upside-down pool on the bottom of an upside-down bowl in some impermeable layer of rocks above it. Probably only a tiny fraction of the oil every produced is in a harvestable form today, that is to say mobile enough to have flowed together but not mobile enough to get squeezed all the way out of the earth. Oil is perhaps the most convenient of the fossil fuels, because it is easily stored and transported in its liquid form. In the United States, oil is mostly used for transportation (**Figure 8-3**).

Perhaps because such special conditions are required to produce harvestable oil, the distribution of oil deposits on earth is very spotty (**Figure 8-5**). Most of the oil is in the Middle East, although there is a smattering scattered throughout the rest of the world. We don't know precisely how much oil there is available to harvest. Economic interests such as British Petroleum, who published the Figure I'm showing you as **Figure 8-5**, have inventoried the proven reserves of oil, stuff that has been discovered and documented already. They say 1150 thousand million barrels of oil, which we can translate into 135 Gton C. The amount of oil that is likely to be ultimately extractable is almost certainly higher than this [Deffeyes, 2001 #2646]. Technological advances in oil extraction can

double the amount of oil you can get from an oil field. More oil fields will certainly be discovered, although there are limits to how much more oil we can expect to discover. Half of the world's oil reserves today are in the 100 biggest oil fields, and most of the largest ones were discovered decades ago.

There is a class of carbon deposits for which the geological conditions were never quite right to produce oil, but from which oil can be extracted if the rocks are mined and cooked. These nontraditional oil sources are called **oil shales** and **oil sands**. The German government during WWII tried desperately to extract oil from coal but never succeeded. The Canadians are mining and processing oil shales now. If future discoveries, technology progress, and all these effects are counted together, the estimates are all over the map, but some of them begin to reach 500 Gton C or so.

How long will the oil last? This is a question that is starting to be discussed in the newspapers. One way to calculate a number with units of time would be

$$Lifetime[yr] = reserves\ size[kg] \cdot \frac{1}{production\ rate[kg]} \left[\frac{yr}{kg} \right]$$

British Petroleum calls these Reserves to Production ratio. The funny thing about this ratio is that its value hasn't changed in 15 years (Figure 8-6). We have 40 years left, and it looks like we always had 40 years left, and we could have 40 years left forever. The clock is ticking very slowly. Actually, what is happening is that oil is being discovered about as quickly as it is being used. The R/P ratio is calculating how many years until the last drop of oil is extracted, assuming that the rate of oil extraction remains constant with time and that no new oil is discovered. The rate of consumption is growing with time, but we have been discovering new oil quickly enough to keep pace with that growing demand.

Another way to estimate the lifetime of the age of oil was developed by geologist M. King Hubbert [Deffeyes, 2001 #2646]. He pointed out that the rate of extraction of a limited natural resource such as oil tends to follow a bell-shaped curve. This is an example of an **empirical** observation; there is no theoretical underpinning to the curve, but it sure does fit nice. The rate of extraction has a spin-up time at the beginning, gets going faster and faster, then slows down as the resource starts becoming depleted. Resource extraction histories are not required to follow this curve. One could imagine an extract-the-juice-from-a-popsicle-on-a-hot-day curve. Starts off slowly at the beginning because the popsicle is too cold to eat, then at the end you have to snarf the last half of the popsicle in one gulp to keep it from hitting the sidewalk (Figure 8-7). The crucial feature of the bell curve is that the maximum rate of extraction, the peak, occurs when half of the resource is still in the ground. The peak of the popsicle-eating curve came at the very end. It could be that global oil will follow a popsicle curve instead of a bell curve, but a bell curve seems like it might be the way to bet. There have been lots of bell curves in the past, for individual oil and coal fields. Figure 8-8 shows a Hubbert's peak for sperm oil harvesting back in the whaling days. Hubbert's theory gains street-cred by the fact that Hubbert used it, in 1956, to forecast the peak extraction of oil in the United

States. He predicted the peak to be sometime between 1965 and 1972, and in fact it came in 1970. The solid line in [Figure 8-9](#) shows the time up to Hubbert's forecast, and the dashed line a prediction like his for the future (benefiting a bit from hindsight, I confess). Hubbert assumed that the future trajectory would follow a bell curve, and he took a guess at the eventual total amount of oil that would be extracted. If you know those two things, a curve like the dashed line just falls out, there's no more wiggle room. That was an incredible shot. The man must have been an amazing miniature golfer.

The forecast for world petroleum depends entirely on what you call oil and how much is out ultimately out there. [Figure 8-10](#) shows world oil production in the past, fit to two Gaussian curves, one for 200 Gton C in oil that will ultimately be extractable, and one for 500 Gton C. We have already extracted 112 Gton C. If 500 Gton C is ultimately right, then the peak comes in in the 2030's sometime. If 200 Gton C is right, consistent with BP, the peak should be now. The significance of the peak for the world economy is that demand and consumption of petroleum is growing exponentially, along with other social factors like population and GDP. If the constant or dwindling oil supply is unable to keep up with the exponentially rising demand, that's what they call shortage. The bottom line is that the time to watch out for is when the oil is half gone; things won't be fine until the last drop.

If the organic carbon in potential oil source rocks becomes buried too deeply in the earth, it will be heated up so much, some of the reduced carbon will convert to **natural gas**, predominantly methane ([Figure 8-11](#)). Methane can also be produced from organic carbon by bacteria. Methane sometimes comes out of oil wells, or it can be extracted from coal beds. Methane is more difficult to transport than liquid petroleum, because it is a gas and must be held under pressure. For this reason, methane associated with oil in remote locations is sometimes just burned, giant dramatic flares into the sky, rather than collected. The ultimate extractable amount of methane is probably even more poorly known than it is for petroleum, because it hasn't been explored for as intensively. Industry estimates are about 100 Gton C as methane in proven reserves. Just like for oil, the ultimately available gas reservoir will certainly be higher than that.

There is also a 363 kg gorilla in the picture, hiding in near-shore ocean sediments. **Methane hydrates** are methane gas molecules frozen into soccer-ball cages of water ice. Water can form these cages around any gas; all the gas is required to do is hold the little soccer-balls open. CO₂ forms hydrates on Mars. It turns out that there are lots and lots of methane hydrate deposits on Earth, in mid-depth sediments of the ocean ([Figure 8-12](#)). I myself have published with Bruce Buffett [Buffett, 2004 #2588] an estimate of 5000 Gton C as methane in clathrates globally, and other estimates are even higher. It difficult to know exactly because the deposits are only a few percent of the solid, they melt on their way back up to the ship. Also there have not been very many sediment cores drilled in these sorts of sediments, because they are geologically unstable.

It takes millions of years to build up the clathrate reservoir. Dead plankton hit the sea floor and mostly degrade in the top 10 cm or so of the sediment. A small amount escapes. Over millions of years, the sediment accumulates, until our dead plankton

reaches a depth of hundreds of meters below the sea floor. At that depth (why that depth? Who knows?) bacteria begin producing methane from the organic matter. The methane can diffuse around, it can react, or if there's enough of it, and it's cold enough, it can freeze into clathrates. The temperature around a given chunk of clathrate increases slowly with time, as more and more sediment accumulates overhead. Eventually, the chunk of clathrate will melt, leaving a layer of bubbles beneath the chunky clathrate layer. This layer of bubbles can be seen in sediments around the world in the reflection of seismic waves, echos of explosions set behind ships. The reflective layer of bubbles tends to parallel the sediment surface, and so is called a **bottom-simulating reflector**.

Methane clathrate deposits seem like the most unlikely and precarious of things. The ices themselves would float in seawater if they were not held down by the sediment above. The bubbles make the deposits even more unstable. They decrease the average density of the sediment column. If the mixture of solid, seawater, and bubbles (called a **slurry**) starts to rise up, the gas bubbles would expand causing the slurry to rise even faster. Some blame the Bermuda triangle on giant ship-swallowing methane burps from clathrate deposits, as bubbles decrease the average density of the fluid the ship was floating in, to something less than the average density of the ship.

Current Energy Consumption

Returning to **Figure 8-1**, we see how different countries get their energy. India and China, the largest of the developing world, use a lot of coal, while Brazil uses mostly petroleum and hydroelectric power. Denmark is remarkable for windmills, which fall into the category "renewable" on this plot. France has invested heavily in nuclear energy. The United States energy sources are reflective of the global average. In part this is because the United States is a big customer for energy, accounting for about a quarter of energy use globally (**Figure 8-13**). If we divide the energy use of a country by the number of people, we get the energy use **per capita**. We see that Americans use five times more energy than the global average citizen of the world, ten times more than the average Chinese or Indian, and twice as much even as the average European or Japanese. Energy use is closely related to economic productivity, which is measured by the total value of everything produced for sale in a country in a year, the gross domestic productivity or **GDP** in units of dollars / year. We can divide energy use by GCP to derive the **energy intensity**, a measure of the energy efficiency of an economy, reflecting waste but also the difference between heavy industry such as steel production versus high-tech industry. We see in **Figure 8-13** that the Europeans and the Japanese can make a buck out of less energy than we can, while the Chinese and Indians are more energy intensive (less efficient).

Carbon emissions from these countries look similar to the energy story (**Figure 8-14**). The United States is responsible for an about a quarter of global CO₂ emissions. Per person, the United State releases five times more than the global average, and more than twice as much as the Europeans or the Japanese, ten times more than the developing world. Per GDP, the U.S. is less than half as efficient as Europe or Japan.

Future Energy Consumption

We have begun to see the factors that control CO₂ emission: population, GDP, energy use per dollar GDP, and carbon release per Watt of energy. CO₂ emission can be estimated by multiplying these factors together; check it out for yourself that the units cancel correctly. Constructing CO₂ emission from these pieces is called the Kaya identity. You can play with the Kaya identity and make your own forecast on-line at <http://geosci.uchicago.edu/~archer/cgimodels/kaya.html>, and see results in **Figure 8-15**. Population has been rising at a rate of about 1.3% per year, but this rate of growth is not expected to continue. Rates of population growth tend to decrease with affluence of societies. Population is forecast to level off at some value, but that value could be 9 billion people or it could be 15 billion people. A typical “business-as-usual” forecast puts the leveling-off point at about 11 billion people. The second factor, \$GDP / person, has been rising throughout the last century at a rate of about 1.6% per year, rising from \$930 in 1990 U.S. dollars in 1900, to \$4800 per person in 2000. The third factor is the energy intensity, how many watts of energy it takes to make a dollar GDP. The energy intensity reflects efficiency but also the difference between heavy and light industry (**Figure 8-13**, bottom). Energy intensity has been dropping by a rate of about 0.55%/year over the past century, but the last few decades have seen a drop by about 1%/yr. According to the model on the web site, the slower rate of decrease results in 825 ppm pCO₂ in 2100, while 1%/yr results in about 700 ppm (**Figure 8-16**). The final factor, the carbon released per energy yield, is the result of efficiency of the power plants but also on the characteristics of the feedstock energy source. Coal contains intrinsically less energy per carbon than does oil or gas, while nuclear and hydroelectric energy release no carbon at all. The IPCC CO₂ emission scenarios, “what-if” stories analogous to Kaya, ranged from 6 to 26 Gton C emission per year in 2100. We will return to this model in **Chapter 9** to see what happens to that CO₂, and in **Chapter 12** to see what it means in terms of energy.

Take-Home Points

Ultimately, the energy available to mankind includes instantaneous solar energy, which is abundant but spread out; stored solar energy in the form of fossil fuels; and stored energy from stellar explosions in the form of radioactive uranium deposits.

Of the fossil fuels, coal is the most abundant. Oil may run out in the coming decades, and the peak rate of oil extraction may be upon us even now.

We can project energy demand in the future as the product of population, economic growth, and energy efficiency.

Lab

1. Hubbert's Peak. Point your web browser to <http://geosci.uchicago.edu/~archer/cgimodels/hubbert.html>.

a. You will see two different data sets to plot against, along with the three parameters (knobs) that control the shape of the Hubbert curve. Nothing fancy here, we're just matching the curve to the data by eye. First start out with the U.S. oil production. The page comes up with some values for the curve that look pretty good to me, but you should try varying the numbers in the box to see how tightly those values are constrained. In particular, there may be combinations of values that could be changed together to fit the data nearly nearly as well, but with different values. How much wiggle room is there for U.S. oil production.

b. Now switch to global oil production with the pull-down menu. When do you forecast the peak of oil extraction? How much wiggle-room is there this time?

2. The Kaya Identity. Point your web browser to <http://geosci.uchicago.edu/~archer/cgimodels/kaya.html>.

a. Find the plots for GDP per capita, energy intensity, and carbon efficiency, and compare the model hind-cast (the solid line) with the data (plusses). How well constrained are the growth rates by the past data? Of course, the future may not follow the dictates of the past; this is not a mechanistic prediction but just a blind extrapolation. Using the past as a fallible guide, however, take a guess at what the range of possibilities is for each of the input values.

b. How much carbon is mankind predicted to emit by the end of the century? Using the uncertainty ranges you made in 2a, what is the highest and lowest plausible carbon emission for 2100?

3. IPCC CO₂ Emission Scenarios. Open a new browser window for the ISAM carbon cycle model, at <http://geosci.uchicago.edu/~archer/PS134/cgimodels/isam.html>. This page shows the results of IPCC carbon emission scenarios made by more sophisticated crystal balls than our simple Kaya identity (and then offers to run the results of the scenarios through a carbon cycle model). On the Kaya page, try to reproduce the year-2100 carbon emissions from scenarios A (Business-as-usual), B (BAU with carbon stabilization), C (slow growth) and F (gonzo emissions). What input parameters are required?

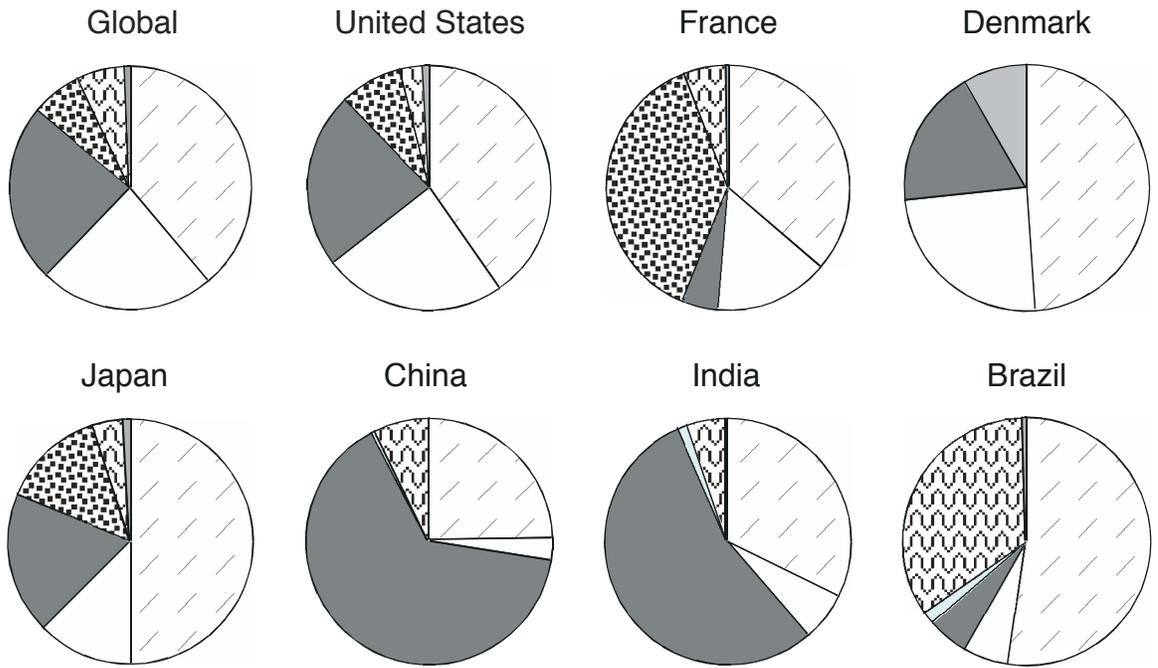


Figure 8-1

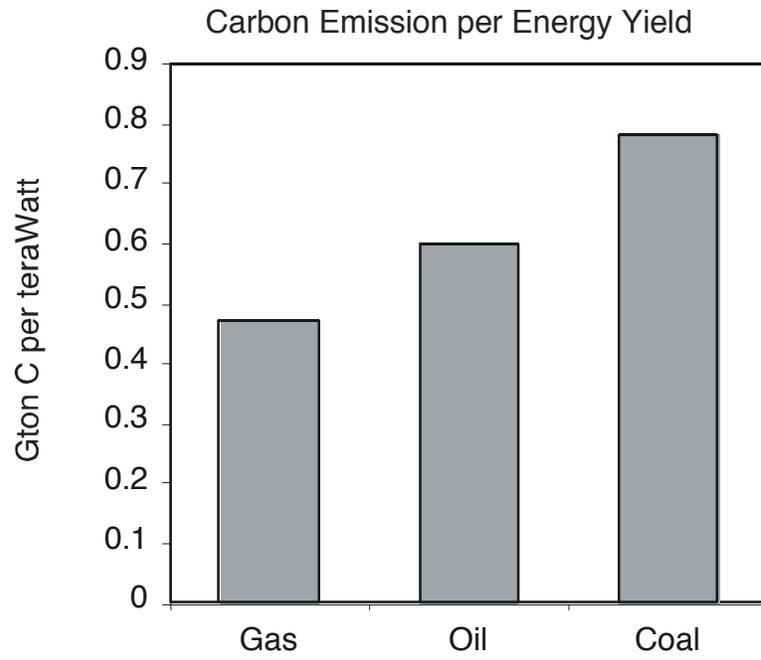


Figure 8-2

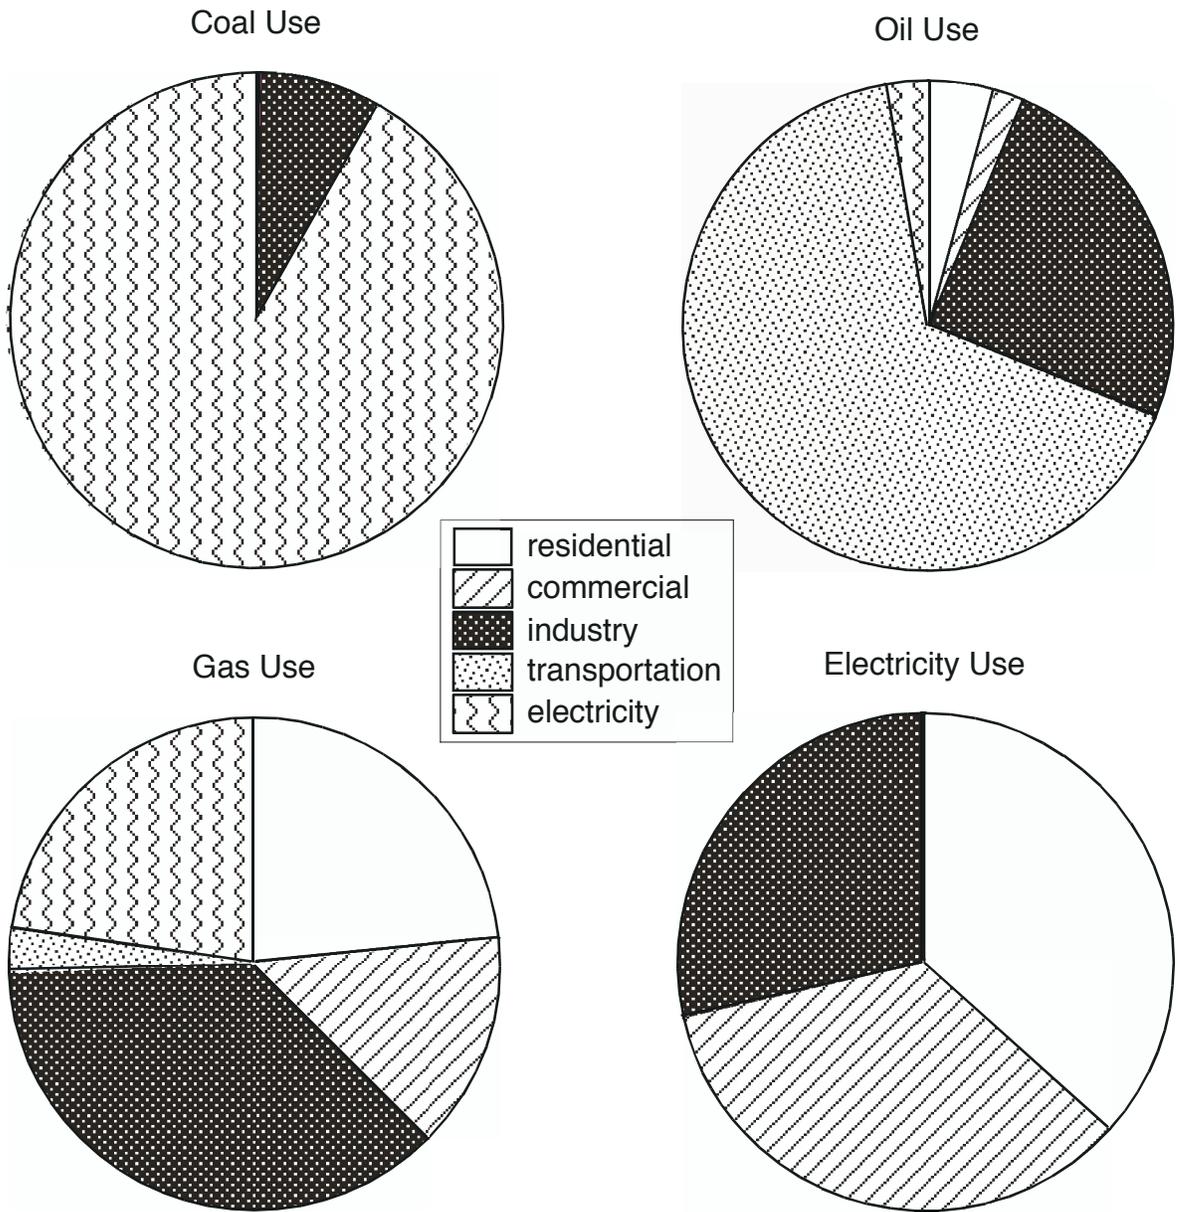


Figure 8-3

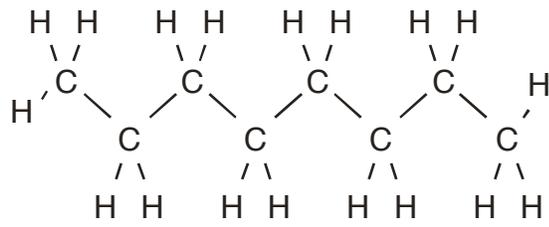


Figure 8-4

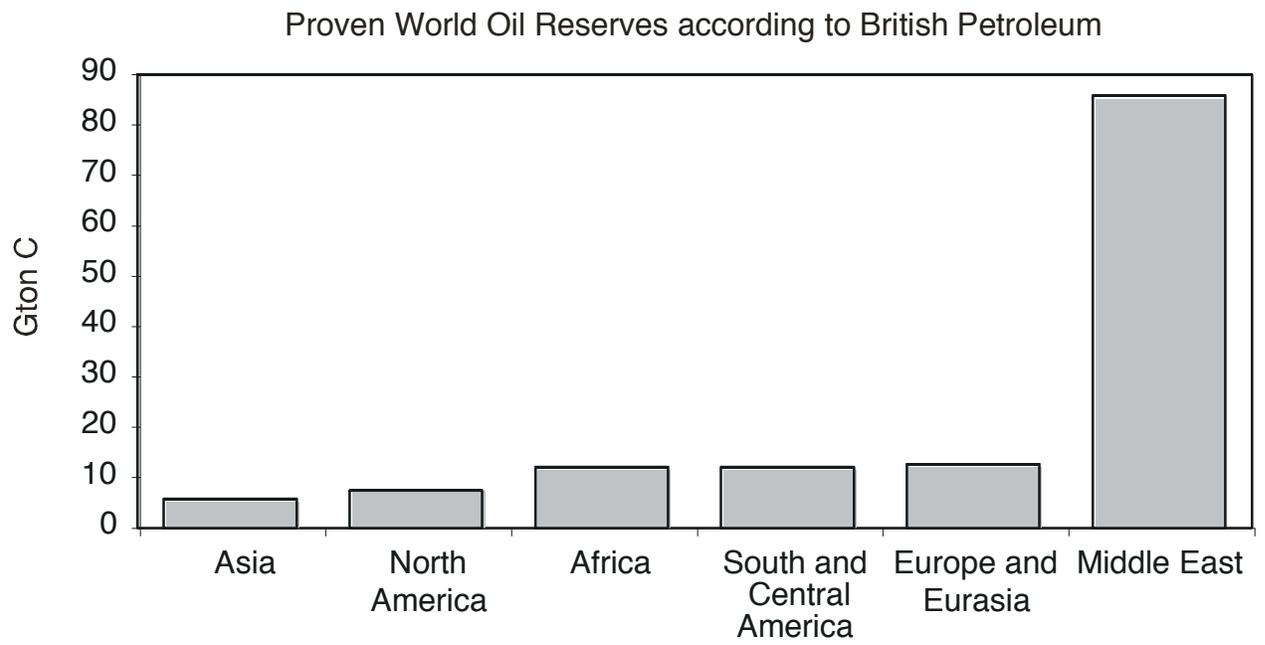


Figure 8-5

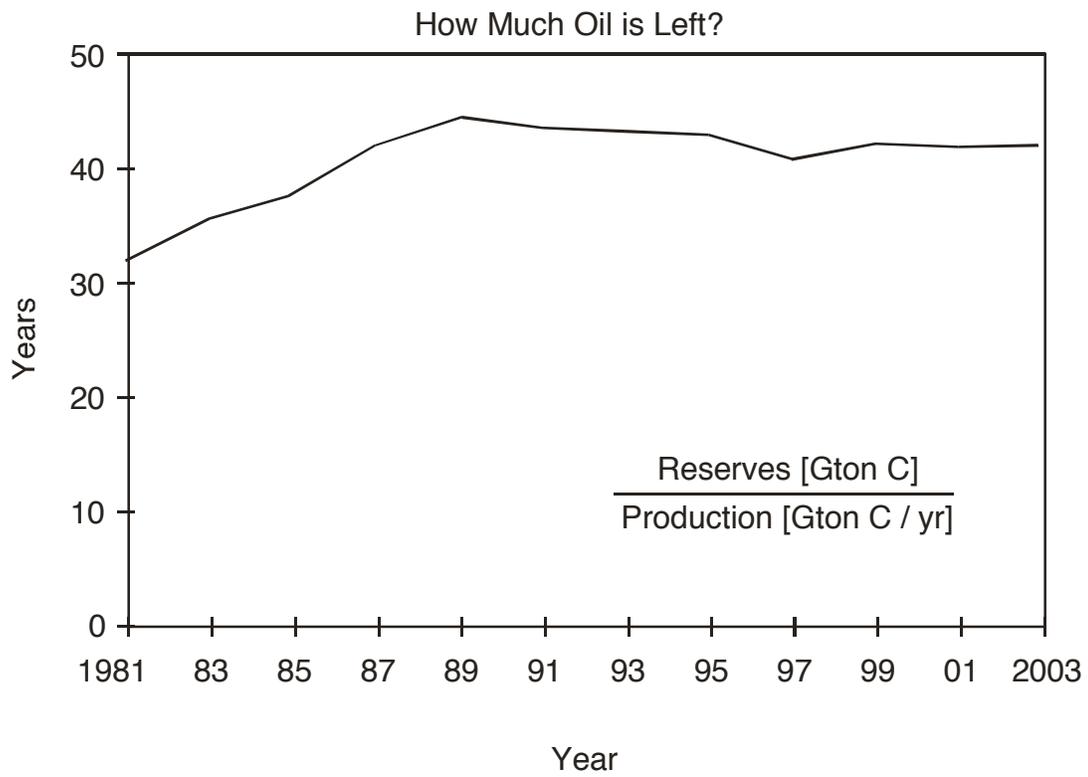


Figure 8-6

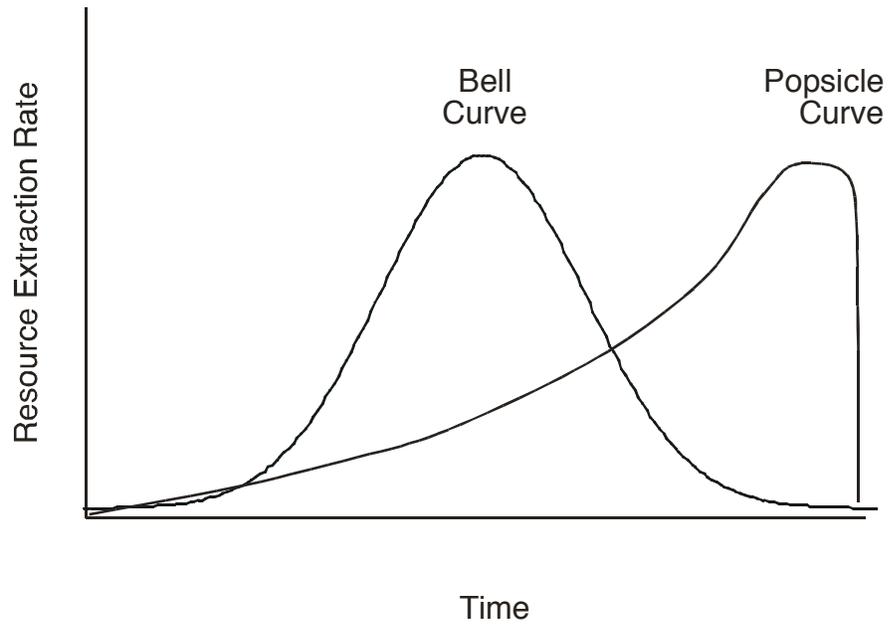


Figure 8-7

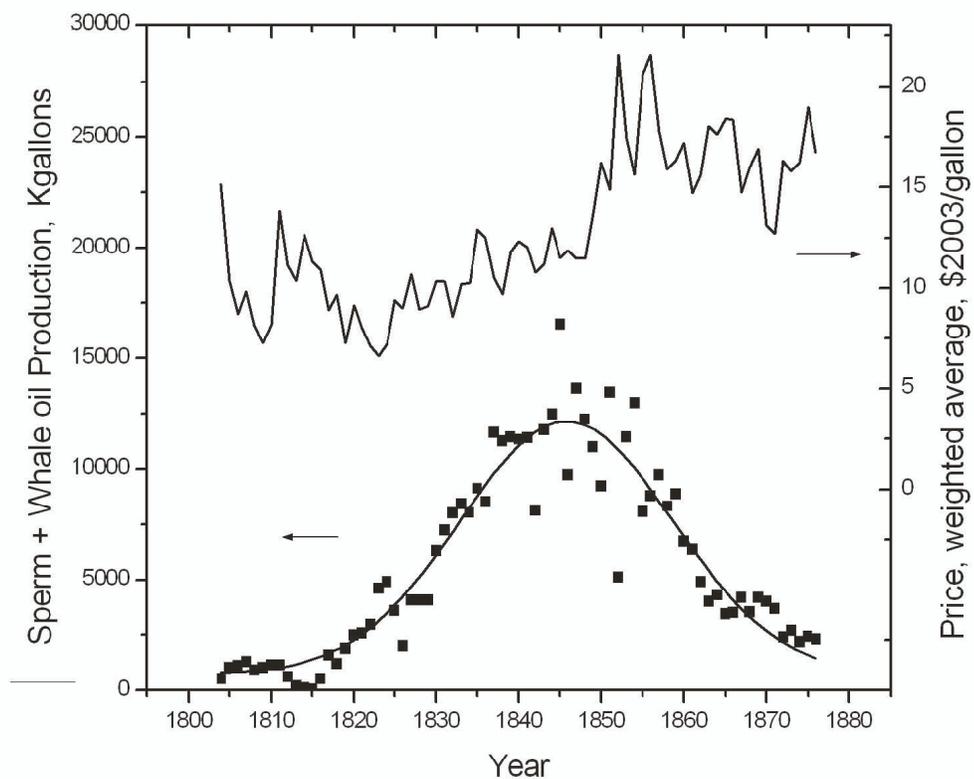


Figure 8-8

By Ugo Bardi
 Dipartimento di Chimica -
 Università di Firenze
www.aspoitalia.net/aspoengli
 sh
bardi@unifi.it

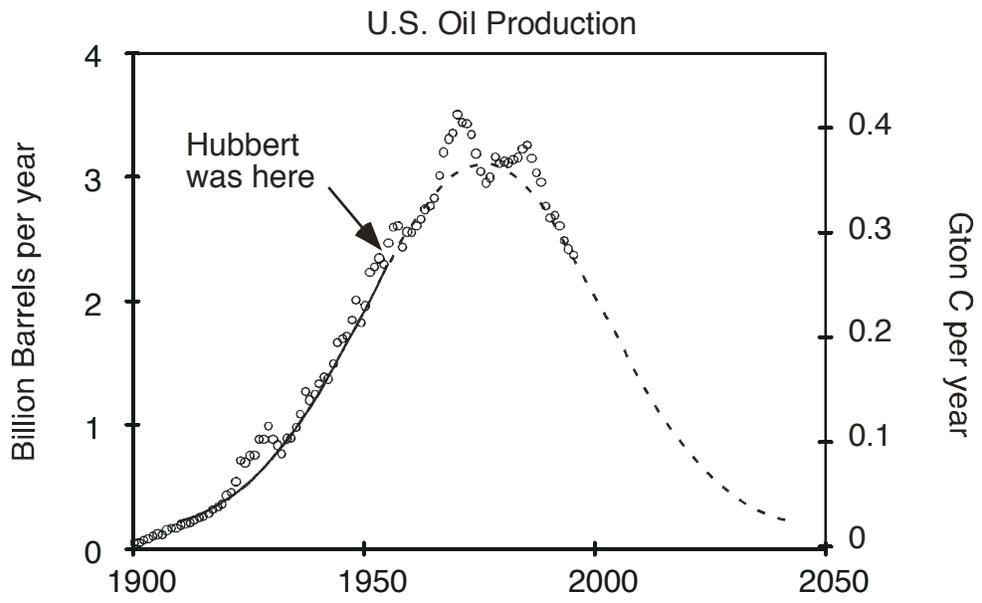


Figure 8-9

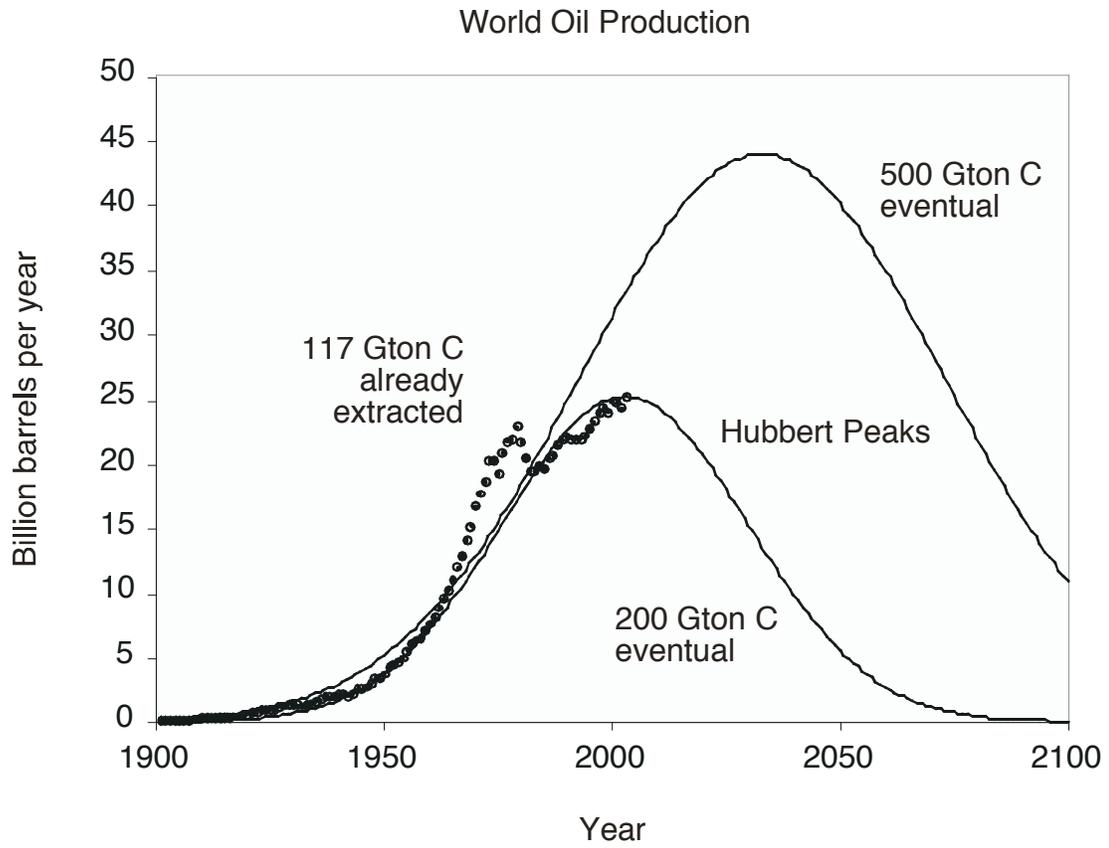


Figure 8-10

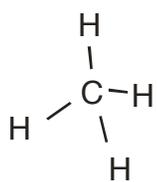


Figure 8-11

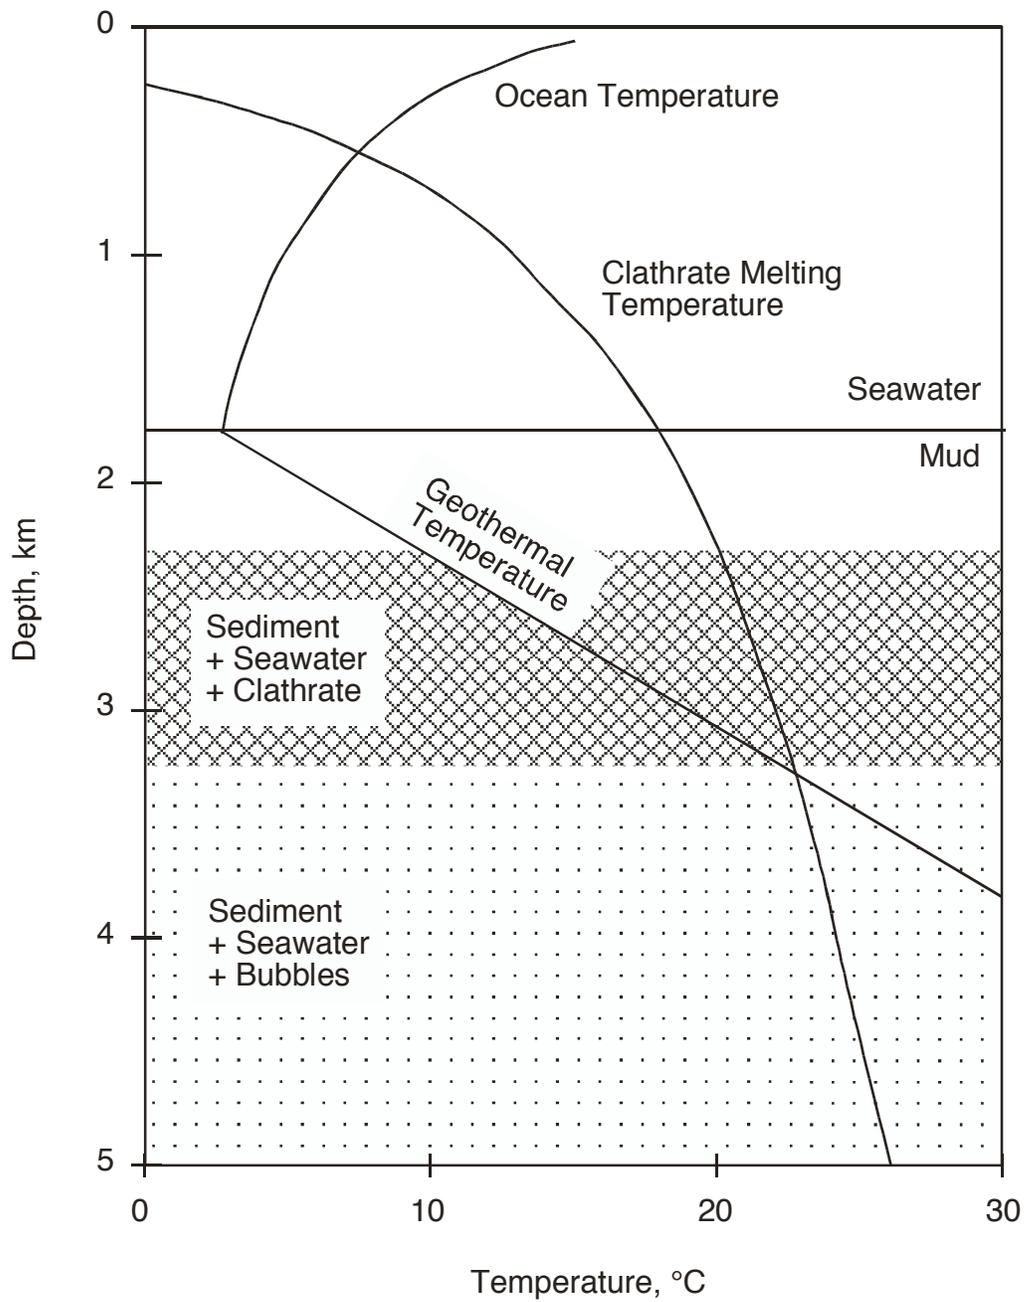


Figure 8-12

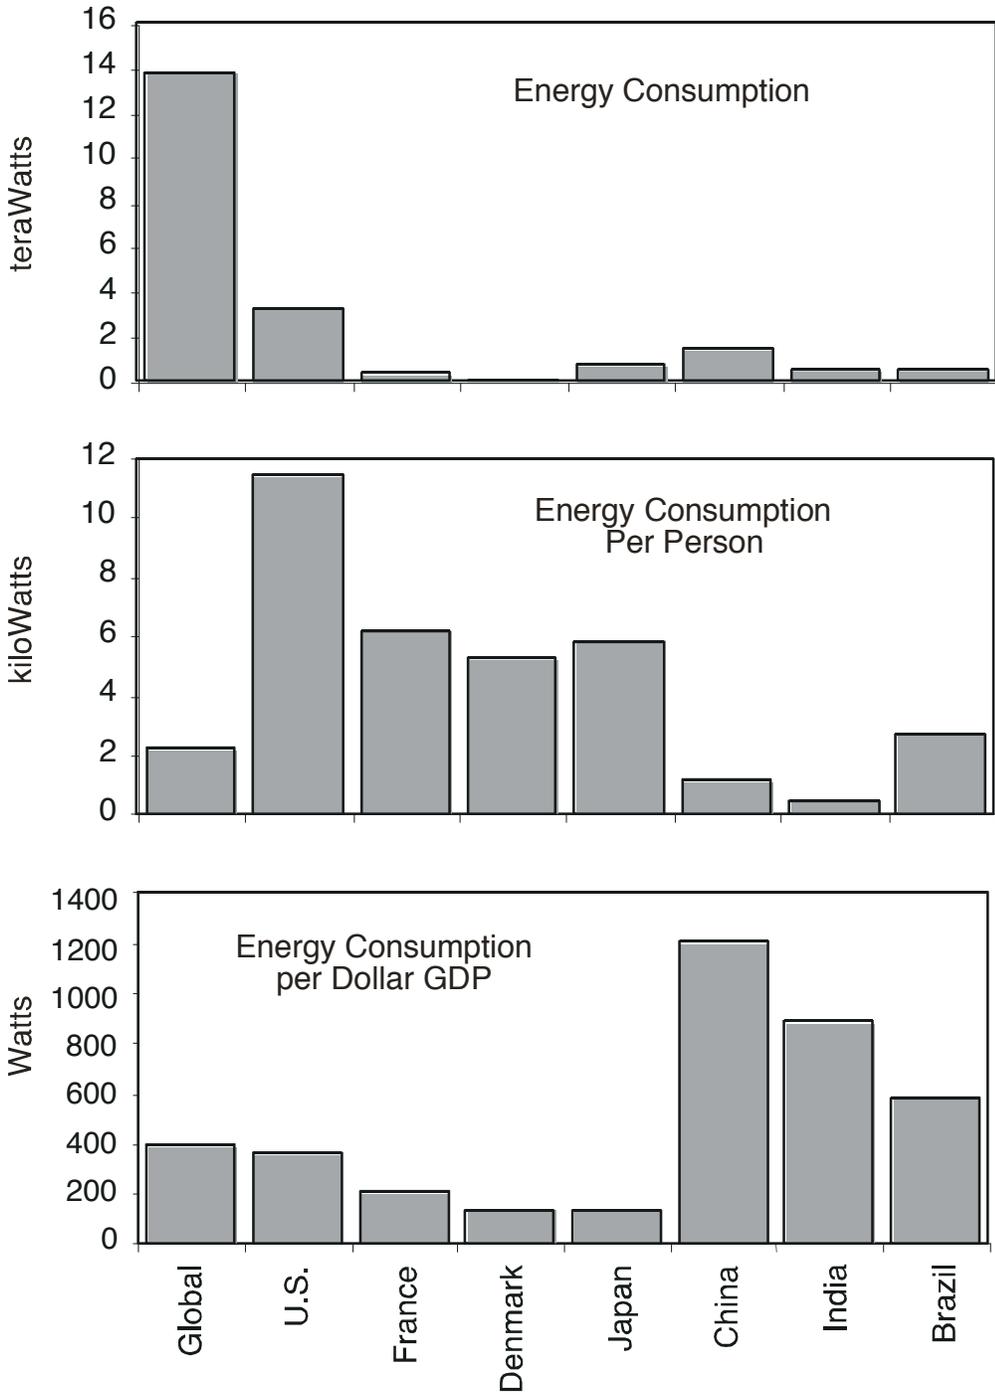


Figure 8-13

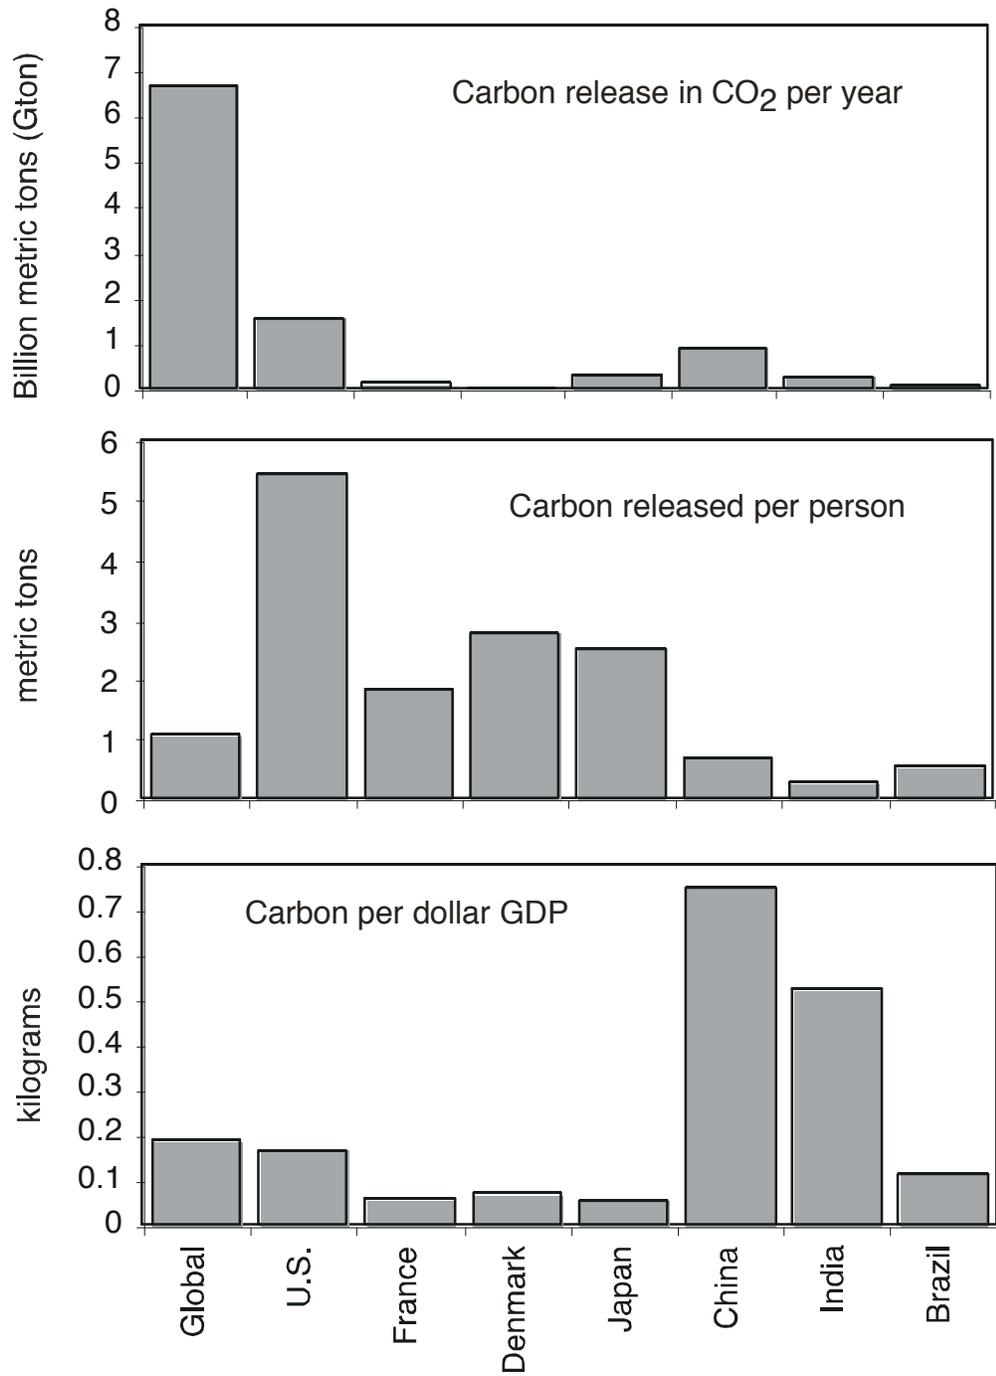


Figure 8-14

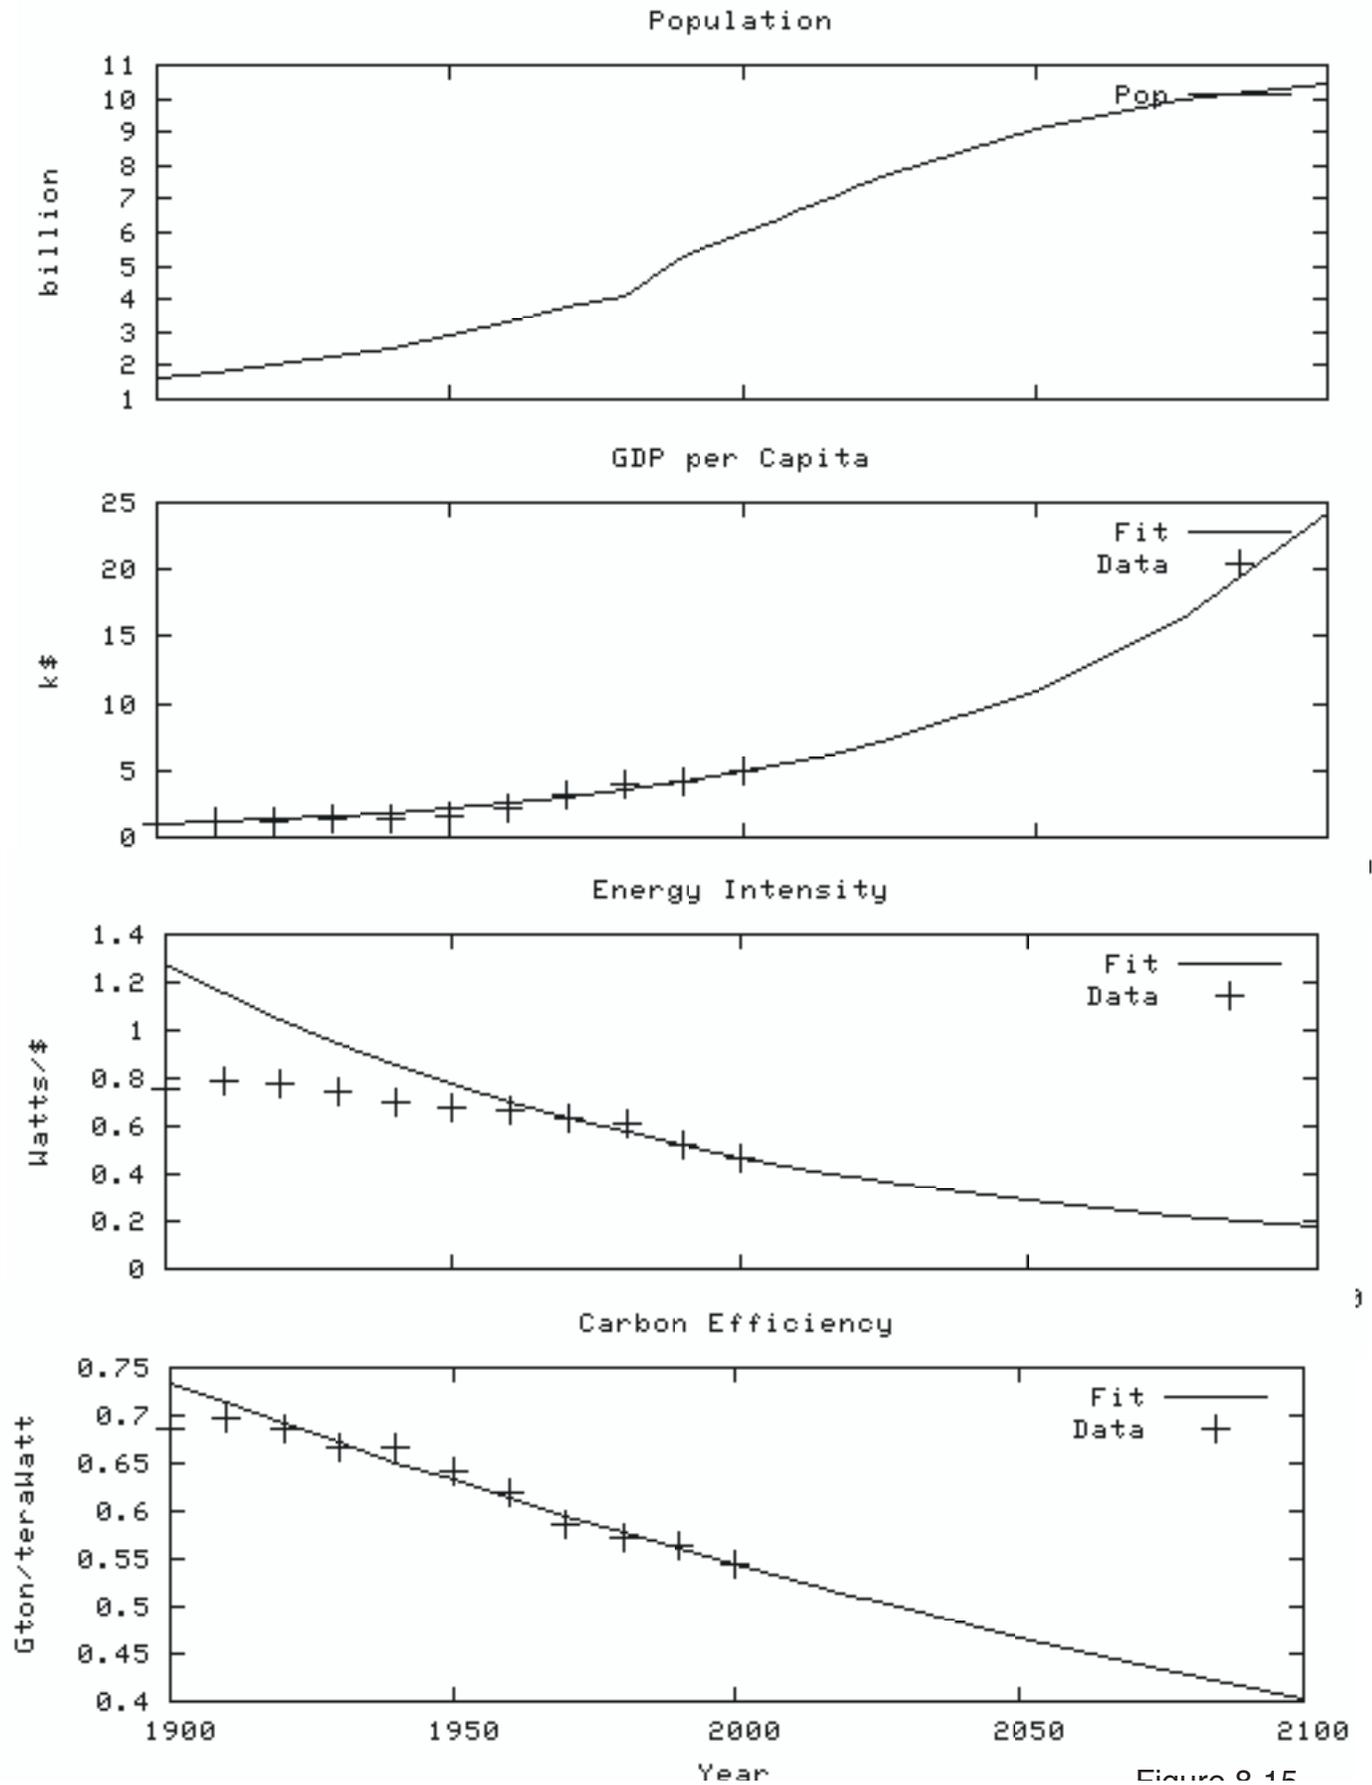


Figure 9.15

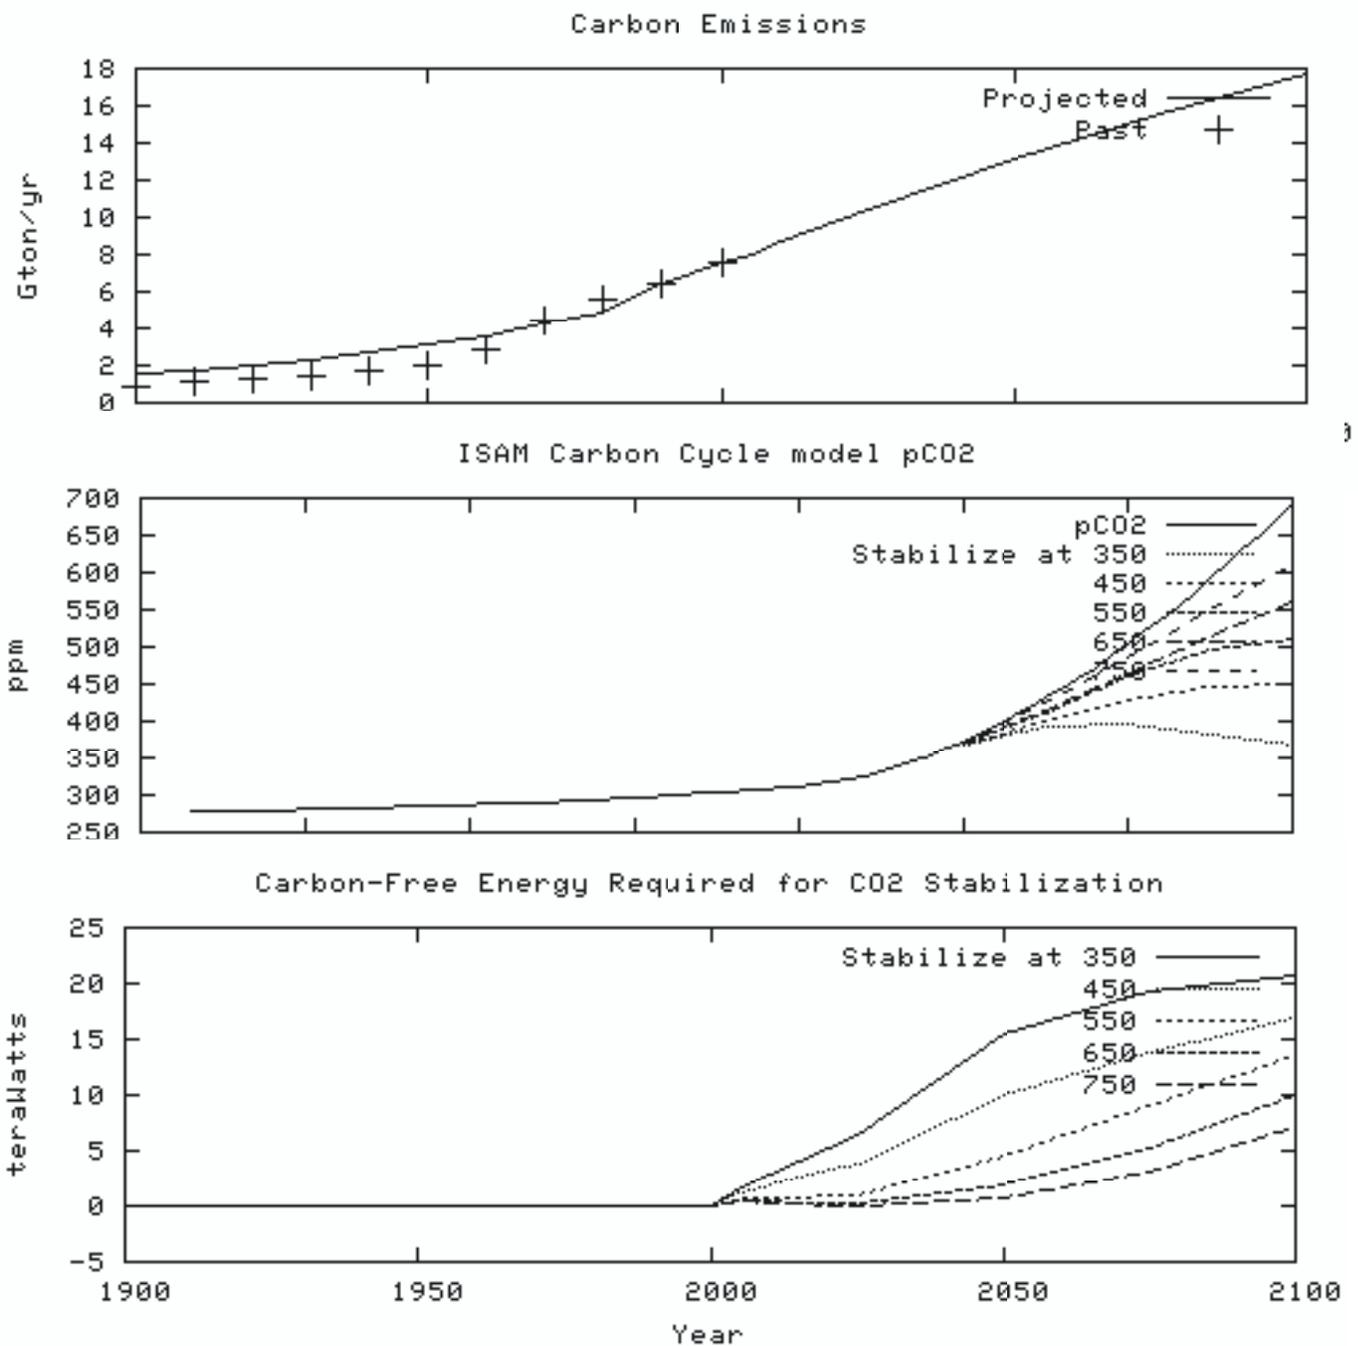


Figure 8-16

Chapter 9. The Perturbed Carbon Cycle

Abstract

We describe three geochemical global change issues here. Ozone in the stratosphere is beneficial because it shields us from ultraviolet sunlight, but it is depleted by chlorofluorocarbons. Ozone in the troposphere is harmful to lungs and to plants, but it is produced by automobile emissions. Neither problem has much to do with global warming but we present it to allay widespread confusion about this.

Methane is produced by livestock and rice paddies as well as natural sources. The lifetime of methane in the atmosphere is about 8 years. The concentration of methane in the atmosphere can be calculated from the source flux multiplied by the lifetime.

CO₂ is accumulating more slowly in the atmosphere than we are releasing it, because of natural uptake by the terrestrial biosphere and the ocean. Ocean uptake is easier to quantify than terrestrial, because chemical properties in the ocean are less patchy than they are on land. CO₂ has a long lifetime in the atmosphere, and will influence the climate for hundreds of thousands of years into the future. Long-term projections of carbon use, compared with the demands of stabilizing the CO₂ concentration of the atmosphere, indicate that a major new source of carbon-free energy will be required in the coming century, with the capacity to produce as much energy as we are currently consuming or even more.

Ozone

Let's get one issue out in the open and out of the way before we begin thinking about global warming. The ozone hole is not the same as global warming. Ozone is a reactive oxygen molecule comprised of three oxygen atoms. Ozone in the stratosphere is produced as O₂ molecules are zapped by energetic UV-C ultraviolet light. Ozone itself absorbs UV light that is less energetic (UV-B) but more abundant. Stratospheric ozone filters UV-B radiation that might otherwise reach the surface, causing skin cancers and sunburn. Stratospheric ozone is depleted somewhat by chlorofluorocarbons, inert chlorine-bearing compounds that are used in refrigerators and air conditioners. The ozone hole is a phenomenon found in Southern spring over Antarctica. The chlorine from the chlorofluorocarbons is transformed chemically, by the presence of nitric acid clouds in the frigid air of the Antarctic stratosphere, into a form which absolutely goes to town on ozone as soon as the sun rises. The ozone hole was unexpected, a "surprise", first observed in measurements of ozone concentration made by hand, on the ground. After these measurements, old satellite ozone records were examined, and it turned out that the satellite had been seeing the ozone hole for several years, but a data quality algorithm had been programmed to throw out any data which violated common sense, and out with the bad data went the ozone hole. This is a measure of how much of a surprise it was. The ozone hole is a terrific example of a **smoking gun**, clear-cut proof of the detrimental effects of chlorofluorocarbons on stratospheric ozone. Before the ozone hole was detected, the models predicted a gradual decrease in stratospheric ozone

globally, superimposed on natural variability; a difficult prediction to prove. As a result of the ozone hole, and the availability of economically and technologically viable alternatives, many countries of the world ratified the Montreal Protocol in 1987, phasing out further production and release of chlorofluorocarbons to the environment. There is an analogy here to CO₂ and global warming, in that the environmental damage predicted from the pollutant is a slow and gradual thing.

Ozone is also produced by reaction of gases from industrial activity, mostly automobile exhaust, in surface urban smog. The ingredients for ozone production in urban air are evaporated organic carbon compounds, like evaporated gasoline or molecules that are emitted naturally from trees and plants, nitrogen oxide compounds that are emitted from automobiles, and sunlight. When ozone concentrations exceed a toxic limit, asthma sufferers begin to feel discomfort, and plant leaves get “burned” and scarred.

Ozone is confusing to the student of the environment because ozone in the stratosphere is a good thing but industrial activity was acting to deplete it, while ozone in the troposphere is a bad thing and industrial activity tends to produce it.

Methane

Methane is a greenhouse gas, 20 times more powerful than CO₂ at current concentrations (see Chapter 3). Methane has natural sources as well as additional anthropogenic sources to the atmosphere. Once released to the atmosphere, methane reacts slowly with activated oxygen compounds to oxidize back to CO₂. The reactive oxygen compounds are produced by sunlight. In the absence of sunlight methane and O₂ gas coexist in ice core bubbles for hundreds of thousands of years with no reaction. Put it in the sunlight and it slowly burns up.

The lifetime of a methane molecule in the atmosphere is about 8 years. If the methane release rate is held steady for awhile, longer than 8 years, the concentration of methane in the atmosphere will settle on a value at which methane degradation equals methane release. We can use the units to guide us to calculating the steady-state methane concentration

$$\text{standing stock [GtonC]} = \text{release} \left[\frac{\text{GtonC}}{\text{yr}} \right] \cdot \text{lifetime [yr]}$$

This relation tells us that the methane concentration in the atmosphere is linearly related to the methane source to the atmosphere, as long as the lifetime of methane stays the same. Methane is consumed by reactive oxygen compounds in the atmosphere, in particular a molecule called OH radical. OH radical is related to ozone, so one could imagine a change in OH radical driven by the change in ozone chemistry of the atmosphere, or that all the OH radicals produced by the ozone chemistry might get used up reacting with methane, if the methane concentration goes up too high. So far, however, the methane lifetime seems to be holding steady.

One of the natural sources of methane to the atmosphere is the degradation of organic carbon in freshwater swamps. Organic carbon reacts most energetically with O_2 , as we have discussed in Chapter 7. In seawater, after the O_2 is gone, there is a substitute in sulfate ion, SO_4^{2-} , which reacts with organic carbon to produce hydrogen sulfide, H_2S . After oxygen and sulfate are depleted, organic carbon is fermented by bacteria, releasing as one of its products methane. This is how the methane is produced that freezes into clathrate deposits below the sea floor. Sulfate is part of sea salt, which there isn't much of in fresh water. Methane is therefore produced much shallower in the sediment column in fresh water. If you step in mucky swampy freshwater mud you may see bubbles of methane rising up around your legs. Methane is sometimes referred to as swamp gas for this reason, and is one of the usual suspects blamed for UFO sightings. Bubbles of ancient atmosphere preserved in ice cores tell us that the methane concentration has fluctuated with climate state over the past 400,000 years (the longest ice core yet available), with lower methane concentrations during colder, drier climate stages (Figure 7-3). This is interpreted to be the result of the changing abundance of swamps.

Anthropogenic sources of methane include production in the guts of ruminant animals, and release as leakage by the fossil fuel industry. Rice paddies often provide ideal anoxic freshwater environments for methane production. The methane concentration has doubled over its pre-anthropogenic concentration (Figure 9-1), and is responsible for a quarter of anthropogenic greenhouse heat trapping (Figure 9-2).

CO_2

There are two main anthropogenic sources of CO_2 to the atmosphere. One is deforestation. A heavily-wooded forest holds more carbon per area than does a plowed agricultural field. The world's forests started to feel the ax thousands of years ago with the development of agriculture and the growing of the human population. Most of the temperate latitudes have been cut long since, and the tropics are currently being cut. The year 1750 has been taken as the beginning of the anthropogenic CO_2 rise, although it has been argued by Ruddiman [Ruddiman, 2003 #2616] that both CO_2 and methane started rising from their "natural" trajectories thousands of years ago. The rise in atmospheric CO_2 after 1750 (Figure 9-1) is clearly the result of deforestation, perhaps in the New World (the "pioneer effect"). Today, visible deforestation is mostly to be found in the tropics, and accounts for about 2 Gton C / yr release of CO_2 to the atmosphere.

The other main anthropogenic CO_2 source is of course the combustion of fossil fuels discussed in the last chapter. Fossil fuel combustion releases 5 Gton C / yr, rising exponentially, driven by population growth and economic growth, rising in spite of increases in energy and carbon fuel efficiency.

Combining the ocean and tropical deforestation, mankind is releasing carbon to the atmosphere at a rate of about 7 Gton C / yr. The atmospheric CO_2 inventory is rising at a rate of about 3 Gton C / yr (Figure 9-3). Where is the other 4 Gton C / yr? There are two main natural **sinks** for CO_2 that are operating today; one is the oceans, and the other is the terrestrial biosphere.

Carbon uptake by the terrestrial biosphere on land is difficult to measure. In comparison with the ocean, the distribution of carbon on land is very spotty. Recall from Chapter 6 that most of the carbon on land is in the soil, rather than in the trees where we could see it. In soils, the amount of carbon depends on the recent history of the land: fires, agriculture, erosion, and so on. It is difficult to know precisely how much carbon there is on land because the measurements vary all over the place; you would have to make a gazillion measurements in order to average out all the noise of natural variations. As a result of this, it would be possible to increase the amount of carbon on land, a little bit here or there, in a way that would be entirely invisible to direct measurements. The land is playing two roles in the carbon budget story, one as visible deforestation source and another as a potential invisible carbon uptake sink.

One way to estimate the invisible terrestrial uptake is by putting together the rest of the carbon budget and assign the terrestrial biosphere whatever is left over. The chemical properties of seawater vary more smoothly than they do on land, so our estimates of ocean uptake are better than they are for the land. Another is to measure CO₂ concentrations in the atmosphere, in the winds as they blow across the land, to see if CO₂ is going into or coming out of a given forest. This doesn't sound easy, but it can be done. There is a global network of "CO₂ observatories" where precise CO₂ concentration measurements are made daily, for uses such as this.

There are several reasons why the land may be willing to take up CO₂ as the atmospheric CO₂ concentration rises. One is that with warming, there will be a longer growing season. This has been observed in many climate and botanical records. With warming, some tundra areas become amenable to conquest by forests. Rising CO₂ in the atmosphere may also directly encourage plants to grow faster, by a process known as **CO₂ fertilization**. Plants run their photosynthetic machinery inside waxy walls on the surfaces of leaves. Gases are exchanged with the outside atmosphere through adjustable vents called **stomata**. When the leaf needs CO₂ for photosynthesis, the stomata open. The cost of opening stomata, though, is loss of water. So if CO₂ concentrations were higher in the outside atmosphere, plants could get the CO₂ they need without opening their stomata as much. They could be stingier with their water. There is no doubt that this is a real effect; CO₂ concentrations in greenhouses are typically higher than in the outside atmosphere, one of the ways that greenhouses are good for plants. However, in the real world, plant growth is very often limited by something else other than water stress, such as fertilizers like nitrogen or phosphorus. Scientists do CO₂ fertilization experiments in natural settings by pumping CO₂ continuously into the air. When the wind changes, they adjust the location of the CO₂ vent, so that the target grove is always downwind from a CO₂ source. These experiments go on for years! What they tend to find is initial growth spurt from CO₂ fertilization, followed by a leveling off at something like the initial rates.

There is another process which may affect CO₂ storage on land, which is a temperature sensitivity to respiration, the process that converts soil organic carbon back into CO₂. Soil respiration really gets going as it gets warmer. Think of a ham sandwich, half of which is safely stowed in the refrigerator while the other half sits on a plate in the

sun. Which half will stay tasty longer? For this reason, there is very little organic matter in tropical soils, while high latitudes may host peat deposits that contain prodigious carbon deposits for thousands of years. Warming and melting and decomposition of high-latitude permafrost may contribute CO₂ to the atmosphere.

The ocean sink depends on ocean circulation, and on the chemical forms that dissolved CO₂ takes in seawater. The ocean covers 70% of the earth's surface, and the length and width of the ocean are huge compared with its depth, which averages about 4 km. The deep ocean is so close to the surface, and yet it is so very far away. The way the ocean circulates, the deep ocean is very cold, and the only place where surface waters are cold enough to mix with the deep ocean is in high latitudes. The ocean surface is huge, but the deep ocean, which is the largest water type in the ocean, only sees the atmosphere through a very small area of sea surface.

The densest water at the sea surface is in the Antarctic and in the North Atlantic, because it is cold there. Surface waters from these locations sink to the deep ocean, filling up the entire deep ocean like a bucket with cold polar water that is only a few degrees warmer than freezing. The cold deep ocean fills up until cold polar waters underlie the warm surface waters in lower latitudes. The warmer waters mix with the cooler, eroding the cold water and making room for more of the coldest water to continue filling the deep sea. As new cold water flows from the high-latitude surface ocean into the abyss, it carries with it atmospheric gases like anthropogenic CO₂, a process known as **ocean ventilation**. It takes centuries for the waters of the deep ocean to travel through this cycle. For this reason, the time scale for getting anthropogenic CO₂ into the deep ocean is centuries.

The shallower ocean has other water masses and circulation modes that are wondrous to learn about and study if one is of a mind to. The zone of the ocean separating the warm from the cold is called the **thermocline**. Thermocline waters may be exposed to the atmosphere in winter, when the sea surface waters are cold. Once a parcel of thermocline water becomes isolated from the sea surface, it follows a trajectory determined by its density and by the rotation of the earth. Thermocline waters ventilate to the atmosphere on a time scale of decades.

The surface ocean water mass is not as large as the deep sea or the thermocline, but it is a respectable carbon reservoir of its own. Turbulence generated by the wind acts to mix the surface ocean down to a typical depth of 100 m. Gases from the atmosphere dissolve in the ocean rather slowly. For an unreactive gas like argon, the concentration of dissolved gas approaches equilibrium with the atmosphere on a time scale of about a month, for a 100-m deep surface ocean layer.

CO₂ differs from argon, however, in that it has chemical equilibrium reactions with water and hydrogen ions. Hydrogen ions are very reactive. If a solution has a high concentration of hydrogen ions, we call it **acidic**. The acidity of a solution is described by a number called the **pH** of the solution, which can be calculated as

$$\text{pH} = -\log_{10} [\text{H}^+]$$

The \log_{10} is the base-10 logarithm, meaning that if $x = 10^y$, then $\log_{10} x = y$. The hydrogen ion concentration is denoted by the square brackets, and usually ranges in seawater from 10^{-7} to $10^{-8.3}$ moles of H^+ per liter of solution. The pH of seawater therefore ranges from 7 to 8.3. Note that the more acidic the solution, the lower the pH of the solution. Ads for shampoo used to claim "low pH" as though pH were some toxic ingredient. I guess it sounded better than calling the shampoo "strongly acidic".

When CO_2 dissolves in water, it reacts with water to form **carbonic acid**, H_2CO_3 .



Carbonic acid loses a hydrogen ion (that's what acids do, in general, is release hydrogen ions) to form **bicarbonate ion**, HCO_3^-



A second hydrogen ion can be released to form **carbonate ion**, $CO_3^{=}$



The concentrations of carbonic acid, bicarbonate, and carbonate ions control the acidity of the ocean, just as they control the acidity of our blood and cell plasma.

These chemical reactions are fast enough that the distribution of chemical species will always be in their lowest-energy distribution. Many chemical reactions that we will encounter will not be in equilibrium, so we should enjoy this equilibrium system while we have it. Equilibrium reactions can be easily and very precisely predicted. In a qualitative way, we can use an idea known as **le Chatelier's** principle to take a stab at the behavior of the system. Le Chatelier's principal states that an addition or removal of a species on one side of the chemical equilibrium will cause the reaction to run in the direction to compensate for that change. Take some of something out, the equilibrium will put some of the something back. Add more something, the equilibrium will remove some of the something. Le Chatelier's principal is treacherous for students of the carbon system in seawater, though, because we have an innate human tendency, I have found, to ignore the hydrogen ions. They are such tiny things, after all. There are far fewer hydrogen ions in seawater than there are of the dissolved carbon species. What this means, however, is that a small change in the concentrations of the carbonate species might make a huge change in the hydrogen ion concentration. The safest assumption to make is that the carbon species have to get along together without counting on sloughing off too many hydrogen ions. Hydrogen ion is such a tiny slush fund, it might as well not exist. We can combine reactions 9-1, 9-2, and 9-3 into a single reaction, in such a way that we don't allow any production or consumption of hydrogen ions



To this reaction we can apply le Chatelier with impunity. If we were to add CO_2 to this system, the equilibrium would compensate somewhat by shifting to the right, consuming

some of the CO_2 by reacting it with $\text{CO}_3^{=}$. Seawater has the capacity to absorb or release more CO_2 than it would if CO_2 had no pH chemistry, because of the other carbon reservoirs HCO_3^- and $\text{CO}_3^{=}$. It is like sitting at a card game, with a rich uncle sitting behind you covering your losses and taking splits of your winnings. The pH reactions of H_2CO_3 , HCO_3^- , and $\text{CO}_3^{=}$ are called a **buffer**, because any changes to the chemistry tend to be resisted or buffered by the chemical reactions. The CO_2 concentration of seawater is buffered by its pH chemistry. Another way of saying this is that the seawater has greater capacity to hold seawater than it would if CO_2 were not buffered. The strength of the buffer is about a factor of 10, meaning that seawater has a capacity to hold 10 times as much CO_2 as it would if there were no buffer.

As the CO_2 concentration in the atmosphere increases, and CO_2 invades the ocean, the concentration of carbonate ion goes down, according to the equilibrium chemical reaction 9-4. Carbonate ion is our anti- CO_2 , reacting with new CO_2 , hiding it away as bicarbonate; this is the action of the buffer. As the carbonate ion becomes depleted, so is its ability to buffer CO_2 . Ocean uptake of new CO_2 would decrease as a result of this.

The future of ocean uptake of fossil fuel CO_2 may also be affected by changes in the circulation of the ocean. Surface warming is expected to be most intense in high latitudes because of the ice-albedo feedback (Chapter 6). If the high latitudes warm, the overall circulation of the subsurface ocean may decrease. The circulation of the ocean may **stagnate**, slowing uptake of CO_2 .

Biology in the ocean acts to decrease the CO_2 concentration of surface waters, by converting CO_2 into organic carbon via photosynthesis (Chapter 7). Dead phytoplankton sink from surface waters, exporting their carbon to the deep sea. This processes has been termed the **biological pump**. If all the life in the ocean were killed, that is if the biological pump were stopped, then the CO_2 concentration of the atmosphere would rise. If the biological pump were stimulated to work harder, it could decrease the CO_2 concentration of the atmosphere. One proposal for dealing with global warming is to fertilize the ocean with iron. Iron concentrations are extremely low in remote parts of the ocean, far away from iron deposition from dust and iron bleeding from surface sediments. Supplying iron to the phytoplankton has been shown to stimulate phytoplankton growth in the Southern Ocean around Antarctica for example. The problem is that it takes hundreds of years for the ocean and the atmosphere to negotiate what the atmospheric CO_2 concentration should be; it's slow, recall, because the ocean circulation is so slow. Model studies have shown that fertilizing the Southern Ocean for hundreds of years might bring the CO_2 concentration of the atmosphere down, but fertilizing for a few decades has very little impact.

Decreasing carbonate ion may also be detrimental to coral reef and other organisms that produce limestone, CaCO_3 , from calcium ion, Ca^{2+} , and carbonate ion. It's like pouring vinegar on limestone steps; you will see bubbles as the acid of the vinegar converts the CaCO_3 to CO_2 . Fossil fuel CO_2 is itself an acid, and drives CaCO_3 to dissolve. Note the counter-intuitive reverse behavior; one might have thought that adding CO_2 to the oceans would lead to an increase in the amount of carbon that winds up as

CaCO_3 , the rising tide floats all boats. The response is backward because CO_2 is an acid and CaCO_3 reacts with acid. Fish and other aquatic organisms react poorly to the acidity and higher CO_2 concentrations resulting from fossil fuel CO_2 release. This danger is called the **acid ocean**.

Uptake of CO_2 into the oceans has been estimated by a number of different independent methods. These include measurements of the chemical concentrations throughout the world ocean, and modeling the circulation and carbon cycle. Other chemicals serve as tracers for how the ocean circulates. These include radioactive elements produced naturally by cosmic rays, or in nuclear bomb tests in the 1960's, and industrial chemicals like chlorofluorocarbons. Another distinction between dissolution of CO_2 in the ocean versus uptake of CO_2 by photosynthesis on land is that photosynthesis releases oxygen, while dissolution in water does not. So you can measure the change in CO_2 along with the change in atmospheric O_2 concentration to figure out what fraction of the missing CO_2 is going into the ocean versus into the terrestrial biosphere by photosynthesis. There are discrepancies between the different estimates and uncertainties associated with each method, but in general they all point to about a 50:50 split. The ocean gets about 2 Gton C per year of the anthropogenic CO_2 and the terrestrial biosphere gets 2 Gton C per year. The terrestrial uptake of new CO_2 is thought to be taking place in high northern latitudes, perhaps into the great forests of Canada and Siberia.

What about carbon uptake by the natural world in the future? There is certainly room for some change in the size of the terrestrial biosphere in the coming centuries, but which way it ultimately goes depends on how people decide to use the land surface, in addition to biological factors such as CO_2 fertilization and increases in soil respiration. So far, the amount of carbon released to the biosphere has been smaller than the size of the terrestrial biosphere, especially if we consider soil carbon. So it's not outrageous to think about the terrestrial biosphere as a potential major player in the carbon budget. However, if we really go to town on the world's coal or methane clathrate deposits, the amount of CO_2 we could release would be several times larger than the carbon stored in the terrestrial biosphere. It would be difficult to imagine the terrestrial biosphere saving the day in that case.

After hundreds of years, most of the fossil fuel CO_2 will dissolve in the oceans, lowering the pH of the ocean and the concentration of carbonate ion. How does the Earth proceed from there? On time scales of thousands of years, the pH and carbonate ion concentration of the ocean are controlled by the dissolution of limestone, CaCO_3 . Limestone deposits continually dissolve on land, a chemical reaction we have already defined as weathering, releasing Ca^{2+} and carbon to the ocean. Plankton make shells of CaCO_3 at a rate which exceeds the supply rate of dissolved CaCO_3 by a factor of 10 or so. The ocean can only afford to bury a tenth of the tiny shells, if it wants the CaCO_3 budget of the ocean to be in balance. The effect of the CO_2 acid invasion of the ocean will be to make it harder for little shells to be buried. In fact, if we ultimately release 1000 Gton C or so, there will be net dissolution of CaCO_3 from the sea floor. Mankind will have reversed the net sedimentation of the ocean! Engineers in Chicago early in the last century reversed the flow direction of the Chicago river. That was impressive in its time,

but it was nothing compared to this. Weathering will exceed burial, so dissolved CaCO_3 will accumulate in the ocean. This process will restore the $\text{CO}_3^{=}$ concentration, and the pH of the ocean, back to their previous natural states, and at the same time, drawing down atmospheric CO_2 somewhat. This process will take thousands of years. On time scales of hundreds of thousands of years, the silicate weathering thermostat, defined and described in Chapter 7, will act to pull CO_2 down the rest of the way toward the pre-anthropogenic value (Figure 9-4). The bottom line is that about 15-30% of the CO_2 released by burning fossil fuel will still be in the atmosphere in 1000 years, and 7% will remain after 100,000 years. Truly, global warming is forever.

The bottom line is that of the 7 Gton C / yr that mankind is releasing to the atmosphere today, 4 Gton C / yr is going away as quickly as we release it. This leads to a simple but powerful conclusion: if we want atmospheric CO_2 to stop going up, tomorrow, we have to reduce our CO_2 emissions from 7 Gton C / yr down to 4 Gton C / yr, say a reduction by 40%. Then the CO_2 concentration in the atmosphere would stop rising, but remain at its current level of 365 ppm. This could continue until the terrestrial biosphere and the ocean equilibrated at this new level, “filled up” with the new higher CO_2 . No one has any idea how long it would take for the terrestrial biosphere to saturate or fill up, but the ocean would take several centuries. If emissions were stopped after that, the atmospheric CO_2 concentration would remain at 365 ppm for thousands of years.

The aim of the Kyoto Protocol, the international agreement to reduce CO_2 emissions discussed in Chapter 12, is to reduce emissions to about 6% below the 1990 levels, resulting in emissions that are still very close to 7 Gton / yr. CO_2 emissions under business-as-usual are projected to grow, so the rather modest-sounding 6%-below-1990 target actually amounts to about 30% reductions from the projected 2010 rate. Still, this is just a drop in the bucket of what would be required to truly stabilize the CO_2 concentration of the atmosphere. The Kyoto protocol by itself is not sufficient to end the problem of global warming; it can only be the first step, alas.

Take-homes

The ozone hole is not global warming. They are different issues.

Methane has a short lifetime in the atmosphere.

CO_2 has a long lifetime in the atmosphere. Stabilizing CO_2 in the atmosphere at some "safe level" (whatever that is) will require major new energy initiatives.

Lab

1. Long-Term Fate of Fossil Fuel CO_2 . Use the on-line geologic carbon cycle model at <http://geosci.uchicago.edu/~archer/cgimodels/geocarb.html>. Use the default setup of the model, and notice that the CO_2 weathering rates etc. for the transient state are the same as for the spinup state. So if there were no CO_2 spike at all, there would be no change in anything at year 0. (Go ahead, make sure I'm not lying about this.) Release some CO_2 in a transition spike, 1000 Gton or more or less, and see how long it takes for the CO_2 to

decrease to a plateau. There are two CO₂ plots in the output, one covering 100 thousand years and one covering 2.5 million years. How long does it take for CO₂ to level out after the spike, according to both plots?

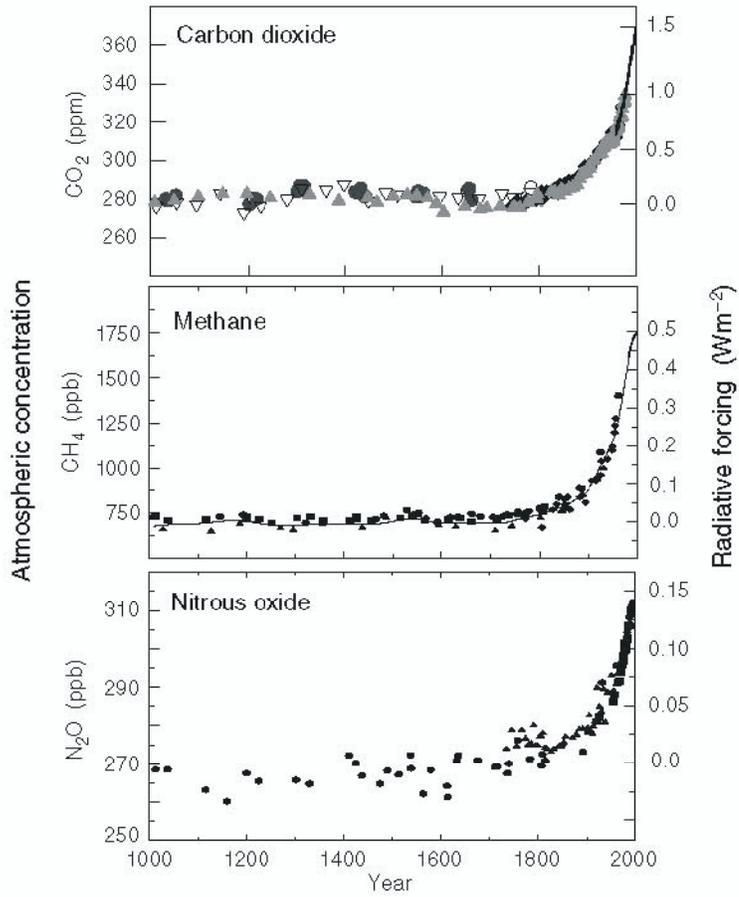
2. Effect of Cutting Carbon Emissions. Look at the on-line ISAM global warming model at <http://geosci.uchicago.edu/~archer/PS134/LabManual/lab.isam.html>.

a. Run the model for the "Business-as-usual" case (Scenario A), and for the cases of halting fossil-fuel and deforestation carbon fluxes now (call it year 2000) and in the year 2050. (The model interpolates between the time points in the input. Zeroing the values in 2075 and 2100, for example, will result in the 2050 emission occurring in 2050, and then a straight line decrease with time to zero in 2075. Therefore you won't be able to shut off emissions abruptly in the year 2050, but the ramp-down will be close enough.) What is the impact of stopping CO₂ emissions now vs. in the year 2050?

3. Climate Sensitivity of this model. Deduce from the above results or new model runs, what is the climate sensitivity, ΔT_{2x} , assumed in this model?

Indicators of the human influence on the atmosphere during the Industrial Era

(a) Global atmospheric concentrations of three well mixed greenhouse gases



(b) Sulphate aerosols deposited in Greenland ice

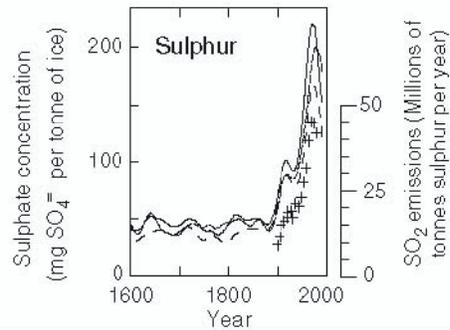


Figure 9-1

The global mean radiative forcing of the climate system for the year 2000, relative to 1750

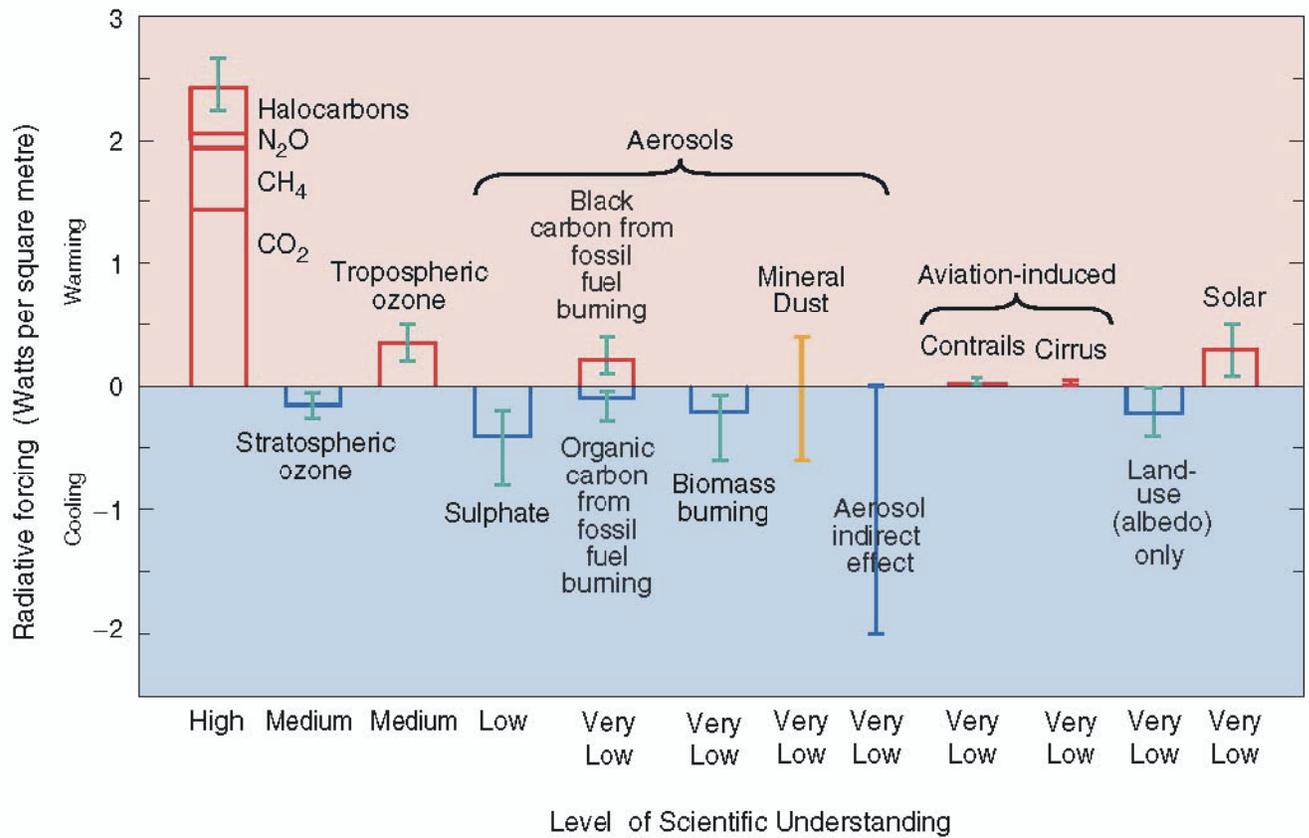


Figure 9-2

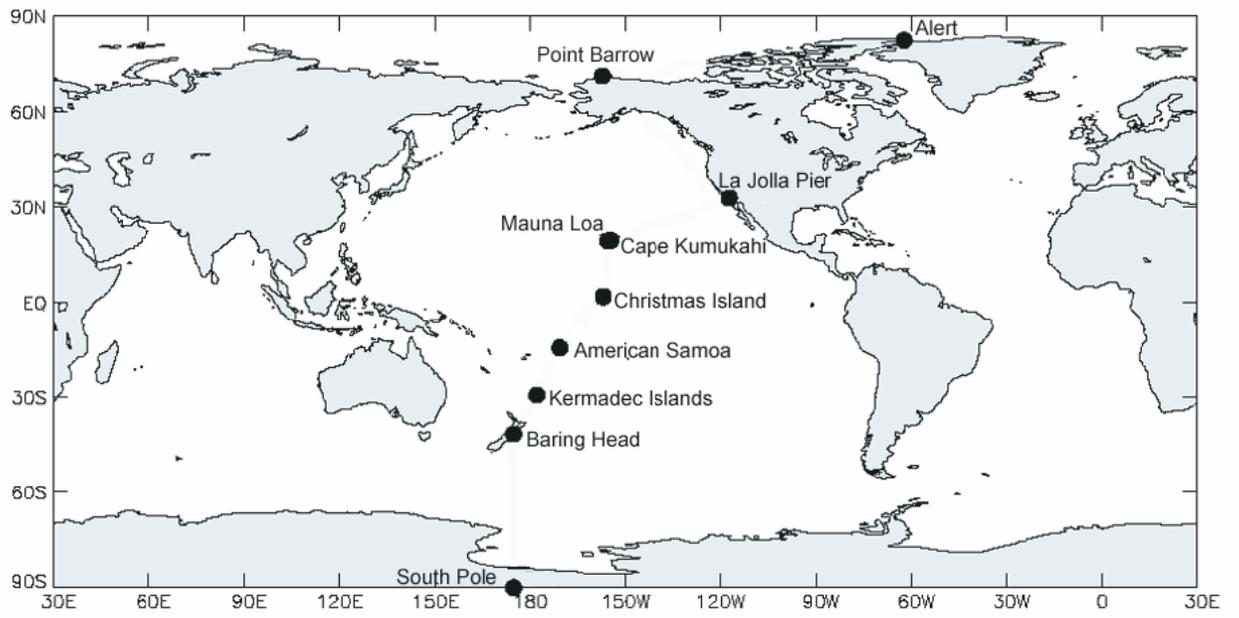


Figure 9-3

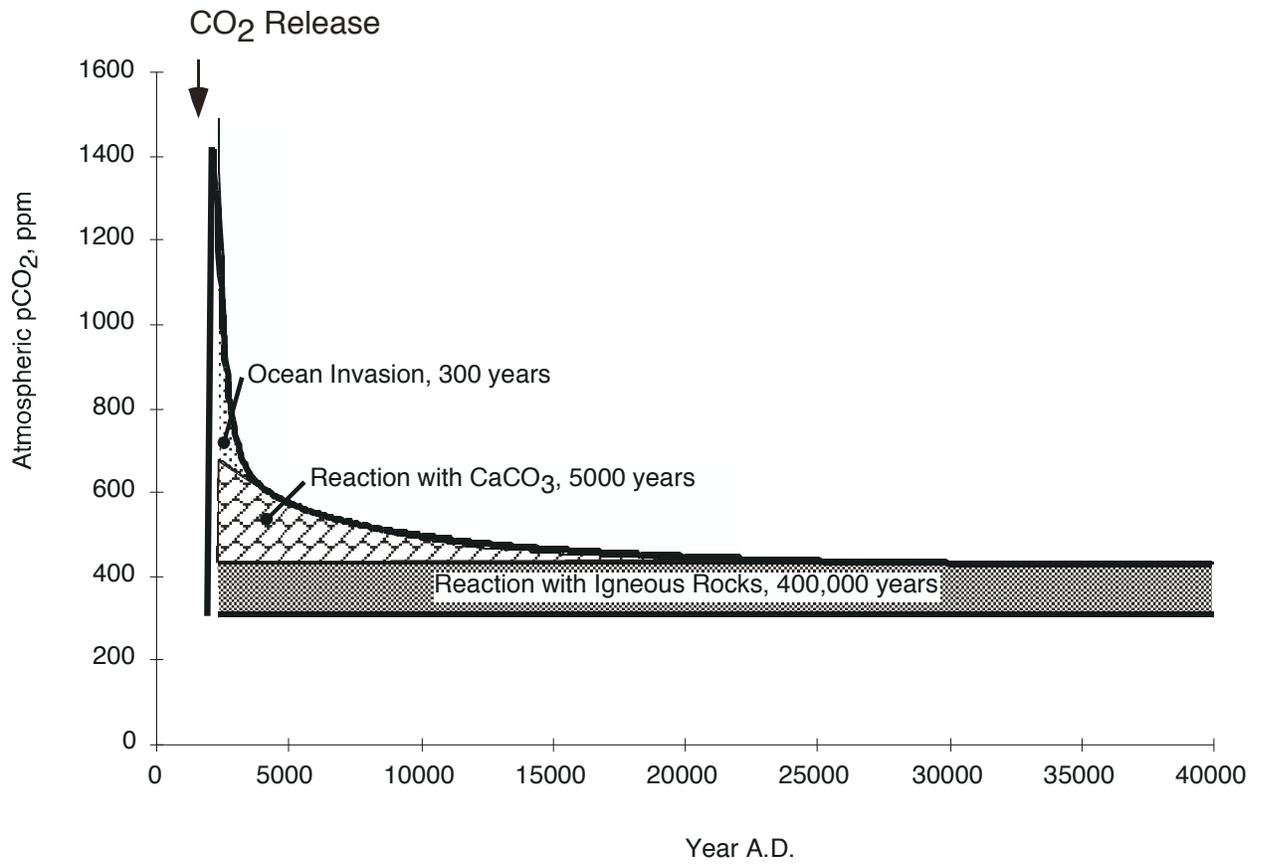


Figure 9-4

Part III. The Forecast

Chapter 10. Is It Reliable?

Abstract

The temperature history of the last century is known from a variety of independent sources which agree with each other pretty well. Reconstructions of the temperature over the last thousand years shows the warming in the late 20th century to be something new, rising faster and higher than has been seen this millennium. Until about 1800, internal climate-system variability and natural forcings such as volcanic eruptions and solar variability are able to account for the reconstructed climate variations. In contrast, natural causes are unable to explain the recent warming; no explanation seems to fit the bill except the rise in greenhouse gas concentrations.

Is the Globe Warming?

How shall we assess the warnings of climate change? We test the forecast. Can it predict past climate changes? Even better, can it make predictions that can afterward be tested? We will begin with the question: How much warming has actually occurred? Answering this question is trickier than one might have thought, because temperatures vary from place to place. The warming signal we are looking for is small compared with the variations in temperature from day to day, season to season. We want to know if the planet surface as a whole is getting warmer, so we will have to average out all the weather. It could be treacherously easy for some bias in the raw data to be undetected or otherwise not corrected for. A computed average temperature for some particular place has to be balanced between daytime and nighttime temperature measurements, for example, and summer versus winter. If we want to compare the average temperature between one year and another, we also have to worry about the possibility that the way temperature is measured might have changed with time.

The Thermometer Record

On land there is a potential source of bias known as the **urban heat island effect**. Sunlight hitting pavement is converted into warmer temperatures, what climatologists call **sensible heat**, “sensible” because you can sense it, like by burning yourself, and you can measure it with a thermometer. Some of the sunlight hitting vegetation drives water to evaporate, and is carried away as latent heat, which you can’t feel or measure with a thermometer until the water condenses again, which it does, eventually, someplace else. So a vegetated landscape doesn’t warm up as much as a paved landscape, because some the heat is escaping invisibly, or insensibly I suppose, as water vapor. This is an observed phenomenon; cities can be 5° C warmer than their surrounding countryside. Nighttimes in particular are warmer in cities. What if we established a weather station in some idyllic rural area, only to have the landscape urbanize over the decades? The

temperature record would show a warming, which would be real for that location, but will confuse you if you try to infer the global average temperature trend from it. The warming in the city is real, but if the data were biased toward cities, the computed average temperature could be higher than the real average temperature of the earth.

IPCC cites three independent studies of the land temperature data, all of which made an attempt to throw out data from locations that might have seen urbanization (Figure 10-1). The studies all agree that removing urban locations has only a small impact on the global average temperature trend that you calculate from the data. It doesn't matter if you throw out the urban data or not, you get pretty much the same answer. The land and sea records are consistent with each other, in that the sea surface temperature record can be used to drive an atmospheric climate model, and the temperature history of the atmospheric model matches the reconstructed land surface temperature record (Figure 10-2). You can see that the largest correction that had to be applied to the temperature records was the bucket correction for the sea surface temperatures, and the abrupt switch in 1942 between bucket and engine-room temperature measurement.

The sea surface temperature, known affectionately as **SST**, used to be routinely measured by heaving a bucket into the ocean to collect surface water, and sticking a thermometer into the bucket on deck. Some of the buckets were made of wood, others of canvas. Evaporation of water stole some heat from the surface of the bucket, cooling the water somewhat. Canvas buckets cooled more than wooden ones, because wet canvas is a poor insulator. After 1942, mariners began measuring the temperature of surface water as it was pumped into the ship for engine cooling water. This is a more accurate measurement, but there is still the possibility of bias. The waters in the surface ocean sometimes get cooler with depth. A large ship might sample colder subsurface waters, which would be systematically cooler than the real SST. These effects must be corrected for, if we are to calculate a global average value of SST through time.

IPCC concluded in 1995 that “the balance of evidence suggests a discernible human influence on global climate”. This was a change from the previous IPCC report in 1990, based on advances in the characterization and modeling of sulfate aerosols in climate. The 1995 report predicted continued tendency toward warmer than average temperatures, and they were right. 1995 was the warmest year on record at the time, but that record has been matched or broken by 6 of the 9 years since then (Figure 10-2).

Box: Consensus in Science. *All three studies agree; what does that mean? You can imagine there might be a lot of judgment calls, trying to eliminate urbanization from the huge morass of all the weather measurements made in the last hundred years. There is always the possibility of prejudice on the part of the researcher, who may allow his or her judgment on a myriad of minor issues to be swayed by an expected result. Skeptics also accuse the scientific establishment of encouraging the creation of false spurious consensus, by making it easier for the scientist who reinforces the consensus view to get grants funded and papers published.*

I would argue that, in my experience, the way a scientist becomes famous (and let's face it, we're not doing science for the money) is to successfully challenge the consensus

view. Shoot the moon. The person writing the second paper would really like to find something that the first paper did not report. It is easier to get a boring consensus paper published; that is true. Extraordinary claims require extraordinary evidence. But, on balance, my opinion is that the structure of the scientific establishment as it is practiced in university and government laboratories around the world encourages the challenger to the consensus view, rather than stifling dissent unproductively. I therefore conclude that in general, replicate scientific studies do bolster a scientific conclusion. Michael Crichton in his fiction work "Climate of Fear" argues the opposite. You will have to decide for yourself.

Temperatures Measured from Satellites

Satellites provide another source of global temperature trend information. Because of the way that electrons are distributed around the nuclei of oxygen molecules, the rotation of oxygen molecules emits and absorbs energy in the microwave range of light. This is longer wavelength, lower energy light than the infrared light which greenhouse gases absorb and emit, and the primary mechanism for microwave emission is rotation rather than vibration. However, the intensity of emitted light increases with increasing temperature, just as the intensity of IR light goes up with warming. The satellites measure microwave emission in a range of wavelengths, and from these measurements construct temperature estimates for several regions of the atmosphere. (See <http://mtp.jpl.nasa.gov/intro/intro.html>). The lowermost region of the atmosphere from which the satellite estimates the temperature spans from the surface to about 8 km altitude, with an exponential decay, rather similar in shape to the curve describing pressure with altitude (Figure 4-2).

These satellite temperature estimates have been the subject of considerable scientific and political discussion because they disagreed with the surface temperature reconstructions from thermometers over the past few decades, until recently (solid black line in Figure 10-3). Just in the last few years the satellite temperature estimate shot up to where it now agrees with the thermometers, but the discrepancy for the 1990's still needs to be explained. In some measure, the discrepancy has shrunk as the satellite data has been corrected for effects such as the decay of the orbit of the satellite. The satellite record after the corrections is consistent with tropospheric temperatures measured by weather balloons, called **radiosondes**. The IPCC and NRC conclusions are that the satellites and balloons are accurately measuring the temperature of the middle of the troposphere, and the thermometers are accurately measuring the temperature of the surface, and that these temperatures are sometimes legitimately different. Models do not show this; the lapse rate sticks closer to the moist adiabat (Chapter 4). Climate skeptic Bjorn Lomborg wonders what impact the cooler-than-predicted tropospheric temperatures will have on the water-vapor feedback. It's a legitimate question, and a bit of a mystery.

Glaciers

The most visual indication of climate change is the melting of mountain glaciers. With the exception of a few locations in Norway and New Zealand, the vast majority of

the mountain glaciers of the world are melting back at astonishing rates. **Figure 10-4** shows records of shortening of 20 glaciers around the world. Many of the glaciers have been melting back since the 1700's, as a result of natural warming since the Little Ice Age. However, melting has accelerated in the past decades. The snows of Kilimanjaro are forecast to melt entirely by 2030, as are all the glaciers in Glacier National Park in the United States. Note to self: Take my children to Glacier before it's too late.

Temperature Reconstructions for the Deeper Past

We would like to extend the temperature record further back in time, so that we can assess the natural variability of the climate system. The thermometer and satellite temperature records indicate a warming trend since about 1970, but this begs the question: is this warming normal, or is it something new? Temperatures from before the days of thermometers can be estimated by a variety of tricks, measurements of something we can measure, interpreted as indicators of things we would like to know. These alternative measurements are known as **proxy** measurements.

One source of proxy temperature measurements is **tree ring** thickness. Trees grow more quickly, on average, during warm years than cold. In places where winter is much colder than summer, the growth rate difference can be seen by eye as tree rings. (There are no tropical tree rings because there is not much of a seasonal cycle there.) Trees also grow more quickly during warmer years than colder years, so it makes sense that tree ring thickness might carry information about temperature. Long records of tree ring thicknesses can be assembled from collections of many individual trees, by matching the patterns of widths to fit the different trees together in time. The widths of the rings must then be **calibrated** against temperature changes from part of the record where the temperature is known independently. As you can imagine, this gets dicey, because trees grow or don't grow in response to other environmental variables, such as the availability of water, nutrients, and even atmospheric CO₂. Young trees grow more vigorously than old ones. Temperatures vary from place to place, and so calculating a global average from local data introduces another possibility of bias. One of the earlier tree-ring reconstructions showed almost no natural variability between 1000 A.D. and about 1900, followed by a sudden twist toward higher temperatures (the light blue line on **Figure 10-5**). This reconstruction has been dubbed the "**hockey stick**", and it reached t-shirt grade celebrity among the climate change community. Since then, other statistical methods and new data have resulted in somewhat more natural variability throughout the natural part of the past millennium. The long-term temperature swings include the **Little Ice Age**, between about 1500 and 1800, and the **Medieval Warm** period, from 1100 to 1400. There are more trees in the Northern hemisphere, and many of these temperature reconstructions are presented as averages of the Northern hemisphere only. This is significant because it is unclear how global the footprints of the Little Ice Age and the Medieval Warm periods were.

Other proxy measurements for temperature include the temperature as a function of depth in the earth and in ice cores, called **borehole temperatures**. The temperature goes up with depth in the earth or in the ice, because heat is being conducted from the interior

of the earth to the surface. The steeper the temperature change with depth, the more heat is transported. If the temperature at the surface remained constant for a long time, and if the rate of heat flow from the earth's interior were constant, then the temperature profile with depth in the earth or the ice would be a straight, linear increase. If the temperature at the surface varies with time, then the temperature profile will reflect this (Figure 10-6). The surface variation will conduct down into the earth, eventually to be smoothed out to nothing. Borehole temperature records lose the little wiggles in the real temperature record, and memory fades with time so the uncertainty gets greater. Borehole temperatures show warming throughout the past 500 years, consistent with the tree-ring record.

Much of the discussion about interpreting proxy information such as tree rings and borehole temperatures involves the machinery of how the raw data are converted into temperature estimates. The real data contains patterns of variability analogous to the different wavelengths of light that can be decomposed into a spectrum. The different patterns of variability in the proxy data are called **eigenvectors**. Much of the art, and the contention, of reconstructing temperature from noisy proxy data involves choosing which of these patterns in the data, which eigenvectors, are really meaningful, and which are simply noise. Much of the evolution in the tree-ring records over time has involved statistical treatment issues of this kind. Borehole estimates of temperature also depend on how the information is boiled down into estimates of temperature.

Other sources of information about past climates come from ice sheets and mountain glaciers. Glacier lengths have recently been used to construct a third independent “hockey stick” figure like Figure 10-6, in substantial agreement with the results from tree rings and boreholes.

Snow accumulate rate tells us something about temperature, because cold air carries less water vapor, and so snows less. We can also read about temperature in the ratios of different **isotopes** in the water. Isotopes of an element have the same numbers of electrons and protons as each other, and hence pretty much the same chemistry. That’s why they are both the same element. The isotopic sisters have different numbers of neutrons, however, so they have different masses. A water molecule comprised of a heavier isotope of oxygen or hydrogen will tend to rain out earlier than would a water molecule made of the lighter isotope. As the temperature over an ice field decreases, the water vapor that remains in the atmosphere to snow down on the ice becomes increasingly depleted in the heavy isotopes in the water vapor. For this reason, there is a systematic relationship today between the temperature and the relative abundances of the heavy and light isotopes of oxygen and hydrogen in water, and the temperature at which the snow fell. Isotopic data are also affected by conditions where the water vapor originally evaporated, and along its trajectory toward the accumulating snow field.

What we can conclude from eyeballing the raw data is that climate has varied in the past millennium, but that there is an ongoing temperature rise in the past centuries unlike any that has occurred in 1000 years. The next questions are called **detection and attribution** of anthropogenic climate change. The ironclad method would be to set up

two Earths, add CO₂ to one of them, and compare their temperature evolutions. Short of that, we can wait for natural experiments like volcanic eruptions, to test our ideas about climate sensitivity. But for CO₂ proper, the first thing to try is to compare data with models. Does rising CO₂ make sense as an explanation for warming? Is it a better explanation than natural variability or natural forcings?

Past Forcings

The first step is to ask about natural processes that drive the climate, natural climate **forcings**. Different climate forcings can be compared with each other in terms of their energy impact on the earth in watts/m². One W/m² in solar output, for example, is roughly equivalent to one W/m² of infrared light trapped by CO₂, or at least these are about the same in models. We have already discussed the climate sensitivity to doubling CO₂, called ΔT_{2x} . A middle-of-the-road estimate for ΔT_{2x} is 3° C. If we wish to think about other forcings than CO₂, we can express the climate sensitivity as temperature change per W/m² of heat forcing. A typical estimate $\Delta T/(W/m^2)$ would be 0.75° / (W/m²).

Solar intensity varies by 0.2-0.5 W/m² over the 11-year sunspot cycle (Figure 10-7). Solar intensity in the past is estimated by the accumulation rate of the isotope of beryllium that has 10 protons + neutrons in its nucleus, called ¹⁰Be. The idea is that a brighter sun is a better shield to cosmic rays. Cosmic rays produce ¹⁰Be. So ¹⁰Be tells us about the strength of the sun. Who made that story up, you may wonder. The story would have no credibility by itself until it is tested, which was done in this case using the most recent data, where both ¹⁰Be and solar intensity are known and can be seen to correlate nicely enough. The time-history of solar forcing variability is to drift up and down on timescales of about 100 years.

Volcanic eruptions inject particles into the atmosphere, sulfate aerosols and dust. Particles injected into the troposphere last a few weeks before they are rained out. It doesn't rain in the stratosphere, so particles spend several years floating around before they eventually sink out. The Mt. Pinatubo eruption in 1991 cooled the planet by 0.5° for several years, a result consistent with a middle-of-the-road climate sensitivity ΔT_{2x} of about 3° C, by the way. Volcanic climate forcings look like randomly spaced spikes of random intensity in Figure 10-7.

To the two natural climate forcings we add two anthropogenic ones, greenhouse gases and aerosols. Neither was important before about 1750. Greenhouse gases warm, and sulfate aerosols cool. Historically, greenhouse gases were added to models first, then sulfate aerosols. Models with greenhouse gas forcing only tend to overpredict the observed warming. Emission of sulfur dioxide from burning dirty coal is masking some of the warming effect of the CO₂ emission from that same coal. A fundamental difference between aerosol and CO₂ induced climate change is that aerosols are cleaned out of the atmosphere in a few weeks, while CO₂ accumulates.

Model Data Comparison

Now that we have some handle on the variability in the forcing that the real climate was subjected to, we can compare the real climate record with model behavior. The most obvious target to try to match on the millennial timescale is the Little Ice Age and the Medieval Warm periods. There is some uncertainty about the amount of temperature change, and the geographical extent, of these events. The solar forcing between 1400-1800 A.D. was a bit on the low side, relative to the present-day and to medieval times. There is quite a bit of uncertainty on both sides of the question, the data, and the forcing for the model, so it's difficult to be definitive on whether the models are successful or not at reproducing the data. If anything, the models may underestimate the natural variability on these time scales. This could be the result of some internal mode of variability that the real climate possesses, but which the models do not do so well, probably having to do with the ocean (which circulates on a time scale of 1000 years). There is some evidence for a 1500-year oscillation in climate records from the glacial time tens of thousands of years ago, which models in general have difficulty reproducing. Alternatively, it could be that we underestimate the long-term variability of solar energy output. The little ice age in particular coincided with a time when there were no sunspots at all for hundreds of years, called the Maunder Minimum.

We are on solid ground comparing models with data for the past century. We can first ask, are the models too tippy, compared to the real world, or perhaps not tippy enough? We can't make a prediction about what happens to a canoe in the rapids if our model is a rowboat. [Figure 10-9](#) compares the unforced variability of the model temperatures with the unforced variability of the real climate. [Figure 10-9](#) is a spectrum, just like the light spectra we looked at earlier, showing the amplitude (for light, that's the brightness) at each different wavelength. The black line is the spectrum of the real climate, with climate effects of volcanic eruptions and things like that taken out. (One might reasonably worry about bias or errors here, by the way. This sounds like a tricky correction.) Some models have higher variability than the real world, and some less. Take a look at the scale on the y-axis. It is logarithmic, so it kind of hides the fact that the model variability spans a factor of 10. Models vary quite a bit in their responses, some apparently more than reality and some less. However, the point is that none of the models are able to explain the recent uptick in temperature as a reasonably likely random event, forced by nothing other than the internal mood of climate. Compare unforced evolution of several models with the recent observed warming in [Figure 10-10](#). The models do not typically spike like that by themselves (long model run on the left, and recent data on the right). Since they are as tippy as the real climate, more or less, we suppose that the real climate wouldn't go off on a spike like that either.

To sum up: The temperature of the earth is warmer than it has been in a millennium. Natural variability of the models is similar to that of the unforced real world, and is unlikely to be the explanation for the recent temperature rise. Among the forcings, the natural ones are unlikely to explain the uptick ([Figure 10-11](#)). Neither can anthropogenic forcings explain the entire 20th century temperature record. A temperature rise between 1900 and 1940 can be explained mostly by an increase in the intensity of the sun. The

temperature rise since 1960 requires greenhouse gases as an explanation. Note that the various model runs on this plot are different runs of the same model, what we called in Chapter 5 an ensemble of model runs, not results from a bunch of different models. However, the computational “experiment” has been redone using other models, however, with substantially the same result. (I’ll remake the plot when I get to it.) The Hadley Centre model, shown in [Figure 10-11](#), has a climate sensitivity ΔT_{2x} that is in the middle of the range of climate model results, neither notably high nor extremely low. The simplest explanation for the multiply-reconstructed temperature record is that models have the temperature sensitivity to energy forcing about right, with no special immunity to CO₂ forcing as opposed to natural sources of variability.

Vostok Record

The final bit of evidence in support of the role of CO₂ as an effective greenhouse gas must be the history of CO₂ concentration and temperature from the Vostok ice core in Antarctica. We already looked at this record in [Figure 7-3](#), and so we won’t reprint it in this chapter. The record goes back 400 kyr, through four glacial / interglacial cycles. The temperature record in [Figure 7-3](#) is the local temperature in Antarctica, derived from isotopes of oxygen and hydrogen in the water. The relevance of this record to our discussion of the global warming forecast is the tight correlation between CO₂ and Antarctic temperature. Every little wiggle lines up between the two records.

Take-home Points

The past few years have been warmer than the thermometer records of the past 140 years, and warmer than reconstructed records over the past 1000 years.

Climate varies naturally on all different time scales, even under steady forcing. Models capture this natural variability, more or less, and lead us to conclude that the recent warming is unlikely to be caused by natural variability

Climate is also driven by natural and human-induced changes in radiative forcing, including volcanic eruptions, solar variability, greenhouse gases, and human-released aerosols. Models cannot explain the recent warming without including anthropogenic forcing.

Lab

Point your web browser to <http://geosci.uchicago.edu/~archer/PS134/lab.hadley.html>. In this page we are browsing through temperature output data from the Hadley Centre coupled climate model. The Hadley Center is the climate-modeling arm of the UKMO (United Kingdom Met. Office), located in Reading, England. It is one of the world’s major general-circulation modeling groups, and has played a particularly important role in the conclusions of the current IPCC report. This model is widely regarded as having perhaps the most sophisticated and accurate physical parameterizations. The web page above has pointers to movies of model output and a data browser.

1. Data Reduction.

(a) **Seasonal cycle.** First look at movies of temperature output every month for 5 years beginning in 1860, 1995, and 2095. Most of what you see is the seasonal cycle with year-to-year variability. By eye, try to see the temperature change between the beginning and the end of the simulation.

(b) **Climatology.** Each of these movies has 12 frames, one for each month of the year. The frame for each month contains the average of 20 years of the temperature during that month in the model. This is called a climatology. There is a climatology movie for the beginning, present, and end of the simulation. Describe how the movies look different from the raw movies in 1a. Is it easier or more difficult to detect warming between the beginning and end of the simulation, than with the raw data in 1a?

(c). **Trends.** Now let's get rid of the effects of the seasonal cycle in our movies by looking at a series of Augusts or Februaries from the beginning to the end of the model run. We are seeing variation from year to year, but no seasonal cycle. Is it easier or more difficult to detect a warming trend than for the other types of movies?

(d). **Anomalies.** When we look at a map of temperature for a month, say, August, most of the color variation we see comes from the warm equator and cold poles. Eyeballing global warming on such a map involves trying to find warmings of maybe only 1°C (needle), on a map that has 50°C of variation on it (haystack). We can subtract out the haystack, leaving only the needle, by subtracting the climatology from the beginning of the run from the values throughout the run. Then if temperature never changed, the anomaly would be 0°C. Where is the warming trend most pronounced, in August and February?

2. Statistics. We also have a data browser, with pointer on the web page. Select a map of temperature from the simulation, as raw data or climatology, and point and click on the map to see the temperature values.

(a) Bring up the raw temperature map for the first August or February in the simulation. Choose some location of interest to you. The world is your oyster. Click on that location find the temperature there, and record it in a table with columns "year" and "T". Do this, for the same location, for the next 9 images, giving us 20 data points.

(b) Compute the mean of the 20 data points as

$$\frac{\sum T}{n}$$

where the capital sigma (Σ) means "take the sum of", in this case of all 20 temperature values, and n is 20.

(c) Add a new column to your table, and fill it with the difference between the temperature each year and the mean temperature. This is the deviation or anomaly.

(d) Add a fourth column that has the square of the deviation.

(e) Compute the mean of the square of the deviation.

(f) Take the square root of the mean of the square of the deviation. This quantity is called the root mean square or RMS of the deviation, also called the standard deviation, abbreviated as σ . Statistically, about 2/3 of the temperatures in the table should be within $\pm 1 \sigma$. Is this true in your case? 95% of the numbers should be within $\pm 2 \sigma$. How many observations ought to be outside of 2σ ? How many do you find?

(g) Move to the last 20 years of the simulation, same month. Start a new table and record 20 data points from the same location. If the new mean value were the same as the old mean value, you would expect 5% or 1 in 20 of the data points to be outside of 2σ . How many of the data points actually are outside of 2σ ?

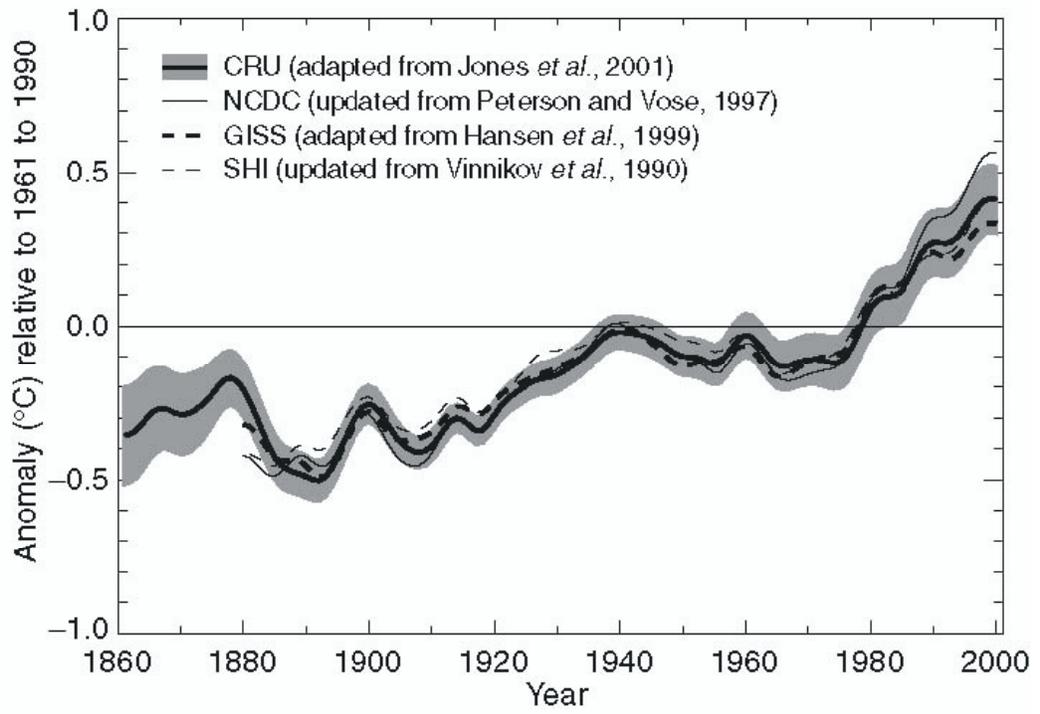


Figure 10-1

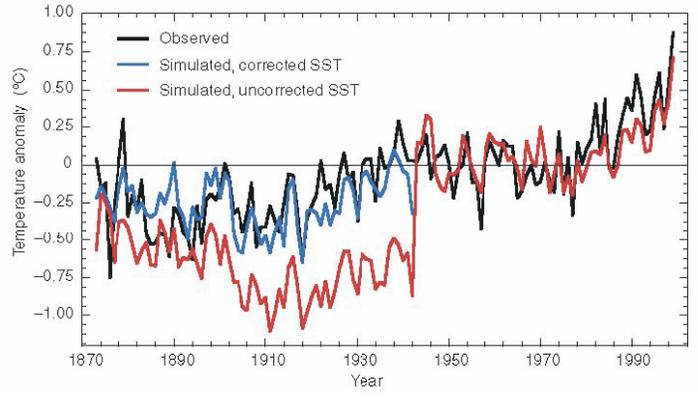
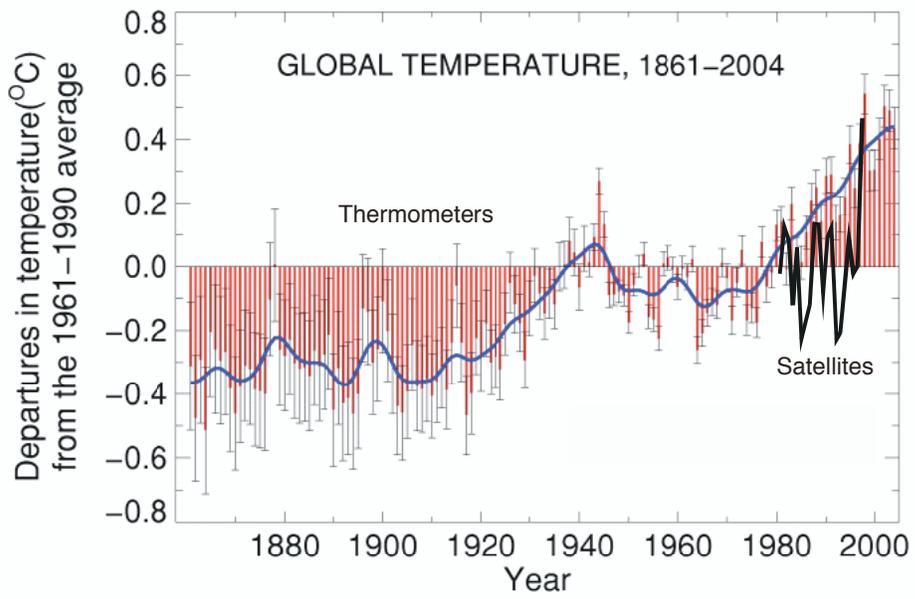


Figure 10-2



MSU Satellite 2LT.D (from Christy *et al.*, 2000)

Figure 10-3

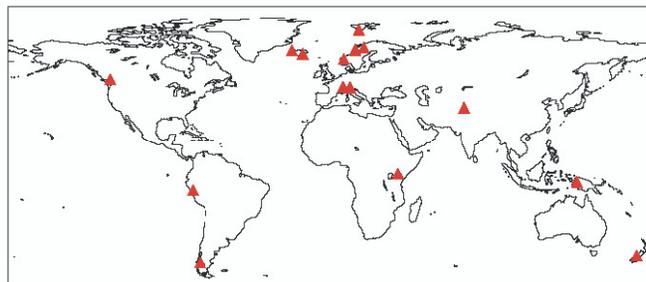
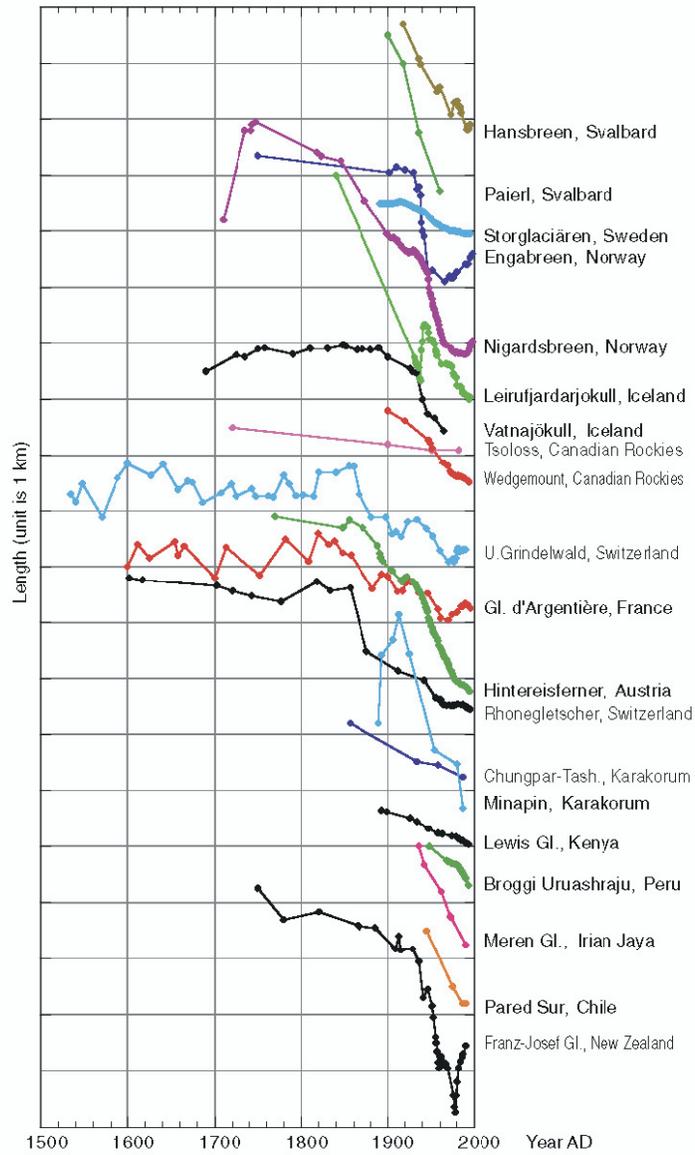


Figure 10-4

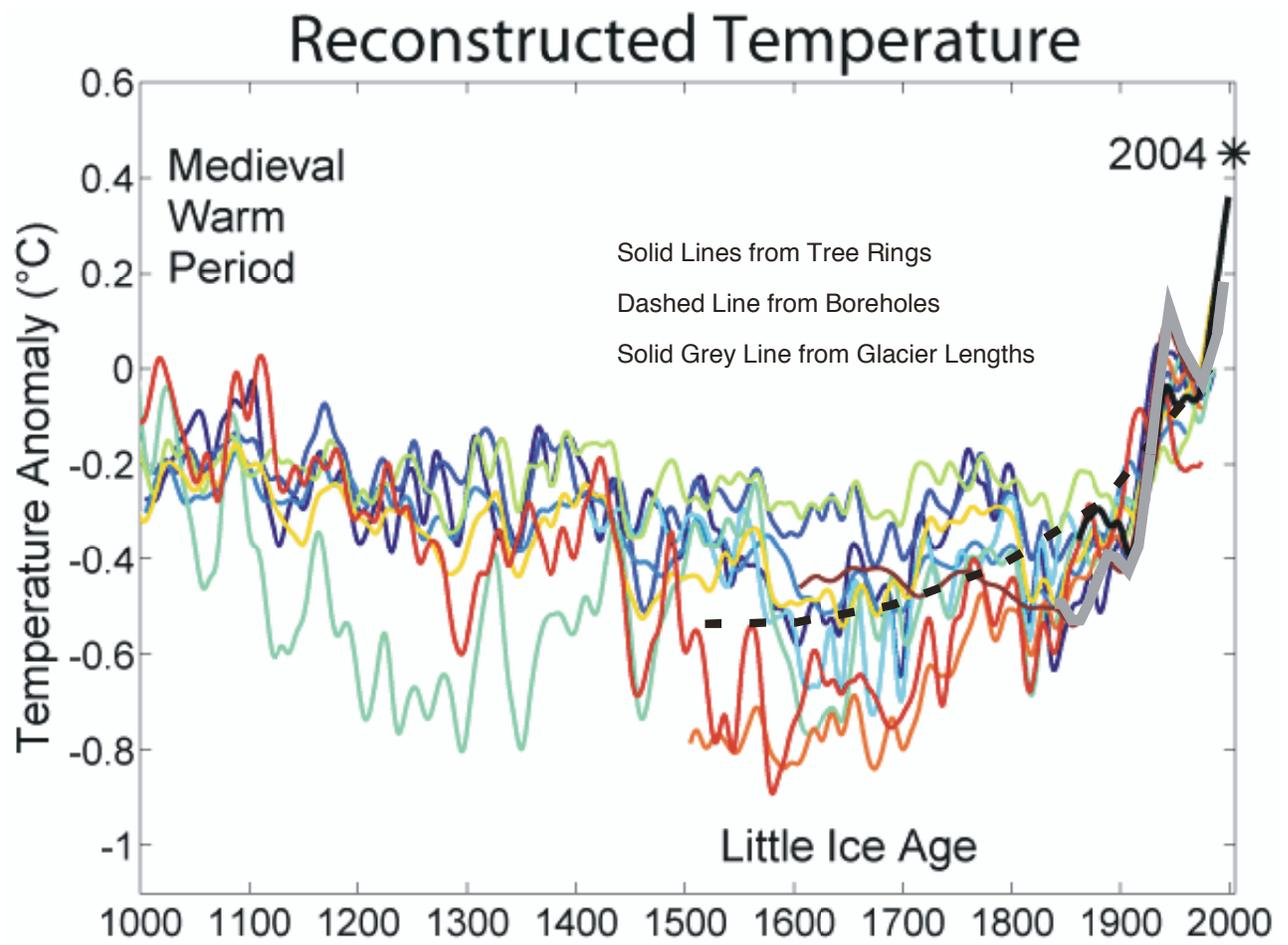


Figure 10-5

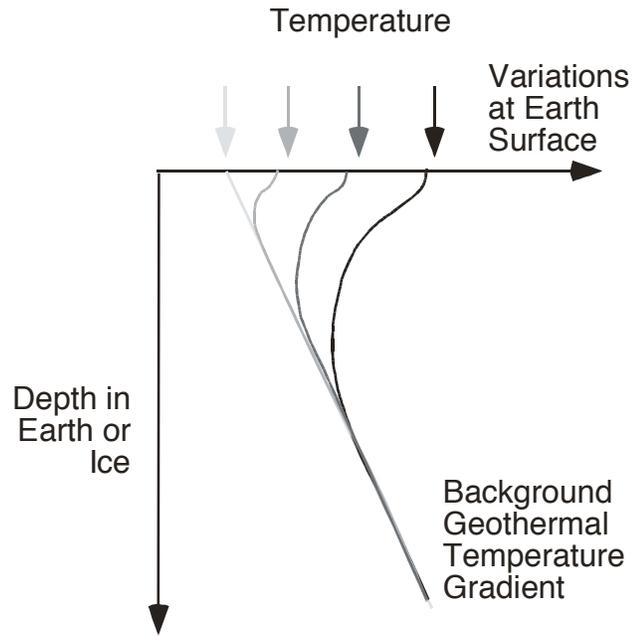


Figure 10-6

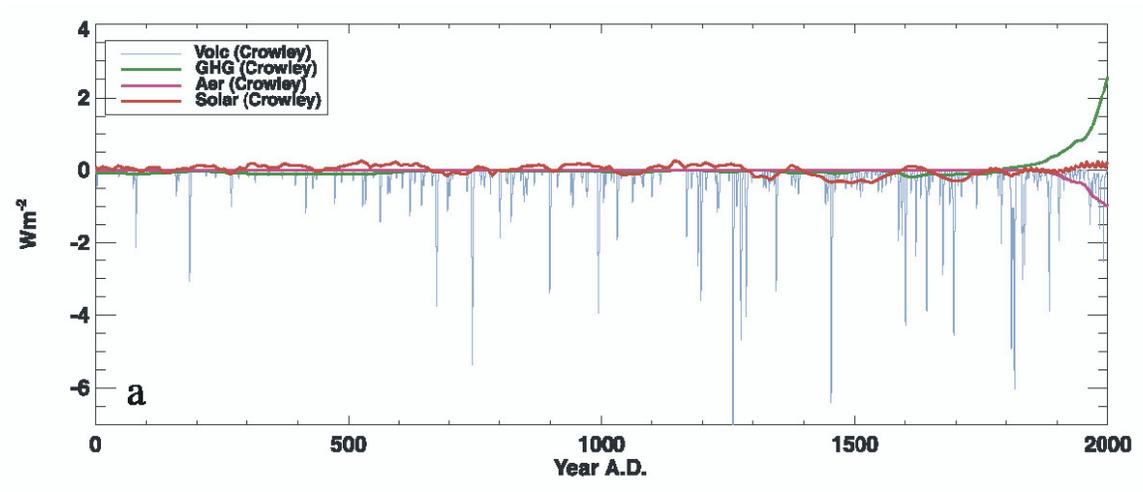


Figure 10-7

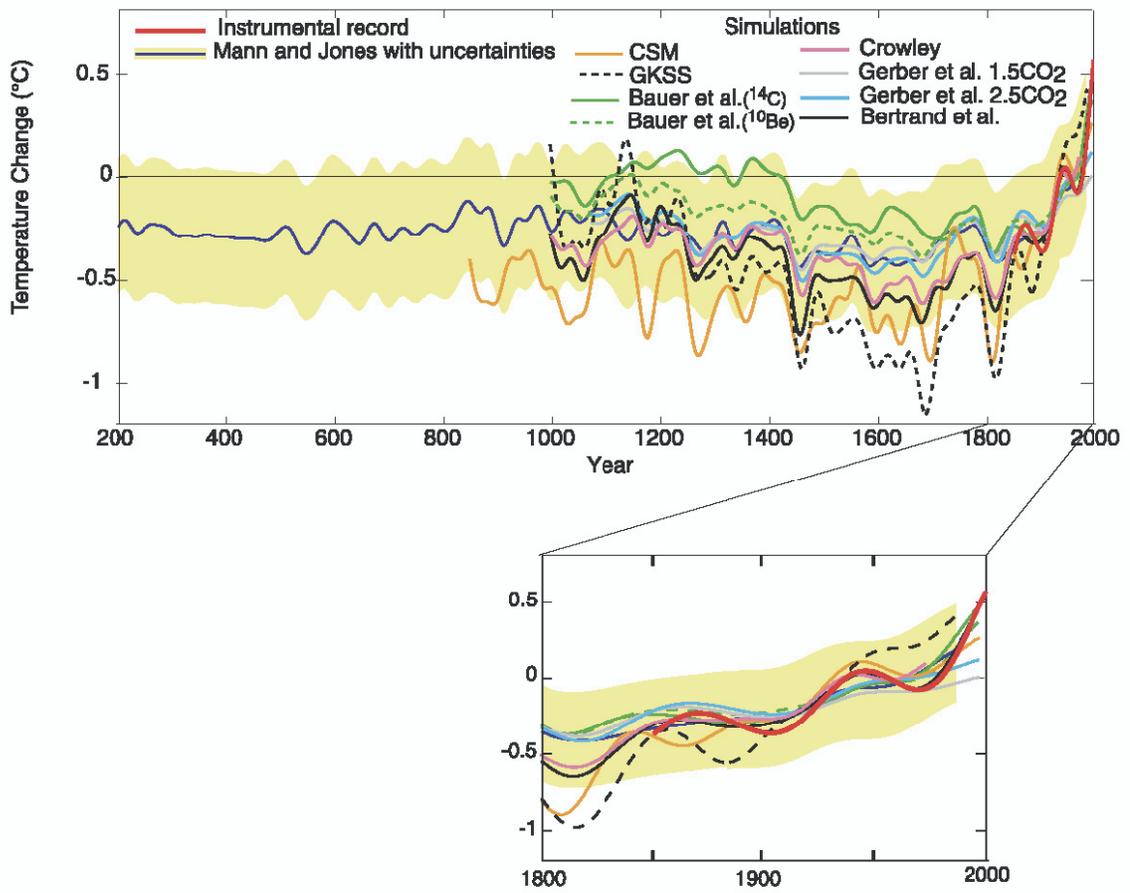


Figure 10-8

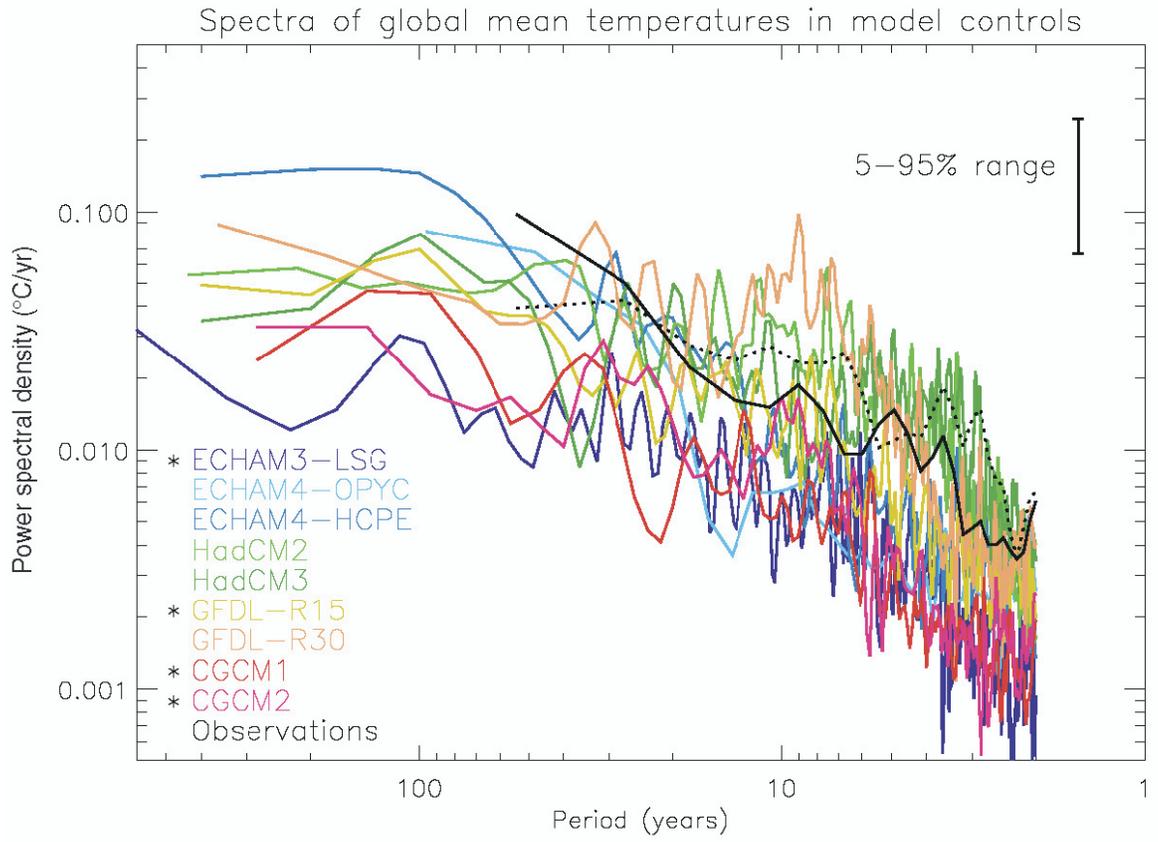


Figure 10-9

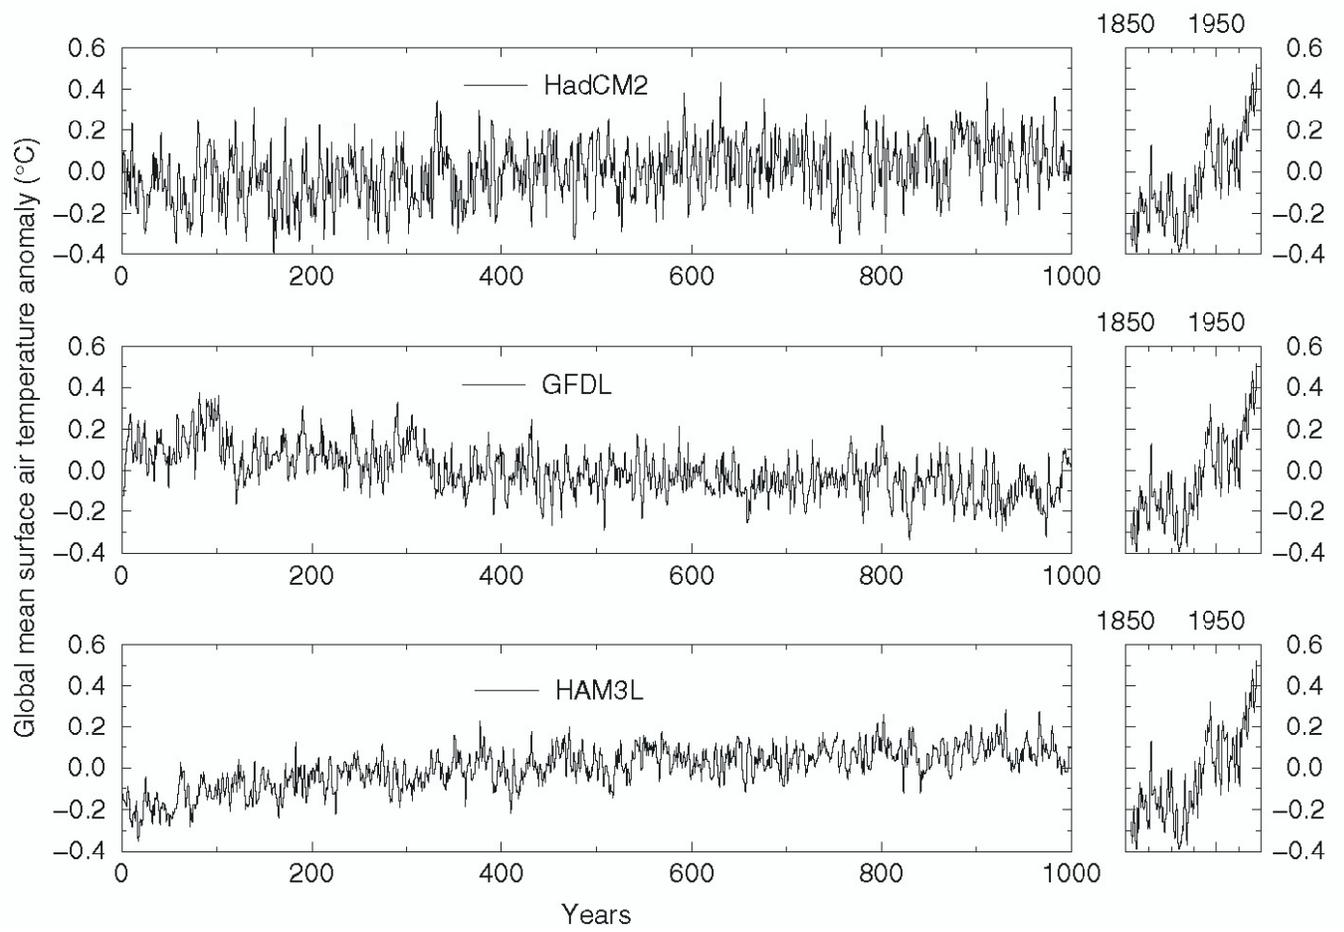


Figure 10-10

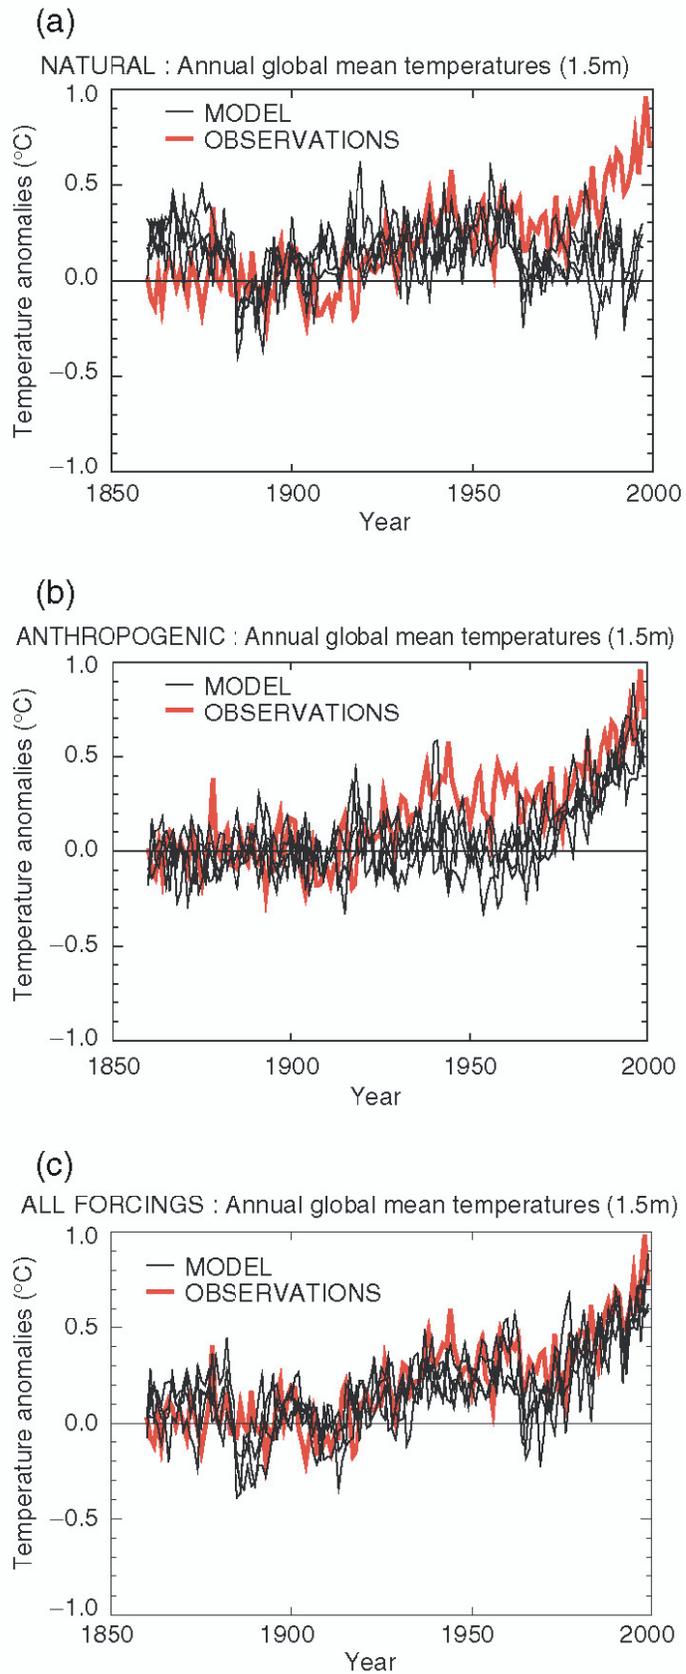


Figure 10-11

Chapter 11. The Forecast

Abstract

There are three sources of uncertainty in the forecast of global mean temperature in the year 2100 (an arbitrary but widely-used benchmark year). One is the equilibrium climate sensitivity, ΔT_{2x} of the real climate. Another is the amount of time-lag in the real climate as the oceans slowly warm and as various feedbacks such as with water vapor come to equilibrium. The third is the amount of greenhouse gas that will be released. In general, we expect 2-6° C warming by 2100 if mankind follows the Business-as-Usual scenario. This can be compared with the socially significant ~1° C temperature variations through the last 1000 years, or the 6° C mean temperature change of massive glacial / interglacial climate cycle.

Rainfall is expected to increase overall, and increase in equatorial and high-latitude regions. The subsidence regions at 30° N and S are expected to dry, as are continental interiors.

Storms are widely feared to increase, but there is not much solid evidence that this is happening already. A 50-year increase in insurance payouts and damages from various forms of weather-related natural disasters is provocative but not well attributed.

The IPCC scenarios can be viewed as best-case scenarios, in that they are smooth with no unexpected surprises. The glacial climate was subjected to savage climate swings which we of the Holocene have largely been spared. It is difficult to predict whether greenhouse warming could push our climate into some rearrangement such as these abrupt climate changes in the past.

Temperature Forecast

We will begin with estimates of the climate sensitivity, ΔT_{2x} . The IPCC Third Assessment Report (2001) reports that the mean of the climate sensitivity, ΔT_{2x} , of 15 climate models, is 3.5° C. The **standard deviation** of the climate sensitivities, denoted by the greek letter sigma (σ), is 0.9° C. The meaning of the standard deviation of a variable quantity like ΔT_{2x} is that 63% of the ΔT_{2x} estimates should fall within 1 σ of the mean (in this case $\pm 0.9^\circ$ C), and 95% of the estimates would fall with 2 σ ($\pm 1.8^\circ$ C). The statistics imply that with 95% certainty, a new estimate of ΔT_{2x} , such as from some new model, should fall between 1.9° and 4.1° C. If we take a really wild leap and assume that the real climate can be represented by the models, then this would be a 95% confidence range for the real climate as well. Of course, there could be systematic biases between models and the real climate, but perhaps we alleviated some of our concern about this possibility by comparing climate models with real climate data, in Chapter 10. The range in model results is taken as a measure of the uncertainty in the real climate

forecast, which makes sense in that this is the range of answers we get. A true uncertainty analysis would be to take the uncertainties in all the bits and pieces that go into a climate forecast, cloud droplet formation processes, mixing rates in the ocean, and everything, and run a climate model many times using all different sets of these input parameters. There is a project to do this sort of analysis using donated “screen-saver” computer time. The URL for this project is www.climateprediction.net. Perhaps you would like to contribute to the effort yourself.

Uncertainty in ΔT_{2x} is only the beginning of the uncertainty in the temperature forecast for the coming century. ΔT_{2x} gauges the equilibrium response to doubling CO_2 , but it takes quite some time for the climate system to reach equilibrium, because it takes a long time to change the temperature of the ocean. Predicting the temperature in the 2100 requires modeling what is called the **transient response**. IPCC uses a standard benchmark for comparing the transient responses of models, which they call Transient Climate Response (**TCR**) and define as the model temperature at doubled CO_2 concentration when the CO_2 has been rising at a rate of 1% per year (**Figure 11-1**). Real atmospheric $p\text{CO}_2$ itself is not rising at 1% per year, but the idea is to raise CO_2 a bit on the fast side, to account for the greenhouse forcing from methane and other greenhouse gases. Even granting this, 1%/year is “on the high side” (IPCC quote) of the rate of change of radiative forcing today. Bjorn Lomborg, author of “A skeptical environmentalist” complains that a comparison of a given year, say 2050, in such a simulation will be warmer than the real 2050, because in the real 2050 the radiative forcing won’t be as high as that. TCR is more appropriate to compare with the radiation = doubled CO_2 year, whatever year that turns out to be. The average TCR values from 20 models is 1.8°C with a standard deviation of 0.4°C .

The intensity of future climate change also depends on how much carbon is released. The IPCC scenario release estimates for the year 2100 span between about 6 Gton C / yr for the slow-growth scenario C to 26 Gton C / yr for the fast-growth scenario F (Chapter 9). The uncertainty in our forecast for the temperature in the year 2100 derives from two sources, which are about equal. One is the uncertainty in model temperature response to a given amount of CO_2 , and the other is uncertainty in what our CO_2 emissions will be in the future.

Transient climate runs forced by the presumably more realistic IPCC business-as-usual scenario predict temperatures 2-5°C warmer by the year 2100 (**Figure 11-2**). A few degrees C does not sound like very much, given that the daily cycle of temperature is greater than that, to say nothing of the seasonal cycle of temperature. One point of comparison to this figure is the amount of cooling during the last ice age, which is estimated to be about 6°C globally. Another is to the temperature changes during the definitely noticeable Little Ice Age and Medieval Warm periods, cooling or warming by perhaps 0.5 to 1°C . These climate intervals were not the end of the world, but they definitely sufficed to rearrange civilizations. A third point of reference is that we are currently 1°C warmer than the little ice age several hundred years ago.

The distribution of the forecast temperature change is not uniform geographically or in time (Figure 11-2). The high latitudes warm more than low latitudes, by a factor of 3 or 4, because of the ice-albedo feedback. Land tends to warm more than water, because evaporation is always an option to carry away heat from the water, but the land may dry out. Wintertime temperatures, and nighttime temperatures, that is to say the minimum or coldest temperatures, tend to warm more than summer daytime temperatures. There will be more days of extreme heat, and fewer days of extreme cold. Projections of mortality from heat waves do show an increase in heat-related death, but these projections are offset by projected decrease in cold-related deaths.

The warming due to CO₂ is offset by cooling by sulfate aerosols, but the distribution of the temperature change from aerosols is not the same. The CO₂ radiative effect is pretty much everywhere the same, proportional to local temperature, because CO₂ is a well-mixed gas in the atmosphere. Aerosols, on the other hand, last only a few weeks in the atmosphere before they are removed from the atmosphere as acid rain. The cooling effect of the aerosols is therefore concentrated near the sources of their release, predominantly in the industrialized northern hemisphere.

Much of the warming to date and projected for the future is expected to express itself through modes of temperature variation that already exist in the global climate system. Most people have heard of the el Niño climate oscillation, originating in the equatorial Pacific but affecting rainfall and temperatures worldwide. For details about el Niño, ask any cab driver. Changes in the frequency and duration of the warm phase of the oscillation, have already been detected, and the global warmed world is expected to have more warm el Niña states than the cooler preanthropogenic world. Another observed climate oscillation is called the North Atlantic Oscillation and/or the Arctic Oscillation. These cycles last about a decade typically, and have a huge impact on the wintertime climates of Britain and Northern Europe. Much of the warming that has already been detected has the same pattern as the NAO/AO. The climate system rings like a bell, and when it is pushed, the response tends to express itself via these ringing patterns.

The warming from rising greenhouse forcing takes some time; it doesn't happen immediately. This is partly because the ocean stores a lot of heat. A rise in greenhouse gas concentration decreases the outgoing flux of heat, so that there today is a net imbalance of energy, of about 0.75 watt / m². The excess heat is absorbed into the ocean (which is measured to be warming, consistent with this heating term). As the ocean warms, it allows the climate of the earth to warm, increasing the energy flux back to space, until the energy budget of the earth reapproaches a state of balance.

As the temperature of the earth warms, the feedback mechanisms we explored in Chapter 6 begin to act, including the water vapor feedback, and the ice albedo feedback. Their effect on the energy budget is to make it harder for the poor old earth to balance her energy books. An increase in temperature of the earth increases the water vapor concentration, for example, so the outgoing energy does not go up as quickly as it would if there were no feedback. These feedbacks also slow down the approach to the new climate equilibrium.

The amount of time it will take to balance the energy budget depends on two things, therefore. One is the heat uptake by the ocean, and the other is the strength of the feedbacks such as water vapor. One estimate of the equilibration time for climate is about 60 years. If the real feedbacks turn out to be on the strong side, this will have two consequences. One is that the ultimate temperature increase will be greater, and the other is that it will take longer to reach the ultimate temperature increase. The implication of this is a prediction that about 40% of the warming that will occur from the CO₂ already released, what is called **committed warming**, has yet to take place. We have paid for 1° C warming, but we have so far received only 0.6° C.

Rainfall Forecast

Warm air holds more water vapor than cool air, so the global rate of rainfall is expected to increase with warming (Figure 11-2). An increase in rainfall overall sounds like a good thing, in a time when fresh water availability is a problem in many parts of the world. However, for the rainfall forecast the devil is in the distribution, even more than it is for temperature (Figure 11-2). Part of the pattern of rainfall change has the appearance of the Hadley circulation. Air rises convectively at the equator, where the solar heating is most intense. The rising air cools, condensing water into rain. Air subsides in the subtropics, bringing bone-dry air down to the surface and creating areas of minimal rainfall at about 30° latitude North and South. The anthropogenic impact is expected to increase this pattern, intensifying rainfall at the equator and further drying out the desert regions in the subtropics. Meanwhile, precipitation in the highest latitudes is expected to increase because of the increase in temperature (warm air carries more water vapor). This increase in high-latitude precipitation is a major player in the forecast for large ice sheets, especially the one on Antarctica.

Several aspects of the rainfall forecast point toward an increased danger of drought. The dry subsidence areas of the earth, the global desert belts at 30° N and S latitudes, are expected to dry further. Continental interiors are expected to get more than their share of both drying and warming. The variability from one year to the next of the monsoon in India is expected to increase. Today, winter snow pack in mountains supplies summer melt water to surrounding lowland areas, but this will decrease with warming.

Storms and Extreme Weather Events

There are several reasons to look for an increase in storminess and extreme weather events in a warmed world, but the observational evidence for this is patchy and models are not very good at forecasting these trends.

There is a correlation in present water patterns between the average amount of rainfall and the variability in rainfall; more rain begets more hard rain. For this reason, the frequency of extreme rain events is expected to increase with warming. This is one of the more reliable predictions. Rising sea levels will amplify the destructive power of storm surges, in which low pressure in an atmospheric storm pulls up the surface of the ocean. This trend is apparent in data from the real climate, and seems like a strong bet for the future as well.

Rising sea surface temperatures could conceivably fuel an increase in the number or the intensity of hurricanes and tropical storms, because warm surface waters are the source of energy for these storms. However, hurricanes do not appear to be any more frequent than they used to be (Figure 11-4). There is some evidence for an increase in the intensity of storms outside of the tropics, but only in the Northern Hemisphere, not in the Southern Hemisphere. Models do not have a strong statement to make in this regard either. The spacing of the grid points in a gridded climate model does not do justice to the physics that govern storms.

The really provocative observation is a massive increase in the cost of weather-related damages in the past 50 years (Figure 11-5). The number of weather-related events has increased by a factor of five since the 1950's. The number of non-weather related disasters has gone up also, but only about half as quickly as weather-related claims. Insurance payments for weather-related property destruction have increased by a factor of about thirteen in the past 50 years. A trend of rising affluence and coastal construction can explain at least part of the increase in payouts, perhaps all of it. Some of the damage is due to floods, which are probably caused by land-use changes, cutting forests, draining wetlands, and building in flood plains, as much as by climate variability. A number of sources, including Working Group II of the IPCC (the human impacts group), blame the increase in insurance payouts on climate change, but the observational support for that conclusion is pretty thin.

Ice and Sea Level

For our purposes we can divide the frozen water on the planet into several types which behave differently. **Sea ice** forms from seawater freezing at the surface, and may be augmented by snow that lands on it. **Glaciers** are formed from compacted snow on a mountaintop, and flow down mountains. **Ice sheets** are larger than glaciers, and are found today in Greenland and Antarctica. **Ice shelves** are former ice sheets that have flowed out onto water, and are currently floating, still attached to the ice sheet. Glaciers and sea ice are expected to melt with warming, indeed have already begun doing so (Figure 10-4 from Chapter 10). Ice cover in the Arctic has decreased to the point that the mythical Northwest Passage may finally become a reality, at least in summer.

The fate of the major ice sheets depends on the balance between snowfall and melting. For the **Greenland ice sheet**, the air around the base of the ice sheet is warm enough that a temperature increase leads to a significant increase in melting. Warming of about 3°C over Greenland over present-day temperatures is thought to be enough to ultimately melt the Greenland ice sheet. Recall that warming is more intense in high latitudes than the global average warming from increasing CO₂, but also that warming is more intense in winter than in summer. The ice sheet is mostly sensitive to summer temperature, when melting takes place. The question is how quickly this can happen. The most conservative answer is that it would take a thousand years or so. The Greenland ice sheet could increase sea level by up to seven meters if it were to melt.

The **Antarctic ice sheet** is colder than the Greenland ice sheet, and an increase in temperature is not predicted to have as strong an impact as it will for Greenland.

Snowfall is expected to increase in the Antarctic interior, because more water vapor will be carried in the warmer atmosphere. Overall, the Antarctic ice sheet is expected to increase in size over the next century; this is consistent with present-day observations. A piece the Antarctic ice sheet has the potential to melt down catastrophically, however. This is called the **West Antarctic Ice Sheet**. The base of this ice sheet rests on rocks below sea level. Portions of the ice itself are flowing at speeds of hundreds of meters per year in rivers of flow called **ice streams**. There is geological evidence to suspect that the West Antarctic ice sheet may have melted in the past, over time periods of centuries, for example during the transition from glacial to interglacial climate around 12,000 years ago. The general consensus is that the West Antarctic ice sheet is not likely to melt down significantly in the coming century, but if it were to melt, it could increase sea level by about five meters.

Sea level is expected to rise with global warming, for two reasons. One is the thermal expansion of the water in the ocean. It takes a thousand years for water to cycle through the deep ocean, so this is how long it will take to warm the entirety of the ocean. Sea level from thermal expansion will also take this long. Depending on the amount of warming, and the circulation and mixing in the ocean, the IPCC forecast calls for 0.1 to 0.4 meters of sea level rise due to thermal expansion alone by the year 2100. Models suggest that a long-term CO₂ doubling might eventually, after a thousand years or so, raise sea level by 0.5 to 2.0 meters.

The other component of sea level rise is from melting ice, but only ice that was not previously floating. Sea ice and ice shelves, when they float in water, displace their weight in water. This is called **Archimedes' principle**. Archimedes is said to have discovered this property while floating in a bathtub, and was so excited that he ran down the streets of Athens naked yelling "Eureka!" (I found it!). When floating ice melts, its water exactly fills the hole that the floating ice had occupied. Land ice, glaciers and ice sheets, do increase sea level when the melt. The melting of land ice also affects the gravitational field in the area, and astonishingly, the gravitational field has an even stronger effect on local sea level than does the fact that the ocean has more water in it after the ice melts. The melting of an ice sheet actually depresses sea level in its immediate previous vicinity, but raises sea level everywhere else.

The forecast for the next century is for 0.1 to 0.7 m of sea level rise, approximately half due to expansion and half due to melting of glaciers and ice sheets. One meter of sea level rise inundates about 10¹¹ meters of land, about 0.07% of the land area of the earth. The roughly half meter of projected sea level rise is small compared to the range of tidal variation in most places, so if sea level were raised by a half meter today, it's not clear that it would be an immediate disaster right away. The impacts would be focused at particular times, such as storm surges, and particular places, like Bangladesh, the Nile delta (Egypt), the Mississippi delta (New Orleans), New York, Miami, and small islands in the Pacific Ocean.

Sea level at any given location on the earth depends on the total volume of water in the ocean, of course, but also on local geological movement of the ground, up or down.

The largest vertical motions on the surface of the earth today are the result of melting of ice sheets from glacial time. Remove all that weight and the crust will float higher relative to the surrounding crust (Eureka!). There is also vertical motion in other parts of the world which is comparable to the projected rate of sea level rise. Bangladesh is a sad example; the local sinking or subsidence rate of the land there is about 0.5 cm / year, comparable to the projected rate of sea level rise, doubling the inundation that Bangladesh would experience naturally.

It may be possible to defend the land against the invasion of the ocean. Much of Holland is located below sea level, and the ocean is kept at bay with dikes and pumps. This could presumably work for Manhattan and New Orleans as well. For Bangladesh and Egypt, where the coastlines are longer, it may be less tractable. A sea level increase of 1 m could inundate about 10% of the farmland in these countries, potentially displacing millions of subsistence farmers.

There are indications from recent events and from the geologic record to worry a bit that the consensus view about slow ice sheet melting might be overly optimistic. Ice shelves have recently demonstrated a dramatic ability to self-destruct catastrophically. A piece of the **Larsen ice shelf** on the Antarctic peninsula, the size of Rhode Island, was there during one pass of a satellite, and had exploded into what insiders call "blue slurpie" for the next pass of the satellite several days later ([Figure 11-6](#)). This explosion can be explained as the result of meltwater at the surface creating crevasses, canyons in the ice, filled with water, cutting the ice shelf into blocks. If the blocks are taller than they are wide, they will have a tendency to tip sideways, expelling the rest of the blocks out to sea and provoking them to tip as well. Kablooey.

During the transition from glacial to interglacial climate 14,000 years ago, sea level rose 5 cm / year, 10 times faster than our forecast, for 400 years. It is not clear exactly where this water came from, it could have been real-time melting or it could have been trapped meltwater in a glacial lake, dammed by ice. A bit earlier than that, during the depths of the glacial climate, sediments of the North Atlantic Ocean tell us about catastrophic ice collapses during glacial time, called **Heinrich events**. The sediments contain layers of rocks (**Heinrich layers**) that are too big to be carried that far out by wind and could only have been carried in floating ice, enormous armadas of icebergs spewing out to melt in the north Atlantic for several centuries. The rate of sea level rise associated with Heinrich events has been estimated to have been 1 - 10 cm / year for hundreds of years. The mechanism by which these icebergs were launched remains a mystery, but launched they were, as ice not as any dammed meltwater. Perhaps some mechanism such as the Larsen ice shelf collapse was responsible. The significance of iceberg discharge is that it is an extremely efficient method of melting ice. The melting of an ice shelf does not contribute to sea level rise, because it was floating before it melted. But the concern is that the ice shelf might have acted to stopper the flow of more ice from the ice sheet uphill from it, and with its demise, the runoff of ice into the ocean could accelerate. It takes more than an ice shelf to create a Heinrich layer. Several of the ice streams draining the West Antarctic ice sheet run into the Ross ice shelf, which is itself within a few degrees of melting like the Larsen ice shelf did.

The long-term geologic record is also somewhat alarming (Figure 11-7). The temperature of the earth was perhaps 4-5°C warmer than today 40 million years ago when the Antarctic ice sheet first made an appearance. Sea level was about 70 meters higher than today. During glacial climate, temperature was about 6°C colder than today, and sea level 120 meters lower. If we connect the dots, we see a huge eventual change in sea level might accompany a long-term increase in Earth's temperature. The question is how quickly the ice melts.

Biological Impacts

The climate of the earth has been warmer than it is projected to get in response to global warming, and life on Earth continued. The transition from preanthropogenic climate to our future greenhouse climate will be fast, but fast transitions have been documented in the natural climate in the past. This is not to say that the natural world will not suffer adverse effects from global warming. There have been mass extinctions in the geologic record documented as well, and we would not want to emulate that.

One piece of the natural world that appears to be particularly vulnerable is **coral reefs**. Reefs are built by sedentary animals, the corals, that house symbiotic algae, plants which can perform photosynthesis to aid the nutrition of the corals. Coral reef ecosystems are among the most diverse on Earth. Reefs of one type or another have been found throughout the fossil record, yet coral reef systems today appear to be rather fragile. If the water is not clear enough, the symbiotic algae will not get enough light, and the coral will suffer. Corals build the reef from CaCO_3 , the production of which becomes more of a challenge as the seawater becomes acidified by rising CO_2 concentrations. Overfishing threatens the coral community, in particular fishing by means such as dynamite or cyanide poison as practiced in some parts of the world. Corals have been ravaged by diseases and by invasive species. In addition to all of these struggles, corals are vulnerable to increases in temperature (Figure 11-8). When corals are stressed, they respond by expelling their symbiotic algae. This is called coral **bleaching** because of their loss of color. Bleaching may be a mechanism for the coral to try to find new symbiotic algae that are more suited to the conditions they find themselves in, but it is a measure of desperation. Bleaching is often followed by the death of the coral. Figure 11-8 shows a record of temperatures in Tahiti, with arrows indicating times of coral bleaching events. The correlation between heat spikes and bleaching is very clear. Projected warming looks like it will have catastrophic effect on corals.

There is some concern that rising temperatures will widen the area where tropical diseases such as malaria will spread. One can play down this threat by pointing out that the extent of malaria depends on two factors, one temperature but the other the health care system. This would tend to minimize the impact of malaria viability spread into the rich industrialized world. On the other hand, people who have been living with endemic malaria for generations may be less susceptible to malaria than would people whose genetic history had historically been free of malaria for climatic reasons.

Polar bears?

Human Impacts

Water availability may be the most important impact of global warming on human welfare. Farmers are making increasing use of non-renewable or only slowly-renewable "fossil water" in **aquifers**, stores of water beneath the earth's surface that can be harvested by pumping. As water is extracted from these aquifers, the depth in the soil where the water is found, or **water table**, has been observed to get deeper. Eventually these aquifers will become depleted, leaving unsustainable demands for water on the hydrological system. Lester Brown in his book Plan B points to the Northern Plains of China as an example of a place where aquifer depletion is limiting agricultural productivity. Eventual inevitable fossil aquifer water depletion around the world, driven by increasing population and water consumption rates, will only aggravate the climate change impacts of rearranged rainfall distribution and increased risk of drought.

The impacts of climate change to agriculture seem rather large. There are projected changes in crop yield, divided by crop and by region, of order 10 to 50%, in response to future climate change predictions for the coming century (IPCC <http://>). In general, negative impacts are stronger in the tropics and in the developing world, while the higher latitude temperature countries suffer less or may even benefit. There are several factors at work here. One is the increase in temperature. Corn yields are sensitive to heat spikes; just a few can have impact the bottom line for the entire season. Water availability is also a primary controller of crop yield. Dry regions are expected to be even drier with global warming, decreasing crop yields in these areas. On the other hand, increasing CO₂ concentration in the air is expected to boost agricultural production by the CO₂ fertilization effect ([Chapter 9](#)). Higher CO₂ levels help plants deal with water stress by allowing them to get the CO₂ they need for photosynthesis without opening their stomata as much. Another factor that is important in the forecast is adaptation. Perhaps if corn no longer grows in a region because of more frequent heat waves, the farmer will plant wheat, or some genetically modified Frankencorn that is more heat tolerant. In general, the story of agriculture in the past century has been explosive, with production growing even faster than population, at ever diminishing costs. This agricultural revolution is in large part due to industrial production fertilizers like nitrate and ammonia from atmospheric nitrogen, an energy intensive process. The fact that the past trajectory has been wildly explosive, probably unforecastable, makes it difficult to forecast the future with much confidence, and it tends to make any worries about a collapse in food production seem a little silly. On the other hand, we have evidence from the archeological record of collapses of entire civilizations in the past, driven largely by regional drought associated with climate changes smaller than what is projected for the future.

Abrupt Climate Change

In one way, the IPCC forecasts represent the best case scenario, because they are smooth. There are no surprises. The climate record from Greenland shows a period of stability throughout the Holocene, the past 10,000 years when civilization and agriculture developed ([Figure 11-9](#)). Prior to this, during the last ice age 20 to 80 kyr ago, the

climate of Greenland was not so stable. The record is dominated by what are known as **Dansgaard-Oeschger events**, thousand-year cycles of temperature change. The events are marked by sudden warmings of 10-15° C within a few years, followed by more gradual coolings. The Dansgaard-Oeschger events come in bundles called **Bond cycles** with the cold extremes getting colder each time until finally during the coldest of the cold a Heinrich layer is deposited on the North Atlantic sea floor, indicative of the collapse of an ice sheet in northern Canada. After this event, the Dansgaard-Oeschger events become milder again, going into another mild phase.

The mechanism for this climate instability is thought to be based on the circulation of the North Atlantic. A warm phase of the cycle has heat carried to the North, enough to maintain ice-free conditions across the central North Atlantic like we have today. The warm surface waters in high latitude cool and sink, making way for more warm surface waters. During the cold phase the overturning circulation in the North Atlantic ceases, and ice covers the ocean. The albedo of the ice tends to cool the northern high latitudes, and air temperatures over ice plummet to values usually found only in continental interiors.

The overturning circulation in the North Atlantic appears to be vulnerable to changes in fresh water input. Fresh water is less dense than salty water, and so may never reach sufficient density to sink as it cools down. If it doesn't sink, and if it's cold enough, it will freeze. Several of the climate swings during the glacial period, such as the Younger Dryas period marked YD in **Figure 11-9**, were provoked by the sudden release of freshwater from Glacial Lake Agassiz, just north of the Great Lakes today, by the breaching of an ice dam. The Heinrich events influenced climate at least in part by the action of their fresh meltwater. The link to global warming is that models show an increase in precipitation in the high latitudes, resulting in a slowing of the overturning circulation in the North Atlantic. Catastrophic melting of the Greenland ice sheet could have a similar result.

If we were to draw a parallel between past and future, the result could be a paradoxical cooling in the northern high latitudes. However, if sea ice plays a major role in amplifying the temperature swings of the past, and if the base temperature of the earth is sufficiently higher in the future than it was in the past (pCO₂ was 200 ppm during glacial time, remember), it could be that the climate is incapable repeating the savage climate swings of its glacial past.

In a more general sense, the abrupt climate change record of the past serves as a warning because climate swings such as this were not anticipated in the past record, and would be impossible to predict for the future. Models tend to underpredict the swiftness and the severity of these events, presumably because positive feedbacks in the real world are difficult to capture in the models. An abrupt climate shift would be a surprise, like the ozone hole. Once the climate rearranged into some different circulation mode, it could remain locked in that mode for centuries.

Clathrates?

Lab

Point your web browser at <http://geosci.uchicago.edu/~archer/PS134/lab.hadley.html> . Use the movies and the data browser to answer the following questions.

(a) Estimate a global average temperature increase from the beginning to the end of the simulation. Document your estimate by recording the spot temperatures or climatological values that went into it.

(b) Where is warming most intense, the winter high latitude, the summer high latitude, or the tropics? What is the predicted warming in each case? Again, document your method of coming to this answer and the numbers that went into it.

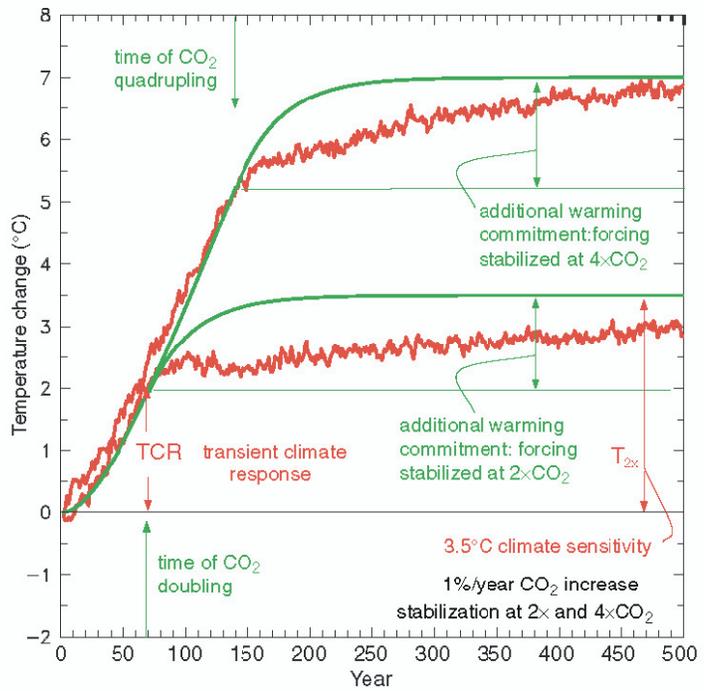


Figure 11-1

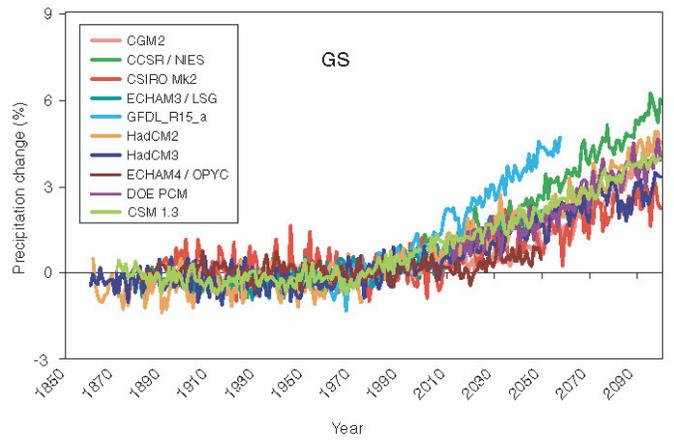
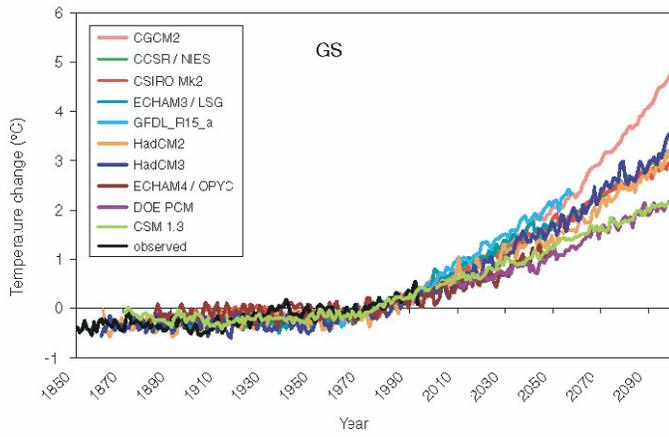
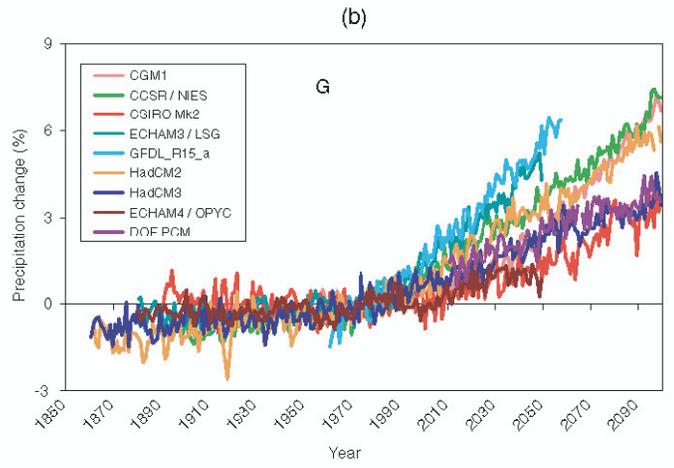
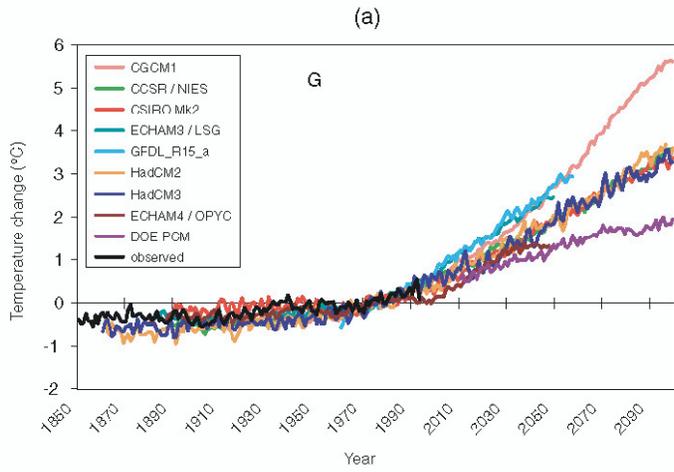


Figure 11-2

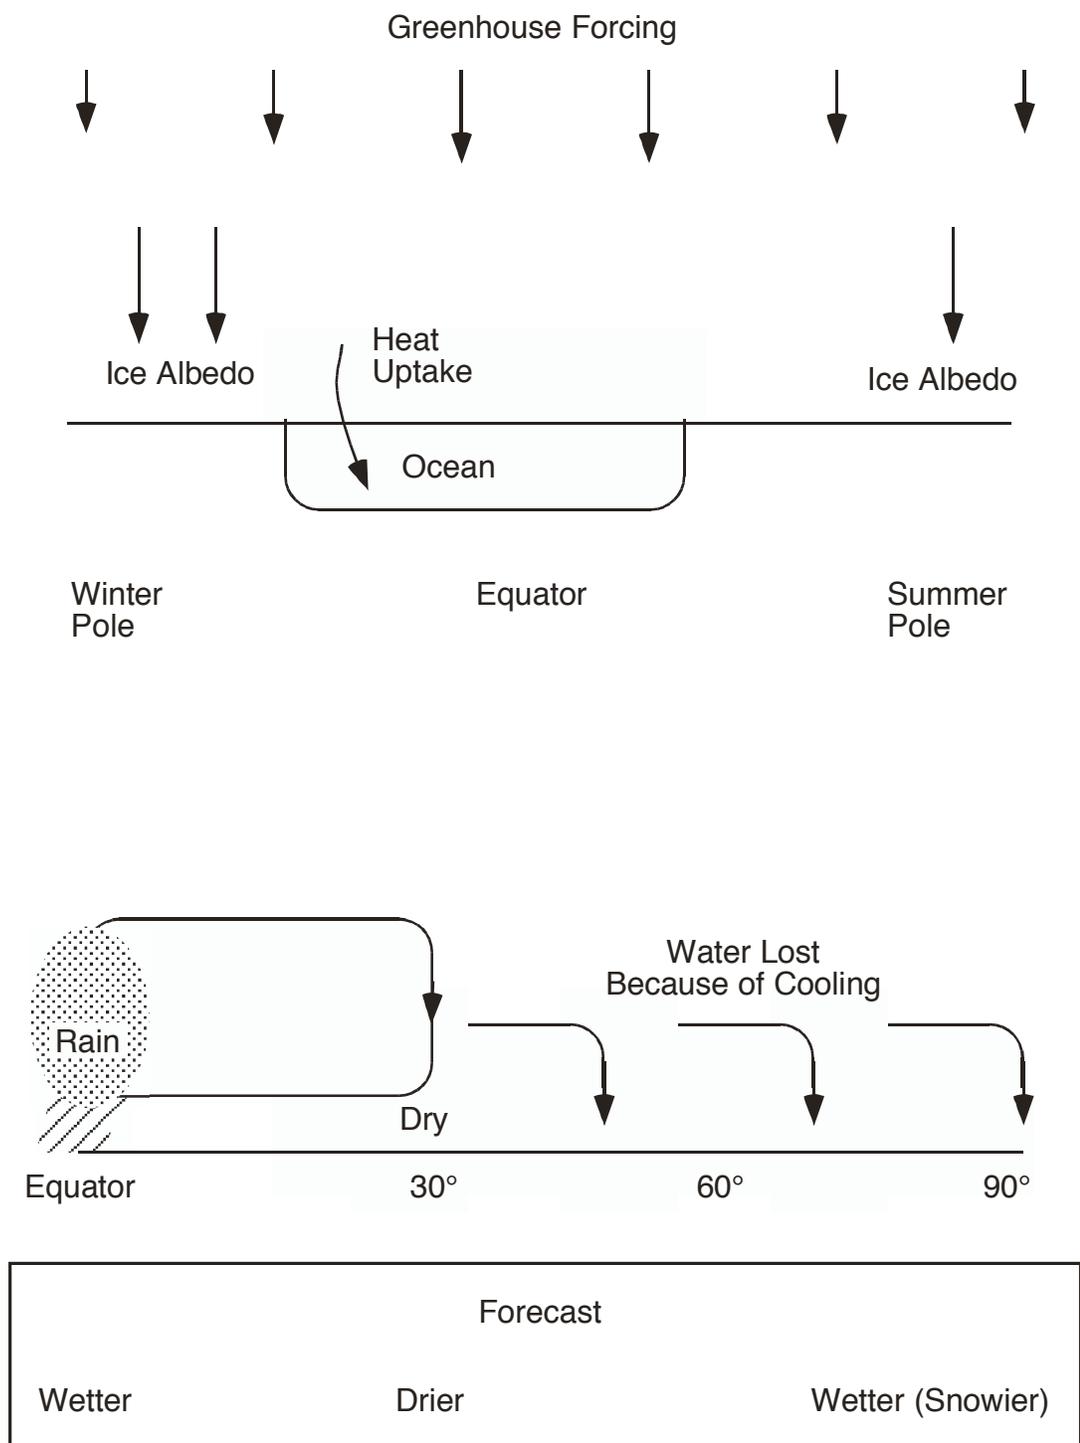


Figure 11-4

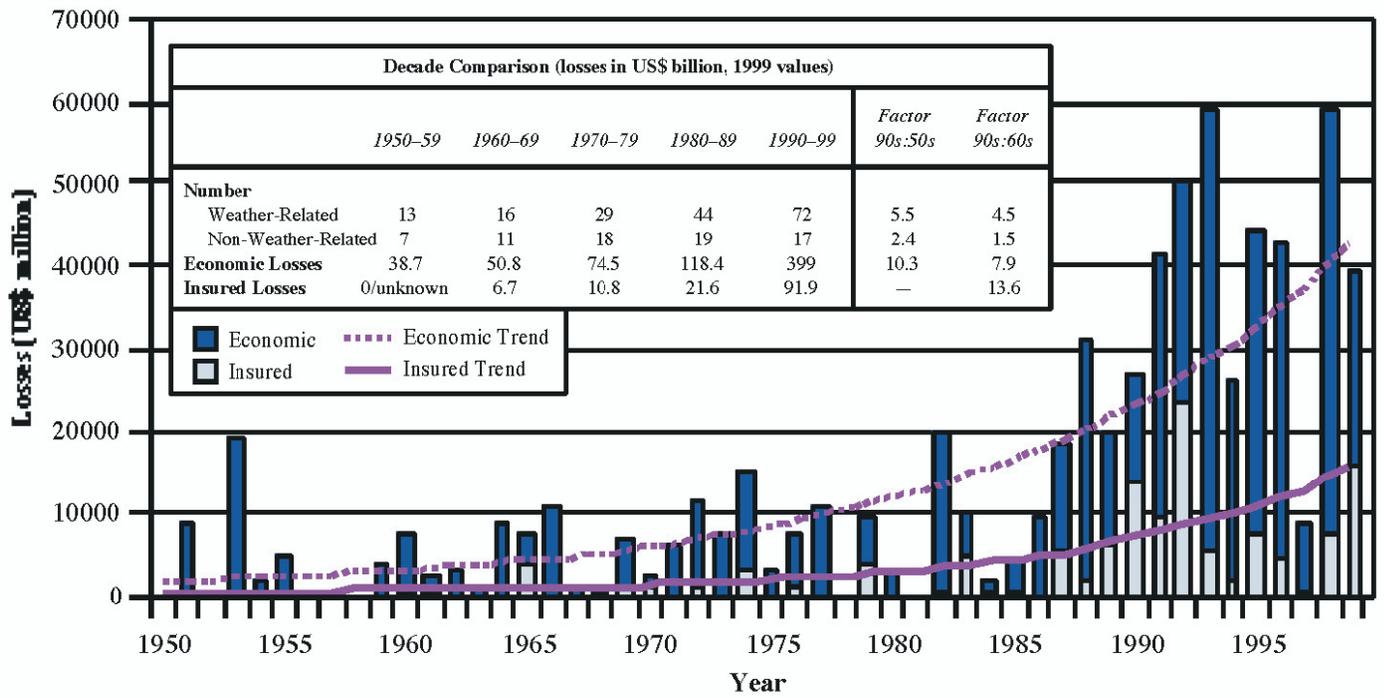


Figure 11-4

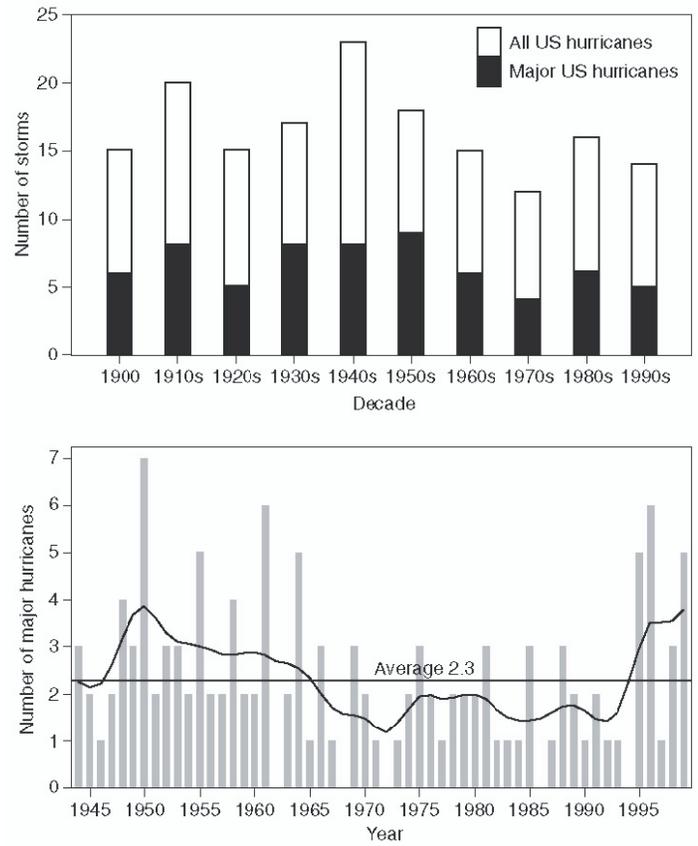


Figure 11-5

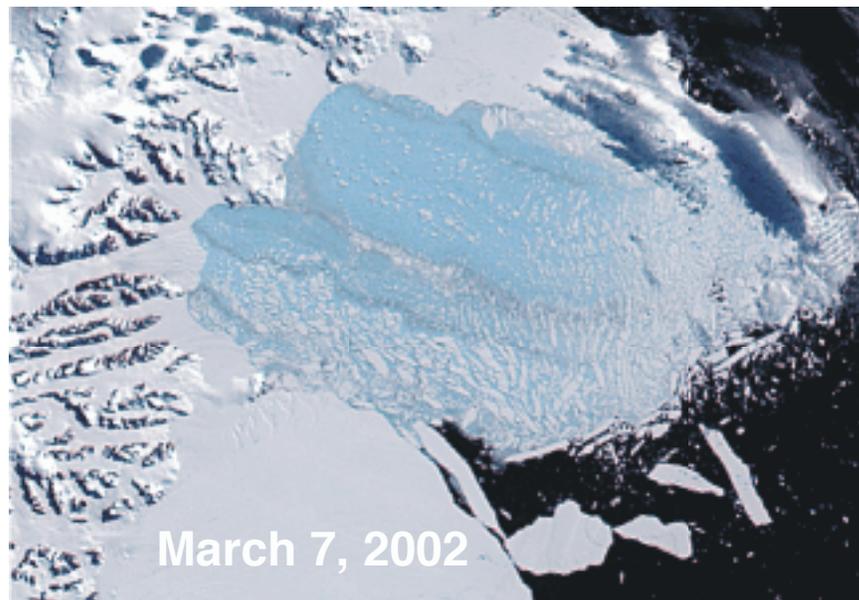


Figure 11-6

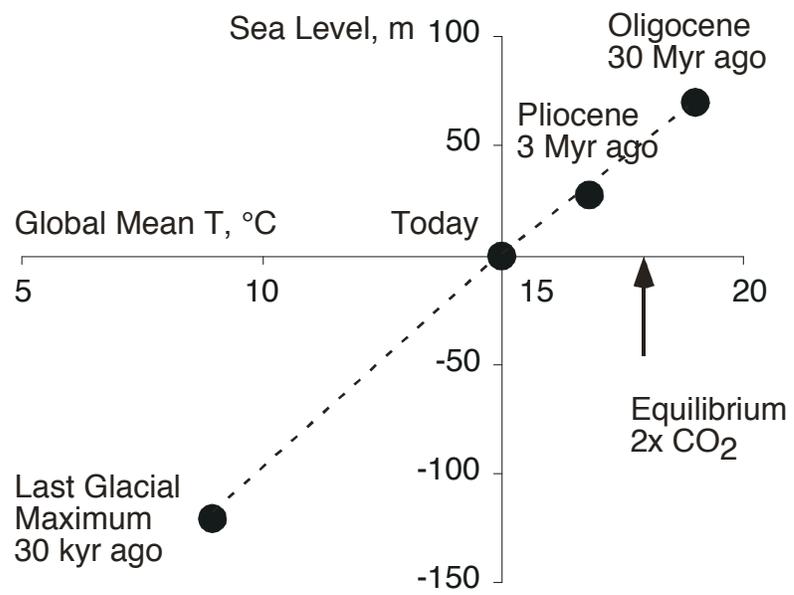


Figure 11-7

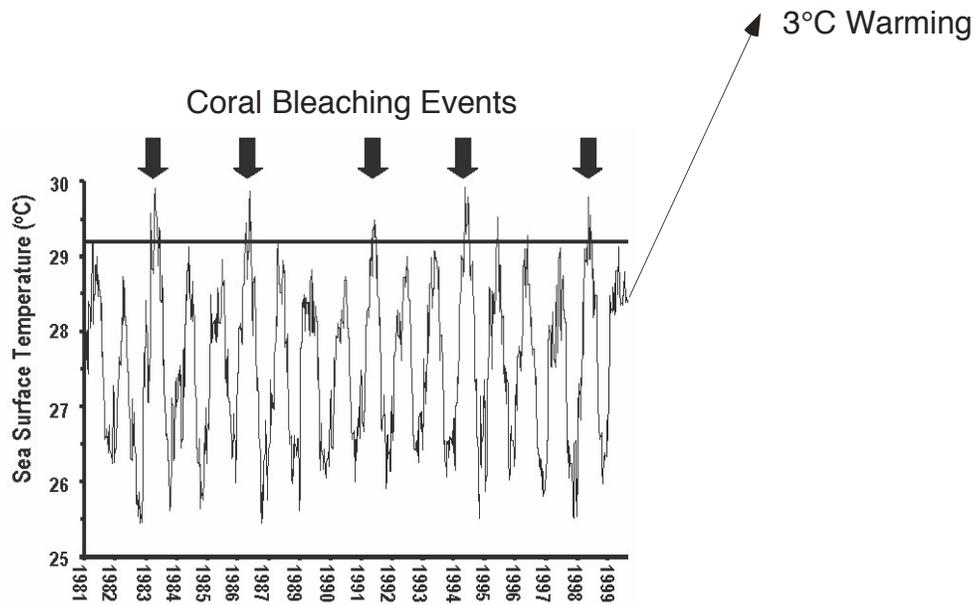


Figure 11-8

Development of Agriculture

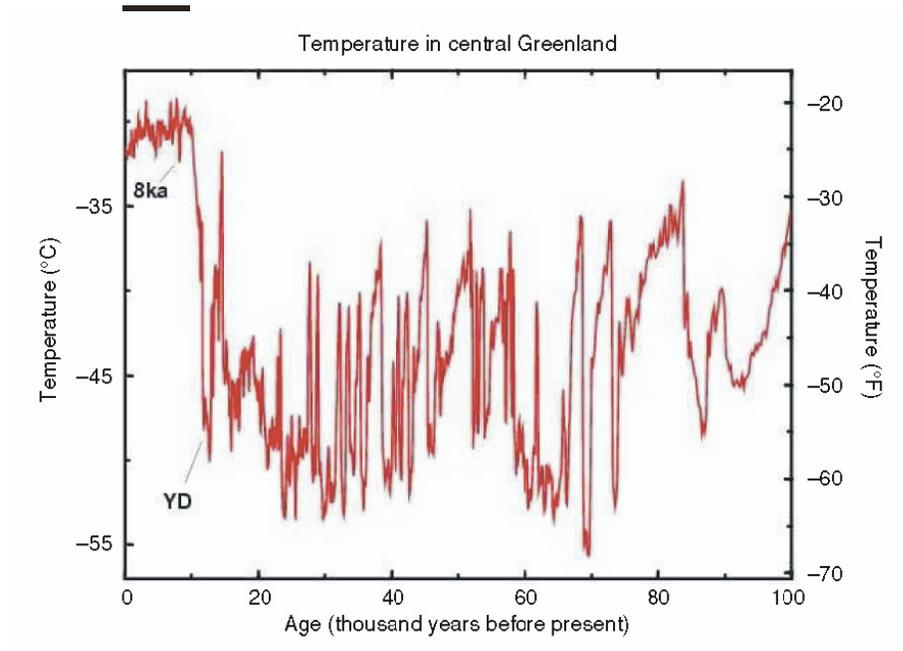


Figure 11-9

Chapter 12. Decisions, Decisions

Abstract

Climate change is the largest-scale problem that mankind has ever faced, global in scope and extending across unimaginable lengths of time. This makes negotiations complicated, because there will be winners and losers no matter how things go.

International negotiations began with the establishment of a framework within which to discuss, called the Framework Convention on Climate Change. The FCCC spawned the scientific body IPCC, and an agreement called the Kyoto Protocol. The Kyoto Protocol can only be the first step in truly preventing climate change.

Economically, the annual costs of CO₂ emission limitation as specified by Kyoto are projected to be comparable to the costs of climate change, but this comparison neglects the long duration of climate change, the issue of fairness, and values of human lives and the natural world that are impossible to quantify monetarily.

Another way to approach the problem is to determine a danger level based on temperature or sea level, and ask how that danger level could be avoided. No single energy source or strategy seems sufficient to solve the problem by itself, but by using a combination of new energy sources and strategies, it may be possible to avoid dangerous climate change.

Global Warming is a Large-Scale Problem

Global warming is a difficult problem to cope with politically because its footprint is so large. It is a global problem, subject to an effect known as the **Tragedy of the Commons**. An example of this effect is of a field for grazing sheep used in common by all residents of a village. Let's say that each new sheep added to the flock decreases the harvestable weight of each of the other cute little sheep by 10%. A farmer with one sheep can add a second, doubling his number of sheep. But the result is that all sheep, including the farmer's, lose 10% of their weight. Each will weight 0.9 times a prior sheep, so the farmer will now own 1.8 original-sized-sheep units. He is still ahead of when he only had one sheep, while all the rest of the farmers lose 10% of their prior sheep stock. The point is that each farmer's personal best interest is served by grazing as many sheep as he can afford. This results in a tendency to overgraze the field. From each farmer's point of view, the other villagers are all also overgrazing the common yard, so his own personal sacrifice, to preserve the common, would be in vain anyhow, and would put his family at a disadvantage. The end result is far more sheep than can be most efficiently supported on the common, decreasing the net yield of sheep in total. Everyone might be better off if they backed off a little bit and cooperated.

There are two potential approaches to preserving the common space. One is to divide it up and assign property rights to individuals. In this case, each individual would have the incentive to preserve the field by not overgrazing, because that individual would directly benefit from careful stewardship. For our example of land stewardship, private ownership is not a universal environmental panacea, because some element of commonality remains. The aquifer holding water underneath the ground may be over-tapped, or soils may be depleted throughout the lifetime of the individual (a sort of commons over time). For fisheries, the ownership society system also has holes in it because fish swim around. For the case of global warming, dividing the atmosphere up into parcels is impossible by the nature of the problem. The other potential solution is for some form of collective self-regulation. Ideally the individuals recognize that regulations applied fairly to all will benefit each individual, relative to the case of the tragedy of the commons.

The decisions made by a single individual are often irrational, and therefore not really well described by economic theory. The collective decisions made by whole societies of individuals, in contrast, often exhibit great sagacity. Economics aficionados speak with reverence of the **Market**, Adam Smith's invisible hand. However, the market has its blind spots, leading to effects like the tragedy of the commons. In economics, a cost which is not paid by the decision maker is called an **external** cost. An example of an external cost associated with driving to work is the traffic. One more car will tend to slow down all the other cars on the road, costing other drivers their time. The cost of climate change is not only paid by people who are responsible, but by everybody, soon and far into the future. Our sheep farmer made the other farmers pay part of the cost of his new sheep. If true costs are left external, then the all-wise Market does not take them into account, and commons tragedies are the result. An external cost can be **internalized** by means of taxes or regulations. The idea is to make the market aware of the true cost of this path versus that path.

One way to harness the wisdom of the Market for preventing global warming is a scheme called cap-and-trade. A regulatory agency allocates permits for emission of CO₂, with the total number of permits they issue totaling some lower overall rate of emission than business-as-usual. If a company or a country is able to cut its emissions even lower than its allocation, it has the right to sell its permit to another company for whom it would be more expensive to cut its own emissions. The price of the permits is set by supply and demand, same as any other commodity. The law of supply and demand is a negative feedback system, just like the many natural examples in Chapter 6. In this way, the market finds the most efficient, that is to say cheapest, means of limiting global CO₂ emissions to the cap value. A cap and trade scheme for sulfur emissions (generating acid rain) from power plants has apparently worked well. The difficulty with cap and trade for a regional problem like acid rain is that it may allow the all the pollution to concentrate in one region, which may be too much for that region. CO₂ is a truly global problem, in contrast, and it makes no difference where the CO₂ is released, so CO₂ emissions would be an ideal problem for limitation by cap-and-trade.

The Market may influence the global warming debate through the insurance industry. We saw in **Chapter 11** the huge increase in insurance payout to cover atmospheric-related natural disasters. It is not at all clear whether this trend is due to climate or to social factors, but the insurance industry is certainly concerned about the possibility of climate change affecting their bottom line. Insurance companies have been responsible for other preventative measures such as requirements for hard hats on construction sites.

Pollution problems get more difficult to solve, politically, as their footprints expand from regional to global, because of the commons effect. Local issues like contaminated drinking water or urban smog are clearly in the interests of local individuals to solve. You fix it same as you would fix the roof over your head. A clear link exists between the costs of cleaning up and the benefits. Larger-scale, regional problems tend to run into us-versus-them issues. Why should the state of Ohio sacrifice their right to burn high-sulfur coal when the costs of that consumption are paid by the folks in New York? Why should power plants in Chicago pay to clean up their mercury emissions when most of the mercury emission to the atmosphere comes from China? At the most difficult, worst end of this problem, global warming is the most global of issues, the most complicated type to solve.

Global warming is also a difficult issue to address because it is long-term. Values need to be related not only from one place to another, but across time. The way that economists deal with values across time is via a construct known as a **discount rate**. The idea is based on a question: which would an individual prefer, a cost today or a cost in the future? A cost in the future, certainly, would be the natural response. A rational reason to justify that choice is that one could invest money today and use the interest to help pay the costs in the future. If we were to assume an interest rate of say 3% per year, we could pay a cost of \$100 in 100 years by investing \$5 today. It would be worth paying \$4 today to avoid a \$100 cost in 100 years, but it would not be worth paying \$6 today. The bottom line of this idea is that costs tend to shrink with time.

Of course, the idea of anybody actually going out and opening a bank account with \$5 in it, to save for that cost in 100 years, seems rather quaint. However, economic theory is more than just a prescription for how one could rationally make the most of one's assets; it is also a description of how money tends to flow in our society. The discount rate theory illustrates why financial decisions made by the Market tend to be somewhat short-sighted. It's a description of the way money flows. Water flows downhill. Money flows short-term.

Negotiations

The process of international negotiations to limit CO₂ emission began with a document called the **Framework Convention on Climate Change** or **FCCC** that was drafted during the Earth Summit meeting in Rio de Janeiro in 1992. The objective of the FCCC was to achieve "stabilization of greenhouse gas concentrations in the atmosphere at a level that would prevent dangerous anthropogenic interference with the climate system." FCCC deferred defining what the dangerous concentration would be, a discussion that continues to this day.

FCCC established a procedure for negotiations by setting up a series of meetings between countries, called **Conference of Parties** or **COP** meetings. FCCC established an organizational body for summarizing scientific results, our familiar **Intergovernmental Panel on Climate Change** or **IPCC**. FCCC also established a procedure for creating amendments to itself, the most famous and significant of which was drafted during the COP-3 meeting in 1997 in Kyoto, Japan, called the **Kyoto protocol**.

The role of IPCC is to assess “the scientific, technical and socio-economic information relevant to understanding the scientific basis of risk of human-induced climate change, its potential impacts and options for adaptation and mitigation”. IPCC scientists are drawn from government and academia, and are divided into three working groups. Working Group I deals with the scientific basis for the climate change forecast. WG II deals with the impacts of climate change on the natural and human world. WG III assesses options for limiting greenhouse gas emissions or otherwise avoiding climate change. The most recent reports from all three are available on line at <http://www.ipcc.ch/> .

IPCC produces reports consisting of chapters and various levels of summaries. The chapters present no previously unpublished research, but are based only on published, peer-reviewed scientific papers. Of course, not every paper published in the scientific literature is guaranteed to be correct in every way, but the practice of sending manuscripts out to rival scientists for review makes sure that most of them are at least plausible. If a scientist disagrees with a paper that has already been published, he or she is free to write another paper attempting to disprove the first. Papers are sometimes retracted by their authors if they are proven wrong, or sometimes a disagreement just stands in the literature, requiring judgment on the part the readership to decide which is right and wrong. Once an IPCC chapter is drafted, summarizing all of the published papers, the chapter is subject to another round of reviews by scientific experts and government workers. The chapters are grouped together into books which also come with technical summaries and summaries for policymakers. These summaries are subjected to line-by-line review, voting, and approval by the entire IPCC Panel. This sounds like a long and painful process!

The conclusion of the first, 1990, IPCC report was that it was not yet possible to detect human-induced global warming. This was largely because of the mismatch between CO₂-only model results and the thermometer record (**Figure 10-11b**). By 1995 it had been worked out that the warm temperatures between 1900 and about 1940 could be attributed to increased solar intensity, while warming from 1970 to the present was clearly due to rising CO₂ levels (**Figure 10-11c**). The 1995 report issued a now-famous statement that “the balance of evidence suggests a discernable human influence on global climate”, thus providing the impetus for drafting the Kyoto Protocol in 1997.

The Kyoto protocol attempts to limit global CO₂ emissions to about 6% below 1990 emission levels, by the year 2010. The countries participating in the negotiations (the “parties”) are divided into industrialized and developing nations. The agreement limits the emissions of the industrialized nations, but does not put limits on emissions from the

developing world. The agreement allows for cap-and-trade arrangements, allowing countries to buy and sell emission credits. Developed countries are allowed to invest in carbon-reduction measures in developing countries, and get carbon emission credit for this. Carbon uptake by re-forestation is allowed, and counts toward meeting emission targets. This is some discussion of penalties for exceeding emissions, including a factor of 1.3 “make-up fee” for emissions that exceed the targets. However, the agreement has no real teeth; it remains to be negotiated whether the consequences for failure are truly legally binding.

The treaty came into force when countries accounting for 55% of CO₂ emissions globally agree to its terms. The United States withdrew from negotiations under the Bush administration in 2001, presenting a fairly sizable obstacle to achieving this target, seeing as how the United States accounts for 25% of CO₂ emissions globally. In spite of this setback, the agreement came into force when Russia agreed to the Kyoto protocol in 2005.

The Kyoto protocol is only the first step toward a solution to the climate change problem. The target CO₂ emission rate from Kyoto is about 6% below 1990 levels. You may recall from Chapter 9 that the uptake rate by the atmosphere and ocean combined is about 4 Gton C / year, about 40% below 1990 levels. The actual projected climate impact of the Kyoto protocol by itself are discouragingly small, because of the small target cut, and because there are no limits to emissions by developing countries.

A major objection raised by the two “dropouts” from the Kyoto protocol, the United States and Australia, is the exemption of the developing world from CO₂ emission limitations. The Kyoto protocol has taken the approach of allocating acceptable CO₂ emissions for each country in the industrialized world according to how much CO₂ they currently emit. By far the greatest fraction of the carbon emissions, historically and today, comes from the industrial world. Capping future emissions according to historical emissions for the developing countries would therefore carve into law the imbalance between industrialized and developing nations. We have seen in Chapter 9 the intimate relationship between CO₂ emission, energy, and economic activity. Within our current technological stage of development, CO₂ emission is the key to national wealth and comfort. The developing world understandably wants the benefits of CO₂ emissions also, and can claim that capping their emissions according to current emissions would be unfair. This is the rationale for exempting developing nations from emissions caps for the time being. In this strategy the Kyoto Protocol follows the example of the Montreal protocol banning chlorofluorocarbon emission, which initially applied only to developed nations and then phased in to restrictions for everyone. However, because of the large numbers of people in the developing world, in particular India and China, the developing world is projected to account for a much larger fraction of CO₂ emission in the future. For this reason, it will be essential for continued negotiations beyond Kyoto to bring the developing world on board.

Economics

One way to approach the decision of whether to limit CO₂ emissions is based on economics. Which path costs less, to prevent climate change aggressively or rather to endure it? A few caveats are in order here before we proceed. Economic behavior is in general not as easy to forecast as climate can be forecast. Physical sciences are based on a solid foundation of natural laws such as Newton's equations of motion, conservation of energy, and so on. It is not always possible to derive the physics of climate solely from these natural laws. Higher-level processes like turbulence and clouds are called “emergent behavior”, surprisingly difficult to build from the fundamental laws of physics. But at least the foundation exists. Economics and the other social sciences are fuzzier. Models of economic trends tend to be empirical, eyeball curve fits. It is not that physical scientists are smarter than social scientists, it is simply that physical sciences are easier, and more tractable.

Economic trends are strongly impacted by technological progress, which is impossible to forecast, and by social and political fancies. One could even argue that economics is pathologically unforecastable because it is self-aware. If I came up with a model that could successfully forecast stock prices, I could use it to make money, until everyone else discovered my methods. At that point, the stock prices would reflect everyone's knowledge of the model forecast, and the future trajectories of the stock prices would alter to evade my forecast. I wouldn't make money any more. Not even on my darkest days do I imagine that the physical world that I model for a living is going to change its behavior for the purpose of making my model wrong!

In general the costs of climate change are predicted to be non-linear with temperature increase. A small increase in temperature, say 1° C or less, will probably not do too much damage, and may even be beneficial overall. As the temperature rises to 2° or 3°C or higher, however, the economic impact starts to rise sharply. This is true for many aspects of climate change: farming, sea level rise, health impacts, and the possibility of abrupt climate change. An economist at Yale, William Nordhaus, has constructed an economic model to gauge the costs of climate change versus the costs of decreasing CO₂ emissions. Nordhaus' book can be downloaded from <http://www.econ.yale.edu/~nordhaus/homepage/homepage.htm> .

His costs are presented as percentage changes in global domestic product, GDP. We have to watch out here because the standard practice among economists is to not write out the units of per year. The units of GDP are \$/yr. The growth rate of GDP is calculated as GDP change per year, which works out to \$/yr/yr or \$/yr². The growth rate of the GDP is typically a few percent each year. This sounds small, a few pennies out of a dollar (specifically, more dollars per year the second year than the first), but yet somehow people's livelihood and sense of prosperity depends extremely sensitively that few percent. If the economy grows by 5% each year everyone gets rich and shoe-shine boys trade stock tips. The great depression decreased global productivity by about 25%. The reunification of Germany put a drain of about 5% on the GDP of West Germany, producing noticeable malaise among West Germans.

Another way to evaluate how big a percent is, however, is to note that economies in general tend to grow by a few % per year, paced perhaps by the growth in population. A cost of a few % of our GDP would just set our level of wealth back by a year or two. A 1% cost only affects the growth rate of the GDP by 1%, changing, say, a growth rate of 3% per year to 2.97% per year. In the long run, this does not seem like a long-term catastrophe to a naïve non-economist like your author. To an Earth scientist, any positive exponential growth like this is an explosion, whether it explodes quickly or slow, it cannot persist forever.

Costs of Climate Change. The projected agricultural costs seem rather small, perhaps 1% of GDP, because agriculture itself is only 2.5% of GDP or so, globally. Of course, this is a special 2.5%, different from say shoes or music videos because one cannot eat shoes or music videos. The costs of sea level rise are projected to be of order 0.1% of GDP. Storm surges currently cost about this much every year. It is not unreasonable to imagine doubling this cost by increasing sea level, just to get an idea of the scale of how big it could be. Health impacts are projected to be worth 0.5% of GDP globally to avoid, with higher impacts in tropical developing countries in Africa and Asia, where it can reach 3%. Of course, this statistic requires coming up with a way to express the worth of a human life in dollars.

The largest impact in Nordhaus' tabulation was a willingness to pay to avoid catastrophic climate change, which he defines as equivalent to the Great Depression, a 25% reduction in GDP, that lasts forever. He gauged the probability of an event like this by surveying experts in climate science, to derive the collective opinion of the group that there is a 1% probability of catastrophe if the globe warms by 2.5°C. If climate warms to 6°C, the probability, pulled out of the air by the panel of experts, rises to about 7%. When the probabilities and downside costs are combined in a way that recognizes that people are willing to pay a little extra in order to reduce risk, Nordhaus derives that it would be worth about paying about 0.5% of GDP globally as insurance to prevent temperatures from reaching 2.5°C, or about 5% of GDP to prevent a temperature rise of 6°C, on the basis of catastrophic climate change alone. Overall, Nordhaus estimates that the cost of a 2.5° warming would be about 1.5-2% of GDP.

Costs of Kyoto. The costs of climate change on the one hand can be compared with projections of the cost of compliance with Kyoto on the other. Of course these numbers are speculative as well. The costs of limiting CO₂ emission will depend on technology development, very difficult to forecast. Ultimately only time will tell. These are guesses, but it is interesting to look at the numbers anyway.

The forecast is that compliance for Kyoto by all nations including the U.S. would cost about 1% of GDP each year, according to the on-line book published in 2001, and 7% in a paper in Science published in 2002. The Science article is the more recent of the two, and the cost went up apparently because the predicted business-as-usual baseline had gotten higher. (A change in our expectations appears as a cost.) Interestingly, the cost to the United States is much higher than the cost to other countries, because the carbon emissions are rising much more slowly in other developed countries than it is in the

States. U.S. participation would actually be a boon to the rest of the industrialized world, because we have such a ravenous appetite for fossil fuels that we would be forced to purchase CO₂ emission credits from other countries. Without U.S. participation, the price of those credits would be much lower. The model forecast a price of \$150 per ton of CO₂ in the year 2025 with United States participation, and \$40 without.

The bottom line is that compliance with Kyoto could cost a few percent of GDP every year, while climate change could cost a few percent of GDP every year. Don't forget that Kyoto will not prevent global warming, it is only the first step, and further restrictions will cost more. One could therefore uncharitably characterize Kyoto as "inefficient and expensive".

It may still however be better than the alternative. The costs of climate change will persist as long as the CO₂ persists, say 300 years for the bulk of it. If you compute the total costs added up over the long-term Kyoto seems like a bargain, because we pay for a few decades and benefit for hundreds of years. Ah, but don't forget that discount rate. We don't care about costs or benefits in the future; we want the money now.

Economics is an awkward tool for the climate change decision, because many aspects of the natural world are simply not economically fungible. The economic value of coral reefs could be tabulated as the tourist dollars they generate, same as Euro-Disney, but somehow something is missing in this tabulation. If coral reefs cease to exist because it gets too hot, we cannot purchase more coral reefs for any amount of money.

The bulk economic framework also hides an issue of fairness, in that people who benefit from business-as-usual are not the same people who pay for business-as-usual, say by living in an increasingly drought-prone world far in the future. I imagine that giving up the institution of slavery in the American South entailed a cost, too, for somebody. Ultimately the abolition of slavery was not an economic decision, but an ethical one.

Although the Kyoto Protocol will not by itself prevent climate change, it would provide the benefit of guiding investment toward decreasing CO₂ emission. Coal power plants, for example, can be built using newer technology called "integrated gasification combined cycle", enabling future CO₂ capture for sequestration into the geological repositories. If CO₂ is released to the atmosphere, IGCC plants cost about 20% more than traditional plants. If CO₂ sequestration starts to be common practice, however, IGCC plants would be cheaper overall than traditional plants. In the regulated United States energy market, regulations often require power companies to use the cheapest technologies. In general this regulation makes sense, if we want to limit the power of monopoly power companies to spend money wantonly, passing the costs on to their captive customers. But in this case, the cost of climate change are external to the accounting, and so the regulation may not be choosing the cheapest option. Once a power plant is built, it has a life span of fifty years or so. The ultimate costs of decreasing CO₂ emissions would be lowered by starting now, rather than continuing to invest in high-emission technology.

Dangerous Anthropogenic Interference

Another way to approach the problem is to ask, What would it take to avoid dangerous climate change? This was how the Framework Convention on Climate Change was constructed, to “prevent dangerous anthropogenic interference with the climate system”. IPCC didn’t specify what dangerous interference is, leaving that topic to future discussion which is still happening.

One way to define dangerous interference is to base it on the range of climate variability that the earth has experienced in the last million years or so. If we change our climate to that of the age of dinosaurs, that sounds dangerous. Let’s not do that. That’s the idea. Global mean temperature may have been 1.5°C warmer during the last interglacial period, 120 kyr ago, or the last time the earth’s orbit resembled the way it is today, a time called stage 11, 400 kyr ago. Based on that, the German Advisory Council on Global Change (<http://>) argues that 2°C is a maximum tolerable global increase in temperature.

Another way to define dangerous is to watch sea level. Jim Hansen (<http://>) argues that 2 meters of sea level rise sounds dangerous to him. The IPCC forecast for sea level rise is about half a meter in the next century, detrimental but not yet dangerous according to Hansen. But 2 meters, you’re into some serious Florida. That sounds like a big deal. The longer-term trends from the past ([Figure 10-7](#)) seem to predict tens of meters of sea level change with long-term doubling of CO₂. The question is how quickly the ice melts down. Heinrich events (Chapter 11) suggest that ice can melt pretty quickly if it is a mind to, although we’re not too clear on how. Hansen argues that the past would suggest the possibility that 1°C temperature increase might exceed his 2 m sea level danger limit. The planet has already warmed by 0.6°C, and another 0.4°C is on the way, just from the CO₂ that we’ve released already. So we are apparently just now crossing the danger limit, according to Hansen.

The other issue is how fast the temperature change occurs. FCCC says that warming should be slow enough for ecosystems to adapt naturally, with no threat to food production, “and to enable economic development to proceed in a sustainable manner”. In other words, the transition to a new climate regime should be slow enough for natural and human systems to adapt. Here again history can be our guide. The trick is, short warming spikes heat up more quickly than long-term trends, so we have to compare temperature trends computed over similar lengths of time. The current rate of warming globally over the past decade is 0.2°C per decade; higher than this in some locations regionally. When we compute the rate of change over multiple decades, the recent rate drops to 0.1°C per decade on average. Historically, warming at 0.1°C per decade is a fairly unusual warming. Based on this, the German Advisory Council (<http://>) concludes that 0.2°C per decade seems about the upper safe limit to the long-term temperature rise.

Now we need to translate from units of °C that we consider danger, to atmospheric CO₂ concentration that would give us that temperature, then to CO₂ emissions to stay within that atmospheric CO₂, then finally to energy policy to keep within those carbon emissions. Beginning with a target temperature that we consider to be dangerous, we use ΔT_{2x} to calculate the atmospheric pCO₂ value you need to get there (Chapter 3). According to equation 3-1, a temperature change of 1°C is already in the pike from our

pCO₂ concentration of 360 ppm; 2°C would be generated by pCO₂ of about 450 ppm. 560 ppm would generate about 3°C (Figure 3-7).

Translating from atmospheric CO₂ concentration to CO₂ emissions requires a model of the uptake of CO₂ into the ocean and the terrestrial biosphere, such as the ISAM carbon cycle model (<http://>). In order to compare the responses of different models, they must each be subjected to the same trajectory of atmospheric CO₂ rise, because the amount of CO₂ the ocean can absorb depends on how much it has already absorbed (its history). IPCC in 1992 constructed for this purpose a series of what-if scenarios for the trajectory of atmospheric CO₂ through the year 2100, called CO₂ stabilization scenarios. These are arbitrary but reasonable-looking trajectories for the CO₂ rise in the atmosphere, ending up at a stable concentration of 350, 450, 550, 650, or 750 ppm.

The on-line model for the Kaya identity (<http://geosci.uchicago.edu/~archer/cgimodels/kaya.html>) has taken advantage of the CO₂ stabilization scenarios to estimate how much carbon-based energy must be replaced by some new form of carbon-free energy, if we are to achieve atmospheric stabilization. Recall from Chapter 8 that the Kaya identity model makes a forecast of energy use and carbon emissions to the year 2100 as a function of population size, economic growth, energy efficiency, and energy sources. Then we compare the carbon emission rate from the Kaya model, say for the year 2100 in Figure 8-16 it is about 700 ppm. The ISAM model tells us that about 2.5 Gton C can be emitted in 2100 for the CO₂ to follow the stabilization scenario for 450 ppm. Kaya wants 18 Gton C/yr, but to stabilize at 450 ppm, we can only emit 2.5 Gton C/yr, so we gotta get rid of 15.5 Gton C/yr.

Coal could generate 17 terawatts of energy from this carbon emission, so to replace coal we need about 17 terawatts of carbon-free energy (Figure 8-16). Coal would be the most sensible carbon fuel to replace, and it may be all the fossil carbon that's left toward the end of the century, if we don't start digging into like clathrates or oil shales. How much is 17 terawatts of energy? For comparison, mankind is currently consuming about 13 TW of energy per year. The sun bathes the Earth in energy at a rate of about 173,000 terawatts. Photosynthesis captures about 100 terawatts.

Alternatives

One obvious possibility is conservation. The United States uses twice the energy per person as is used in Europe or Japan (Figure 8-13). In part this is because the human terrain of the United States developed in a more automotive way, while public transit and trains are better developed in Europe. In addition, Americans drive larger cars and live in larger houses. I personally like life in Europe. I would far prefer to take a train to work, in which I can read or watch people, rather than curse at other drivers in a private car. Certainly it makes more sense, in a global sense, for Americans to learn to live like Europeans rather than Europeans and other Earthlings to learn to live like Americans.

On the other hand, it seems unlikely that conservation alone will solve the world's energy problems. Most of the world lives at a standard far below that of Europe or the United States. If the citizens of the developing world ultimately consumed as much as

the European model, the footprint of humanity on the earth would increase by about a factor of five. The Kaya projection of energy use in the future already assumes increases in energy and carbon efficiency, and it still predicts a thirst for energy that results in higher CO₂ concentrations than the stabilization scenarios.

None of our traditional sources of carbon-free energy seems capable of making up the shortfall for the coming century by itself [Hoffert, 2002 #2647; Pacala, 2004 #2648]. The first option that seems to come to people's minds is **nuclear energy**, because it generates no carbon, and we know it works. The French are using nuclear for more than a third of their energy needs. However if we want to generate 17 terawatts of nuclear energy, using present-day standard 1000 Mwat reactors, it would require 17,000 new reactors within 100 years, for an average reactor construction rate of one new reactor every second day, continuously, for 100 years. This seems like a lot of nuclear reactors, or else development of reactors much larger than today's.

It is a close call whether or not there would be enough minable uranium to support this level of nuclear energy production. The global inventory of minable uranium today would be sufficient to generate 10 terawatts of energy at our current efficiency for about 10 years before it would be exhausted ([source](#)). However, uranium today is rather inexpensive, and we only bother to mine the purest uranium deposits, because production of nuclear energy is limited by supply of reactors, not by supply of uranium. We could move to lower-purity uranium deposits. Ultimately it may be possible to extract uranium from seawater, where it is present in huge amounts but at very low concentrations. In addition, it is theoretically possible to stretch the energy yield of natural uranium by converting it to plutonium in what is known as a **breeder nuclear reactor**. This increases the energy yield from the uranium by a factor of 50, but the downside is that it is very easy to produce nuclear weapons from plutonium. For this reason, breeder nuclear reactors are not used today.

Opposition to nuclear energy arises because of the potential for catastrophic accidents. The explosion at the Chernobyl power plant in the former U.S.S.R. could have been much worse, if the core had melted down to the water table, contaminating the groundwater supply to Ukrainian agriculture and the city of Kiev. Nuclear energy has the potential to poison large stretches of land to human use essentially forever. The other objection to nuclear energy is the problem of waste storage. Breeder reactors have the advantage of incinerating the long-lived radioactive wastes.

Windmills are becoming economically competitive with new traditional power plants. The trick word in that sentence was "new". An already-existing coal power plant can be run very inexpensively; nothing can compete with that. Windmills supply 7% of the energy needs of Denmark with no adverse impact on the beauty of the landscape of Denmark in my opinion. Wind energy currently accounts for 0.3% of energy globally and is growing at 30%/yr. At this rate wind could supply 10% of the world's energy within the next couple of decades. Wind power may not be able to generate tens of terawatts, but scaling up current wind energy production by a factor of 50 would generate a few terawatts of energy.

It may be possible to extract energy from winds at high elevations in the atmosphere. Winds get faster with altitude in the troposphere, peaking in the jet stream winds in mid-latitudes of both hemispheres. The air up there is at lower pressure than down here, and so therefore has less density to drive wind blades, but still the power density is much higher at high altitude than it is at the earth's surface. Wind energy collectors could function like giant kites tethered to the ground by electrically conducting cable, remaining aloft passively by riding the wind. It has been proposed that high-altitude windmills could scale up to tens of terawatts of energy (<http://www.skywindpower.com>).

Solar cells, or **photovoltaics**, generate electricity from sunlight. Photovoltaics are currently rather expensive to produce, but if solar energy generation were scaled up by a factor of 700 from present-day rates, they could generate several terawatts of energy. Similarly to surface windmills, photovoltaics seem unlikely to solve the entire energy problem but they can certainly help. One new idea is to build solar cells on the moon (<http://>), beaming energy back to the earth as microwaves. The moon would be an ideal location for solar cells, in that it never rains there, there are no clouds, no birds to poop on the solar cells. The cells could be produced from lunar soil, reducing the amount of mass that would have to be lifted by rockets. Transmission of energy via microwaves to Earth seems feasible, safe, and relatively efficient technologically. An array of solar cells in orbit around the earth would also share the benefits of a clean environment and solar flux that is unimpeded by the atmosphere. Solar cells in orbit would have to be produced on Earth and lifted to orbit.

Both surface wind and solar power share a technical difficulty that energy must be stored, because periods of maximum power demand may not coincide with times of maximum windiness or sunshine. One way to deal with this problem is to globalize the power grid, so that electricity could be efficiently transported around the world. Electrical energy can also be used to generate **hydrogen**. You have probably heard the hype about the hydrogen economy. Hydrogen itself is not a primary energy source. There are no minable sources of energy as hydrogen like the energy in fossil fuel deposits. Hydrogen can be produced from other sources of energy, for example electricity can be used to break water into components hydrogen and oxygen gas, a process called **hydrolysis**, or produced chemically from coal in a process called **gasification**.

Hydrogen could be used to store energy for use in transportation. Because hydrogen is a gas, it must be contained under pressure, as opposed to gasoline, which is liquid at room temperatures, or propane which can be liquefied at moderate pressure. Hydrogen is also a very flammable gas, and burns invisibly unless the flames happen to hit some other material that emits light. Propane, in contrast, can be stored at lower pressure because it forms a liquid at room temperature. Propane is already being used to power some buses, as well as things like forklifts. Hydrogen is more explosive than propane, but this added danger is offset by the tendency of hydrogen to rise up, like helium would, escaping from a crashed vehicle.

Another option for carbon-free power is to use fossil fuels and then dispose of the carbon, called **CO₂ sequestration**. CO₂ could be captured from the exhaust stream of a power plant, or coal could be treated with steam to release hydrogen gas and CO₂. The CO₂ could be injected into the earth into geological formations that must be porous enough for CO₂ to flow away from the injection site, but isolated from the surface so that the CO₂ doesn't escape back to the surface. The largest type of geological formation that would fit the bill is called saline aquifers. These deposits contain water in their pore spaces, but the water has salt dissolved in it, so the assumption is that this water is no good for anybody and we might as well inject CO₂ into it. Methane gas has remained stable in deep Earth reservoirs for hundreds of millions of years, so the idea is in principal possible. These aquifers are thought to have the capacity to store 10,000 Gton of C as CO₂.

Scientists are also discussing the possibility of sequestering CO₂ in the deep ocean. CO₂ released into the atmosphere will ultimately mostly dissolve in the ocean anyway; 75% of the carbon is in the ocean when the air and water reach equilibrium after hundreds of years. One could envision direct injection into the ocean as simply bypassing the transit of CO₂ through the atmosphere. CO₂ released in the ocean will equilibrate with the atmosphere toward the same equilibrium point; 25% of ocean-released CO₂ would escape to the atmosphere after hundreds of years. The immediate difficulty with ocean injection is that the CO₂ is very concentrated right near the injection site, acidifying the water ([Chapter 9](#)) and killing marine life.

You may have read about the idea of fertilizing the plankton in the ocean to take up CO₂ as a means of carbon sequestration. The idea is superficially attractive in that plankton in many parts of the ocean, in particular in the cold surface waters around Antarctica, are starving for tiny amounts of the element iron. One atom of iron would allow plankton to take up 100,000 atoms of carbon (more or less). Iron is typically supplied to the surface ocean in dust blown through the atmosphere, but the Southern Ocean is remote enough that not much dust blows down there. When oceanographers add iron to surface water, the plankton bloom. The idea is that the plankton would grow, die, then sink to the deep ocean, thus taking charge of the jobs of extracting CO₂ from the surface ocean, ultimately from the atmosphere, and transporting it to the deep ocean for us. The problem is that models of the carbon cycle in the ocean and atmosphere predict that even if the entire Southern Ocean could be successfully fertilized, the effect on atmospheric CO₂ in the coming century would be discouragingly small. The reason is that it takes hundreds of years for the atmosphere and ocean to reach equilibrium. Ocean fertilization could have a larger impact on atmospheric CO₂ if we were willing to wait 500 years. The impact in the next decades is too small to be worth it.

The last possibility to mention is idea of deliberately altering the climate in such a way as to counteract the warming effects of rising CO₂ concentrations. One possibility is to deliberately inject sulfate aerosols into the stratosphere. Sulfate aerosols cool by scattering light. Particles in the stratosphere remain there for several years, as opposed to aerosols in the troposphere where rain scrubs them out in a few weeks. The cooling effect of volcanic eruptions such as Mt. Pinatubo has been clearly documented. The

aerosols could be shot up to the upper atmosphere in 1-ton shells fired by large cannons. Relative to most of the alternatives, cooling the earth by this method seems inexpensive and benign. My own personal objection to this idea is that the warming effects from CO₂ will last for centuries, even millennia, committing mankind to taking an active role in climate maintenance essentially forever for a few decades of carelessness. Other proposals for climate engineering include the placement of a large reflective object in space, in orbit around the earth or at the stable “Lagrange” point between the earth and the sun where objects can sit indefinitely. The space-based option would be more costly, and would require active removal if its effects were eventually judged to be detrimental, but it would not require ongoing participation in order to continue working.

Summary

In summary, we appear to have multiple options for dealing with climate change within the context of continued economic growth, but none of the options are trivial or without obstacles. Socially, politically, and technologically, the issue of climate change poses a challenge to mankind, on a larger scale than humankind has ever had to face before.

Take-home points

Human-induced climate change is an example of the tragedy of the commons effect. Economists refer to climate change as an external cost, which can cause the market to make poor decisions. Economic forces also keep the market focused on short-term profit and cost, disregarding costs that come far in the future.

International negotiations under the auspices of the United Nations have resulted in a treaty called the Kyoto Protocol to limit CO₂ emissions. The emissions cuts mandated by the Kyoto Protocol are small compared to what would be needed to avoid climate change. The treaty also only imposes cuts on industrial countries, who are currently emitting the most CO₂. For these reasons, the Kyoto Protocol can be only the first step toward preventing anthropogenic climate change.

The ultimate goal of the negotiations is to prevent “dangerous interference” with the climate system. The definition of “dangerous” within this context might be in terms of temperature, or in terms of sea level change.

Lab

1. Compound Interest. Using equation 4-1 in Chapter 4, and assuming an interest rate of 3%/year, how much would an investment of \$1.00 made today be worth in the year 2100? What if the interest rate were 5%/year?

2. ΔT_{2x} . Using equation 3-1, calculate the temperature that would result from increasing atmospheric pCO₂ from pre-anthropogenic value of 280 ppm to 1000 and 2000 ppm. The temperature during the Cretaceous period might have been 6°C warmer than today.

Using a ΔT_{2x} value of 3°C , what must the $p\text{CO}_2$ have been at that time? How does your answer change if ΔT_{2x} is 4°C ?

2. Carbon-Free Energy. The Kaya Identity web page actually runs a carbon cycle model to predict the atmospheric $p\text{CO}_2$ response to its predicted carbon emissions. You will learn more about this model in the next chapter, but for now let's take its output at face value. The Kaya web page then computes how much coal energy would have to be replaced by carbon-free energy, if we wish to stabilize atmospheric CO_2 at some concentration (choices are 350, 450, 550, 650, and 750 ppm). Using the default web page settings, which are something like a business-as-usual scenario, find from the plot the amount of energy in terawatts required to stabilize CO_2 at 450 ppm.

a. If a typical nuclear reactor generates 1000 megawatts of energy, how many power plants would be required by the year 2100? (the prefix tera means 10^{12} , while mega means 10^6). How many power plants would this require?

b. A modern windmill generates about 1 megawatt of energy; let's say that future ones will generate 10 megawatts per tower. How many of these would be required to meet our energy needs by 2100? The radius of the earth is 6.4×10^6 meters. What is its surface area? Land occupies about 30% of the surface of the earth; what area of land is there? Let's assume that windmills could be placed at a density of four windmills per square kilometer. What fraction of the earth's land surface would be required to supply this much wind energy?

4. Carbon Stabilization and Kyoto. How much would we have to cut emissions to stop atmospheric $p\text{CO}_2$ from rising beyond current levels? You could just randomly plug numbers into the model, or you could do it a smarter way: Run the model for BAU (climate geekspeak for business-as-usual) and determine from the printed output below the plots what the rate of CO_2 uptake is by the ocean today (call it year 2000). You'll have to take the difference in the cumulative ocean inventory between two adjacent time points, and divide by the number of years between those two points (five), to get Gtons of C uptake per year. Plug these numbers into the model and run it, to see if that really stops $p\text{CO}_2$ from going up. What percentage change in emission is this? The Kyoto Protocol aimed to limit emissions to some percentage below 1990 levels. The Japanese and Europeans argued at the high end, for something like 9% reductions below 1990. The Americans and the Saudis were at the low end, say 2% below 1990. Compare this range of emission with your answer