6

Conceptions and Misconceptions of Pressure

> *"I'm becoming frightened, Doctor," she said, in a low voice. She smiled as she spoke, but I could see that she was seriously worried. "Do you know what the barometer reads? Twenty-eight seventy! And the wind is growing stronger every minute."*

Charles Nordhoff and James Hall surely would not have included this scene in their rattling good tale, *The Hurricane,* had they doubted that a mere barometer reading would convey an apprehension of impending disaster. For who, however rudimentary his knowledge of meteorology, does not associate the rise and fall of barometric pressure with changes in the weather? Indeed, there is no physical quantity the variation of which in time and space is more important to weather forecasting. What, then, is pressure? Although it is true that sea level atmospheric pressure is a measure of the weight of the atmosphere, this is not sufficient to explain all phenomena associated with atmospheric pressure and its variations. Moreover, a fluid's pressure and its weight are two different concepts; they are not synonymous, for if they were then one of them could—and should—be discarded, at no loss to physical understanding and at a gain in economy.

Like temperature, pressure in its everyday sense is a prescientific concept, and you do not have to know what it is to experience it. This is obvious to anyone who dives under water and feels his surroundings squeeze him more tightly the deeper he descends. This experience of pressure is as old as humankind. Yet only comparatively recently did it become evident that the atmosphere, like water, is also a fluid capable of exerting a pressure. This realization was long in coming because we do not ordinarily sense atmospheric pressure. Awareness of atmospheric pressure comes with the ability to measure it, and instruments to do so were not invented until the seventeenth century.

DEMONSTRATING ATMOSPHERIC PRESSURE

There are so many demonstrations of atmospheric pressure that it would be futile to labor mightily on devising a new one or varying somehow an old one. Instead, I shall describe the one I consider to be the best of the lot. It was reported by Pieter Visscher of the University of Alabama in a short note to the *American Journal of Physics* (1979, Vol. 47, p. 1015).

Very little is needed for this demonstration: a hotplate, an empty soft (or hard) drink can, and a shallow pan filled with water. Put a small amount of water (about a spoonful) into the can and place it on the hotplate. The water

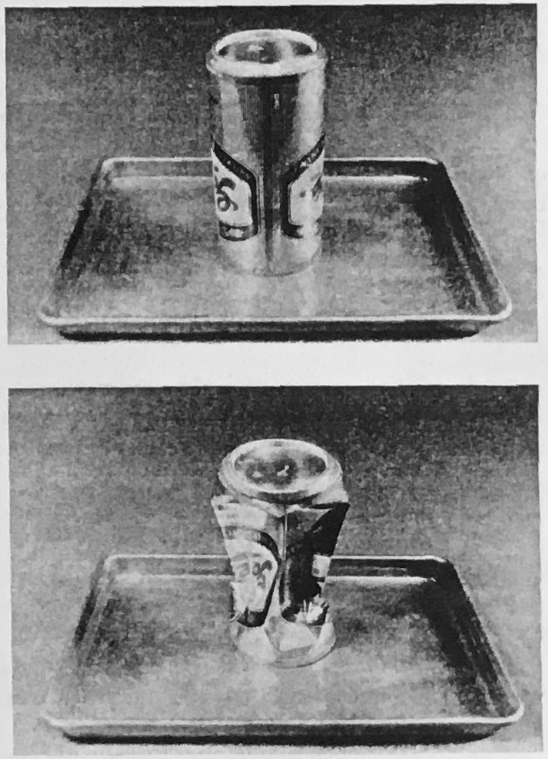

Figure 6.1 When a can filled with hot water vapor is placed in a shallow pan filled with cooler water, the can is crushed because of the pressure difference between its inside and its outside. That outside is atmospheric pressure. That inside is initially atmospheric pressure but drops as the water vapor condenses onto the cooled walls of the can.

temperature will rise quickly and the can will fill with water vapor. Unlike the gaseous constituents of the atmosphere, water vapor is readily condensed (this is the distinction between a gas and a vapor). To do so just snatch the can from the hotplate and plunge it mouth first into the pan. In a few seconds, if all goes well, the can will be crushed (see Fig. 6.1) accompanied by a dull pop. Very dramatic. It sometimes takes more than one try depending on how successful you are at finding something with which to grasp the can.

Tongs of the proper span would be best, but these never seem to be available so I often improvise them using whatever is at hand. As a consequence, I sometimes drop the can or it ends up on its side in the water—a fizzle. But never mind, the demonstration is easily repeated, and even the most fumbling efforts are usually met with success. Be sure to use a *shallow* pan. If you do not, all that is likely to happen is that the can will sink into the water without being crushed.

In Chapter 2 I discussed the concept of saturation vapor pressure—the pressure of a vapor in equilibrium with its liquid—and in Chapter 5 I emphasized how rapidly it increases with increasing temperature. For example, in going from 40° to 100°C (the normal boiling point) the saturation vapor pressure of water increases by more than a factor of ten. It is this large change that crushes the can.

The can sizzling on the hotplate was filled mostly with water vapor at a pressure about that of the surroundings (at the boiling point the saturation vapor pressure is that of the surroundings). When this hot can was plunged suddenly into room-temperature water it cooled, hence the rate of evaporation from its inner surface decreased. The rate of condensation onto it, however, was slower to respond. So there was *net* condensation of water vapor for a moment, and the pressure inside the can consequently dropped to well below that of the surroundings. Thin aluminum cans are not able to withstand such large pressure differences so they collapse. You may wonder why the pressure difference doesn't merely force water into the can, like mercury in a barometer. With a much stouter can—one made of cast iron, for example—this is what would happen. And it would also happen even with a thin aluminum can were it not for the inertia of the water: before the water has sufficient time to flow in, the can collapses.

This demonstration of the *existence* of atmospheric pressure is unambiguous. It is only when we *interpret* what we observe that the ambiguities and contradictions arise.

It is often asserted that atmospheric pressure is just the weight of the atmosphere above a unit area. While this is true in the absence of vertical acceleration, it does not come to grips with the concept of pressure in a fundamental way. For gas pressure and weight are, in general, independent. To convince yourself of this consider the following thought experiment. Fill a sealed container with a gas, which could be air but need not be. At a given location the weight of this gas—that is, the force exerted on it by the gravitational attraction of the earth—is fixed. Yet it is a matter of experience that the gas pressure would rise if the container were heated. And if the gas were taken beyond the pull of the earth's gravitational field it would become weightless but not pressureless. Thus the relationship between pressure and weight is not unique. Moreover, the weight interpretation of pressure does little to help us understand the demonstration. There is more to the pressure concept than that it is an alternative guise for weight.

INTERPRETING PRESSURE

To our coarse senses all matter is *continuous:* we can discern no gaps, no lumps. This is particularly true of the common fluids air and water. Despite this, the hypothesis that matter is ultimately *discrete,* the atomic (or molecular) hypothesis, has had many adherents almost since the beginning of recorded thought. And its appeal is not to be wondered at, for we too are atoms. This does not mean that we are very small: the root meaning of atom, from the Greek *atomos,* is indivisible. Neither we nor atoms can be divided, at least not without losing our salient characteristics: cut a man in two and he just won't be the same.

Although the existence of atoms was hotly debated for centuries, it was only during the first decade of this century that enough evidence had been amassed to silence the critics, among whom were counted some of the most eminent scientists of the day. Today, anyone who seriously denies the existence of atoms and molecules (although not strictly indivisible, a molecule is a more or less stable aggregation of atoms) is classed with the flat-earth believers. But in one sense the critics of the molecular hypothesis were right: it is sometimes neither necessary nor even desirable to take account of the discreteness of matter. For many purposes it is more expedient to proceed *as if* matter were continuous.

Few people have not at least heard of molecules. Unfortunately, what they have heard is often so fragmentary, so superficial, so distorted, that it would have been better had they not. A little bit of knowledge is sometimes less than none at all. For example, if one adopts the continuum interpretation of matter, then one has little choice but to consider the pressure within a gas as resulting from adjacent portions of a continuous medium pushing against one another. If one now switches to the molecular interpretation it is only natural to carry over this notion, that is, to consider gas pressure as resulting from gas molecules pushing against one another. As attractive and widespread as this notion is, it is nevertheless false. The error arises in going from one interpretation to another. Translation errors, sometimes the cause of embarrassment, are not uncommon in going from one language to another. Concepts in one language do not necessarily have exact counterparts in another. So also is it with different languages for describing the physical world. Let us explore how one such language, that of atoms and molecules, describes gas pressure.

MOLECULAR INTERPRETATION OF GAS PRESSURE

Even after one adopts the molecular viewpoint, one is confronted with two aspects of the concept of gas pressure: as a *measurable* quantity and as a *property* of the gas. I shall consider each of these in turn.

Consider a gas in a rigid container. In even a tiny volume of this gas there is an enormous number of molecules: in one cubic millimeter, more than a million billion of them (under normal conditions). Molecules that collide with the walls of the container will suffer a change in ***momentum***. Momentum, which is the product of mass and velocity, is a directed quantity. Thus even if the speed of a molecule—how fast, but not where, it is going—does not change upon collision, its momentum does: before collision it was moving ***toward*** the wall, after collision it is moving ***away*** from it.

Newton's second law states that a body will not change its momentum unless a force acts on it. And from his third law—to every action there is an equal and opposite reaction—it follows that if the wall exerts a force on a molecule that collides with it, each such molecule exerts an equal and opposite force on the wall (this is painfully obvious to anyone who has ever run headlong into a wall). Because of their great number and speed (at room temperature the average speed of an air molecule is about that of a rifle bullet) molecules collide with the wall at a tremendous rate: in a millionth of a second, each square centimeter of wall is struck by billions upon billions of molecules. And each of them contributes its mite to the pressure, the total force acting on a unit area of the wall. This is the force that causes the liquid in barometers to rise and that caused the can to collapse: the total force due to molecular collisions with the inside of the can was overwhelmed by that due to collisions with the outside.

Note that there is an important distinction between the target (i.e., a surface) and the projectiles (i.e., gas molecules) impinging on it. The average separation between molecules in a gas (at normal temperatures and pressures) is about ten times that in a solid or liquid. The density of solids and liquids is therefore about a thousand times greater than that of gases. As a consequence, the rate at which gas molecules collide with one another cannot be anywhere near the rate at which they collide with a liquid or solid surface. Indeed, to good approximation air may be treated as an ***ideal*** gas, another name for which is a ***collisionless*** gas. The very name gives the game away. It is not that collisions between gas molecules do not occur, or even that they are not important in determining ***some*** properties of gases, it is just that they are not the determinants of such quantities as gas pressure. As far as pressure is concerned, the molecules might just as well be collisionless—with one another, of course. And the same holds true for temperature. Two gases with greatly different rates of intermolecular collision (e.g., two gases with greatly different densities) can be at the same temperature.

When gas pressure is measured, it is always by virtue of molecules interacting with a solid or liquid surface (e.g., a barometer). If we were concerned only with the measurability of pressure then we would need say no more than this. But what meaning is to be attached to pressure construed as a property of the gas? That is, how should we interpret pressure within a gas as opposed to at its interface with a liquid or solid?

The dimensions of pressure are force per unit area, as are those of ***momentum flux,*** the rate at which momentum is transported across a unit area. This is not just a coincidence. Recall that pressure was associated with the rate at which gas molecules transfer momentum to unit area of a material surface at the ***boundary*** of a gas. Now imagine a point ***within*** a gas. Molecules in the neighborhood of this point are whizzing about in all directions, but their total momentum is zero if the gas is at rest: for every molecule with a given momentum there is another with the opposite momentum. Contrary to what you might expect, however, the total momentum flux is not zero. Consider an arbitrary direction, call it the positive direction (the opposite direction is the negative direction). Each molecule has a component of momentum, either positive or negative, along this direction. Moving molecules carry their properties, one of which is momentum, with them wherever they go; that is, molecules transport momentum. The reason that the total momentum flux is not zero, even though the total momentum is, is that positive momentum transported into a region is equivalent to negative momentum transported out. A simple analogy comes to mind: getting rid of a debt (an outflow of negative money) is equivalent to a raise in salary (an inflow of positive money). Lest I mislead anyone, I must emphasize that I have in mind momentum transported across a unit area of an imaginary ***flat*** surface in a gas. The net rate at which momentum is transported across a ***closed*** surface in a gas is zero if there are no forces acting on the molecules.

According to the molecular interpretation, the pressure in a gas is the total rate at which its molecules transport their momentum, along a given direction, across a unit area perpendicular to that direction. Regardless of the direction chosen, the corresponding momentum flux is the same. Consequently, the pressure acts equally in all directions. Pressure has essentially nothing to do with the rate at which gas molecules collide with one another. This rate is small compared with the rate at which they collide with a solid or liquid surface, and it is by means of the latter that pressure is measured.

How does the molecular interpretation elucidate observations of atmospheric pressure, such as its decrease with height? This is usually explained by saying that the weight of the atmosphere above a given elevation decreases with increasing elevation. This is true enough, but the observation has an alternative explanation. At each successive elevation lower pressure implies lower momentum flux. This means that for any horizontal layer of atmosphere more vertical momentum goes in at its bottom than comes out at its top. What happened to this momentum? According to Newton's second law, momentum changes (e.g., decreases) only if caused to do so by a force. In this instance, the force is that of gravity. And weight is, after all, merely the force of gravity. So the two interpretations—microscopic and macroscopic—are in harmony. But they can be made to clash unnecessarily by mixing concepts appropriate to one with those appropriate to the other. Consistency is what is wanted. Once you have adopted a particular viewpoint, stick with it. If you must change, beware of translation errors.

The choice of interpretation is to some extent a matter of taste. More often it is guided by considerations of expediency. It would be foolish, for example, to try to do a molecule by molecule forecast of the weather. For this purpose the view that matter is continuous is the most sensible one. For other purposes, however, understanding is acquired only by recognizing the ultimate discreteness of matter.

14

Multiple Scattering at the Breakfast Table

> *white . . . is not a mere absence of colour; it is a shining and affirmative thing, as fierce as red, as definite as black.*
> G. K. Chesterton

The scene shown in Figure 14.1 is familiar enough. In it are the kinds of things likely to be found on your breakfast table: salt, sugar, and a glass of milk. Although dissimilar in many ways, they all are white. Why this is so is the subject of this chapter. And I can think of no better way to begin than with a simple demonstration.

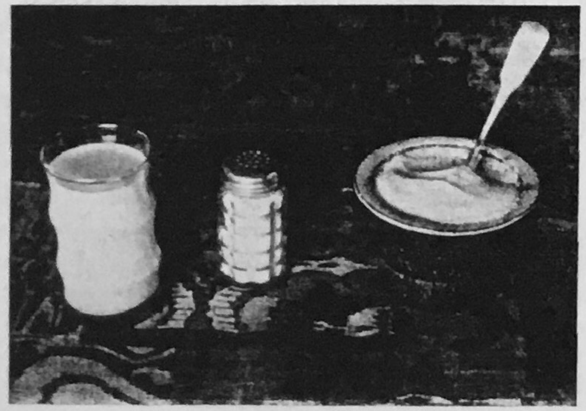

Figure 14.1 Multiple scattering at the breakfast table.

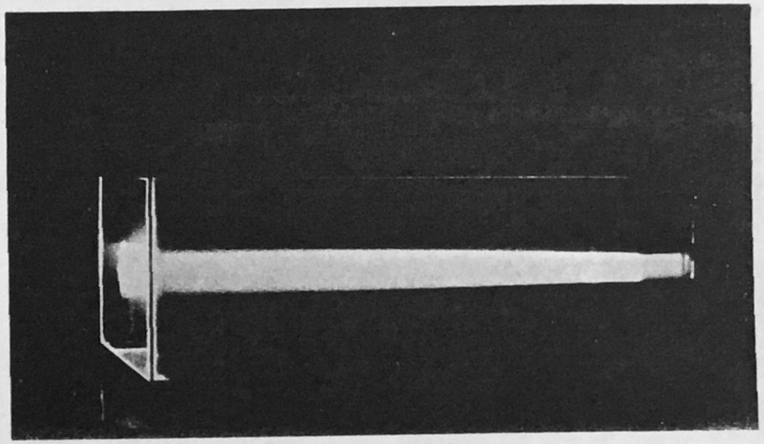

Figure 14.2 Single scattering by a suspension of fat globules (milk) in water.

SCATTERING IN AN AQUARIUM

Scattering of light by particles and by molecules was discussed in Chapters 11, 12, and 13. Without saying so, in these chapters I had in mind what is called *single* scattering in contradistinction to ***multiple*** scattering. The difference between the two kinds of scattering can be demonstrated using a slide projector, an aquarium filled with clean water, and some milk.

A slide projector in which an opaque slide with a hole in it is inserted provides a collimated beam of light. You cannot see this beam from the side unless it traverses a medium that scatters some of the light from the beam toward your eyes (this, by the way, is why movie scenes in which combating spaceships fire brilliant laser beams at one another are absurd: in space, where there is little to scatter light, you would see a beam only if it were pointed directly at you).

If the projector beam is shone into an aquarium filled with clean water you may, in a very dark room, barely be able to see the beam. To make it more evident add a few drops of milk to the water (Fig. 14.2). Milk contains tiny globules of fat which scatter light from the beam, and it is attenuated along the direction it propagates. Its margin is well defined as a consequence of single scattering: a photon (see the previous chapter) must be scattered at least once for us to see the beam (to be precise, we do not see the beam but rather the light removed from it), but its likelihood of being scattered more than once is small. To increase this likelihood add more milk to the water; that is, increase the concentration of scatterers. The result is shown in Figure 14.3. Note that light now comes from beyond the limits of the beam where previously none had been evident (Fig. 14.2). This is because photons scattered by particles in the beam are scattered again by particles outside it and thence to our eyes. Multiple scattering increases the number of ways in which a photon may reach our eyes.

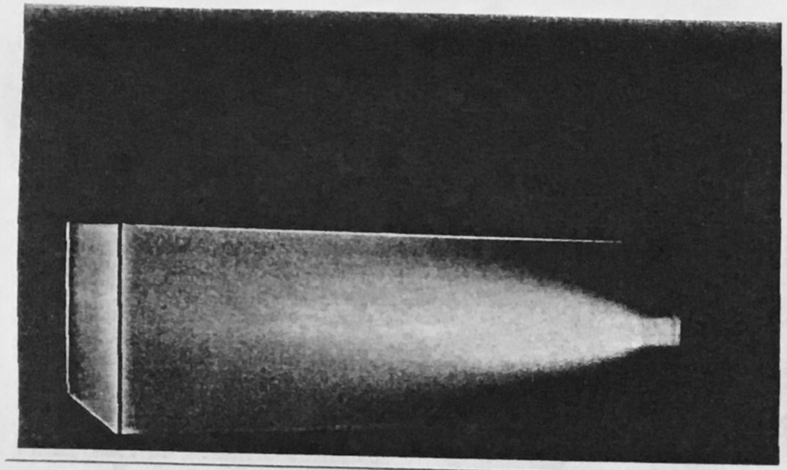

Figure 14.3 Multiple scattering by a suspension of fat globules (milk) in water.

A PILE OF PLATES

Plates cut from clear plastic sheets provide raw material for demonstrating other aspects of multiple scattering. One such plate, on a black background, is weakly reflecting; but a pile of them is strongly reflecting (Fig. 14.4). Here is an example of an ensemble—a pile of transparent plates—with optical properties quite different from those of its individual members. As plates are successively added to the pile it reflects more light, but beyond a certain number, in which instance the pile is said to be *optically thick,* each additional plate yields an ever smaller increment. So it is also with any optically thick scattering medium, such as a cloud. A single water droplet does not scatter much light, but because of multiple scattering an optically thick cloud of such droplets reflects much of the visible light incident on it, which is obvious to anyone flying over a thick layer of clouds.

For a collection of scatterers (e.g., a pile of plates, a cloud) to be bright and white upon illumination by white light it must not only be optically thick but its members must only *weakly* absorb such light. Consider a pile of plates in photon language. Most of the photons incident on a single plate are transmitted to the underlying black surface where they are absorbed. But with two plates in a pile some of the photons transmitted by the first are reflected by the second. Each plate added to the pile increases the probability that a photon eventually gets to your eyes, so with enough plates the pile is white. Again, as with the aquarium demonstration, multiple reflection (scattering) increases the number of ways—reflected once, twice, and so on—in which a photon can get to your eyes. But multiple scattering takes as well as gives: the greater the number of scatterers the greater the chance that a photon will be absorbed; multiple scattering not only increases the pathways by which incident photons can reemerge from a medium, it exposes them to a greater hazard of being absorbed. I shall return to this later.

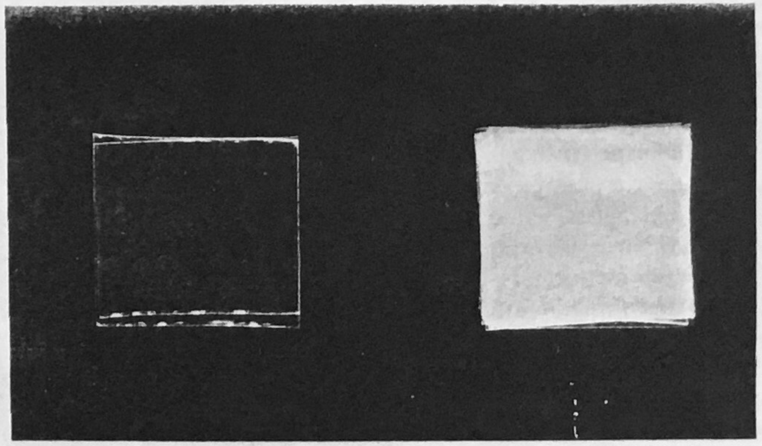

Figure 14.4 A pile of transparent plates on a black background.

WHITE CLOUDS: THE CONVENTIONAL WISDOM

It has been stated countless times that clouds are white because they are composed of droplets sufficiently large that they scatter visible light of all wavelengths about equally. Indeed, this notion is so widespread and has so many adherents that it now transcends science and has become an article of faith; to believe it is a mark of piety. I can therefore anticipate the great howling that will be raised in response to my assertion that this explanation is demonstrably false. At the very least, those who advance it fail to distinguish between a *necessary* condition and one that is merely *sufficient*.

For a cloud to be white (upon illumination by white light) it is sufficient that its droplets scatter visible light of all wavelengths about equally—but it is not necessary. If it were, a glass of milk would be blue. For milk is a suspension of particles that, unlike cloud droplets, scatter blue light more than red light; this was why milk was used for the green flash demonstration discussed in the previous chapter. Yet a glass of ordinary milk is white, not blue. Why? The single-scattering characteristics of particles in milk are really quite irrelevant to the appearance of a glass of milk. What is relevant is that it is optically thick—a small glass of milk looks no different from a large glass—and its particles are very weakly absorbing. Although incident photons corresponding to the color red may have of Chapter 2. It is not the amount of water in a cloud that makes it appear so different from the patch of sky in which it was born but rather the state of aggregation of this water. Why? The answer lies in the concept of coherence, which I discussed briefly in the previous chapter.

A group of water molecules when randomly separated scatter incoherently; when these same molecules condense into a droplet (i.e., they all are part of the same entity rather than independent agents) they scatter coherently. And coherent scattering can be much more intense, all else being equal, than incoherent scattering. Suppose that a group of hare-brained people is trying to push a stranded car. If they push incoherently (i.e., each individual pushes randomly when and in whatever direction he chooses) the car is not likely to do more than rock back and forth, possibly inching forward laboriously. Now suppose that a leader emerges from the group and urges everyone to push coherently: One, two, three, Push! In no time at all the car will be on its way. I have implicitly assumed that the pushing is done coherently *in phase:* coherence by itself is not sufficient unless directed to the same end (I shall have more to say about coherent scattering in Chapter 18). Just as coherent (in phase) pushing yields a much greater effect than incoherent pushing, so also is coherent scattering vastly more intense even though the scatterers are the same in both instances.

While I am on the subject of coherence and incoherence I should note that although a single cloud droplet is a coherent scatterer, a group of them scatters incoherently. If you are reeling after this apparent contradiction, I'll back up a bit. A group of N isolated water molecules in the atmosphere scatters incoherently: scattering by N is N times scattering by one. But when these same N molecules condense to form a droplet, scattering by them is much greater. Now consider a group of water droplets. If they are as far apart as those in atmospheric clouds, scattering by N droplets is N times scattering by one. That is, the droplets considered as a group are incoherent scatterers, whereas each droplet itself is a group of coherent scatterers. The multiple scattering I have discussed in this chapter (and will discuss further in subsequent chapters) is incoherent multiple scattering. Now let us imagine that a cloud of water droplets is compressed until it forms a continuous film. Again, what is observed is predominantly coherent scattering, more intense than incoherent scattering by the cloud but confined mostly to two directions. For example, sunlight reflected by a film of water is much brighter than that reflected by a cloud, at least a thousand times greater, possibly much more depending on the angle of incidence of the sunlight (to convince yourself of this compare the reflection of the sun in a pond with light from a cloud). But this reflected light is concentrated in a single direction, that given by the law of reflection, whereas the cloud scatters light in all directions. Light from the fogged mirror of Chapter 7 has about the same brightness in all directions; the unfogged mirror is much brighter, but only in one direction.

In going from a group of independent molecules to independent droplets to a water film, nothing changes except the arrangement of the molecules. Chemically they are the same. Yet what is observed is markedly different depending on the degree of coherence—the extent to which the molecules stick together, both figuratively and literally—of their scattering.

ANOTHER PILE OF PLATES

As long as we have plastic plates at hand we might as well use them to demonstrate yet another salient characteristic of multiple scattering: a little bit of absorption goes a long way. In the following section the relevance of this demonstration to the atmosphere is discussed.

I implied that the plates shown in Figure 14.4 are weakly absorbing. And so they are. But they are not nonabsorbing. To show this, take a single plate and lay it on a white background. Then add plates. With each additional plate the pile gets darker (Fig. 14.6). Multiple reflection exposes photons to multiple opportunities to be absorbed by the plates. The piles of plates in Figures 14.4 and 14.6 do not look the same because of different contrast and because they were photographed under different illumination and at different exposure times. Yet the two piles have the same reflectance (fraction of incident light reflected), which is independent of the underlying surface if the pile is optically thick. You can verify this with a pile of plates that straddles the boundary between black and white pieces of paper.

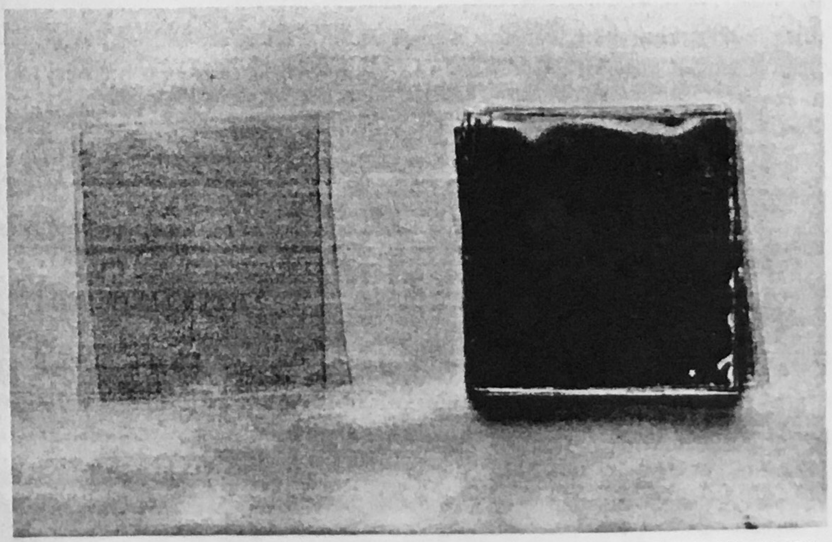

Figure 14.6 The same pile of plates as in Figure 14.4 but on a white background.

to rattle around more in a glass of milk before re-emerging than those corresponding to blue, almost all of them ultimately re-emerge having escaped absorption. It is these two characteristics—optically thick and weakly absorbing—that unite the dissimilar objects with which I began this chapter. Regardless of their single-scattering characteristics, if enough weakly absorbing particles are heaped into a pile it will be white if the source of illumination is white. Besides the examples of milk, salt, and sugar in Figure 14.1 there is snow, flour, powdered glass, white sand—the list is endless. Taken individually these particles may have vastly different scattering properties; collectively, however, they are nearly identical.

Every glass of milk refutes the conventional explanation of why clouds are white. In fairness to those who espouse it, I might add that the scattering characteristics of individual particles in optically *thin* clouds are relevant to their appearance. For example, thin cirrus clouds are white—although not nearly as bright as cumulus clouds—because their particles are large compared with visible wavelengths; in contrast, noctilucent clouds—very tenuous high-altitude clouds—are often described as bluish, which is offered as evidence of the smallness of their particles. But I think it also fair to say that people who ask why clouds are white do not have in mind wispy cirrus clouds but, rather, proper clouds: big, fluffy, towering cumuli.

As a harsh critic of what I consider to be superficial explanations, I would be wise to forestall criticisms of mine. It might be—indeed, has been—asserted that reflection and scattering are fundamentally different processes, in which instance the two experiments described in preceding paragraphs are unrelated: the first demonstrates multiple scattering whereas the second demonstrates multiple reflection. Such logic-chopping would make a medieval theologian blush. I argued in the previous chapter that the term scattering embraces reflection as well as refraction. We continue to use these terms separately for historical reasons (i.e., intellectual inertia) and because they are convenient ways of summarizing what is observed. But other than metaphorically, light does not bounce off of a window pane. What we call the light reflected by glass is the sum of all the light scattered by its constituent molecules; if we ignore these details it is for convenience not by necessity.

CLOUDS AND SNOW

You may have noticed that clouds are often not as bright as snow (Fig. 14.5). This is a variation on the black cloud theme, which was discussed in Chapter 11. Of course, to compare fairly a cloud with a snowpack the conditions of illumination and viewing must be identical, and this is not always easy to obtain. Nevertheless, I have yet to see clouds brighter than the brightest snow. This is not because cloud droplets are more absorbing than ice grains in snow. Indeed, the reverse is true: the ratio of incident light scattered to that absorbed by a single cloud droplet is *greater* than that by single ice grains in snow because the grains are much larger (see the next chapter for more on this).

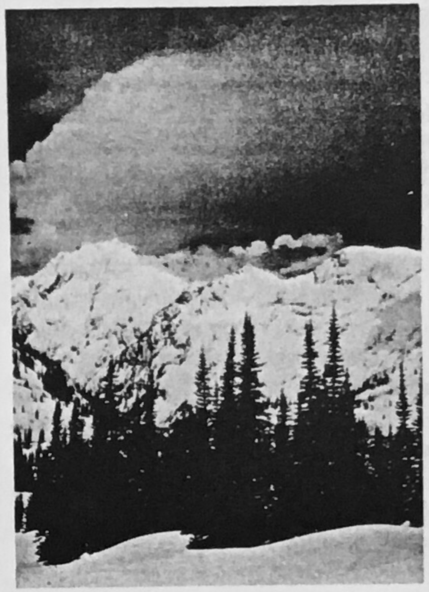

Figure 14.5 All else being equal, snowpacks are usually brighter than clouds, as in this winter scene near Alta, Utah.

Both clouds and snowpacks are multiple-scattering media. The difference between them lies in their different *optical* thicknesses, that is, their thicknesses measured in units of mean free paths (see Chapter 16). Snowpacks are usually optically thicker than clouds. Indeed, snow on the ground is often effectively infinitely optically thick. Except for very shallow snowpacks, the addition of another layer does not sensibly change the fraction of the incident light it reflects. You may have observed this many times. Snow a few inches deep is indistinguishable from that a few feet deep. Clouds are darker than snow not because they absorb more light—they absorb less—but because they transmit more of it to their surroundings.

AN APPARENT PARADOX

Before discussing yet another aspect of multiple scattering, there is one other matter to be disposed of. It may have occurred to you that every cloud presents what at first sight is a paradox. Water molecules, like all the other molecular constituents of the atmosphere, scatter light. But when a given amount of water vapor condenses to form droplets the resulting cloud scatters much more than the water molecules did. This was evident in the cloud bottle demonstration iting beam—the (greatest) angle it makes with the incident beam—changes as the angle of incidence changes; this will be more apparent if you project the exiting beam onto a screen. All the incident beams must be parallel to one another, which can be achieved by fixing the laser in place and moving the beaker along a line perpendicular to the beam. Move the beaker so that the beam makes all angles of incidence with the beaker, from zero degrees (normal incidence) to 90 degrees (glancing incidence). As the angle of incidence is increased, the deviation angle of the exiting beam will first decrease uniformly then increase uniformly (see Fig. 21.1). This shows that there is an angle of incidence such that the deviation angle of the exiting beam is a minimum; this is the primary rainbow angle.

Light internally reflected more than once before exiting also gives rise to rainbow angles, an infinity of them. Each successive bow, however, is less bright than its predecessors. Despite travelers' tales to the contrary, natural rainbows of orders higher than the second are rare—I have yet to see a photograph of one—because they fade into the background illumination. Recently, however, what appears to be a reliable observation of a tertiary rainbow was reported by D. E. Pledgley (*Weather*, December 1986, p. 401). Unlike the primary and secondary bows, the tertiary bow is seen looking toward the sun, about 40 degrees from it. Bows of even higher order may be observed in the laboratory, where background illumination can be eliminated. To learn more about them I highly recommend Jearl Walker's article in *American Journal of Physics* (1976, Vol. 44, p. 421).

AN INDOOR RAINBOW

The laser demonstration described in the previous section shows that at special angles—rainbow angles—light scattered by a water drop is much more intense than at other angles. But a rainbow is more than just an intense arc of light. Because the angle of minimum deviation depends on wavelength, light of different colors is separated in natural rainbows. Indeed, "rainbow" evokes images of a profusion of colors. A polychromatic source of light is therefore required for a rainbow—as that word is commonly understood—rather than a nearly monochromatic source such as a laser.

Glass globes filled with water and illuminated by slide projectors (or sunlight) have been used to demonstrate rainbows. In my experience the rainbows that result are adequate but not very impressive (for a different viewpoint, see the letter by Arvid Skartveit and Frank Cleveland in the October 1985 issue of *Weatherwise*). This is because only a small fraction of the incident light contributes to the various rainbows, which are formed by light that undergoes one or more internal reflections; most of the incident light is transmitted without internal reflection.

A colleague of mine at the University of Arizona, Donald Huffman, devised a clever way of greatly enhancing the brightness of the rainbow obtained with a water-filled globe: he deposited, in a vacuum, a mirror coating on part of the globe. The rainbow angles are not changed, but the brightness at these angles is greatly increased: light that would have been lost to the rainbow is directed into it by the mirror.

Although Huffman's variation on a common rainbow demonstration is a great improvement, it requires apparatus that is not readily available. In a discussion about this with Dennis Thomson, one of my colleagues at Penn State, we hit on the idea of wrapping a flexible mirror around a beaker. The mirror was a sheet of aluminized plastic film (not aluminum foil, which is likely to be too crinkly). The beaker shown in Figure 21.3 has such a mirror covering part of its surface; this is why you see no light transmitted through the beaker. I also tried painting part of the beaker with aluminized paint; the results were less satisfactory, although someone who takes more care might have more success than I did. In any event, aluminized film is almost as common as aluminized paint.

You need not cover an entire half of the beaker with aluminized film. The ray diagram in Figure 21.1 will guide you in determining what fraction of the beaker's surface needs a mirror. For example, if you are interested in angles near the rainbow angle, only a small amount of the surface needs a mirror.

Figure 21.3 shows what is obtained using a water-filled beaker with a mirror covering part of its surface. The source of illumination is a slide projector in which an opaque slide with a slit in it is inserted. Note how bright it is inside the bow. This illustrates a point Alistair Fraser often makes: the rainbow is merely the bright edge of a disc of light.

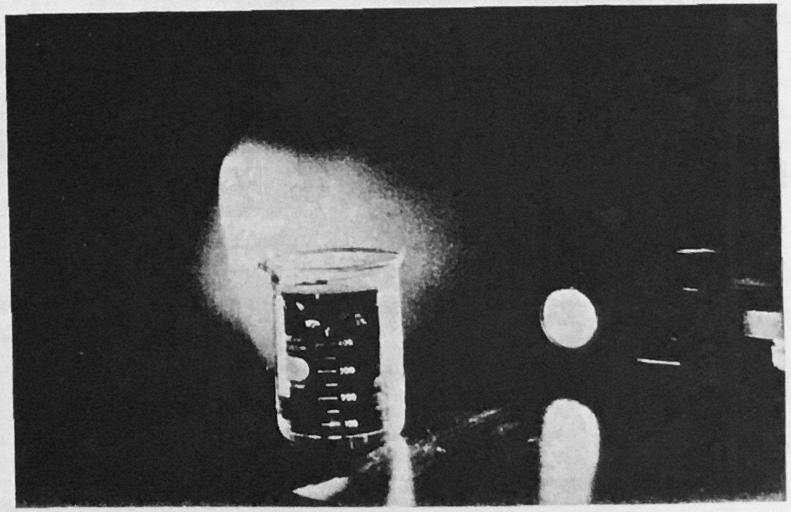

Figure 21.3 An indoor rainbow—without rain.

SUPERNUMERARY BOWS

Geometrical optics applied to large water drops predicts the existence of rainbows, their angular position, and their color separation. But it fails completely to account for supernumerary bows, even though Nature has no trouble making them. Here is one example of the inadequacy of geometrical optics. This is hardly cause for hand-wringing. Geometrical optics is so simple that it is somewhat of a surprise that it works at all. It would be unreasonable to expect it to describe all observable features of rainbows.

Supernumerary bows, as their name implies, are those in excess of the expected number. To understand them requires us to come to grips with the wave nature of light, to which geometrical optics is oblivious.

Alistair Fraser's article in the December 1983 issue of *Weatherwise* is the best exposition on supernumerary bows I have ever seen. It would be pointless for me to try to do better, so I have merely lifted one of his diagrams (Fig. 21.4), which shows the interference pattern produced when waves, rather than rays, are incident on a sphere.

In natural rainbows only a few supernumerary bows may be seen, three, perhaps four, at most. If you lust for more, they are easy enough to obtain in the laboratory. All you need is a hypodermic syringe or a hobbyist's oiler and a laser. If you use a hypodermic needle, you'll have to remove its sharp tip with a grinding wheel.

Fill the syringe with water and clamp it to a stand. With a bit of practice you'll be able to suspend drops from the tip of the needle. Transparent liquids other than water, such as syrups or oils, may be used.

In a darkened room illuminate a suspended drop, as large as possible, with

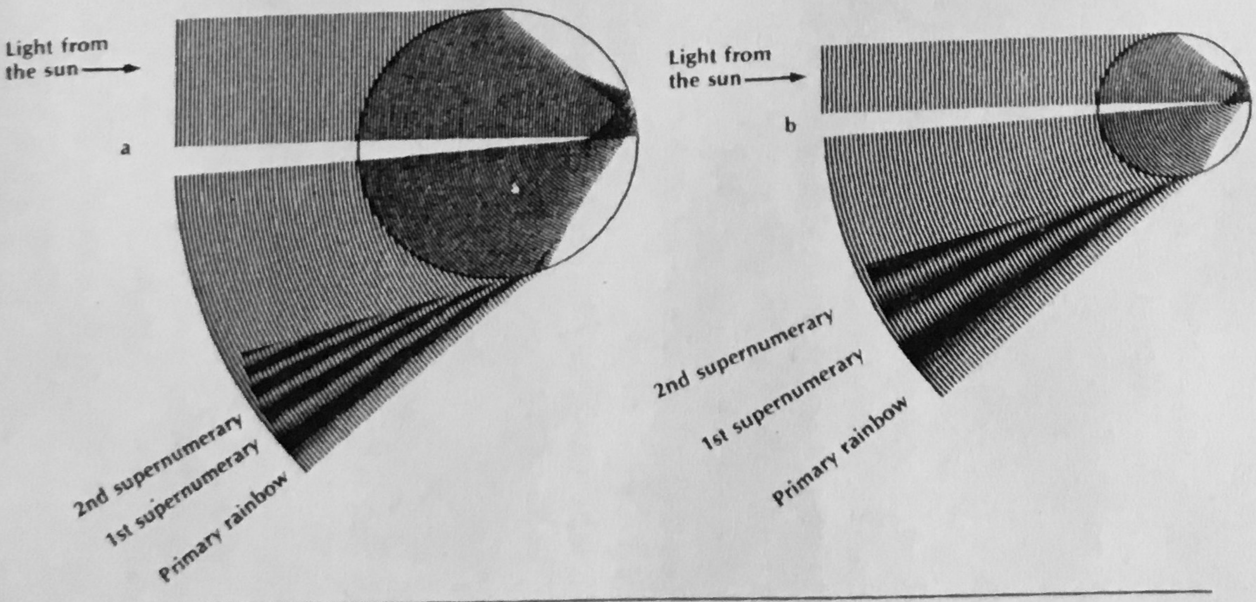

Figure 21.4 This diagram is the same as that in Figure 21.1 except that rays have been replaced with waves. Although rays cannot interfere, waves can, which is what gives rise to supernumerary bows. Figure courtesy of Alistair Fraser.

so by tracing rays, which lead us to the physical origins of these splashes. But we cannot trace rays that will lead us to the origins of the moon illusion. Whenever we see the moon on the horizon, we interpret it as being larger than it really is, and no amount of learning can disabuse us of this. Intellectually, we can accept that the moon is not objectively larger, but it will always appear larger to us. An enlarged moon is a creation of the mind; a mirage is a creation of the atmosphere.

Although I began this chapter with tales of summertime mirages, you don't have to wait until summer; you can see highway mirages all year round, even in winter when the sides of the road are piled high with snow. High temperatures do not give rise to mirages; high temperature gradients—the rates at which temperature changes with height—do.

7

The Greenhouse Effect Revisited

By the influence of the increasing percentage
of carbonic acid in the atmosphere, we may
hope to enjoy ages with more equable and
better climates, especially as regards the
colder regions of the earth
SVANTE ARRHENIUS: *Worlds in the Making* (1908)

Several years ago, I wrote a series of articles in which I rashly predicted that the greenhouse effect might give way to nuclear winter as the global catastrophe most favored by headline writers. Subsequent events have shown otherwise. Nuclear winter is fading into obscurity, whereas only a hermit snug in a cave is spared news of the greenhouse effect. Newspapers hardly let an issue go by without warnings of impending global warming. At breakfast, I learn that Cleveland will become an oasis, whereas the verdant and fertile fields of Iowa will be transformed into an American Empty Quarter.

Concern accelerated during the drought of 1988, taken by some to be a portent of the grim future awaiting us. As the drought stretched on, a trickle of articles on the greenhouse effect turned into a stream, then a raging river, finally a flood of biblical magnitude. It may have crested by now. I even sense a waning, but I have been wrong before.

Although a failure as a prognosticator, I may have been one of the first to observe a new malady, the Greenhouse Effect Anxiety Syndrome (GREAS). Not long ago, a man called my department. He wanted to talk to "a meteorologist," and so was passed on to me by a receptionist with a malicious sense of humor. He did not ask any questions of substance; he just wanted to talk to a sympathetic listener, to vent his anxiety about the impending doom he had read about in a newspaper.

Many of the writers who are giving night sweats to people like him don't bother to learn much about their subject or to inject any freshness into their writing. They rely on stock formulas, reminding us ad nauseam that "the earth radiates heat, this heat is trapped by our atmosphere, which acts like a blanket." These are shaky metaphors at best, shopworn from overuse. We can understand the greenhouse effect without invoking notions of radiation trapping and atmospheric blankets. With this end in mind, I offer the following experiments and meditations on them as a small step toward grasping some of the subtleties of the greenhouse effect.

Cooling Curves

For my experiments, I used a 250 ml glass flask filled with water heated to about 75°C (167°F). Every 10 or 15 minutes, I measured the temperature of the water, thereby generating a cooling curve. According to Newton's law of cooling, the difference between the temperature of a cooling body and that of its surroundings decreases exponentially with time. Thus a plot of temperature on a logarithmic scale versus time on a linear one should yield a straight line. Its slope is inversely proportional to the *cooling time,* which I define as the time required for the temperature difference to drop to half its initial value. The larger the cooling time, the more slowly the flask cools.

Figure 7.1 shows two cooling curves. For a bare flask, the cooling time is about 50 minutes. When the flask is wrapped in aluminum foil, the cooling time increases to 78 minutes. To explain why requires some backtracking.

All bodies emit electromagnetic radiation of all wavelengths at all times. Yet at normal temperatures, say between about 250 and 330°K, the emission spectrum of most terrestrial objects peaks at an infrared wavelength of around 10 μm; shortward of 4 μm and longward of 20 μm they emit much less, but not an unmeasurable amount. Emission is merely small enough that we need not consider radiation outside this range if our interests are limited to energy transfer.

Water in the flask cooled because of energy exchange between it and its surroundings. For the moment, we are concerned only with the radiative part of this exchange: the flask radiates to its surroundings and they radiate to it. The amount radiated by each depends on their temperatures. But this is not all. Radiant emission also depends on the radiator's composition.

Figure 7.1
These curves were obtained by allowing a 250 ml flask filled with heated water to cool in still air. The temperature difference between the water and that of its surroundings is plotted on a logarithmic scale. Data points indicated by crosses are for a bare flask; those indicated by circles are for a flask wrapped in aluminum foil. Because the infrared emissivity of aluminum is less than that of glass, the foil-wrapped flask cools more slowly.

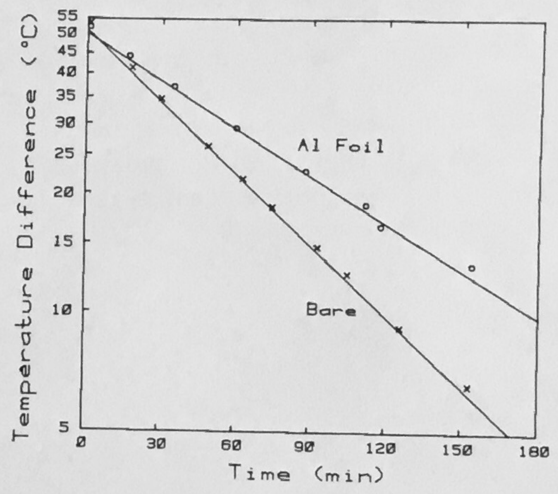

The flask experiments suggest that emission by the foil was less than that by the glass. The foil is so thin and its thermal conductivity so large that it provides little insulation in the sense usually meant—suppression of energy transfer by conduction. But aluminum foil does provide insulation of a different sort, radiative insulation, which brings us to the concept of *emissivity*.

Emissivity

A blackbody is one that absorbs completely all radiation incident on it regardless of wavelength, direction of incidence, and even state of polarization. The emission spectrum of such a hypothetical body depends only on its temperature. At any wavelength, radiant emission by a real body is less than that by a blackbody, both at the same temperature. One way to understand why is to recognize that emission is the reverse of absorption. Suppose that radiation of a given wavelength is incident at a given direction on a body. We may look upon this radiation as a stream of incoming photons, some of which are absorbed by the body. We further imagine that we can film this process (of course, we cannot literally photograph photon paths) and obtain a motion picture of photon absorption. When we run the film backwards, absorbed photons become emitted ones.

The fraction of the incident radiation absorbed by a body is called its absorptivity. Its emissivity is its emission rate relative to that of a blackbody at the same temperature. Because of the film-reversing arguments I made previously, emissivity equals absorptivity. This is called Kirchhoff's law and is strictly valid only for a specified wavelength, direction, and polarization of the radiation. Equality of absorptivity and emissivity is necessary for thermal equilibrium. For example, a body placed in an opaque cavity is bathed in blackbody (or equilibrium) radiation. When the body's temperature becomes constant and equal to that of the cavity walls, the rate of emission by it must equal the rate of absorption.

All radiation incident on a body is not necessarily absorbed; some is reflected and some transmitted. The fraction reflected is the reflectivity; that transmitted is the transmissivity. Because incident radiation can be only absorbed, reflected, or transmitted, the sum of the absorptivity, reflectivity, and transmissivity must be one. If the object is opaque, its transmissivity is negligible, so its absorptivity is one minus its reflectivity. From Kirchhoff's law, the absorptivity is equal to the emissivity, hence that of an opaque, highly reflecting body is low.

Although our eyes cannot tell us, aluminum foil is highly reflecting not only at visible wavelengths but well down into the infrared and beyond. It also is opaque in the infrared. Thus we conclude that it has a low infrared emissivity.

Radiative exchange between the flask and its surroundings depends on the emissivities of both. To good approximation, the surroundings in which I did my experiments are black (i.e., have an emissivity close to one). The net radiative exchange between the flask and its surroundings is proportional to the difference between the fourth power of the absolute temperature of the flask and that of its surroundings, where the proportionality factor contains the emissivity of the flask (strictly, its emissivity averaged over its emission spectrum). Thus, the lower the emissivity of the flask, the less the radiative exchange with its surroundings and the more slowly the water in it cools. A lower infrared emissivity for aluminum than for glass is at least consistent with the different measured cooling times.

Now some of you might say, Bosh! The aluminum foil keeps the flask warmer because it reflects the infrared radiation emitted by the glass. To test this hypothesis, I painted one side of some aluminum foil black, then wrapped it around the flask with the black side inward. The cooling time was only a few percent lower than that for the flask wrapped with unpainted foil, an insignificant difference given the accuracy and reproducibility of my measurements.

The reason for the irrelevance of the radiative characteristics of the foil's inner surface is that the foil and the flask are at nearly the same temperature. The net radiative exchange between the glass and the inner surface of the foil depends on the temperature difference between them. Because this difference is small, so is the net radiative exchange regardless of the emissivities of glass and foil.

In refuting that reflection by the foil is what causes a flask wrapped with it to cool more slowly than a bare one, I was deceptive: I said that I painted the inside of the foil black. Black to my eyes, that is, but neither you nor I see in the infrared. Our eyes are powerless to tell us if black paint at visible wavelengths is also black (i.e., has an emissivity close to one) in the infrared. Indeed, it takes a conscious act of will to convince oneself that an object perceived as black to visible radiation may not be black (i.e., have a low reflectivity) to infrared radiation, and conversely. One way to drive this home is with a simple experiment.

I obtained two cooling curves, one for a flask wrapped in foil painted black and one for a flask wrapped in foil painted white (see Figure 7.2). The cooling time for the white flask was 52.6 minutes, that for the black flask was 50.8 minutes, and that for the bare flask was 50 minutes. Thus the infrared radiative characteristics of glass, black paint, and white paint are nearly identical.

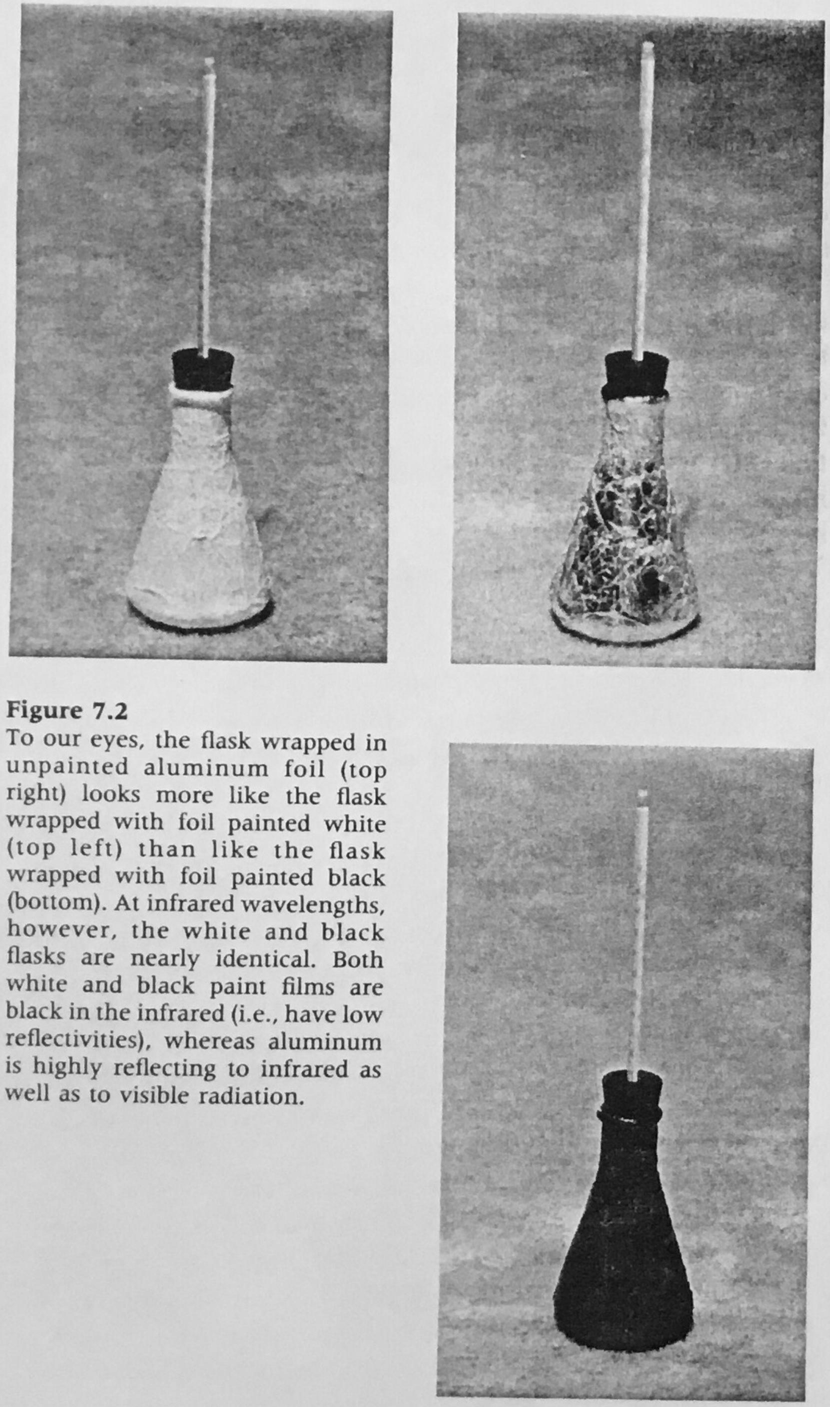

Figure 7.2
To our eyes, the flask wrapped in unpainted aluminum foil (top right) looks more like the flask wrapped with foil painted white (top left) than like the flask wrapped with foil painted black (bottom). At infrared wavelengths, however, the white and black flasks are nearly identical. Both white and black paint films are black in the infrared (i.e., have low reflectivities), whereas aluminum is highly reflecting to infrared as well as to visible radiation.

Engineering students are taught that covering a metallic furnace pipe with insulation can increase rather than decrease the rate of energy transfer from it to the surroundings. This was discussed as long ago as 1933 in *American Journal of Physics* by William Schreiver, who cited experiments done even earlier (1920) at the University of Illinois. To make his point, Schreiver did the same kind of cooling-curve experiments as I have done.

If you were to ask anyone how to decrease the rate of cooling of an object, you would likely be told to wrap it with insulation (wool, for example). This suggestion is so much in accord with common sense that to hint that it might not yield the desired results borders on heresy. Ever the heretic, I wrapped a swatch of thick wool cloth around my foil-covered flask filled with hot water, then measured its cooling time. The cooling time increased by less than 2 percent. Although wool giveth, it also taketh away: it is a much better conductive insulator than aluminum foil, but a much worse radiative insulator. So the net effect of dressing the flask in a warm wool coat is almost nil.

Lest I mislead you, I must point out that the radiative characteristics of a cooling (or warming) body become immaterial when other modes of energy transfer (e.g., convection) predominate. To demonstrate this, I placed a foil-covered flask about one-half meter from a low-speed fan tilted downward. I estimate the wind speed to have been a few knots. The cooling time was 23 minutes. Then I painted the foil black and repeated the experiment; the cooling time was similar, about 25 minutes. In still air, however, the two cooling times differed by more than 50 percent.

These results also demonstrate why *the greenhouse effect* when applied to the atmosphere is somewhat of a misnomer. Greenhouses exchange energy with their surroundings both radiatively and convectively. In even a light wind, convection dominates and the emissivity of the glass becomes immaterial.

Do Good Absorbers Go to Heaven?

Kirchhoff's law is often transformed into a slogan: "A good absorber is a good emitter." Whenever I see or hear this, I imagine patting an absorber on the head and murmuring approvingly, "You're a good little absorber." I was pleased to discover that my distaste for this kind of anthropomorphism was shared by John Henry Poynting, whose name—if not career—is well known to all students of electromagnetic theory, having been enshrined in the Poynting vector (sometimes maligned as the "disapoynting" vector), which specifies the magnitude and direction of the flow of electromagnetic energy (including that which we call light). In an obituary notice about Poynting published in *Nature* in 1914, Sir Joseph Larmor opines that "his rebellion against an excessive anthropomorphism which had begun to cling around the notion of natural laws, as if they were really legal enactments to be obeyed or disobeyed by inert matter almost as if it possessed will-power and could exercise choice, some substances being praised as good radiators while others stigmatised as bad . . . Poynting's revolt against this kind of attitude to laws of nature, though doubtless more than half humorous, was in itself wholesome."

At the very least, the notion that a good absorber is a good emitter is misleading. Even if it does not mislead (my experience is that it does), the overused word *good* is meaningless unless accompanied by a criterion. Let us consider each of these points in turn.

The rate of absorption by a given object depends on its radiation environment, whereas its rate of emission depends solely on its temperature. We can imagine taking an internally heated blackbody (i.e., a "good absorber") into space. It emits infrared radiation but absorbs much less because it isn't bathed in much. Yet we can imagine exposing this same body to an intense source of infrared radiation, in which instance it would be a "good absorber" but not nearly so "good" an emitter.

What is good (or bad) depends on our aims. Suppose that we wanted to keep a sandwich warm in cooler surroundings. Because of our experiments with the flask, we would elect to wrap the sandwich in aluminum foil. Aluminum is not a "good emitter," but its "bad" emissive properties make it good for keeping the sandwich warm. To keep the sandwich cool in warmer surroundings, we could again wrap it in aluminum foil. Here are two examples in which what we might unthinkingly refer to as a bad emitter is good for our purposes: keeping a sandwich cool in warmer surroundings or warm in cooler surroundings. But we haven't exhausted the possibilities.

Our aim might be to keep an object cool in radiatively cooler surroundings. One summer in Arizona, I painted the roof of my cottage with aluminum paint so that it would reflect solar radiation, thereby giving some inexpensive cooling. In retrospect, I should have used white paint—assuming that the reflectivities for solar radiation of the two paints are about the same—because its infrared emissivity is higher than aluminum's.

To say that the surroundings on a hot Arizona day are *radiatively* cooler than a roof, which at first glance might strike you as absurd, does not imply that the temperature of the air is lower than that of the roof. My roof viewed the sky, the radiative temperature of which is that of a blackbody that emits the same amount of radiation. As a rough estimate, the radiative temperature of the clear sky is about 250°K. This will vary up or down depending on the absolute humidity and actual air temperature, by which I mean a vertically averaged temperature, not that of the air adjacent to the roof.

A final possibility is that we want to keep an object warm in radiatively warmer surroundings. An example is a transatlantic hot-air balloon. At night, the balloon must be kept warm so that it doesn't sink. Since temperature usually decreases with height and water is nearly a blackbody, the balloon's temperature can be less than that of the sea below and still be positively buoyant. To keep the air in the balloon warm, its bottom half should have a high infrared emissivity, whereas that of the top half should be low because its surroundings are radiatively cooler.

I am told that the top halves of transatlantic balloons are aluminized and their bottom halves are painted black. The reflecting top keeps the balloon from getting too cold at night and too warm during the day. (A white top would be even better during the day, but not at night.) My guess is that it is unnecessary to paint the bottom of the balloon black since it, like a glass flask or white paint, is probably already black in the infrared. The black paint is more psychologically than physically beneficial.

Emissivity and Climate Change

Many textbooks convey that emissivity is a surface property, which indeed it may be considered for opaque bodies. Both a flask and aluminum foil are sufficiently opaque that their emissivities are surface properties: slice them thinner (within limits), and their emissivities would not change.

The atmosphere is different from aluminum foil or a glass slab. Unlike the latter, the atmosphere has no palpable surface. Nevertheless, the atmosphere has an emissivity, the steady increase of which is what fuels global warming.

Another difference between the atmosphere and glass or foil is that the atmosphere's infrared reflectivity is negligible, aluminum is highly reflecting from visible to radio wavelengths, and glasses have spectral regions of fairly high infrared reflectivities. Whereas the emissivity of aluminum foil is one minus its reflectivity, that of the atmosphere is one minus its transmissivity. To say that the emissivity of the atmosphere is increasing because of increasing concentrations of carbon dioxide (as well as water vapor, methane, and other gases) is therefore the same as saying that its transmissivity is decreasing or that it is becoming more opaque.

On average, the amount of surface radiation absorbed by the atmosphere must equal that emitted by it, so the temperatures of both are linked. Since the surface is heated both by solar radiation and by atmospheric infrared radiation, the average temperature of the surface is higher than that of the atmosphere.

The solar radiation absorbed by our planet must be balanced on average by infrared radiation emitted upward by the atmosphere together with that emitted by the surface and transmitted by the atmosphere. Thus if the absorbed solar radiation is fixed, an increased atmospheric emissivity must be accompanied by temperature changes to keep the infrared flux constant. Its two components depend on atmospheric emissivity in opposite ways: emission by the atmosphere goes up with emissivity, and transmission of surface radiation by the atmosphere goes down. Since the atmospheric temperature is less than that of the surface, the temperature of the latter must increase to compensate for an increased emissivity (i.e., decreased transmissivity).

Emissivity enters the greenhouse effect by another route. A possible consequence of global warming is a greater water vapor flux into the atmosphere from the oceans, hence more cloud cover accompanied by a lower net solar radiation flux to the planet. Although more clouds would indeed reflect more solar radiation, thereby cooling the planet (a negative feedback), the emissivity of clouds is greater than that of clear air (a positive feedback). Moreover, these feedbacks depend not only on how much cloud cover is increased, but where it occurs. For simplicity, divide the troposphere (which contains almost all of the mass of the atmosphere) into two parts, lower and upper. The infrared emissivity of the lower troposphere is larger than that of the upper. Clouds added to the lower troposphere (which is already nearly black) would therefore have less effect than clouds (e.g., cirrus clouds) added to the upper troposphere.

Which of the two feedbacks—negative because of increased solar reflection by clouds or positive because of decreased infrared transmission by clouds—will dominate is uncertain. The current conventional wisdom favors net warming from increased cloud cover, but this could change after the latest results from computer models are rushed to press.

Figure 10.3
At the same latitude (45°N in this diagram), the sun is more nearly overhead at solar noon in June than in December because the earth's rotation axis is tilted with respect to the plane of its orbit around the sun.

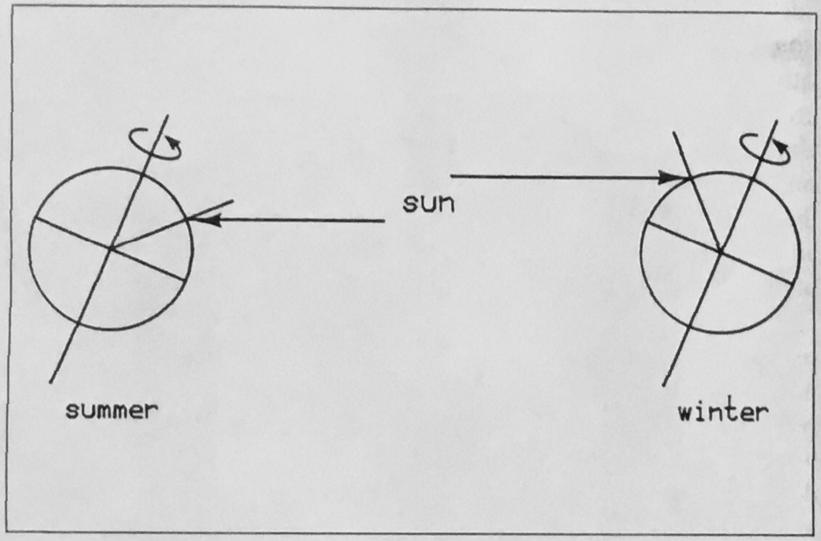

21 June), the rays of the sun at solar noon (when the sun is highest in the sky) are directed about 22 degrees from the zenith, whereas at the winter solstice (about 22 December), the rays are directed about 68 degrees from the zenith. As a consequence, a given patch of the earth receives considerably more solar radiation at the summer than at the winter solstice.

One wonders how Wilbur Glen Voliva would have explained both latitudinal and seasonal variations in temperatures over a flat earth.

✵11

Temperature Inversions Have Cold Bottoms

Dust in the air suspended
Marks the place where a story ended
T. S. Eliot

It is well known—to those who know it well—that Los Angeles owes its infamous smog to its large number of automobiles and plentiful sunshine for transforming their emissions into noxious substances. It may be less well known that there is a third ingredient in the recipe for smog: the variation of temperature with height in the atmosphere. While Morris Neiburger was senior meteorologist at the Air Pollution Foundation in Los Angeles, he was subjected to a stream of suggestions for modifying the meteorological conditions of the Los Angeles Basin. These suggestions reflected the character of the area: unconventional, imaginative—but unrealistic. Before discussing Neiburger's analysis of the feasibility of one of these meteorological solutions to the smog problem—penetrating or dissipating the inversion—we need to understand what an inversion is and how it is related to the dispersal of pollution.

Demonstrating Atmospheric Stability

A striking demonstration of the connection between atmospheric temperature profiles and pollution was devised by Hans Neuberger and George Nicholas and is included in their *Manual of Lecture Demonstrations*. Place a tall, clear plastic cylinder, tightly sealed at the bottom, in a coffee can partly filled with ice and water. Gently blow smoke—you will need a fair amount—into a tube that passes through a small hole to the bottom of the cylinder. The result is shown in Figure 11.1. Note that the smoke lies

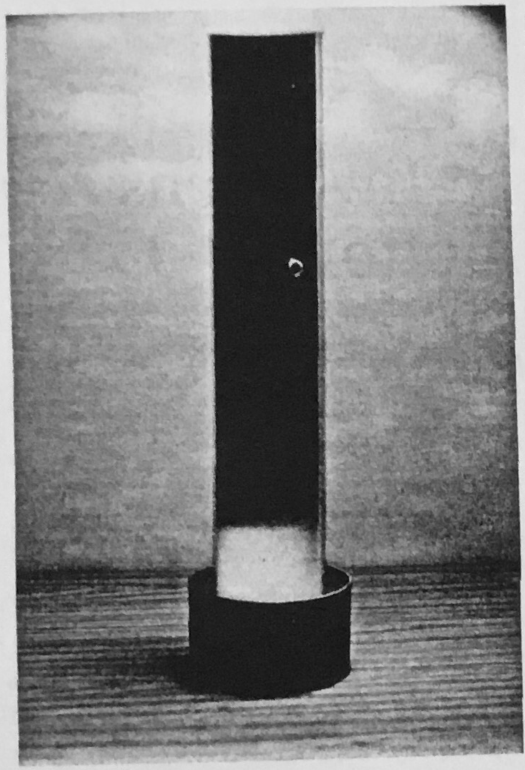

Figure 11.1
The bottom of this cylinder is in cold water. As evidenced by the smoke, there is little vertical mixing because warm air overlies cold, denser air. Photograph by Gail Brown.

quietly at the bottom of the cylinder; it does not rise much even if you jostle the cylinder a bit. But now replace the ice water with boiling water. Within a minute the air will suddenly come alive with curling, turbulent motion, ending in thoroughly mixed smoke (see Figure 11.2). Merely changing the temperature at the bottom from low to high dramatically transformed the state of the air in the cylinder: in the first instance it was *stable*; in the second it was *unstable*.

Atmospheric stability is closely connected with buoyancy. Bubbles rise in beer because of buoyancy. Buoyancy also can be demonstrated with a pan of sand and two ping pong balls, one filled with lead shot. If both balls are placed on the surface of the sand and the pan is agitated vigorously, the heavy ball disappears while the light one remains on top. This is shown in Figure 11.3. If the light ball is buried, it will eventually rise to the surface as the sand is agitated; it may be observed to suddenly break through the surface and float on the top. The light ball is positively buoyant: it is less dense than the sand; the heavy ball is negatively buoyant: it is more dense than the sand.

Figure 11.2
If the cold water in the can shown in Figure 11.1 is replaced by hot water, the air becomes unstable resulting in rapid vertical mixing (left) until the smoke is completely dispersed throughout the cylinder (right). Photograph by Gail Brown.

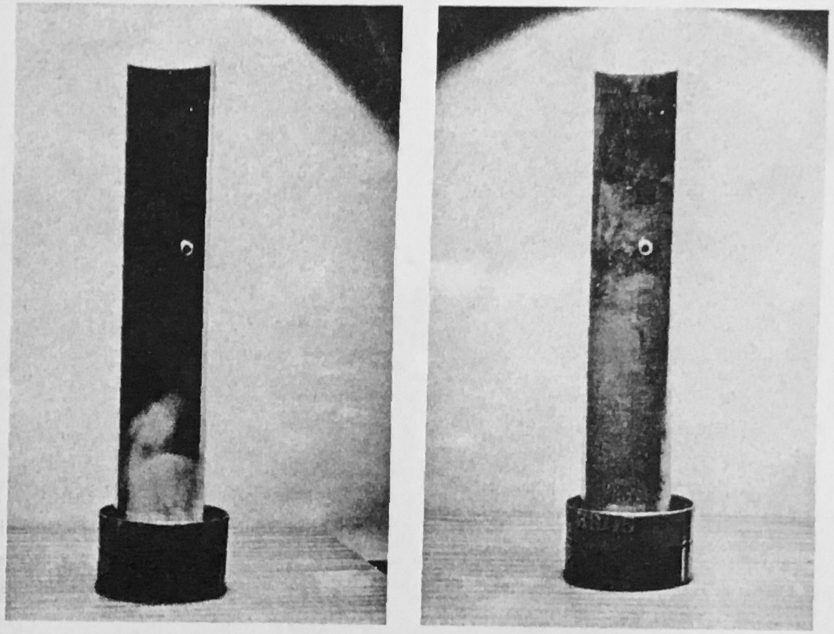

Robert Hooke, a seventeenth-century English physicist, used a dish of agitated sand to demonstrate that the ease with which liquids and gases flow—their fluidity—implies that matter is in continual motion. For if you make a hole in the sand with your finger, the hole immediately fills in again. But Hooke was also aware that the behavior of heavy and light objects in agitated sand demonstrates buoyancy. Sir William Bragg, in his superb book *Concerning the Nature of Things*, quotes Hooke as follows: "For by this means [shaking], each sand becomes to have a vibrative or dancing motion, so as no other heavier body can rest on it, unless sustein'd by some other on either side: nor will it suffer any body to be beneath it, unless it be heavier than itself."

Figure 11.3
When the sand in the pan is agitated, the light ball remains on the surface (top) while the heavy ball begins to sink (bottom) and eventually disappears (next page).

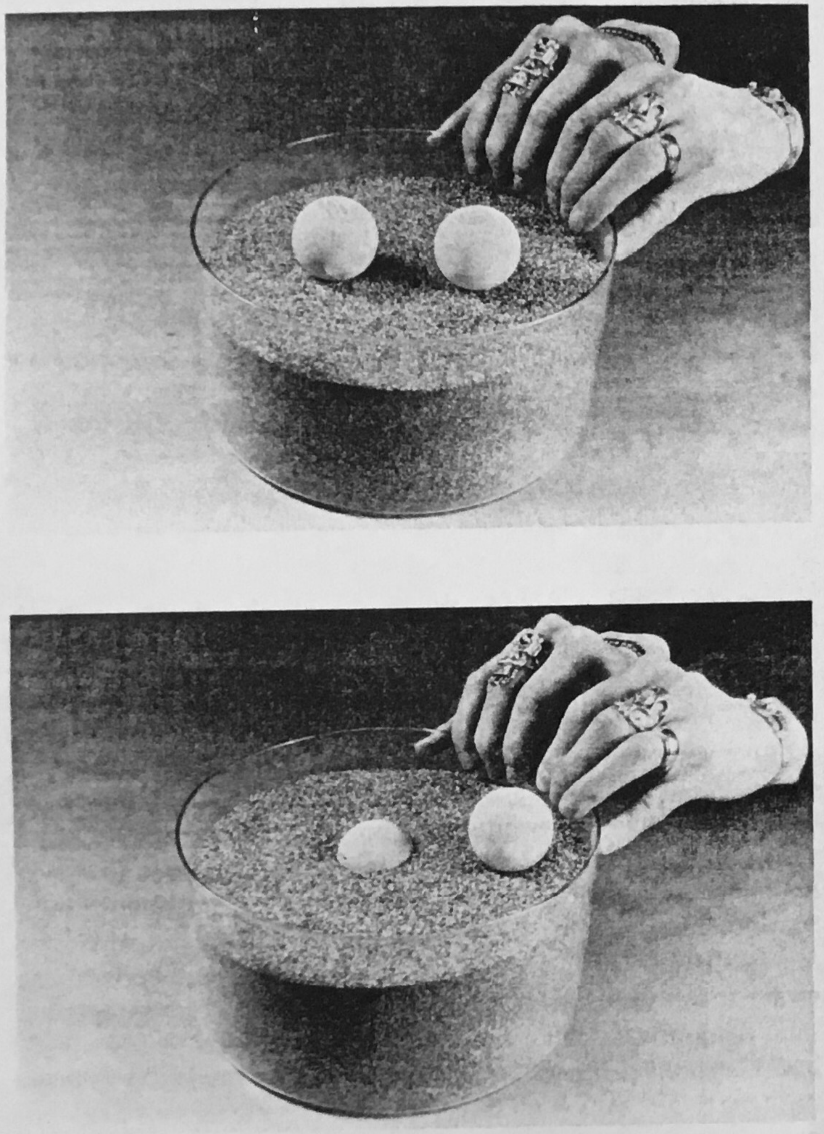

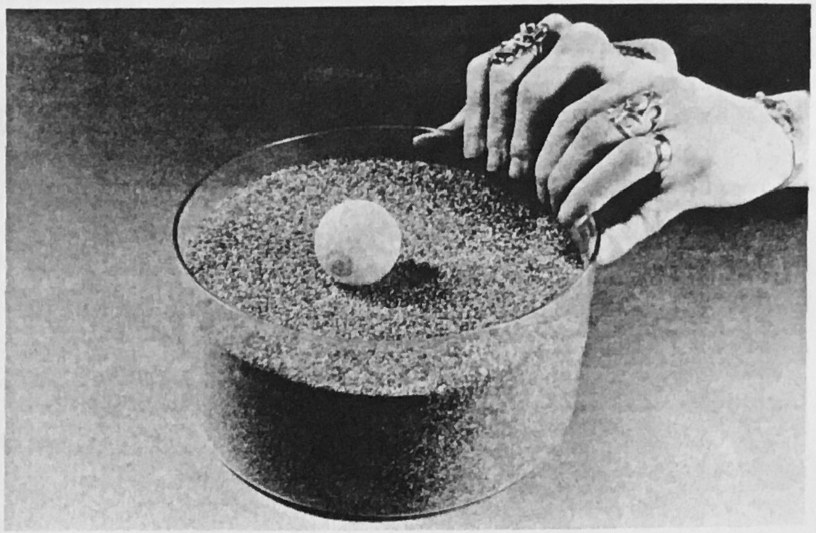

Temperature Differences

Cold air is denser than warm air at the same pressure. When the bottom of the cylinder was colder than its top, warm air overlaid cold air. This is stable: if warm air were to sink into the cold air it would be buoyed upward, back whence it came; on the other hand, if cold air were to rise into the warm air it would be heavier than its surroundings and would sink. The reverse—cold air above warm—is unstable: cold air sinking into warm air would find itself colder, hence denser, than its surroundings, and would continue to sink; warm air rising into cold air would be lighter than its surroundings and would continue to be impelled upward by buoyancy. Mixing of layers of air at different temperatures is therefore suppressed if the temperature increases with height. Conversely, mixing is promoted if the temperature decreases with height. Rising air cools, however, and it is the *difference* in temperatures between this air and its surroundings that determines if it will continue to rise or will sink. So we must ask, at what rate must the temperature decrease with height if the air is to be unstable?

Consider a parcel of air, which we may imagine to be enclosed within an elastic balloon (Figure 11.4). If the balloon is impelled upward by some means, it will expand: pressure decreases with

Figure 11.4
Two hypothetical vertical temperature profiles (solid lines). The dashed line shows the dry adiabatic lapse rate.

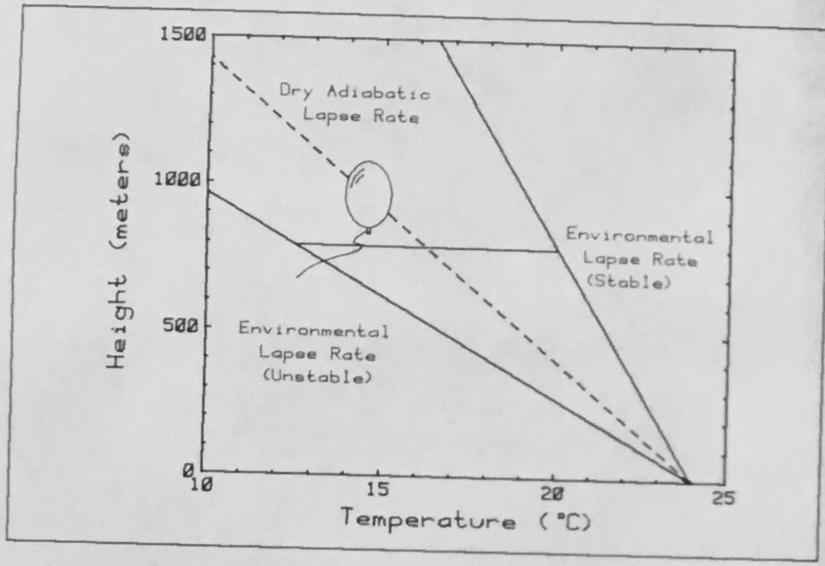

height, and at each instant the balloon increases its volume so that the pressure inside balances that outside. But it takes energy for the balloon to expand against its surroundings, just as it takes energy to lift a weight against gravity. The energy required for expansion is taken from the thermal energy of the air in the balloon, so it cools. There may be water vapor in the parcel, but we assume that it does not condense. And we also assume that the parcel is neither heated nor cooled by its surroundings because of temperature differences between the two. With these restrictions, the rate at which the temperature decreases is called the *dry adiabatic lapse rate*: dry because there is no condensation, adiabatic because there is no heat transfer, and lapse because there is a decrease.

Now suppose that the temperature of the atmosphere decreases with height according to the dry adiabatic lapse rate, which is about 10°C per kilometer. At each height, a rising parcel of air—we may imagine it to be in a balloon carried aloft on the tail of a kite—has the same temperature and pressure, hence density, as that of its environment. If the balloon is released at any height, it will neither sink nor rise.

If you move away from a warm stove, the temperature around you decreases. Similarly, the temperature usually decreases with height in the atmosphere, the earth's surface warmed by the sun taking the place of the stove. But the rate of temperature decrease need not be, indeed rarely is, the dry adiabatic lapse rate: it may be greater or less. Suppose that the temperature of the balloon's surroundings decreases more rapidly than 10°C per kilometer; this is shown by the leftmost curve in Figure 11.4. At any point in the ballon's ascent, its temperature, which follows the dry adiabatic lapse rate, is greater than that of its surroundings; if it were released at, say, 1,000 meters, it would continue to rise. In this instance the atmosphere is unstable. But if the temperature decreases less rapidly than 10°C per kilometer (shown by the rightmost curve in Figure 11.4), the temperature of the balloon is always less than that of its surroundings. If released it would therefore sink. In this instance the atmosphere is stable. So the dry adiabatic lapse rate merely defines the boundary between a stable and an unstable atmosphere. If we replace the imaginary balloon with real parcels of polluted air, it follows that, in the context of pollution dispersal, instability is desirable whereas stability is not. We are now better able to understand the meteorological factor in Los Angeles smog.

Los Angeles Smog

The average variation of temperature with height at 7 A.M. in September at Long Beach, California, as taken from Morris Neiburger's article "Weather Modification and Smog," is shown in Figure 11.5; the dashed line indicates the dry adiabatic lapse rate. Below about 475 meters the atmosphere is stable, but only slightly. Between 475 and 1,055 meters, however, temperature *increases* with height, which indicates a very stable atmosphere. This is called a temperature *inversion*, and polluted air under an inversion is trapped there; the temperature inversion acts as a lid on the upward dispersal of pollution. To eliminate the smog, one need merely eliminate the inversion by, for example, heating all the air below its top. Neiburger calculated that to do this over the Los Angeles Basin would require the energy obtained from burning (at 100 percent efficiency) 1.27 million tons of oil, the amount of crude oil processed in 12 days by all the refineries in

Figure 11.5
The average temperature profile at 7 A.M. in September at Long Beach, California (from M. Neiburger, "Weather Modification and Smog," *Science*, Vol. 126, 1957, p. 637). The dashed line indicates the dry adiabatic lapse rate.

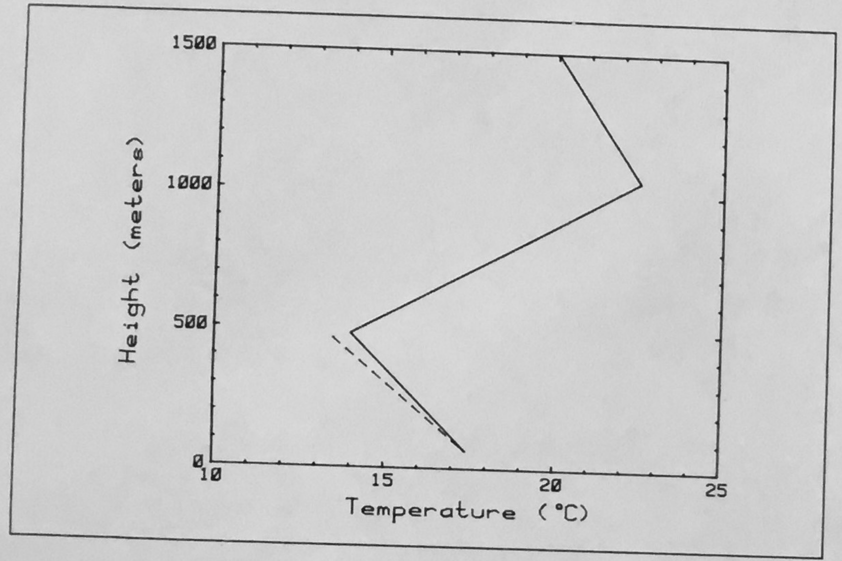

the Los Angeles Basin. And the air would not just sit still while being heated but would rise and be replaced by colder air, thereby reestablishing the inversion.

The prevailing temperature inversion in the Los Angeles Basin, particularly during the warmer months, is an inevitable consequence of the general circulation of the atmosphere. To eliminate the inversion would require either modifying the general circulation or moving the entire basin to a more meteorologically benign location. But neither solution has yet found advocates among even the most fanciful of Southern Californians.

Excuse My Dust

My most vivid—and unpleasant—memories of the consequences of temperature inversions are from my student days in Arizona. I lived on the desert, about 15 miles from Tucson. I would ride my bicycle to the university early in the morning, following less-traveled dirt roads as far as possible to avoid automobiles. Other than the large black dog who would wait patiently for me in hopes of catching me as if I were a jackrabbit (I always won the race by a margin of the dog's hot breath), what I dreaded most was that a car would precede me in the early morning hours. Because of overnight radiational cooling of the ground in the clear desert air, the air near the surface would be colder than that aloft. Dust raised by a passing automobile would hang in this stable air, and I would have to eat the gritty stuff for miles. But at the end of the day, after the sun had been baking the ground for hours, air temperature decreased with height, the atmosphere was unstable, and dust raised by passing automobiles was so quickly dispersed that I hardly noticed it.