1 Estimation and Dimensional Analysis

How large a crater does an asteroid make when it impacts the Earth? How much does sea level change as global temperature changes? What is the average distance between bacterial cells in the ocean? Simple questions such as these frequently give us insight into more complicated ones, such as how often do large asteroids collide with the Earth, and can bacteria communicate with each other in the oceans? These are complicated questions, and to get accurate answers often involves using complicated computer simulations. However, by simplifying the problem we can often get a good estimate of the answer and a better understanding of what factors are important to the problem. This improved understanding can then help guide a more detailed analysis of the problem. Two techniques we can use to simplify complicated problems and gain intuition about them are *back-of-the-envelope calculations* and *dimensional analysis*.

Back-of-the-envelope calculations are quick, rough-and-ready estimates that help us get a feeling for the magnitudes of quantities in a problem.[1] Instead of trying to get an exact, quantitative solution to a problem, we aim to get an answer that is within, say, a factor of 10 (i.e., within an order of magnitude) of the exact one. To do this we make grand assumptions and gross approximations, all the time keeping in mind how much of an error we might be introducing. Back-of-the-envelope calculations also help us to understand which variables and processes are important in a problem and which ones we can ignore because, quantitatively, they make only a small contribution to the final answer.

Dimensional analysis is another useful tool we can use to simplify a problem and understand its structure. Unlike back-of-the-envelope calculations, which provide us with a quantitative feeling for a problem, dimensional analysis helps us reduce the number of variables we have to consider by examining the structure of the problem. We will rely on both techniques throughout this book.

1.1 Making Estimates on the Back of the Envelope

One of the first steps we have to take when tackling a scientific question is to understand it. What are the variables we need to consider? What equations do we need? Are there

[1] The myth of back-of-the-envelope calculations is that one should need a piece of paper no bigger than the size of the back of an envelope to do them. In reality, one sometimes needs a little more than that. However, the name conjures the right spirit, to use intuition and approximations to make the calculation as simple as you can, but not too simple!

assumptions we can use that will make the problem easier? Can we make an initial, rough estimate of the answer? This sort of understanding is needed whether we are tackling a complicated research problem, or a problem in a textbook. When we first start working on a new problem, we might feel unsure of how to proceed to a solution, particularly if the problem is in an area we are unfamiliar with. Our initial impulse is often to list all the variables and processes we think might be important and see if something leaps out at us. Back-of-the-envelope calculations can help us reduce this list by determining which variables and processes play quantitatively important roles in the problem.

To make good back-of-the-envelope calculations we need to be comfortable making good estimates of numbers. A good estimate is one that is likely within an order of magnitude of the actual value. We might wonder how we know this if we do not know the actual value. We do not. Like a painter who roughly sketches a scene, trying different arrangements and perspectives before undertaking the actual painting, or a writer trying different outlines before writing a book, we use back-of-the-envelope calculations to help us build a broad understanding of the problem we are tackling. We want to learn which variables and processes might be important for a more detailed investigation. For that we need good quantitative estimates. Estimating that the Earth is 2 km in diameter, or that a microbial cell is 1 m in diameter, will definitely lead us into trouble. But estimating that the diameter of the Earth is 12000 km, or that a microbial cell is 1 μm in diameter is acceptable. An actual microbial cell may be 2 μm in diameter, but this is only a factor of 2 different from our estimate. A more accurate value for the equatorial diameter of the Earth is 12756.28 km (Henderson and Henderson, 2009), so our estimate is only 6% off from the original and is far easier to remember. The idea is to develop a feeling for the magnitude of numbers, to build an intuition for the sizes of objects and rates of processes.

How accurate do we need our estimates to be? We may be tempted to give our answers to many decimal places or significant figures, but we should resist this because we are making only rough estimates. For example, using our estimate for the diameter of the Earth, we can estimate its surface area using $A = \pi d^2 \approx 3 \times (12 \times 10^6)^2 \approx 4.4 \times 10^{14}$ m^2 (a more accurate value is 5×10^{14} m^2, so our estimate is about 12% lower than the accurate value, good enough for a back-of-the-envelope calculation). Doing the calculation on a calculator yielded $A = 4.523889 \times 10^{14}$ m^2, but all the digits after the first or second significant figure are meaningless because we used an estimate of the diameter that differed from an accurate value by 6%. Keep in mind that the aim of a back-of-the-envelope calculation is to obtain a rough estimate, not a highly precise one, and a good rule of thumb is to keep only the first two or three significant figures when making an estimate—this also reduces the number of digits you have to write down and so minimizes the chances of copying a number incorrectly.

Our first back-of-the-envelope calculations will demonstrate how they can help us visualize the scales and magnitudes of quantities in a problem. In science, we frequently come across numbers that are either much larger or much smaller than those we experience in our daily lives. This can make them hard to visualize or think about clearly. For example, in the oceans bacteria are responsible for much of the natural cycling of elements such as carbon and nitrogen, and bacterial abundances in the surface waters are typically 10^5–10^6 cells cm^{-3}. But does this mean that the cells are crowded in the water and almost touching

Cambridge University Press
978-1-107-11748-8 — Mathematical Methods in the Earth and Environmental Sciences
Adrian Burd
Excerpt
More Information

each other? Or are they well separated? Having a good feeling or intuition for this helps us understand processes such as the ability of bacteria to take up nutrients, or to detect chemical signals that indicate the presence of food. We will return to this problem a bit later.

One simple technique that can help us visualize very large or small numbers is to compare them with similar quantities that we might be more familiar with. As an example, let us think about visualizing the Gulf Stream, which is a large, surface current in the North Atlantic Ocean that transports water and heat northward from the subtropics to more temperate latitudes. The transport of water in the Gulf Stream increases from approximately 3×10^7 m^3 s^{-1} near Florida, to approximately 1.5×10^8 m^3 s^{-1} near Newfoundland (Henderson and Henderson, 2009). These numbers are large, and it is hard to visualize a flow of hundreds of millions of cubic meters per second; we are probably not even used to visualizing volumes of water in units of cubic meters.

To put the flow of the Gulf Stream in perspective, we can compare it with something more familiar, but what should we choose? We experience the flow of water from a tap (or faucet) whenever we wash our hands, so we have an intuitive feeling for that. The idea is then to think, "How many taps would have to be turned on to obtain a total flow equivalent to that of the Gulf Stream?" However, the flow from a single tap is too small to make a meaningful comparison—we would end up with numbers as large as the ones we had trouble visualizing in the first place. Comparing the flow of the Gulf Stream to something that is larger and that we have seen for ourselves might make more sense. One possibility is to use the flow of a large river, such as the Amazon, instead of a tap. This has the advantage of having a much larger flow rate than a tap, and we stand a good chance of having seen a large river personally, or in movies, so we can visualize what it is like.

Exercise 1.1.1 What is the typical flow speed of a medium to large river? This question is intentionally vague to encourage you to use your experience. When you walk by a large river, is it flowing faster than your walking speed? Would you have to sprint to keep up, or could you amble along at a leisurely walking pace? You then have to ask how fast you walk!

Exercise 1.1.2 Taking the average width near the river mouth to be 20 km, and the average depth of 10 m, use your answer from Exercise 1.1.1 to estimate the discharge (in m^3 s^{-1}) of the Amazon River. Compare your answer with the number given in Table 1.1. If your answer is more than an order of magnitude different from that in the table, determine which of your estimated numbers could be improved.

Now that we have an estimate for the discharge of the Amazon River, we can compare it with the flow of the Gulf Stream. By simple comparison, the flow of the Gulf Stream is between 150 and 750 Amazon Rivers, or approximately between 2000 and 9000 Mississippi Rivers,[2] while the Amazon itself is equivalent to more than 10 Mississippi Rivers. In making this calculation we have effectively come up with our own "unit"—one Amazon River's worth of flow—for visualizing the transport of water on the scales

[2] The flow of water in the Gulf Stream is 50 times greater than the combined discharge of all the rivers that flow into the Atlantic Ocean.

Table 1.1 Approximate average discharge (volume rate of flow) of some major rivers in the world.

River	Average discharge ($m^3\ s^{-1}$)
Amazon	2.0×10^5
Congo	4.0×10^4
Ganges/Brahmaputra	3.8×10^4
Orinoco	3.6×10^4
Yangtze	3.0×10^4
Mississippi	1.7×10^4

Henderson and Henderson (2009).

of ocean currents. To do so, we came up with a quantity that is a few orders of magnitude different from the one we are interested in—it would be inappropriate to use the same scale for a small stream, for example. The point is, we can come to grips with quantities that are far larger or smaller than those we experience every day by comparing them with things that are more familiar to us.

Exercise 1.1.3 Estimate the amount of milk (or your favorite drink) you drink in a week, and use this to estimate how long would it take you to consume 1 m^3 of it.

Exercise 1.1.4 What is the volume of water in a standard Olympic-sized swimming pool (50 m long, 25 m wide, and 3 m in depth)?

Exercise 1.1.5 How long would it take you to fly a distance equivalent to the diameter of the Earth?

Exercise 1.1.6 How many times could the Earth fit in the distance between the Earth and the Moon, and between the Earth and the Sun?

Exercise 1.1.7 Rainfall rates in a hurricane can be as high as 3.5 cm per hour within 56 km of the center of the hurricane. If that rainfall rate occurred uniformly over a circle of radius 56 km for 1 hour, how many Olympic-sized swimming pools would this fill?

Back-of-the-envelope estimates frequently involve more detailed calculations, but we need to always keep in mind that we are seeking an estimate, an answer that is likely accurate to within a factor of about ten. To do this, we sometimes need to know good estimates to certain numbers (e.g., the diameter of the Earth) and we need to learn how to make judicious approximations. Let us look at a more involved example.

Example 1.1 The average concentration of gold in seawater is approximately 100 fmol kg^{-1} of seawater (Falkner and Edmond, 1990; Henderson and Henderson, 2009). This is a very small number, but we can visualize it by recasting this number in terms of something more familiar. For example, if a gold ring contains 4 g of gold, how many rings could one make using all of the gold in the world's oceans?[3]

[3] The extraction of gold from seawater has actually been put forward as a serious business proposition several times.

We rarely come across a femto (see Appendix A) of anything in our daily lives, so it is hard to visualize what 100 fmol (i.e., 100×10^{-15} moles) of gold looks like. Instead of thinking about such a low concentration, we might ask what is the mass of gold in 1 kg of seawater. This raises another question: what does 1 kg of seawater look like? The density of seawater is approximately[4] 1000 kg m^{-3}, so 1 kg of water occupies 10^{-3} m^3 = 1 L, or the equivalent of a milk carton. The atomic weight of gold is 197 g mol^{-1}, so 100 fmol kg^{-1} is the same as

$$100 \times 10^{-15}\,\text{mol} \times 197\,\text{g mol}^{-1} \approx 100 \times 10^{-15} \times 200\,\text{g} = 2 \times 10^{-11}\,\text{g L}^{-1}.$$

Notice that we have approximated 197 by 200 to make our numbers easy, and which incurs an error of only 1.5%. To calculate the total amount of gold in the oceans, we need to know the total volume of the world's oceans. Knowing the radius r of the Earth ($\approx$ 6000 km) we can calculate its surface area: $4\pi r^2 \approx 12 \times (6 \times 10^3)^2 \approx 12 \times 36 \times 10^6 = 432 \times 10^6$ km^2, assuming $\pi \approx 3$. The average depth of the oceans is 4 km and they cover approximately 70% of the Earth's surface (Henderson and Henderson, 2009), so we can calculate the volume of the oceans; $\approx 1.2\times10^9$ km^3 or $\approx 1.2\times10^{18}$ m^3 or $\approx 1.2\times10^{21}$ L. Now, we estimated earlier that 1 L of seawater contained about 2×10^{-11} g of gold, so the total amount of gold in the oceans is approximately 2×10^{10} g of gold, enough to make about 5×10^9 rings.

Let us return to the problem of the bacteria in the ocean that we described earlier in this section. Typical abundances of bacteria in seawater are 10^6 cm^{-3}. How can we determine if the cells are crowded together or not? One approach is to think of the distance between the cells in units of the typical size of a cell. First, we have to estimate a typical distance between cells. One way to do this is to assume that the cells are uniformly distributed in the 1 cm^3, so that the typical distance between them will be

$$l = \left(\frac{1}{10^6\,\text{cm}^{-3}}\right)^{1/3} = 10^{-2}\,\text{cm} = 100\,\mu\text{m}.$$

A typical diameter of a bacterial cell is about 1 μm, so this means that we could fit 100 bacterial cells between each bacterium. From the perspective of an individual bacterium, that is quite a low density of cells and has implications for the mechanisms that bacteria use to detect chemical signals and survive in the oceans.

The real power of the back-of-the-envelope calculation appears when we want to obtain quick, approximate answers to complicated problems. This can be useful if we want to know whether or not a problem is worth pursuing in more detail, or whether it is a small (though possibly interesting) effect in the big scheme of things.

Example 1.2 Between 1900 and 2010, Greenland lost an estimated 9×10^{12} tonnes of ice. We might wonder how much of this ice contributed to global sea level rise. To figure this out, we can estimate the rise in global sea level if all this melting ice contributed to sea level rise. First, we need to determine the volume of ice that has been lost. We can use the

[4] The density of seawater varies with temperature, salt content (i.e., salinity), and pressure. The average density of seawater at the surface is 1025 kg m^{-3}. So our estimate introduces an error of approximately 2%.

fact that 1 kg of water occupies a volume of 1 L—bearing in mind our approximation from Example 1.1. So, 1 m^3 of water has a mass of 1 tonne and 9×10^{12} tonnes of ice occupies 9×10^{12} m^3, or 9×10^3 km^3.

To obtain the rise in sea level this would cause, we need to make some simplifications about the shape of the oceans. As we move offshore, the depth of the ocean generally increases relatively slowly until we reach what is called the *shelf break*, where the depth increases more rapidly from an average of about 130 m down to the *abyssal plain* at a depth of about 4000 m. The shallow coastal regions make up less than 10% of the total area of the oceans. So, we can approximate an ocean basin as being a straight-walled container with sides 4000 m tall. We will also assume that the melting ice gets uniformly distributed throughout all the world's oceans, so we can combine them into a single ocean. To get the change in sea level height, we simply divide the volume added to the oceans from the melting ice by the total surface area of the oceans. We have already estimated that the surface area of the Earth is about 4.4×10^{14} m^2, so knowing that the oceans cover approximately 70% of the Earth's surface, we can estimate the area of the oceans as approximately 3×10^7 km^2 to get approximately 30 mm.

It is always a good idea to perform a "sanity check" after doing such a calculation, just to make sure that our approximations are reasonable. Over the twentieth century, global sea levels rose approximately 19 cm (Jevrejeva et al., 2008), so we estimate that about 15% of this came from Greenland losing ice.

Exercise 1.1.8 The Greenland ice sheet contains approximately 2.8×10^6 km^3 of ice. Estimate the mass of this ice sheet and compare it with the 9×10^{12} tonnes that was lost between 1900 and 2010. Estimate the rise in sea level if *all* of the Greenland ice sheet were to melt and flow into the oceans.

Solving back-of-the-envelope calculations can often involve many steps, and sometimes we get stuck and cannot readily see what the next step in the calculation should be. One tactic to use to get unstuck is to examine the units of the quantities we need to calculate and see if that provides enough information to move ahead. To illustrate this, consider the following question: atmospheric carbon dioxide concentrations are increasing and values are often given in units of parts per million (ppm). But at the same time, we hear that humans emit several gigatonnes (10^{15} tonnes) of carbon into the atmosphere per year (Le Quéré et al., 2016). How many gigatonnes of carbon emitted yields a 1 ppm change in atmospheric CO_2 concentration?

The first sticking point we have here is one of units: what is meant by parts per million? Parts per million by mass? By volume? This is quite an abused notation, and we have to be careful that we understand how it is being used in the context of the question. In atmospheric sciences, these units are really mole-fractions—that is, 1 ppm is really shorthand for "1 mole of specific stuff for every million moles of all the stuff combined." It just so happens that for gases, actually for ideal gases, the mole-fraction is the same as the volume fraction (ppmv) because of the ideal gas laws,[5] and atmospheric gases at room temperature and surface pressures behave almost like ideal gases.

[5] An ideal gas is an idealized gas of particles that only interact through collisions, with no forces of attraction or repulsion between them, and the collisions are "perfectly elastic," which means that none of the kinetic energy

Because 1 ppmv is a mole-fraction, we need to know how many moles of gas there are in total in the atmosphere in order to know how many moles of CO_2 are present. We could calculate this if we knew the molecular weight (grams per mole) of air and the total mass of the atmosphere. To tackle the first part we need to know the composition of air (approximately 79% N_2 and 21% O_2) and the molecular weights of the components of air (the molecular weight of N_2 is 28 and that of O_2 is 32). If we assume that the atmosphere is well mixed so that the composition of the air is everywhere the same, then the molecular weight of air is approximately

$$\begin{aligned}\left(0.79\frac{\text{mol N}_2}{\text{mol air}}\times 28\frac{\text{g N}_2}{\text{mol N}_2}\right)+\left(0.21\frac{\text{mol O}_2}{\text{mol air}}\times 32\frac{\text{g O}_2}{\text{mol O}_2}\right)\\ = 22.12\frac{\text{g N}_2}{\text{mol air}}+6.72\frac{\text{g O}_2}{\text{mol air}}\\ = 28.84\frac{\text{g}}{\text{mol air}},\end{aligned}$$

or about 29 g per mole.

Next, we need to estimate the total mass or total number of moles of gas in the atmosphere. Calculating the volume or mass of the atmosphere is difficult—the concentration of gases is not uniform with height, and where does the atmosphere end? But we might be able to find a way to estimate the mass of the atmosphere by listing what we know about it. In this way we can see if there are any quantities we know that have units containing mass. We know an average surface temperature, but it is hard to see how knowing something with units of temperature will help us calculate a mass. We need to estimate a mass, so we should try and list relevant variables that have the units of weight or force in it them.[6] How about pressure? Pressure is defined as a force per unit area, and Newton's laws tell us that force is mass multiplied by acceleration. Atmospheric pressure at the surface of the Earth is 1.01×10^5 Pa (N m^{-2}). To get the total force of the mass of the whole atmosphere, we need to estimate the surface area of the Earth, which is about 510×10^6 km^2, and we can look up the acceleration due to gravity (9.81 N kg^{-1}). So, the mass of the atmosphere is the atmospheric pressure at the surface multiplied by the surface area of the Earth and divided by the acceleration due to gravity,

$$\left(1.01\times 10^5\times\frac{1\text{ kg}}{9.81\text{ N}}\right)\times\left(510\times 10^6\text{ km}^2\right)\times\left(\frac{1000\text{ m}}{1\text{ km}}\right)^2\approx 5.2\times 10^{18}\text{ kg}.$$

Combining this with the average molecular weight we estimated earlier, the number of moles of gas in the atmosphere is

$$\left(\frac{5.2\times 10^{21}\text{ g}}{29\text{ g mol}^{-1}}\right)\approx 1.8\times 10^{20}\text{ moles}.$$

Our next step is to determine how much of this is in the form of CO_2. Because 1 ppm is 1 part in 10^6, 1 ppm CO_2 in the atmosphere is 1.8×10^{14} moles CO_2. There is 1 carbon atom in each CO_2 molecule, so 1 ppm of CO_2 corresponds to 1.8×10^{14} moles of

of the particles motion is converted to other forms of energy. The ideal gas law implies that equal volumes of any ideal gas held at the same temperature and pressure contain the same number of molecules.

6 Recall that Newton's laws tell us that a force is a mass multiplied by an acceleration.

carbon contained in CO_2 molecules in the atmosphere. The end result is that the mass of 1 ppm of C is

$$(1.8 \times 10^{14}\ \text{moles}) \times \left(12 \frac{\text{g}}{\text{mole}}\right) \sim 2 \times 10^{15}\ \text{gC} = 2\ \text{Pg C} = 2\ \text{Gt C}.$$

So 1 ppm of CO_2 corresponds to $\approx$ 2 Gt C. Knowing this allows us to quickly convert between the two sets of units when we see them in articles and research papers. It also allows us to ask other interesting questions, such as what is the contribution of fossil fuel burning to the rise in atmospheric CO_2 (see Problem 1.15)?

Back-of-the-envelope calculations can also be useful in determining spatial and temporal scales over which different processes are important. Many processes relevant to the Earth and environmental sciences have characteristic scales that determine how fast they occur and over what distances they work. For example, typical wind speeds over most of the United States vary between 4 and 5 m s^{-1}, but can be greater than 10 m s^{-1}. Open ocean surface currents have typical speeds of 0.1–2 m s^{-1}. So, we might expect the transport of gaseous pollutants in the atmosphere to be approximately 2–40 times faster than the transport of dissolved pollutants in the surface ocean.

Diffusion is an important process in both air and water, and we will meet it often in this book. Diffusion has the effect of smoothing out differences in concentration and is characterized by a quantity called the *diffusion coefficient* ($\mathcal{D}$) which, analogous to a velocity, is a measure of how fast diffusion can spread material. However, whereas velocity is a length divided by a time, the diffusion coefficient is a length squared divided by a time—we can think of it as the square of the distance a particle diffuses divided by the time it takes to diffuse that distance (Berg, 1993). This difference has important consequences for the distances and times over which diffusion is an important process. For example, the diffusion of a small molecule in air is roughly 10^{-5} m^2 s^{-1}, whereas in water it is $\approx 10^{-9}$ m^2 s^{-1} (Denny, 1993). We can use this to estimate the time (t) it will take a small molecule to diffuse a given distance (l), say 1 cm, in air and water:

$$t_{\text{air}} \sim \frac{l^2}{\mathcal{D}_{\text{air}}} = \frac{10^{-4}}{10^{-5}} = 10\,\text{s} \qquad \text{and} \qquad t_{\text{water}} \sim \frac{l^2}{\mathcal{D}_{\text{water}}} = \frac{10^{-4}}{10^{-9}} = 10^5\,\text{s} \sim 1\ \text{day}.$$

So, diffusion is a far slower process in water than in air, all other things being equal. What is more, because the diffusion coefficient is characterized by a length squared, it takes relatively longer to diffuse further distances. For example, to diffuse 10 cm takes 10^3 seconds (approximately 15 minutes) in air, and 10^7 seconds (about 116 days) in water. So, knowing something about the units of the diffusion coefficient and its value allowed us to estimate these diffusion times.[7]

Exercise 1.1.9 Estimate the time it takes a small molecule to diffuse a distance of 1 μm, 10 mm, 1 m, and 10 m in both air and water.

Exercise 1.1.10 Estimate the surface area of the Earth, the total surface area occupied by oceans, the total surface area occupied by land, and the total volume of the oceans.

[7] This is a calculation that is quick and easy to do, and can often be used to impress friends, family, and colleagues.

Exercise 1.1.11 Given that the average concentration of the salt in the oceans is 35 ppt, estimate the total mass of salt in the oceans and compare that to the mass of humanity on planet Earth.

1.2 Scaling

The phenomena we want to understand and explain in the Earth and environmental sciences cover a large range of spatial and temporal scales. At the smallest scales we might want to understand the processes of microbial interactions and how they affect biogeochemistry, or the nucleation of raindrops in the atmosphere. At the opposite end, the largest spatial scales encompass the planet, or large fractions of it. Consequently, it is useful to know if there are some general, unifying frameworks that allow us to understand how the importance of certain processes changes with scale. This is where *scaling arguments* become important.

There are generally two types of scaling that occur, *isometric*, or geometric, scaling and *allometric* scaling. Isometric scaling describes situations where the variables scale geometrically: for example, if you double the length of the side of a cube, the new surface area will be four times the old surface area, and the new volume will be eight times the old volume. In other words, the *shape* of the object stays the same, even though the size has increased. This geometric scaling can help us to understand how the importance of many processes changes with scale. For example, a microbial cell takes up nutrients through the surface of the cell, so all other things being equal, a cell B with twice the diameter of cell A should be able to take up nutrients four times faster than cell A. However, a cell's metabolic rate (a measure of how fast it uses energy) depends on its volume—the larger the cell, the more of it there is that has to be kept going. So, our geometric scaling argument implies that it should be harder for larger cells to obtain sufficient nutrients to support their energy needs than smaller cells, all other things being equal. What is more, this will vary by cell size according to the ratio of the cell surface area to its volume, and if the cells are spherical

$$\frac{\text{area}}{\text{volume}} = \frac{6}{\text{diameter}}.$$

Not all objects have a simple geometry like a sphere, so we might wonder what we use for a typical length scale when the object we are studying is not a sphere. Generally, most objects will have some characteristic length scale that is relevant to the problem and that we can choose to use. For example, if we are interested in relating maximum running speed to body length, we might choose stride length as a measure of length and relate this to body size. We could choose another variable such as leg length, and we could develop a scaling relationship using it, but leg length by itself is not necessarily a good indicator of running speed. A cheetah is about 60–90 cm feet tall at the shoulder, but has a running stride of several meters in length, much longer than a human.

A more fundamental question is whether or not geometric scaling always works. Galileo recognized that geometric scaling arguments often fail, even though data show that a scaling relationship still exists. These relationships, where a scaling relationship is not isometric, are *allometric*. If you watch a King Kong movie, you will see that Kong is just a geometrically scaled ape. However, if you carefully compare a mouse with an elephant, you will notice that the legs of a mouse seem thinner compared to their body size than a simple geometric scaling would suggest. This implies that larger animals require disproportionately thicker legs to support themselves.

Now, given that it is harder to break a thick branch than a thin twig, we might suspect that the diameter of the leg determines how easy it is to break it. The material that animal bone is made from is pretty similar between animals, so we expect the strength to be similar between animals. But what do we mean by "strength"? In this case we mean the strength of the bone to withstand fracture and buckling from just bearing the weight of the animal; in other words, a static measure of strength. The mechanical strength of a cylindrical shape (a good approximation to the shape of a leg bone) is proportional to its cross-sectional area (A)—cylinders with larger diameters are harder to buckle than those with smaller diameters and the same length. So, we expect that a heavier animal would need to have thicker bones, and hence thicker legs, to support its weight, so that $A \propto M$, where M is the animal's mass (if we double the body mass, we might expect to have to double the strength of the bone by doubling its cross-sectional area. The length (l) of a bone should scale with the size (L) of the animal,[8] so $l \propto L \propto M^{1/3}$, where we have assumed that the mass is proportional to the animal's volume. With the length and area of the bone, we can get a scaling for the bone mass, and so an estimate for the animal's skeletal mass (m), $m \propto A \times l \propto M \times M^{1/3} \propto M^{4/3}$, or $m = aM^{4/3}$ where a is a constant. If we take logarithms of this expression, we obtain the equation of a straight line $\log(m) = \log(a) + (4/3)\log(M)$, with a slope of 4/3. If we plot data for various bird species, for example, we find a slope closer to 1.0 than 1.33 (Figure 1.1). This is interesting because Figure 1.1 indicates that there is indeed a nice scaling relationship between skeletal mass and total body mass for birds, but it is not quite the relationship we expected from our geometric scaling argument. This tells us that something else is going on, and our assumptions are incorrect, so this is an example of an allometric scaling. In this case, there are several possibilities. One is that bone size is not determined by the ability of the bone to bear the animal's weight when standing still, but rather bone size is related to the ability of the bone to withstand dynamic processes such as walking (Prothero, 2015). Another possibility is that the structure of bones in a large bird is in some way different from that in smaller birds. So, comparing our scaling argument to data has revealed the assumptions behind our geometric simple scaling to be incorrect and has presented us with some interesting questions.

As another example, consider the scaling of river basins. River networks are formed from small streams that merge into larger rivers that themselves merge into still larger rivers until the final, large river discharges into the oceans (Figure 1.2). We can use two lengths to characterize the shape of the river basin, the length (L) and the average width (W). The area of the river basin is then $A \approx LW$. Observations of river basins show that

$$W \propto L^H, \qquad \text{where} \qquad \frac{1}{2} \lesssim H \lesssim 1$$

[8] This implies that a larger animal has more difficulty supporting its weight than a smaller one, and argues that an animal like King Kong could not exist; its bones would break when it moved. An excellent, if somewhat gruesome, description of this effect is given in Haldane (1945).