Packing for the Trip
III: Light

5

Light is the most amazing thing in the Universe. It would be wonderful to spend the entire semester talking about light. Well, wonderful for me at least! Of course, we can't do that—we've got to get on with our astronomy. But since light is so important to astronomy, we do need to at least go over the essential points; and that will be enough to give us a glimpse of why light is so fascinating and baffling.

5.1 The Basics: What We Know About Light

To begin with, let's look at some of the essential facts about light. This is the basic stuff—nothing weird or strange; and much of it you probably know already. We'll break it down into five points.

1. **Light is pure energy.** That's one thing you can say about light that will always be correct. It's just about the only thing you can say about it that will be—but we'll get to that later.
2. **You can only see light when it enters your eye.** This is very important and somewhat subtle. What it means is that we "see" light differently than we "see" other things. When you look at a frog, the frog does not have to enter your eye in order for you to see it. The frog is out there, and you see it out there. Light is fundamentally different. You can't "see" a beam of light the way you "see" the frog. You can't "see" a beam of light "out there" as it zips past you. The fact is, you can only see light if it is directed into your eye. Your eye is a little light detector, and unless the light goes into your eye you cannot see it.

   "Wait a minute," you say, "I went to a concert last summer and they had lasers, and I saw the laser beams."

   Actually, you didn't. What you saw was dust, or smoke particles in the air that were being lit up by the laser beam. If you shine light onto a frog, you will see the frog. If you shine laser light onto a dust particle, you will see the dust particle. If you shine a laser beam through a room that is filled with dust particles, you will see dust particles along the path of the laser beam. It will *look* like you are "seeing" the beam of light, but you're not really. You are seeing the dust particles where the beam of light is, but that is *not* the same as seeing the beam of light itself. If the air in that room had been completely clean you would not have been able to see where the laser beam was at all. It would have been completely invisible.

3. **There are different kinds of light.** Your eyes are little light detectors, but they detect only a narrow range of all of the different kinds of light that are out there. The different kinds of light are shown in Figure 5.1. Notice first of all that the terms used for the different kinds of light are probably familiar to you: *radio*, *microwave*, *infrared*, *visible*, *ultraviolet*, *X-ray*, *gamma-ray*. The light that you can see with your eye we call *visible light*. But you can see a whole range of different kinds of visible light, which are seen by you as different *colors*. We're not yet ready to appreciate what makes one kind of light different from another, so you can ignore where it says "photon description" and "wave description in Figure 5.1 for now. But we will get to that soon.

   Notice by the way that there is no such thing as "white light." "White" is what it looks like to you when all of the different colors of visible light enter your eye all at once, in more or less the same amounts. (Pure black, of course, is the absence of any light. But an object that looks black is really an object that emits or reflects very little light compared to other objects around it.)

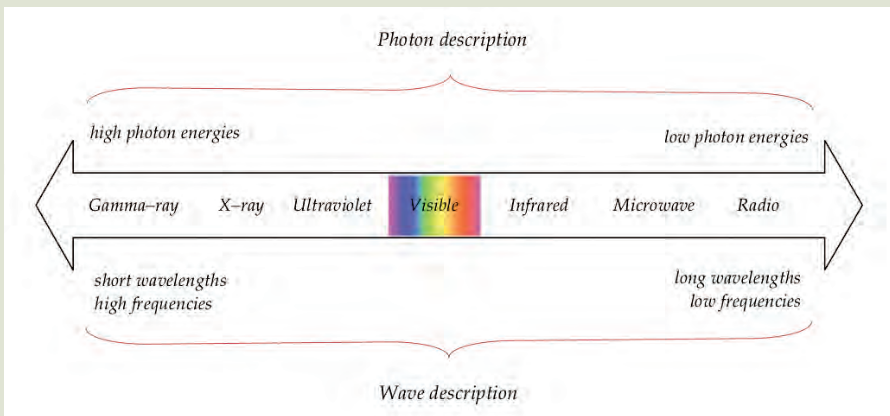

Figure 5.1 The Electromagnetic Spectrum

4. **When light hits an object, there are three things that can happen to it.**

   (1) The light can pass right through the object. This is called *transmission*.
   (2) The light can "bounce off" of the object. This is called *reflection*.
   (3) The light can be "soaked up" by the object. This is called *absorption*.

   There are two things that are very important to realize about these possibilities. First, they are not mutually exclusive. In fact, most objects cause two or more of them to happen at the same time. A window, for example, will *transmit* light—light from outside will come into the room; but it will also *reflect* light—if you stand outside a window looking into the room, you will be able to see your own reflection in the window.

   Second, the degree to which each of the three possibilities occurs depends not only on the object itself, whether it is made of glass or wood or whatever, but also on the type of light. Wood will absorb and/or reflect visible light, but it will transmit radio light—a portable radio works perfectly well inside a log cabin.

5. **There are two ways that we can see things.** We can see things when they *emit* their own light, and we can see things when they *reflect* light that has been emitted by something else. Which is more common do you think? The second, definitely. You could count on one hand the number of different things you can see by emitted light: Flames, light bulbs, stars (including our Sun), and maybe a few other things—glow in the dark stickers and things. Maybe you'd need two hands, but that's it. Everything else is seen by reflected light.

Let's bring all of these ideas together to explore how it is that we see something—a green leaf, for example.

The light that strikes the leaf (assuming that you are seeing that leaf outside on a sunny day) originated in the Sun. As I mentioned before, the light from the Sun is created from the Sun converting its own mass into pure energy, or light. That sunlight travels out in all directions, and some of it—a very tiny fraction of it—reaches the Earth; and a very, very, very tiny fraction of the light that hits the Earth, hits the leaf. Most of the light that hits the leaf is absorbed by the leaf; little if any is transmitted. (Leaves are generally not transparent. You can't see through them.) The light that is absorbed by the leaf provides the leaf with energy which is then used by the plant or tree for various purposes like growing, and photosynthesis, and whatever else it is that plants and trees do. Some of the light, however, is reflected by the leaf. When you see the leaf, it is this reflected light that you are seeing, since that is the only light that can possibly go into your eye. If the leaf looks green, it means that green light is going into your eye; that is, the leaf is reflecting green light. So a "green leaf" is a leaf that absorbs everything *except* green.

This picture is very simplified I'm afraid, and you should know that things can be more complicated. Just because something looks a particular color, that doesn't necessarily mean that it is reflecting light of that color. It might be reflecting other colors that, when combined in your eye, cause you to see that particular color. But the basics of this simplified picture are still useful.

5.2 What *is* Light? A Very Famous Experiment

We've got the basics now. That's good. But before we explore the nature of light any further, we're going to have to take a moment to look at a very famous and very important experiment. It may seem at first as if this experiment has nothing to do with light; but let me assure you, it is *profoundly* relevant and important.

The experiment is a physics experiment, and it is usually referred to by physicists as the *double–slit experiment.* The setup for the experiment is very simple (see Figure 5.2). To begin with, we need a source of something. The "something" can be pretty much anything we like: light, electrons, baseballs, you name it. We'll put this source on the left. Then, on the right, we need a detector that can detect whatever the source is putting out. Finally, midway between the source and the detector we need a screen with two small openings in it. Whatever is coming from the source will be blocked by the screen, except where the two openings are. In other words, the only way for the "something" to get from the source to the detector is to go through the openings. The experiment was first done with light, with thin vertical slits as openings for the light to get through; hence the experiment's name. But since we want to keep this very general, we will refer to them as *windows.*

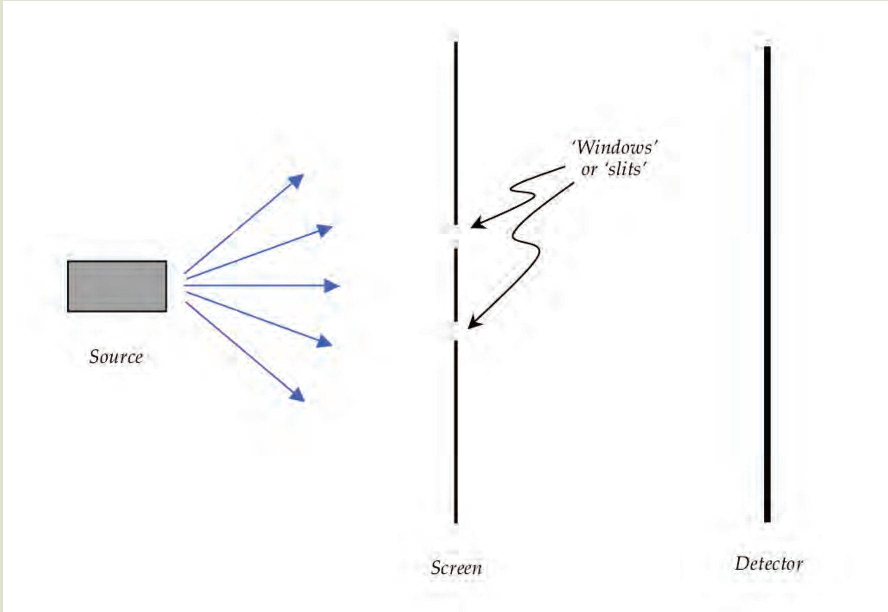

Figure 5.2 The Set Up for the Double–slit Experiment

On the left is the source of whatever you wish to send through the experimental apparatus (see text). The screen blocks whatever is emitted by the source and prevents it from getting through, except where the two 'windows' or 'slits' are located. The detector can be anything that is capable of recognizing or detecting the arrival of whatever is emitted by the source.

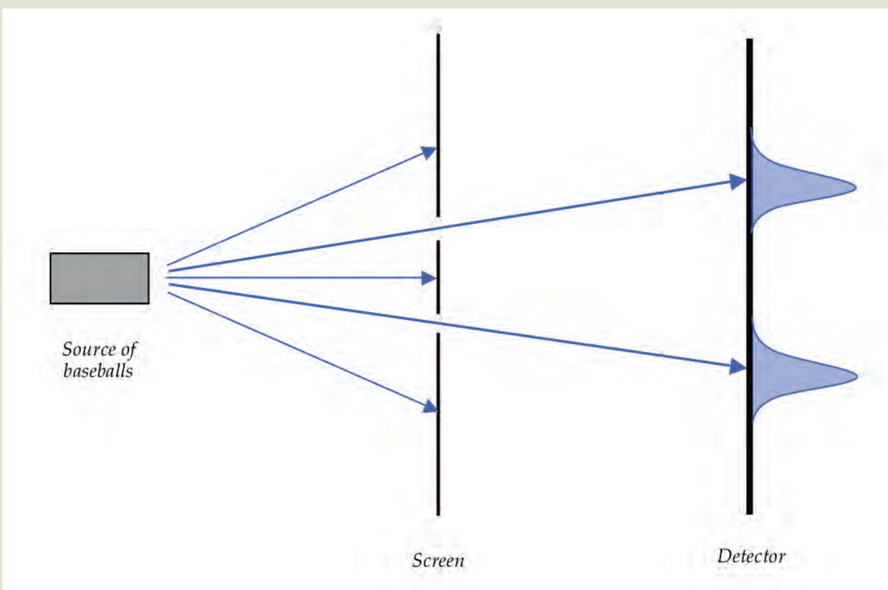

Figure 5.3 The Double–slit Experiment with Baseballs

Each baseball will either be stopped by the screen, which in this case might be a wall or a fence, or it will pass through one of the windows, and so be 'detected' in a direct line from the source to the detector. The curves along the detector are meant to indicate the number of baseballs that hit the detector in a given amount of time—lots of baseballs in line with each window, but few or none anywhere else.

5.2a An Experiment with Baseballs

For the first case, let's imagine that our source is a machine that throws out baseballs. The screen will be a wall with two windows in it, each one large enough for a baseball to get through, but small compared to the size of the screen and the distance from the source. The "detector" can be anything that registers the arrival of a baseball. It could be a wall with built–in electronic sensors that record the precise location and timing of anything that hits the wall; or it could be a line of nets or baskets, each of which collects the baseballs that are thrown into it. It could even be a line of "catchers" with their mitts ready, who will catch any ball that comes their way! Finally, we will assume that the source throws the baseballs out in random directions, so that we never know where any particular baseball is going to end up.

We then turn on our baseball throwing machine, and watch what happens. Clearly, each baseball is either going to go through one of the windows, or not. We will assume that if it goes through a window, it will reach the detector and be recorded as a "hit;" otherwise it won't. So after lots and lots of baseballs have been thrown, our detector will have recorded lots of "hits" (baseball detections) in line with each window, and few if any "hits" elsewhere along the detector (Figure 5.3). Of course, one could imagine that the occasional baseball might hit the edge of one of the windows, and so bounce off to the side, hitting the detector in some unexpected place. But we can assume that this doesn't happen very often, so that most of the "hits" are pretty well lined up with the windows. So far, so good.

5.2b An Experiment with Water Waves

interference
The process of two waves combining or overlapping with each other.

For our second experiment, let's imagine that the entire arrangement is set up in a swimming pool, and that our source is a device that creates ripples on the surface of the pool. These ripples, or waves, will then move across the surface of the water towards the other end of the pool, where there is some kind of mechanism for detecting the waves as they go by. One might imagine a line of toy boats, for example, that bob up and down whenever a ripple passes under them. Midway along the length of the pool is the screen, which in this case will simply be a wall of some sort that the waves cannot get past, with two openings in it (the windows) where the waves *can* get through.

We now push the "ON" button, and the source begins making ripples which move across the water. What happens when the ripples reach the screen? Well, this experiment is going to be considerably more complicated than our experiment with baseballs, since the motion of waves on water is more complicated than the motion of baseballs flying through the air. However, if we look at it carefully, we can work things out pretty well. First of all, the screen will act kind of like a breakwater in a bay. The waves will wash up against it, and bounce back. However, where the windows are the waves will continue on through, creating two waves on the other side. In other words, each window in the screen will act as if it is a new source of waves (Figure 5.4).

So what happens when these two sets of waves reach the detector? Ah! That is where things get interesting! You see, each wave can be thought of as a series of up–and–down motions that travel across the water's surface. When the two waves overlap with each other, their up–and–down motions combine, a process physicists refer to as **interference**. There will be places, for instance, where both waves are causing the water's surface to move up and down at the same time. When that happens, the water's surface will move up and down—a lot! Physicists call that *constructive interference*. But there will also be places where one wave is trying to move the water up, at the same time that the other wave is

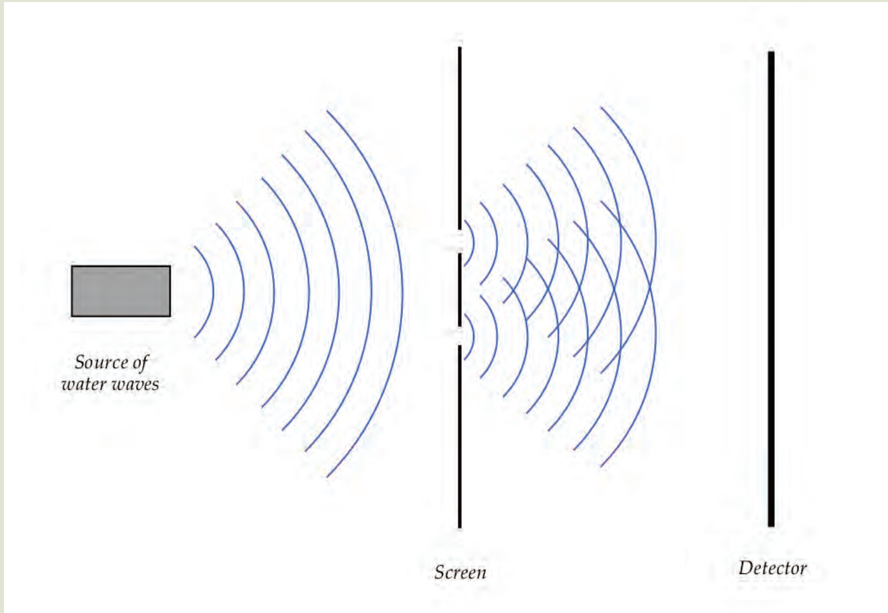

Figure 5.4 The Double–slit Experiment with Water Waves

Each window will act like a new source of water waves, so the net result of the screen is to turn a single source of waves into a double source.

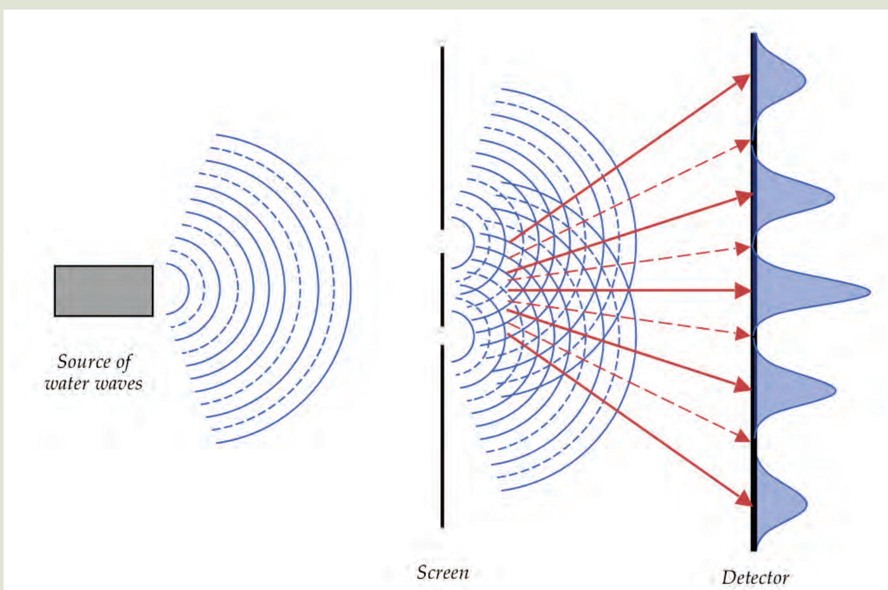

Figure 5.5 An Interference Pattern Produced by Two Overlapping Water Waves

Solid curves are used to indicate "peaks," or high points in the wave; dashed curves indicate "troughs," or low points in the wave. *Constructive interference* occurs wherever peaks overlap with peaks, or troughs overlap with troughs (the solid arrows); *destructive interference* occurs wherever peaks overlap with troughs (the dotted arrows). The curve to the far right indicates the amount of wave motion that would be occurring along the detector, with maximum wave motion occuring at points of constructive interference, and no wave motion occurring at points of destructive interference. The peaks in the curve get smaller further from center because the intensity of the wave decreases as it moves further from the source.

trying to move it down. When that happens, the two waves cancel each other out and the water's surface doesn't move at all, a situation physicists call *destructive interference.*

interference pattern

The pattern of wave–no wave–wave–no wave that results from two waves overlapping or interfering with each other.

Now, you might think that at the detector the waves will be interfering with each other more or less randomly. But that's not what happens! Although it isn't obvious unless you work through the *mathematics,*[1] the places of constructive and destructive interference will not be randomly arranged. On the contrary, a regular pattern of constructive interference, destructive interference, constructive interference, destructive interference, will be seen along the length of the detector. More specifically, what we will see when we look at our line of boats is that the boat that is in the center—that is, the one that is directly opposite the source—will be bobbing up and down a lot (constructive interference). However, on either side of this will be a place where the waves completely cancel each other out (destructive interference), and the boats will not be moving at all. Moving further along, there will again be another place (on either side of center) where the boats are moving up and down a lot; then further out another place (on either side) where the boats are not moving at all; and so on, and so on, along the entire length of the detector. If there is room between the points of constructive and destructive interference for more boats, then those in–between boats will bobbing up and down an in–between amount. The closer they are to the point of constructive interference, the more they will be bobbing up and down; the closer they are to the place where there is destructive interference, the less they will be moving.

In terms of waves, our "wave detector" will be registering a very definite, symmetric pattern of wave—no wave—wave—no wave—wave—no wave… along the length of the detector, a pattern known to physicists as an **interference pattern** (Figure 5.5). Note how different this is than the case with baseballs!

5.2c An Experiment with Light

So far, we have been talking about pretend experiments. No physicist would ever do a double–slit experiment with water waves, unless perhaps they were trying to demonstrate interference for a freshmen physics class; and no one would be interested in a double–slit experiment with baseballs, except maybe a baseball player who needed pitching practice! So for our last variation of this experiment, let's consider a more realistic case, one that has actually been done for actual physics research. Let's consider what would happen if we performed a double–slit experiment with light.

There is no great difficulty in conducting such an experiment, although the required equipment is a bit beyond the means of your typical college physics classroom. The source of light would generally be some kind of inexpensive laser. A light bulb would do, but a laser makes it easier. A thin sheet of metal with a couple of windows or slits in it makes a very nice screen. If you're using visible light, the windows or slits would have to be extremely tiny, so unless you know an excellent machinist you're going to want to buy the screen from a physics equipment supplier. The detector is where things get a little tricky. Finding something that can detect light is of course not a problem in itself, since your eyes can do that. But for the experiment to be of any use, you need a detector that is considerably more sensitive than the human eye. However, electronic light detectors are readily available if you know where to look for them.

Imagine then that we've got all our equipment in place, and we turn on our laser. What do we see at our detector? What we see is a clear pattern of bright spots, separated by dark

1 The mathematics of wave interference is not difficult. It requires only a little trigonometry, and can be found in most introductory physics textbooks.

places. Bright–dark–bright–dark–bright–dark… all along the length of the detector. In other words, we see an interference pattern. So light is a wave!

But not so fast. Your eyes may see an interference pattern, but let's take a look at the readout from that fancy, highly sensitive electronic detector. What you find there is that in the brief instant that it took to go from darkness to light, the detector was busily recording the arrival of zillions of individual little "blips" of light. In other words, what your eye sees as a glowing band of light–dark–light–dark across the detector, is actually the result of a zillion individual little dots of light, arriving one "blip" at a time—exactly like little tiny baseballs.

So what can we conclude? Is light a wave? Yes, it must be, since it produces an interference pattern in a double–slit experiment. Is light made up of individual particles? Yes, it must be, since light can be detected as individual 'blips' of energy. Is light both a wave, and individual particles? No, that's impossible, since a wave is *by definition* spread out in space, whereas individual particles *by definition* are not.

So we have a contradiction? Yes! The nature of light is intrinsically contradictory, at least as far as our current understanding of physics can take us. Physicists deal with this apparent contradiction through a branch of physics known as quantum mechanics, which underlies our understanding of the nature of both light and matter. It is beyond the scope of this book to explain how quantum mechanics works, or how quantum mechanics addresses the paradoxical nature of light; but it is important for you to be aware of the problem. You see, even though astronomers don't generally spend a lot of time worrying about the paradoxes and contradictions of fundamental physical reality, they *do* work with light—a lot. So a basic knowledge of how light behaves under various conditions is essential to astronomy.

5.3 What Is Light?

Because of the difficulties just described, physicists have two *models* for what light actually is. I want you to understand at the outset that these two models are what we call *mutually contradictory*. That is, they contradict each other. They cannot both be true. Nevertheless, physicists—and astronomers—use both models. The two models that we have are (1) that light is a *wave*, and (2) that light is *particles*.

A wave is basically ripples that move along together in some coordinated way. To say that light is a wave, is to say that what a light bulb is doing is creating ripples that travel out through space. To say that light is made of particles is to say that that same light bulb is sending out little 'blips' of light, little particles of light. Now, it doesn't make any sense to say that light is both a wave and particles, since they are contradictory ideas. A wave is spread out over some region of space. Ripples emerging from the light bulb would fill up the room. A particle is a little 'blip' that exists at a particular point in space. Light can't be both, and yet it is both.

You can actually avoid a lot of trouble if you just choose your words carefully. Instead of saying light is a wave in some respects, and it is individual particles in other respects, say that under certain conditions light behaves *as if* it is a wave spread out in space, filling up the entire room; and yet the very next instant, that same light will behave *as if* it is made of little tiny point particles that are *not* spread out in space, and do *not* fill up the room. Physicists have learned to live with this apparent contradiction, and you will have to as well—at least while you are reading this book—because we are going to be using both models interchangeably. The fact is, light is what it is. The difficulty comes when we try to put it in a particular box, or wrap our feeble human minds around it. (It is this difficulty, by the way, that I was thinking of when I said that light is the most amazing thing in the Universe!)

Despite these conceptual difficulties, it would be worthwhile to spend a few pages exploring each model in some detail. If you're going to learn astronomy, you need to understand how physicists talk about light, and what they mean when they refer to a "photon detector," or the "frequency of the light beam."

5.3a Light as Particles

To describe light as particles you have to picture a beam of light as consisting of zillions of little 'blips' of pure energy, 'blips' of light. These 'blips' are called **photons**. A photon can be described as a "particle of light."

photons
Individual 'blips' of pure light energy.

The only difference between one photon and another is the amount of energy that it carries (at least, that's the only difference that will be important for this class), and it is the energy of the photon that determines the type of light that it belongs to, the "label" that we would put on it. There are high energy photons that carry lots of energy, and low energy photons that carry very little energy. The highest energy photons are called *gamma rays*, or *gamma ray photons*; the lowest energy photons are called *radio light* or *radio photons* (see Figure 5.1)[2].

Your eyeball is a light detector that can respond to a certain range—a fairly narrow range, as it turns out—of photon energies. As I mentioned before, this range is called *visible light*. The lowest energy photon that your eye can respond to is associated in your perception with the color "red." A photon with any less energy that this would simply not be detected by your eye. The highest energy photons you can see are recorded by you as "violet."

The number of photons determines how bright the light is. When you look at a very bright light, it just means that lots of photons are entering your eye each second. A faint light source is one that is not emitting very many photons. Your eye needs to detect something like 6–8 photons each second in order for you to "see."

The next question to address is a big one: How do you make a photon? The answer to that seemingly simple question is going to take a lot of writing (or reading on your part), and a lot of explaining; and since it is inseparable from the question of how objects emit light, we will save it until we are ready to address that side of things. First, we need to consider our other model of what light is.

5.3b Light as a Wave

For some reason, basic science books almost always talk about light as a wave and ignore altogether the description of light as photons. It's somewhat unfortunate that the wave model is more common, since it really is less intuitive than the photon model. At least, I think so. To understand how the idea of "light as a wave" works you will need a little background on waves, what they are and how they work. So let's start with that.

As I mentioned earlier, a wave is just ripples traveling along together. All waves—it doesn't matter what kind: ripples on the surface of a pool of water, sound waves, light waves—all waves are characterized by just two things: (1) how close together the ripples are; and (2) how big the ripples are.

2 Note that radio light is sometimes called "radio waves." This does *not* mean that radio light is more "wave like" than "photon like," or that radio light is any more "wave like" than any other kind of light. The name is purely historical, and was assigned at a time when scientists knew comparatively little about light.

wavelength

The spacing between the crests and/or troughs of a wave.

frequency

The number of crests or troughs that pass a given point in space each second. Alternatively, the frequency can be thought of as the number of times that a given location on the wave moves up and down (or back and forth) each second.

How close together the ripples are can be measured in two ways. One way is to simply measure the distance between them. This is called the **wavelength**. The wavelength of a typical wave is shown in Figure 5.6. Wavelengths can be measured in meters, or centimeters, or any other useful unit of distance. The other way is to stand in one place and let the wave pass you, and then count how many ripples pass you every second. This is called the **frequency**. The frequency is usually measured with a unit called *Hertz (Hz)*, where 1 Hz just means "one ripple each second."

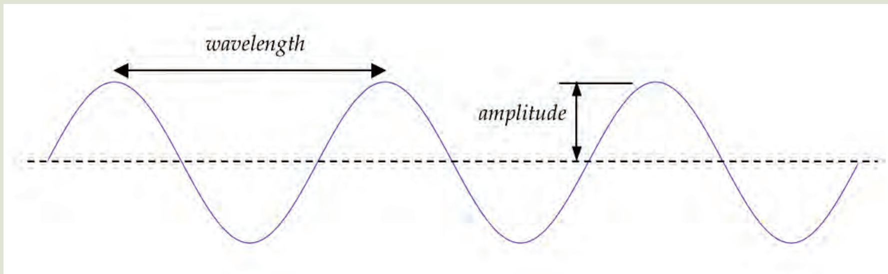

Figure 5.6 A Typical Wave

Notice that wavelength and frequency give the same information—almost. They both are a measure of how close together the ripples are, so for any given wave they are redundant to some extent. If a wave has a *short wavelength*—that is, if the ripples are close together—then a lot of ripples will pass you each second if you stand in one place, meaning that it will have a *high frequency*. If a wave has a *long wavelength*—if the ripples are far apart—then fewer ripples will pass you each second when you stand in one place, so the wave will have a *low frequency*. That is,

short wavelength = high frequency

long wavelength = low frequency

So you can refer to a particular wave as having a "short wavelength" or a "high frequency." It really doesn't make any difference since you're saying the same thing. However, if you are comparing different waves, then things do get a little more complicated. For example, imagine that you have two waves, both of which have a wavelength of 1 meter, meaning that the ripples are 1 meter apart, but which are traveling at different speeds. Let's say one wave is traveling at 10 meters per second (10 m/s). This means that each ripple travels a distance of 10 meters each second; which in turn means that if you stood in one place watching the wave go by, 10 meters worth of wave would pass you each second. Furthermore, since the ripples are 1 meter apart, 10 meters worth of wave will include 10 ripples; which means that 10 ripples would pass you each second. In other words, this wave would have a frequency of 10 Hz. Now imagine that the other wave has a speed of 20 m/s. By the same reasoning, 20 of its ripples would pass you each second, meaning that it would have a frequency of 20 Hz. This shows that two waves can have the same wavelength, but different frequencies, and vice versa; so when you are comparing different waves, frequency and wavelength do not necessarily give the same information.

amplitude

Loosely speaking, the 'height' of the wave. Specifically it is usually defined as the distance from the centerline of the wave to the top of the crests, or the bottom of the troughs.

All of this is actually illustrative of a simple relationship between the wavelength, frequency, and speed of any wave. The most common notation for these quantities is to use the Greek letter λ (pronounced "lambda") to represent the wavelength, a lower case *f* to represent the frequency, and a lower case *v* for the speed. The relationship then looks like this:

$$\lambda \text{ x } f = v$$

Since I wanted this to be a non-mathematical introduction to astronomy, I'm not going to demonstrate any numerical examples here. But if you're mathematically inclined, you can easily show that the numerical values we were tossing around a couple of paragraphs ago are all consistent with this relationship. Note in particular that for a given wave—or more accurately for a given wave *speed, v*—a wave that has a short wavelength (a small value for λ, meaning that the ripples are close together) also has a high frequency (a large value for *f*, meaning that lots of ripples pass you each second), and vice versa. (If you're feeling a little muddled about that last point, hang on; we'll be returning to it in a moment, and explaining it more clearly.) Getting back to our discussion of the general properties of waves, the size of the ripples is referred to as the **amplitude**. The amplitude is usually measured as the distance from the centerline to the top of a crest or to the bottom of a trough (see Figure 5.6), where the centerline is an imaginary line that runs in the direction the wave is traveling, positioned halfway between the top of the crests and the bottom of the troughs. In other words, if you were watching a wave traveling across the surface of a pond, the centerline would simply correspond to where the surface of the water would be if there weren't any wave.

Let's see how these terms and concepts relate to light as a wave. If you look at Figure 5.1 again you'll see that near the lower left it says *short wavelengths* and near the lower right it says *long wavelengths*. This indicates that, if you want to think of light as a wave, the different kinds of light correspond to different wavelengths of light. The longest wavelengths of all are radio light, often known as *radio waves*; whereas the shortest wavelengths correspond to gamma rays.

Now since light always travels the same speed through empty space, our previous mathematical relationship looks like this:

$$\lambda \text{ x } f = c$$

for light traveling through space, where you will recall from before that a lower case c is always used to represent the speed at which light travels. This means that for any type of light, the product of the wavelength of that light multiplied by its frequency must be the same number as it would be for any other type of light. So the product of the wavelength and the frequency of radio light, for instance, must be the same number (the same speed) as the product of the wavelength and the frequency of gamma rays. Since radio light has a much longer wavelength than the gamma rays (its value of l is much larger), it follows that its frequency must be lower than the frequency of gamma rays (its value of f must be smaller). This is why in the lower left of Figure 5.1 where it says short wavelengths, it also says *high frequencies;* and in the lower right, where it says *long wavelengths,* it also says *low frequencies*. So referring again to what we mentioned before,

short wavelength light = high frequency light

long wavelength light = low frequency light

Within the range of light that you can see with your eye—*visible light*—the different wavelengths and/or frequencies correspond to the different *colors* that we see. The longest wavelength/lowest frequency light that you can see with your eyes is seen by you as "red," while the shortest wavelength/highest frequency light is seen as "violet."

The amplitude of a light wave corresponds to how bright the light looks. A bright source of light can be thought of as generating a large amplitude wave; a faint source of light is generating a small amplitude wave.

So how do you make a light wave? Well, that's an interesting question! Before we attempt to answer it, let's first look a little more carefully at water waves, since they are probably the type of wave that is most familiar to all of us. Let's imagine then that you are standing waist deep in a flat pool of water. Let's imagine also that some distance away from you is a ball, floating in the water. Now imagine that you hold your hand on the surface of the water and begin moving it up and down, up and down, over and over. As you wiggle your hand up and down, ripples will begin traveling outward from your hand across the water's surface. When those ripples reach the ball, the ball will start wiggling up and down as well, in response to your hand. The ball will mimic the motion of your hand. If your hand moves up and down a lot or a little, the ball will eventually start moving up and down a lot or a little. If your hand moves up and down quickly or slowly, the ball will start moving up and down quickly or slowly.

Obviously there will be a delay between what your hand does and what the ball does. If you change the motion of your hand it will take some time before the ball changes its motion in response. In fact, if you knew how far away the ball was you could use this time delay to figure out how fast the ripples are moving. This would be the *speed* of the wave, which might be an important thing to know.

Now imagine that the pool of water is invisible, and all you can see is the ball floating apparently in mid-air. You start wiggling your hand up and down, and some short time later the ball starts wiggling up and down, apparently in response to your hand. Even if you couldn't see the water and the ripples traveling across it, you might speculate that some kind of wave was being generated by your hand and traveling through the invisible space to the ball. After all, *something* is making the ball move up and down, and an invisible wave would be a perfectly reasonable explanation.

Now let's turn our attention to a very different scenario. Let's imagine that you have two particles that have electric charge—let's say a proton and an electron; and let's imagine that the electron is stuck inside a long thin strip of metal that is held vertically. Metals have the property of allowing electrons to move through them pretty easily, so this electron is able to move up and down, but it can't move side to side since it is trapped inside the metal strip. Of course, the electron would like to be as close to the proton as possible since it is attracted to the proton by the electromagnetic force. You can therefore expect it to move along the metal strip until it is as close to the proton as it can be.

Now imagine that you lift the proton up high. The electron will then move upward too, until it is level with the proton, assuming the metal strip is long enough for it to go that high. Now move the proton down low, and the electron will move down too, until it is again level with the proton. Now move the proton back up, and the electron will follow it up. Move the proton back down, and the electron will follow it back down. You can see that if you start wiggling the proton up and down, up and down, over and over, the electron will start wiggling up and down, up and down, too, mimicking the motion of the proton.

This situation would look very much like the floating ball responding to the motion of your hand. And just as that cause and effect is due to ripples, or a wave traveling across the water, so it is reasonable to speculate that the proton is generating some kind of ripples, a wave of some sort that is traveling through the intervening space and effecting the electron.

Just as with the ball on the water, there will be a delay between what the proton does and when the electron responds. When you move the proton up, the electron moves up too—but not immediately; there is some time delay. And just as with the ball on the water, you can use this time delay to figure out how fast the wave being generated by the proton is traveling. If you did this, you would find that the wave being generated by the wiggling proton is traveling at exactly the speed that light travels. And you might conclude from this, as the famous physicist James Clerk Maxwell did in the late 1800s, that the wave being generated by the proton *is* light.

This is in fact our wave description of light. The most accurate way to state it is that, in many situations light can be thought of *as if* it is a wave generated by wiggling electrically charged particles. If you wiggle them quickly, you generate high frequency light waves with short wavelengths. If you wiggle them slowly, you generate low frequency light waves with long wavelengths. If you wiggle them a lot, you generate large amplitude light waves—that is, a bright light. If you wiggle them comparatively little, you generate small amplitude light waves that will look like a faint light. Make sense?

5.4 Light from Astronomical Sources

Since this is an astronomy class rather than a physics class, our primary concern is not so much light itself, as how astronomers use light to learn about the Universe. And this begins of course when astronomers point their telescope at some astronomical source of light, and begin detecting and ultimately analyzing that light. To appreciate this process, we need to recognize that different types of objects emit light differently; and by "type" of object, I mean basically whether the object or source of light is a gas or a solid. What about liquids? Well, astronomers don't have to deal with liquids simply because liquids cannot exist in space. If you're an astronomer and you're pointing your telescope at something, you can be sure that what you're looking at is either (1) a gas; (2) a solid or very dense object; or (3) some combination of the two. Now in reality, there is only one way you can combine light from a gaseous source with light from a solid source, and that is if the gas is in front of the solid. Obviously, if the solid is in front of the gas, then you wouldn't be able to see the gas, because the solid object would be in the way. But if the gas is in front of the solid, then it's perfectly possible that you would be able to see the light from the solid object, but it would first have to pass through the gas.

So this then sets us up to explore three possible types of astronomical light sources. What you will find is that astronomers use either the photon/particle model of light or the wave model, depending on which seems to do a better job of explaining how light behaves. It's an odd way of approaching it, but it's the best we can do at this time.

5.4a Light from Individual Atoms: The Emission Spectrum

The difference between a gas and a solid is that in a solid the atoms or molecules are stuck to each other, or "bonded" to use the term chemists prefer; whereas in a gas the individual atoms and molecules are flying about freely, independently of each other. Hence, when you point your telescope at a cloud of gas, you are going to see light that is being emitted by individual atoms. And here again I have to warn you about a difficult topic ahead. Caution! Challenging course work ahead! The full description of how individual atoms emit light is, I would suspect, considered by many first-time astronomy students as the most difficult sub-topic of the entire class. I recommend that you read through the description I'm about

energy levels

The specific amounts of energy that an electron can have when it is bound to an atom or molecule.

to give you once and see if it makes sense. If not, go back and read it again and see if you can pick up a little more. Maybe give it a third reading. If you still haven't got it worked out after that, you may want to just continue with the book and not worry about it too much. You can always come back and look at it again later. If worse comes to worse, you can still understand most of the astronomy we'll be covering (and pass the class!) without it.

In brief, the light that is emitted from individual atoms comes from the electrons; and since light is energy, our starting point will be the simple fact that electrons in atoms have energy. In fact, electrons in atoms can have different amounts of energy. They can have a lot of energy. They can have very little energy. But they can't have zero energy. There is no way I can really explain why, except to say that zero energy is forbidden by the laws of quantum mechanics.

Now it turns out that an electron in an atom cannot have just *any* amount of energy. It can only have very *specific* amounts of energy. A physicist would say that the electron's energy is *quantized*. What do we mean by this? Well, think about a baseball. You can give a baseball energy by throwing it. That's kinetic energy. The harder you throw it, the faster it goes and the more kinetic energy it has. But what if the baseball could only travel at certain very specific speeds? What would that be like?

Let's say you throw the baseball and it travels at some speed—say 20 meters per second (20 m/s). You throw the baseball again, a little harder this time, but instead of moving faster the baseball still goes 20 m/s. You throw it again, still harder—but it still goes 20 m/s. You keep throwing it harder and harder, but it seems to have no effect at all on how fast the baseball moves. Finally, you throw it hard enough, and *now* the baseball goes faster—say 25 m/s. You conclude that the baseball can go 20 m/s or 25 m/s, but it can't go any speed in between. It apparently can't go 22 m/s or 24.5 m/s. This would be an example of quantized speed, and hence quantized energy for a baseball. It would certainly be an odd situation!

But that's what it's like for an electron. An electron can have certain amounts of energy in an atom. But if you try to give it more energy, it won't take it. You can keep trying to give it more and more energy, but it won't take it—it just keeps the same energy that it had. Then finally, if you give it *just the right amount of energy*, it will take the energy and become an electron with more energy.

The specific amounts of energy that an electron can have in a given atom are called **energy levels**, and they are generally represented by diagrams like the one shown in Figure 5.7. So let's imagine we have an electron in an atom, and this electron has the lowest possible energy that is allowed for that atom. This is known as the *ground state* (*Level 1* in Figure 5.7). You want to get the electron up to a higher energy level, so you're going to have to give it some more energy. How do you do this? Well, there are lots of ways. You could hit it with something small, say another electron. But a particularly neat way to do it is to "hit" it with a photon. Photons are after all 'blips' of pure energy; so if you can get the electron to absorb a photon, you will have succeeded in giving it some energy.

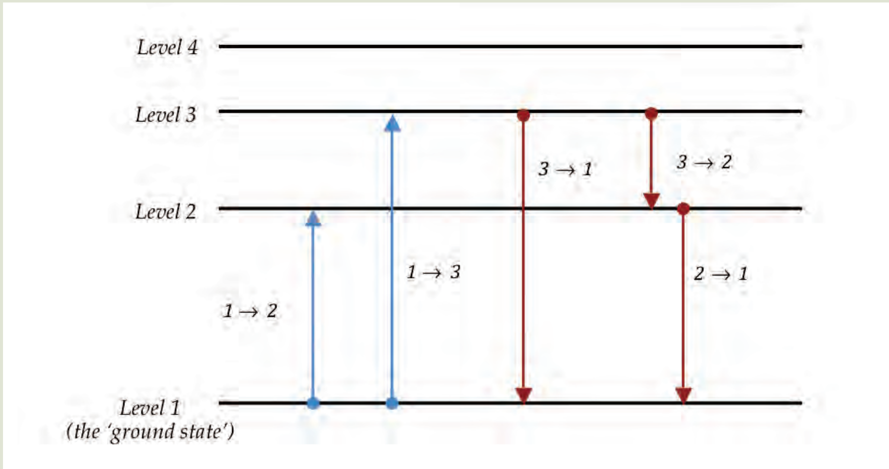

Figure 5.7 Simplified Energy Level Diagram for the Electrons in a Typical Atom

Let's say first that you start hitting the electron with very low energy photons. What happens? Nothing. The electron won't absorb the photons, so the photons will just pass on through. Next you start hitting it with some higher energy photons, but the electron ignores them as well so they pass right on through. You keep increasing the energy of the photons, but the electron just keeps ignoring them; until…

When you hit the electron with a photon that has *just the right amount of energy*—just the amount of energy that it needs to get to the next highest energy level (*Level 2* in the diagram)—*then* the electron absorbs that photon. The photon is gone, it no longer exists, and its energy is given to the electron. The electron is no longer in the ground state; it now has an amount of energy that corresponds to the second energy level (transition *1* → *2* in the diagram).

Now what happens? Well, electrons "want" to have as little energy as possible. So if you wait long enough, the electron is going to give up the energy that you gave it; and when it does, it will "drop" back down to the ground state, the lowest energy state (transition *2* → *1* in the diagram). How does it do this? It emits a photon. It emits a photon with exactly the amount of energy that it needs to get rid of in order to go back down to the ground state. In other words, it emits a photon with exactly the same amount of energy as the photon that it absorbed.

That's how you make a photon! You get an electron up to a high energy level—any level higher than the ground state—and then wait for it to drop back down.

Now, if you're thinking that this is a "Which came first, the chicken or the egg?" problem—we had to have a photon to make a photon—that's not really true. Remember that I said there are lots of ways to give energy to an electron, you don't have to hit it with a photon. But it's a useful option that we'll return to a little later.

What if you hit your ground state electron with a photon that has enough energy to get it up to the third energy level—two levels up from the ground level? What then? Well, *provided the photon has exactly the right amount of energy, which would be the energy difference between level 3 and level 1*, the electron will absorb the photon; the photon will disappear; and the electron will have an amount of energy that corresponds to level 3 (transition *1 → 3* in the diagram). But now the electron has *two* possibilities for what it can do next. It could re-emit a photon that is identical to the one that it just absorbed. If it does, it will be giving up exactly the amount of energy that it just absorbed, and so will "drop" all the way back down to the ground state again (transition *3 → 1* in the diagram).

However, it could also emit a photon with less energy than the one it absorbed—a photon with energy equal to the difference in energy between levels 3 and 2—and "drop" down to level 2 (transition *3 → 2* in the diagram). Then, if you wait long enough, the electron is going to give up a second photon, one whose energy corresponds exactly to the difference in energy between level 2 and the ground state, and return to the ground state (transition *2 → 1* in the diagram).

So if an electron is at level 3, there are two possibilities for what you might see: You might see a single, high-energy photon come out of the atom, or you might get two low energy photons, one after the other. The electron doesn't really have a "choice" in this. It's not like it thinks it through and makes a decision. But one interesting fact about it is that there doesn't seem to be any *law* that tells the electron what to do. There is no law that says: "If this occurs, then the electron will emit one photon; otherwise, it will emit two photons." All we know is that the electron's behavior will follow certain laws of *probability*.

For example, let's say that for a particular atom there is a 50% chance that an electron at leve 3 will drop down to the ground state in one jump and emit a single photon, and a 50% chance that it will take two steps and emit two photons. Even knowing what the probabilities are like this you still have no idea what any particular electron will do. All you can say is that, if you bump the electron up to the third level 100 times, then on average 50 of those times it will drop down to the ground state in one step, and the other 50 times it will take two steps. It's like flipping a coin. You don't know whether it's going to be "heads" or "tails," you just know that if you flip it 100 times, roughly 50 of them will come up "heads" and 50 "tails." That's how probabilities work.

This is a particularly interesting fact, especially when you realize that this is one of the fundamental phenomena of nature. It can be a little unsettling to realize that physics tells us that some of the most fundamental phenomena of nature have no law behind them whatsoever, but are simply statistical flips of a coin. As I've alluded to a couple of times already, the branch of physics that describes all of this is known as *quantum mechanics*. It was because of his objection to this aspect of quantum mechanics that Einstein once wrote in a letter to a colleague that "the *Old One* does not play dice," a comment which is often quoted as, "God does not play dice!"

So photons are created when an electron in an atom drops from a high energy level down to a lower energy level. Because the energy levels in an atom are quantized, it follows that only photons with very specific energies are emitted by the atom. It is also a fact that the specific energy levels that are available to an electron in an atom depend on what type of atom it is; and the energy levels are unique to that type of atom. For example, the hydrogen atom has very specific energy levels that are available to its electron. The helium atom has different energy levels available to its electrons. The carbon atom has still different energy levels. Since the photons are created when electrons jump from one level down to another level, it follows that the specific photons that are emitted by an atom are unique to that particular type of atom. The specific photons that are emitted by a hydrogen atom are different than the specific photons emitted by a helium atom, or by any other atom.

So here is something that you can do: You can take a chamber and fill it with hydrogen atoms. Then, either by sending electricity into it to energize the atoms, or by some other means, you can get the electrons in those hydrogen atoms to jump up to higher energy levels. Then, when they drop down, you can record the photons that they emit. If you have lots of hydrogen atoms doing this continuously, what you will get is a whole bunch of photons from the hydrogen atoms; and those photons will have very specific energies. If you then take those photons and send them through a device like a prism that will split up light according to different colors, or photon energies, then you will be able to see exactly which photon energies are coming from the hydrogen atoms. This pattern of specific colors, or photon energies is called the **emission spectrum** for hydrogen. Each specific color or photon energy is known as an **emission line**. If you then fill your chamber with helium atoms, you will see the emission spectrum for helium. Since the energy levels are unique to each type of atom, the emission spectrum is also unique to each type of atom (see Figure 5.8).

emission spectrum

The type of spectrum emitted by individual atoms or molecules, and therefore the type of spectrum that is emitted by a gas.

emission line

A specific photon energy (or wavelength or frequency) emitted by electrons making a particular energy transition, from a higher energy level to a lower one.

nebula

A cloud of gas out in space somewhere.

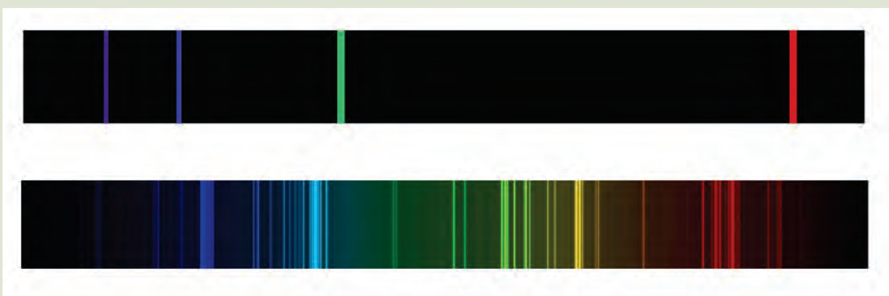

Figure 5.8 Emission Spectra

The top spectrum is for hydrogen atoms; the bottom spectrum is for nitrogen atoms.

"Spectral lines emission," by Stkl, via German Wikipedia (top spectrum source);
"Spectral lines absorption," by Stkl, via German Wikipedia (bottom spectrum source).

What scientists have done over the last century or so is measured and catalogued the emission spectra for every conceivable type of atom and molecule. Molecules also have unique emission spectra, since when two atoms are bonded to each other in a molecule they affect each other's energy levels. This catalog of emission spectra provides one of the most useful tools that astronomers have at their fingertips. Here's how it works.

Let's say you are looking at a cloud of gas—what astronomers call a **nebula**—out in space somewhere. You take the light that is coming from the nebula and split it up according to the different colors, or photon energies. You will then obtain the emission spectrum of that nebula. Now look at the spectrum and see what emission lines are present. If the nebula has hydrogen atoms in it, you will see the emission lines for hydrogen in its spectrum. But unless the nebula is pure hydrogen, you will also see the spectra of other atoms. If there is helium in the nebula, the emission lines for helium will be mixed into its spectrum. If carbon dioxide is there, you'll see the spectrum of that molecule. If there is carbon monoxide, those emission lines will be there too. There may be water; there may be ammonia. You will see the spectra, the emission lines of all of the different atoms and molecules that make up that nebula all mixed together in its emission spectrum.

spectroscopy
The process of splitting up a beam of light to see which photon energies/wavelengths/frequencies are present in the beam.

So by studying the nebula's spectrum, *you can learn what it is made of*. This is how astronomers learn which atoms and molecules are out there, and how abundant each kind is. As I say, this is an *extremely* powerful tool, and it is one that works really, really well. The study of spectra is known as **spectroscopy**; and spectroscopy is a huge branch of astronomy.

An emission spectrum is what you see when you look at the light coming from an individual atom. It is also what you see when you look at the light coming from lots of atoms, as long as they are free to act on their own. In other words, it is what you see when you look at the light coming from a gas, because a gas is the phase of matter in which atoms are free to act on their own.

There! That's most of the difficult part. If what we've talked about in this section makes sense to you, you're in good shape. If it still seems kind of shaky, you might want to go back and re-read this section, since we'll be using these ideas again in a few pages.

5.4b Light from Atoms Working Together: The Continuous Spectrum

continuous spectrum
The type of spectrum emitted by solid or very dense objects, in which the atoms or molecules are tightly bonded to each other.

What happens when atoms are bound together in a solid? Well, it turns out that something very strange, almost magical happens when you force atoms to be so close together that they are bonded tightly to each other to form what we call a solid. In this situation the electrons in a particular atom no longer respond only to that particular atom; they also interact with their neighbors around them. What this does is open up other possible energies for the electron, energies that might not have been available to it if that atom were on its own. So in addition to the energy levels that are available from the particular atom to which it belongs, energy levels that are unique to that particular atom, each electron also has other energy levels available to it because of its interaction with other atoms around it.

Now, you might think this would make things pretty complicated. But this is where the magical thing happens—for it actually makes thing simpler! Through a mathematical argument based on statistics, it is possible to show that when atoms are bound together so tightly as to form a solid, the electrons in those atoms get so many energy levels opened up to them that they essentially can have any energy possible. In other words, the energy levels are no longer quantized; they become *continuous*. The electrons in a solid can take on any energy at all within some range. The range of energies available to the electron is determined by one and only one thing—and this is the second magical thing! The range of energies available to the electron depends only on the temperature of the solid. That is, it has nothing to do with the *type* of atoms or molecules the solid is made of.

Now that is really quite remarkable. What it comes down to is this: If you look at the light coming from a solid object, and you spread the light out according to all of the different colors or photon energies, instead of seeing specific photon energies—like the emission spectrum of a gas—you will see photons of all possible energies, within some range. There will be a continuous spectrum of colors or energies within some range; and the range will have nothing to do with what the solid is made of, but only on its temperature. This type of emission, which is clearly different than the emission from a gas, is called *continuous emission,* or a **continuous spectrum**; or sometimes thermal emission. It is the emission that you see from a solid object (or an object that is sufficiently dense).

Figure 5.9 A Continuous Spectrum

"Spectral lines continuous," by Stkl, via German Wikipedia

If you look at the visible light range of the spectrum from a solid or dense object, instead of individual, specific emission lines, you will see a continuous rainbow of colors. Since this doesn't convey much information, continuous spectra are usually represented schematically by a plot. Along the horizontal axis you put the energy of the photon, and on the vertical axis you plot how much light the object will emit at each photon energy. A plot of this type of emission is shown in Figure 5.9. Thermal emission curves for objects at four different temperatures have been shown. Again, you don't have to know what the objects are made of. It doesn't make any difference. Any object at a temperature of 300 K will emit the same thermal emission as any other object at300 K. Magic!

Now, let me clarify this. What you're seeing here is the amount of light that is emitted *from a given sized region on the surface of an object*. It's usually stated as the amount of light emitted *per square meter*, although you could just as well measure it per square centimeter, or per square foot, or whatever. It's important to realize that if you're going to compare the amount of light from two different objects, you have to compare the light being emitted by equal sized areas on the surfaces of the two objects. In general, the entire surface of an object will be emitting light, depending on the temperature of the surface. If one object is bigger than another, then it obviously has a lot more surface area that is emitting light. For the comparisons in Figure 5.10 then, we'll will assume that the objects are all the same size.

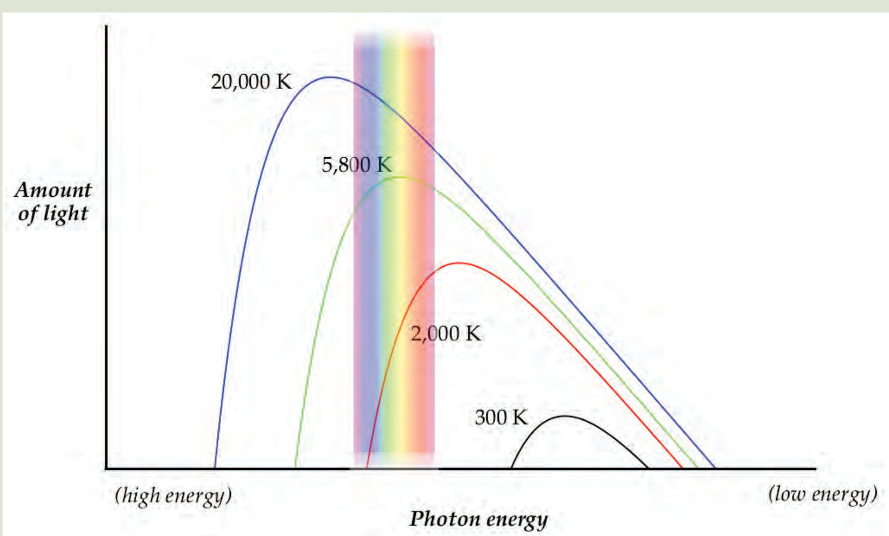

Figure 5.10 Thermal (continuous) Emission Curves for Objects at Different Temperatures

Let's look first at the 300 K object. This is approximately "room temperature." So this could be the light emitted by any object in the room with you right now; or more specifically, the light emitted *per square meter* from the surface of any object in the room. Of course, the objects around you are not all *exactly* 300 K, but they're pretty close.

These objects that are around you right now are emitting light. The chair you're sitting on is emitting light. Your body is emitting light! They are all examples of solid objects that are composed of atoms that have electrons that are continuously being "bumped" up to high energy levels and then dropping down to lower levels, emitting photons while they do. Because the objects are solid, the atoms are packed so tightly together that the electrons have essentially any energy available to them, so that the atoms are emitting photons with all possible energies *within a certain range*. This range depends only on the temperature of each object.

Now because the objects around you right now are not very hot, the photons they are emitting are all fairly low energy photons. Specifically, the photons are entirely in the range known as infrared. So even though the objects are emitting light, it is not light that you can see with your eyes. That's why if you turned out all of the lights in the room, and it was dark where you are sitting, you wouldn't be able to see anything. Even though the objects *reflect* visible light that is generated by light bulbs or the Sun, they only *emit* infrared light. So you can only see them by reflected light.

Now let's imagine that we take one of the objects that is near you right now and we begin heating it up. If we want to be realistic we'd better choose an object that you *can* heat up without destroying it, so we'd better not use you, or this book. How about a metal paper clip that is sitting on your desk? (If there isn't a metal paper clip sitting on your desk, pretend there is.) As you heat up the paper clip, two things will happen. First, it will begin to emit more light at *all* photon energies. The hotter it becomes, the more light it will emit across the board, so to speak. Second, the energy at which the paperclip emits *most* of its photons will shift towards higher energy photons. The hotter it becomes, the more the peak will shift to higher and higher energy. You can see both of these effects illustrated in Figure 5.10 by comparing the four curves that are shown.

Getting back to our paperclip—since it is initially at room temperature, it will only be emitting infrared light to begin with. But because of this shift towards higher energy photons, if you heat the paper clip enough it will reach a temperature where it will begin to emit a significant number of visible light photons (compare the 300 K and 2,000 K curves in Figure 5.10). This means that you can see it! Recall that the lowest energy photons that you can see with your eye appear to you as the color red. So when the paper clip first begins to emit visible light photons, it will begin to glow red. You would say that it is "red hot"!

What would happen if you kept heating up the paper clip? Now, I have no idea what the melting temperature of a paperclip is, so maybe you couldn't heat it up any further. But let's pretend. What if you heated it to, say, 5,800 K? At 5,800 K it would be emitting even more light at all photon energies than it was before. It would be emitting more infrared photons, for instance, than it was at 2,000 K or 300 K. But at 5,800 K the peak would have shifted all the way into the visible range, right smack in the middle of it in fact. In other words, it would then be emitting mostly visible photons. So you would no longer be seeing only red photons coming from it, you would be seeing orange photons, and yellow photons, and green and blue and violet photons; and you would be seeing them in roughly equal quantities. Now, we mentioned before what it looks like to you when your eye takes in all of the colors of visible light, all mixed together in equal amounts—you see white. So the paper clip would no longer be "red hot," it would be "white hot."

Interestingly, at 5,800 K the paper clip would be essentially the same temperature as the surface of the Sun. At this temperature, the light emitted by the paper clip *per square meter*—or since it's a paper clip, we could say *per square millimeter*—would be essentially the same as the light that is emitted by the Sun per square millimeter of its surface. Remember—thermal emission has nothing to do with what the object is made of! It depends only on the temperature.

Finally, what would happen if you made the paperclip *really* hot? Hotter even than the surface of the Sun? What if you heated it to, say, 20,000 K? Well, as you can see in Figure 5.10, at 20,000 K the peak in the curve is actually to the left of the visible light; in other words, the paperclip would be emitting mostly *ultraviolet light*. But notice that this does *not* mean that you can't see it. Even though most of its emitted light would be beyond the range of our eyes, it would still be emitting plenty of visible light. In fact, if you compare one more time the curves in Figure 5.10, you'll see that the more you heat up a given object, the more light it emits at *all* photon energies. So you can't make something invisible by heating it up too much!

5.4c Light from a Solid Seen through a Gas: The Absorption Spectrum

There is one other situation we need to mention to complete the picture of what light looks like when we observe astronomical objects. If you are looking at a solid object, you see a continuous spectrum. If you are looking at a gas, you see an emission spectrum. What if you are looking at a solid object, but you are seeing it *through a cloud of gas*? You might be observing a solid planet, for instance, on the other side of a cloud of gas; or a planet "hidden" by its own atmosphere. Let's apply what we know so far to this situation.

The solid object will be emitting all possible photons within a certain range of photon energies. This is thermal emission, which gives a continuous spectrum. But before those photons get to you they have to pass through the cloud of gas, and the gas is filled with atoms, each of which has one or more electrons; and those electrons have specific energy levels that are available to them, that are unique to that type of atom. Now remember that one way to get the electrons in an atom to jump to higher energy levels is to hit them with photons that have just the right energy.

So if the thermal emission from your solid object is sending photons of all possible energies (within a given range) through that cloud of gas, then some of the photons will have energies that are exactly the energies that the electrons need to jump up to their higher energy levels. When an electron is hit by one of these "just right" photons, it will absorb it. What that means is that as your photons from the solid object pass through the cloud of gas, some of them aren't going to make it. They will get absorbed by the gas atoms. So your nice rainbow of continuous colors, or photon energies, will have gaps in it. These gaps are called **absorption lines**, and the resulting spectrum is known as an **absorption spectrum**. What's more, the energies at which the absorption lines occur are exactly the same as the energies that correspond to the emission lines for that particular kind of gas. To help you see what I mean, Figure 5.11 shows both the emission spectrum and the absorption spectrum for hydrogen.

absorption lines
The spectrum that is seen when a continuous distribution of photon energies (as emitted by a solid object, for example) passes through a gas. The spectrum is a combination of the continuous spectrum emitted by the solid object, and the specific photons that are absorbed by the gas.

absorption spectrum
A specific photon energy (or wavelength or frequency) absorbed by electrons making a particular energy transition, from a lower energy level to a higher one.

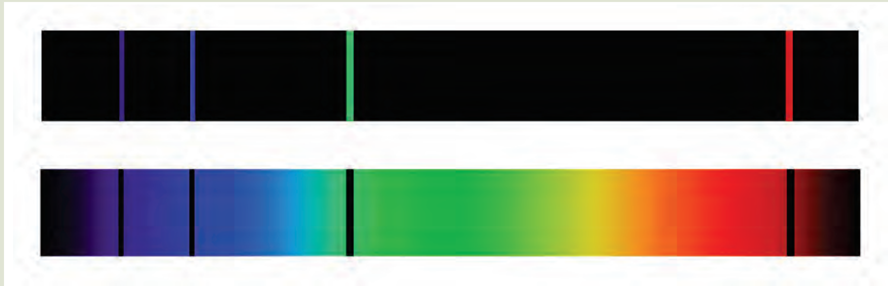

Figure 5.11 The Emission and Absorption Spectra of Hydrogen

"Spectral lines emission," by Stkl, via German Wikipedia (top spectrum source);
"Spectral lines absorption," by Stkl, via German Wikipedia (bottom spectrum source).

Now you might think, "Wait a minute … aren't the electrons in the atoms in that cloud of gas going to re-emit those same photons? Won't I get the continuous rainbow back again anyway?" Well, yes, the electrons will re-emit the photons that they absorb. But here's the thing: Even though the atoms that absorb those photons *will* re-emit them, *they won't necessarily re-emit them in the direction of your telescope*. In other words, for every 100 photons that would have made it to your telescope if they hadn't been absorbed, maybe only 2 of them will be re-emitted in the direction of your telescope. The other 98 will go off in other directions and you'll never see them. So that's why the gaps are there.

What it comes down to is that if you look at a solid object seen through a cloud of gas, you will see what is called an *absorption spectrum*. It will look like a continuous rainbow with gaps in it. The continuous rainbow is due to the thermal emission from the solid object. The gaps are due to the atoms in the gas, and their specific energies are unique to the type(s) of atoms in the gas. So you can use an absorption spectrum to find out what a gas is made of just as you can an emission spectrum. An absorption spectrum is what you see if you look at a planet, if the planet has an atmosphere. It's also what you see when you look at a star, because even though the star itself will radiate thermal emission like a solid object, stars are always surrounded by layers of gas, known as *stellar atmospheres*, that will absorb some of the photons.

5.4d A Quick Summary of a Challenging Topic

The last three sections have described three particular scenarios that astronomers deal with regularly, where each scenario has to do with light being emitted by a specific kind of astronomical source. If we were to reduce these three scenarios into quick bullet points, it would look like this:

- If you look at a cloud of gas, you will see an emission spectrum. You can use this emission spectrum to learn what the gas is made of.
- If you look at a solid (or very dense) object, you will see a continuous spectrum (thermal emission). You can*not* use this to determine what the object is made of, but you *can* use it to figure out the temperature of the object.
- If you see a solid object through a cloud of gas, you will see an absorption spectrum. You can use the background continuous spectrum to determine the temperature of the solid object, and you can use the absorption lines to determine what the cloud of gas is made of.

A visual representation of this summary is shown in Figure 5.12.

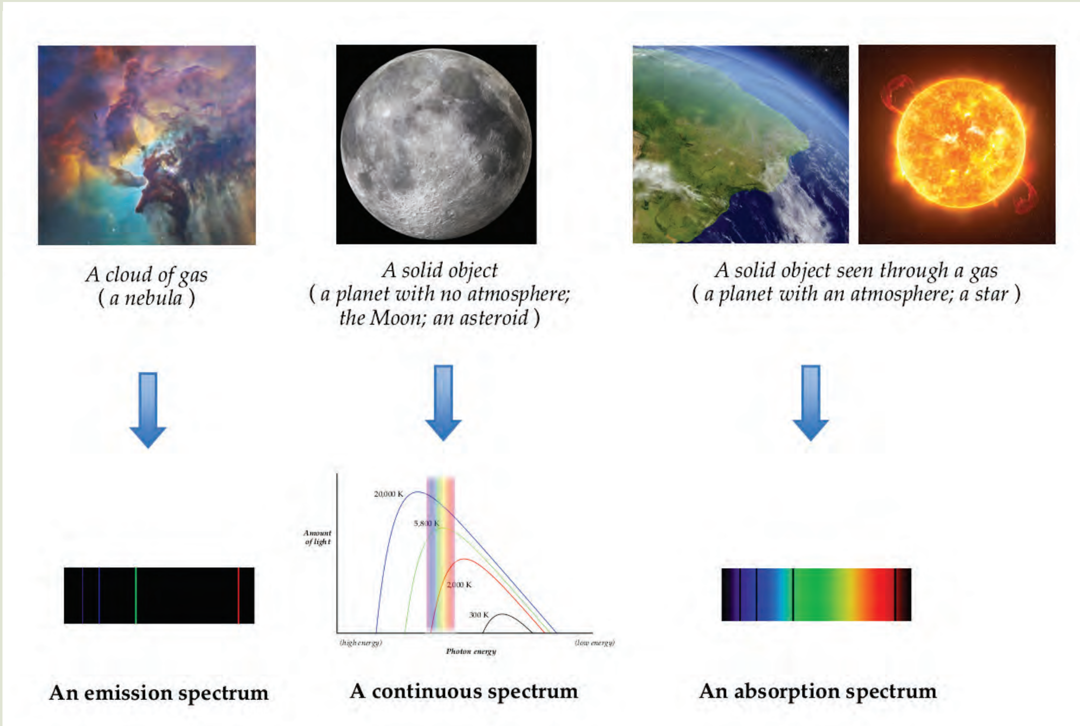

Figure 5.12 The Three Types of Spectra and the Astronomical Sources Associated with Each

Left to right, top : "Spectral lines continuous," by Stkl, via German Wikipedia; Moon, Earth and Sun images: iStock /Getty Images; bottom left: "Spectral lines emission," by Stkl, via German Wikipedia; bottom right: "Spectral lines absorption," by Stkl, via German Wikipedia

And now you can relax. You've made it through the most difficult part of the book! At least I suspect it's the most difficult part for most students. You might want to re-read Section 5.4, just to make sure you've got it. But then, maybe you just did…?

5.5 So What is Light? Particle or Wave?

You have now been introduced to the two models, or descriptions that physicists and astronomers use when they are talking about light. Sometimes light *behaves* as if it is a wave, as if it is composed of ripples traveling through space. At other times light *behaves* as if it is made of little particles, little 'blips' of energy called photons. Which description is correct? *Both, always*. Never mind that it is impossible for anything to be both a wave and a particle. I never said that light was. I said that both descriptions are correct. Light can always be described *as if* it is a wave, or *as if* it is photons.

So if both descriptions are correct, which one should you use? Well, the simplest answer is: Choose whichever description makes more sense to you! If you study light a lot, as astronomers do, you will find that there are some situations where the wave description makes more sense, and there are some situations where the photon description makes more sense. Most scientists therefore choose their description according to the context. That's what I'll be doing throughout the book, so even if you prefer one description over the other (I always prefer the photon description myself) you do need to be familiar enough with both descriptions so that you won't be thrown off by what you're reading.

Although we've hardly scratched the surface, we have said all that we need to say about the nature of light itself. However, one related topic that is worth a brief look is the tools that astronomers use to collect light and study it. So that will be the last tool that we'll look at, which will form the subject of the next chapter: telescopes.

Key Terms

absorption lines, **pg. 89**
absorption spectrum, **pg. 89**
amplitude, **pg. 79**
continuous spectrum, **pg. 86**
emission line, **pg. 85**
emission spectrum, **pg. 85**
energy levels, **pg. 82**
frequency, **pg. 78**
interference, **pg. 73**
interference pattern, **pg. 75**
nebula, **pg. 85**
photons, **pg. 77**
spectroscopy, **pg. 86**
wavelength, **pg. 78**