1 Wetting Your Feet

Most technical education emphasizes exact answers. If you are a physicist, you solve for the energy levels of the hydrogen atom to six decimal places. If you are a chemist, you measure reaction rates and concentrations to two or three decimal places. In this book, you learn complementary skills. You learn that an approximate answer is not merely good enough; it's often more useful than an exact answer. When you approach an unfamiliar problem, you want to learn first the main ideas and the important principles, because these ideas and principles structure your understanding of the problem. It is easier to refine this understanding than to create the refined analysis in one step.

The adjective in the title, **order of magnitude**, reflects our emphasis on approximation. An order of magnitude is a factor of 10. To be "within an order of magnitude," or to estimate a quantity "to order of magnitude," means that your estimate is roughly within a factor of 10 on either side. This chapter introduces the art of determining such approximations.

Writer's block is broken by writing; estimator's block is broken by estimating. So we begin our study of approximation using everyday examples, such as estimating budgets or annual production of diapers. These warmups flex your estimation muscles, which may have lain dormant through many years of traditional education. After the warmup, we introduce a more subtle method: scaling relations (Chapter 2).

Everyday estimations provide practice for our later problems, and also provide a method to sanity check information that you see. Suppose that a newspaper article says that the annual cost of health care in the United States will soon surpass $1 trillion. Whenever you read any such claim, you should automatically think: Does this number seem reasonable? Is it far too small, or far too large? You need methods for such estimations, methods that we develop in several examples. We dedicate the first example to physicists who need employment outside of physics.

1.1 Armored cars

How much money is there in a fully loaded Brinks armored car?

The amount of money depends on the size of the car, the denomination of the bills, the volume of each bill, the amount of air between the bills, and many other factors. The question, at first glance, seems vague. One important skill that you will learn from this text, by practice and example, is what assumptions to make. Because we do not need an exact answer, any reasonable set of assumptions will do. Getting started is more important than dotting every i; make an assumption—*any* assumption—and begin. You can correct the gross lies after you have got a feeling for the problem, and have learned which assumptions are most critical. If you keep silent, rather than tell a gross lie, you never discover anything.

Let's begin with our equality conventions, in ascending order of precision. We use $\propto$ for proportionalities, where the units on the left and right sides of the $\propto$ do not match; for example, Newton's second law could read $F \propto m$. We use $\sim$ for dimensionally correct relations (the units do match), which are often accurate to, say, a factor of 5 in either direction. An example is

$$\text{kinetic energy} \sim Mv^2. \tag{1.1}$$

Like the $\propto$ sign, the $\sim$ sign indicates that we've left out a constant; with $\sim$, the constant is dimensionless. We use $\approx$ to emphasize that the relation is accurate to, say, 20 or 30 percent. Sometimes, $\sim$ relations are also that accurate; the context will make the distinction.

Now we return to the armored car. How much money does it contain? Before you try a systematic method, take a guess. Make it an educated guess if you have some knowledge (perhaps you work for an insurance company, and you happened to write the insurance policy that the armored-car company bought); make it an uneducated guess if you have no knowledge. Then, after you get a more reliable estimate, compare it to your guess: The wonderful learning machine that is your brain magically improves your guesses for the next problem. You train your intuition, and, as we see at the end of this example, you aid your memory. As a pure guess, let's say that the armored car contains $1 million.

Now we introduce a systematic method. A general method in many estimations is to break the problem into pieces that we can handle: We **divide and conquer**. The amount of money is large by everyday standards; the largeness suggests that we break the problem into smaller chunks, which we can estimate more reliably. If we know the volume V of the car, and the volume v of a US bill, then we can count the bills inside the car by dividing the two volumes, $N \sim V/v$. After we count the bills, we can worry about the denominations (divide and conquer again). [We do not want to say that $N \approx V/v$. Our volume estimates may be in error easily by 30 or 40 percent, or only a fraction of the storage space may be occupied by bills. We do not want to commit ourselves.[1]]

We have divided the problem into two simpler subproblems: determining the volume of the car, and determining the volume of a

1. *"Once at a Fourth-of-July celebration, a reporter wondered and later asked why Mr. Murphy (he was always Mr. Murphy even to his closest associates) did not join in the singing of the National Anthem. 'Perhaps he didn't want to commit himself,' the boss's aide explained." From the Introduction by Arthur Mann, to William L. Riordan, Plunkitt of Tammany Hall (New York: E. P. Dutton, 1963), page ix.*

bill. What is the volume of an armored car? The storage space in an armored car has a funny shape, with ledges, corners, nooks, and crannies; no simple formula would tell us the volume, even if we knew the 50-odd measurements. This situation is just the sort for which order-of-magnitude physics is designed; the problem is messy and underspecified. So we **lie skillfully**: We pretend that the storage space is a simple shape with a volume that we can find. In this case, we pretend that it is a rectangular prism (Figure 1.1).

To estimate the volume of the prism, we divide and conquer. We divide estimating the volume into estimating the three dimensions of the prism. The compound structure of the formula

$$V \sim \text{length} \times \text{width} \times \text{height} \tag{1.2}$$

suggests that we divide and conquer. Probably an average-sized person can lie down inside with room to spare, so each dimension is roughly 2 m, and the interior volume is

$$V \sim 2\,\text{m} \times 2\,\text{m} \times 2\,\text{m} \sim 10\,\text{m}^3 = 10^7\,\text{cm}^3. \tag{1.3}$$

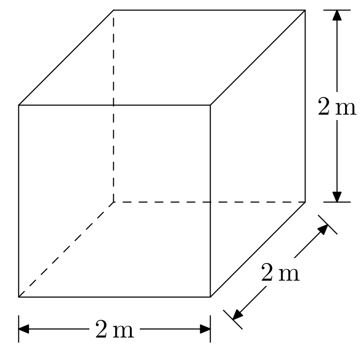

Figure 1.1. *Interior of a Brinks armored car. The actual shape is irregular, but to order of magnitude, the interior is a cube. A person can probably lie down or stand up with room to spare, so we estimate the volume as* $V \sim 2\,\text{m} \times 2\,\text{m} \times 2\,\text{m} \sim 10\,\text{m}^3$.

In this text, $2 \times 2 \times 2$ is almost always 10. We are already working with crude approximations, which we signal by using $\sim$ in $N \sim V/v$, so we do not waste effort in keeping track of a factor of 1.25 (from using 10 instead of 8). We converted the m^3 to cm^3 in anticipation of the dollar-bill-volume calculation: We want to use units that match the volume of a dollar bill, which is certainly much smaller than $1\,\text{m}^3$.

Now we estimate the volume of a dollar bill (the volumes of US denominations are roughly the same). You can lay a ruler next to a dollar bill, or you can just guess that a bill measures 2 or 3 inches by 6 inches, or $6\,\text{cm} \times 15\,\text{cm}$. To develop your feel for sizes, guess first; then, if you feel uneasy, check your answer with a ruler. As your feel for sizes develops, you will need to bring out the ruler less frequently. How thick is the dollar bill? Now we apply another order-of-magnitude technique: **guerrilla warfare**. We take any piece of information that we can get.[2] What's a dollar bill? We lie skillfully and say that a dollar bill is just ordinary paper. How thick is paper? Next to the computer used to compose this textbook is a laser printer; next to the printer is a ream of laser printer paper. The ream (500 sheets) is roughly 5 cm thick, so a sheet of quality paper has thickness 10^{-2} cm. Now we have the pieces to compute the volume of the bill:

$$v \sim 6\,\text{cm} \times 15\,\text{cm} \times 10^{-2}\,\text{cm} \sim 1\,\text{cm}^3. \tag{1.4}$$

The original point of computing the volume of the armored car and the volume of the bill was to find how many bills fit into the car: $N \sim V/v \sim 10^7\,\text{cm}^3/1\,\text{cm}^3 = 10^7$. If the money is in \$20 bills, then the car would contain \$200 million.

2. *"I seen my opportunities and I took 'em."—George Washington Plunkitt, of Tammany Hall, quoted by Riordan [53, page 3].*

The bills could also be \$1 or \$1000 bills, or any of the intermediate sizes. We chose the intermediate size \$20, because it lies nearly halfway between \$1 and \$1000. You naturally object that \$500, not \$20, lies halfway between \$1 and \$1000. We answer that objection shortly. First, we pause to discuss a general method of estimating: **talking to your gut**. You often have to estimate quantities about which you have only meager knowledge. You can then draw from your vast store of implicit knowledge about the world—knowledge that you possess but cannot easily write down. You extract this knowledge by conversing with your gut; you ask that internal sensor concrete questions, and listen to the feelings that it returns. You already carry on such conversations for other aspects of life. In your native language, you have an implicit knowledge of the grammar; an incorrect sentence sounds funny to you, even if you do not know the rule being broken. Here, we have to estimate the denomination of bill carried by the armored car (assuming that it carries mostly one denomination). We ask ourselves, "How does an armored car filled with one-dollar bills sound?" Our gut, which knows the grammar of the world, responds, "It sounds a bit ridiculous. One-dollar bills are not worth so much effort; plus, every automated teller machine dispenses \$20 bills, so a \$20 bill is a more likely denomination." We then ask ourselves, "How about a truck filled with thousand-dollar bills?" and our gut responds, "no, sounds way too big—never even seen a thousand-dollar bill, probably collectors' items, not for general circulation." After this edifying dialogue, we decide to guess a value intermediate between \$1 and \$1000.

We interpret "between" using a logarithmic scale, so we choose a value near the geometric mean, $\sqrt{1 \times 1000} \sim 30$. Interpolating on a logarithmic scale is more appropriate and accurate than is interpolating on a linear scale, because we are going to use the number in a chain of multiplications and divisions. Let's check whether 30 is reasonable, by asking our gut about nearby estimates. It is noncommittal when asked about \$10 or \$100 bills; both sound reasonable. So our estimate of 30 is probably reasonable. Because there are no \$30 bills, we use a nearby actual denomination, \$20.

Assuming \$20 bills, we estimate that the car contains \$200 million, an amount much greater than our initial guess of \$1 million. Such a large discrepancy makes us suspicious of either the guess or this new estimate. We therefore **cross-check** our answer, by estimating the monetary value in another way. By finding another method of solution, we learn more about the domain. If our new estimate agrees with the previous one, then we gain confidence that the first estimate was correct; if the new estimate does not agree, it may help us to find the error in the first estimate.

We estimated the carrying capacity using the available space. How else could we estimate it? The armored car, besides having limited space, cannot carry infinite mass. So we estimate the mass of the bills, instead of their volume. What is the mass of a bill? If we knew the density of a bill, we could determine the mass using the volume computed in (1.4). To find the density, we use the guerrilla method. Money is paper. What is paper? It's wood or fabric, except for many complex processing stages whose analysis is beyond the scope of this book. Here, we just used another order-of-magnitude technique, **punt**: When a process, such as papermaking, looks formidable, forget about it, and hope that you'll be okay anyway. Ignorance is bliss. It's more important to get an estimate; you can correct the egregiously inaccurate assumptions later. How dense is wood? Once again, use the guerrilla method: Wood barely floats, so its density is roughly that of water, $\rho \sim 1\,\mathrm{g\,cm^{-3}}$. A bill, which has volume $v \sim 1\,\mathrm{cm}^3$, has mass $m \sim 1\,\mathrm{g}$. And $10^7\,\mathrm{cm}^3$ of bills would have a mass of $10^7\,\mathrm{g} = 10$ tons.[3]

This cargo is large. [Metric tons are 10^6 g; English tons (may that measure soon perish) are roughly $0.9 \cdot 10^6$ g, which, for our purposes, is also 10^6 g.] What makes 10 tons large? Not the number 10 being large. To see why not, consider these extreme arguments:

- In megatons, the cargo is 10^{-5} megatons, which is a tiny cargo because 10^{-5} is a tiny number.
- In grams, the cargo is 10^7 g, which is a gigantic cargo because 10^7 is a gigantic number.

You might object that these arguments are cheats, because neither grams nor megatons is a reasonable unit in which to measure truck cargo, whereas tons is a reasonable unit. This objection is correct; when you specify a reasonable unit, you implicitly choose a standard of comparison. The moral is this: A quantity with units—such as tons—cannot be large intrinsically. It must be large compared to a quantity with the same units. This argument foreshadows the topic of dimensional analysis, which is the subject of Chapter 3.

So we must compare 10 tons to another mass. We could compare it to the mass of a bacterium, and we would learn that 10 tons is relatively large; but to learn about the cargo capacity of Brinks armored cars, we should compare 10 tons to a mass related to transport. We therefore compare it to the mass limits at railroad crossings and on many bridges, which are typically 2 or 3 tons. Compared to this mass, 10 tons is large. Such an armored car could not drive many places. Perhaps 1 ton of cargo is a more reasonable estimate for the mass, corresponding to 10^6 bills. We can cross-check this cargo estimate using the size of the armored car's engine (which presumably is related to the cargo mass); the engine is roughly the same size as the engine of a medium-sized pickup truck, which can carry 1 or 2 tons of cargo (roughly 20 or 30 book boxes—see Example 4.1). If the

3. *It is unfortunate that* mass *is not a transitive verb in the way that* weigh *is. Otherwise, we could write that the truck masses* 10 tons*. If you have more courage than we have, use this construction anyway, and start a useful trend.*

money is in $20 bills, then the car contains $20 million. Our original, pure-guess estimate of $1 million is still much smaller than this estimate by roughly an order of magnitude, but we have more confidence in this new estimate, which lies roughly halfway between $1 million and $200 million (we find the midpoint on a logarithmic scale). The Reuters newswire of 18 September 1997 has a report on the largest armored car heist in US history; the thieves took $18 million; so our estimate is accurate for a well-stocked car. (Typical heists net between $1 million and $3 million.)

We answered this first question in detail to illustrate a number of order-of-magnitude techniques. We saw the value of lying skillfully—approximating dollar-bill paper as ordinary paper, and ordinary paper as wood. We saw the value of waging guerrilla warfare—using knowledge that wood barely floats to estimate the density of wood. We saw the value of cross-checking—estimating the mass and volume of the cargo—to make sure that we have not committed a gross blunder. And we saw the value of divide and conquer—breaking volume estimations into products of length, width, and thickness. Breaking problems into factors, besides making the estimation possible, has another advantage: It often reduces the error in the estimate. There probably is a general rule about guessing, that the *logarithm* is in error by a reasonably fixed fraction. If we guess a number of the order of 1 billion in one step, we might be in error by, say, a factor of 10. If we factor the 1 billion into four pieces, the estimate of each piece will be in error by a factor of $\gamma = 10^{1/4}$. We then can hope that the errors are uncorrelated, so that they combine as steps in a random walk. Then, the error in the product is $\gamma^{\sqrt{4}} = 10^{1/2}$, which is smaller than the one-shot error of 10. So breaking an estimate into pieces reduces the error, according to this order-of-magnitude analysis of error.

1.2 Cost of lighting Pasadena, California

What is the annual cost of lighting the streets of Pasadena, California?

Astronomers would like this cost to be huge, so that they could argue that street lights should be turned off at night, the better to gaze at heavenly bodies. As in Section 1.1, we guess a cost right away, to train our intuition. So let's guess that lighting costs $1 million annually. This number is unreliable; by talking to our gut, we find that $100,000 sounds okay too, as does $10 million (although $100 million

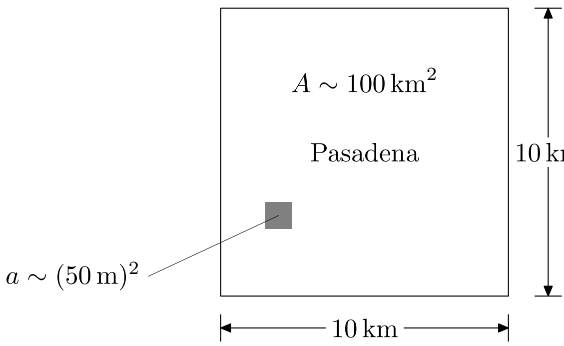

Figure 1.2. *Map of Pasadena, California drawn to order of magnitude. The small shaded box is the area governed by one lamp; the box is not drawn to scale, because if it were, it would be only a few pixels wide. How many such boxes can we fit into the big square? It takes* 10 min *to leave Pasadena by car, so Pasadena has area* $A \sim (10\,\text{km})^2 = 10^8\,\text{m}^2$*. While driving, we pass a lamp every* 3 sec*, so we estimate that there's a lamp every* 50 m*; each lamp covers an area* $a \sim (50\,\text{m})^2$*.*

sounds too high).

The cost is a large number, out of the ordinary range of costs, so it is difficult to estimate in one step (we just tried to guess it, and we're not sure within a factor of 10 what value is correct). So we divide and conquer. First, we estimate the number of lamps; then, we estimate how much it costs to light each lamp.

To estimate the number of lamps (another large, hard-to-guess number), we again divide and conquer: We estimate the area of Pasadena, and divide it by the area that each lamp governs, as shown in Figure 1.2. There is one more factor to consider: the fraction of the land that is lighted (we call this fraction f). In the desert, f is perhaps 0.01; in a typical city, such as Pasadena, f is closer to 1.0. We first assume that $f = 1.0$, to get an initial estimate; then we estimate f, and correct the cost accordingly.

We now estimate the area of Pasadena. What is its shape? We could look at a map, but, as lazy armchair theorists, we lie; we assume that Pasadena is a square. It takes, say, 10 minutes to leave Pasadena by car, perhaps traveling at 1 km/min; Pasadena is roughly 10 km in length. Therefore, Pasadena has area $A \sim 10\,\text{km} \times 10\,\text{km} = 100\,\text{km}^2 = 10^8\,\text{m}^2$. (The true area is $23\,\text{mi}^2$, or $60\,\text{km}^2$.) How much area does each lamp govern? In a car—say, at 1 km/min or $\sim 20\,\text{m}\,\text{s}^{-1}$—it takes 2 or 3 sec to go from lamppost to lamppost, corresponding to a spacing of $\sim 50\,\text{m}$. Therefore, $a \sim (50\,\text{m})^2 \sim 2.5 \cdot 10^3\,\text{m}^2$, and the number of lights is $N \sim A/a \sim 10^8\,\text{m}^2/2.5 \cdot 10^3\,\text{m}^2 \sim 4 \cdot 10^4$.

How much does each lamp cost to operate? We estimate the cost by estimating the energy that they consume in a year and the price per unit of energy (divide and conquer). Energy is power × time. We can estimate power reasonably accurately, because we are familiar with lamps around the home. To estimate a quantity, try to compare it to a related, familiar one. Street lamps shine brighter than a household 100 W bulb, but they are probably more efficient as well, so we guess that each lamp draws $p \sim 300\,\text{W}$. All N lamps consume $P \sim Np \sim 4 \cdot 10^4 \times 300\,\text{W} \sim 1.2 \cdot 10^4\,\text{kW}$. Let's say that the lights are on at night—8 hours per day—or 3000 hours/year. Then, they consume $4 \cdot 10^7$ kW–hour. An electric bill will tell you that electricity costs \$0.08

per kW–hour (if you live in Pasadena), so the annual cost for all the lamps is \$3 million.

Now let's improve this result by estimating the fraction f. What features of Pasadena determine the value of f? To answer this question, consider two extreme cases: the desert and New York city. In the desert, f is small, because the streets are widely separated, and many streets have no lights. In New York city, f is high, because the streets are densely packed, and most streets are filled with street lights. So the relevant factors are the spacing between streets (which we call d), and the fraction of streets that are lighted (which we call f_l). As all pedestrians in New York city know, 10 north–south blocks or 20 east–west blocks make 1 mile (or 1600 m); so $d \sim 100\,\mathrm{m}$. In street layout, Pasadena is closer to New York city than to the desert. So we use $d \sim 100\,\mathrm{m}$ for Pasadena as well. If every street were lighted, what fraction of Pasadena would be lighted? Figure 1.3 shows the computation; the result is $f \sim 0.75$In New York city, $f_{\mathrm{L}} \sim 1$; in Pasadena, $f_{\mathrm{L}} \sim 0.3$ is more appropriate. So $f \sim 0.75 \times 0.3 \sim 0.25$. Our estimate for the annual cost is then \$1 million. Our initial guess is unexpectedly accurate.

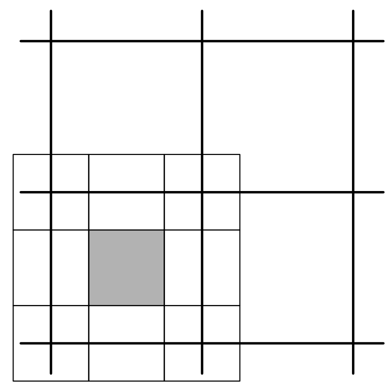

Figure 1.3. Fraction of Pasadena that is lighted. The streets (thick lines) are spaced $d \sim 100\,\mathrm{m}$ apart. Each lamp, spaced $50\,\mathrm{m}$ apart, lights a $50\,\mathrm{m} \times 50\,\mathrm{m}$ area (the eight small, unshaded squares). The only area not lighted is in the center of the block (shaded square); it is one-fourth of the area of the block. So, if every street has lights, $f = 0.75$.

As you practice such estimations, you will be able to write them down compactly, converting units stepwise until you get to your goal (here, \$/year). The cost is

$$\begin{aligned}\text{cost} \sim\ & \underbrace{100\,\mathrm{km}^2}_{A} \times \frac{10^6\,\mathrm{m}^2}{1\,\mathrm{km}^2} \times \underbrace{\frac{1\,\mathrm{lamp}}{2.5\cdot 10^3\,\mathrm{m}^2}}_{a} \times \underbrace{\frac{8\,\mathrm{hrs}}{1\,\mathrm{day}}}_{night} \\ & \times \frac{365\,\mathrm{days}}{1\,\mathrm{year}} \times \underbrace{\frac{\$0.08}{1\,\mathrm{kW\text{–}hour}}}_{price} \times 0.3\,\mathrm{kW} \times 0.25 \\ & \sim \$1\ \text{million}.\end{aligned} \tag{1.5}$$

It is instructive to do the arithmetic without using a calculator. Just as driving to the neighbors' house atrophies your muscles, using calculators for simple arithmetic dulls your mind. You do not develop an innate sense of how large quantities should be, or of when you have made a mistake; you learn only how to punch keys. If you need an answer with 6-digit precision, use a calculator; that's the task for which they are suited. In order-of-magnitude estimates, 1- or 2-digit precision is sufficient; you can easily perform these low-precision calculations mentally.

Will Pasadena astronomers rejoice because this cost is large? A cost has units (here, dollars), so we must compare it to another, relevant cost. In this case, that cost is the budget of Pasadena. If lighting is a significant fraction of the budget, then can we say that the lighting cost is large.

1.3 Pasadena's budget

What fraction of Pasadena's budget is alloted to street lighting?

We just estimated the cost of lighting; now we need to estimate Pasadena's budget. First, however, we make the initial guess. It would be ridiculous if such a trivial service as street lighting consumed as much as 10 percent of the city's budget. The city still has road construction, police, city hall, and schools to support. 1 percent is a more reasonable guess. The budget should be roughly \$100 million.

Now that we've guessed the budget, how can we estimate it? The budget is the amount spent. This money must come from somewhere (or, at least, most of it must): Even the US government is moderately subject to the rule that income $\approx$ spending. So we can estimate spending by estimating income. Most US cities and towns bring in income from property taxes. We estimate the city's income by estimating the property tax per person, and multiplying the tax by the city's population.

Each person pays property taxes either directly (if she owns land) or indirectly (if she rents from someone who does own land). A typical monthly rent per person (for a two-person apartment) is \$500 in Pasadena (the apartments-for-rent section of a local newspaper will tell you the rent in your area), or \$6000 per year. (Places with fine weather and less smog, such as the San Francisco area, have higher monthly rents, roughly \$1500 per person.) According to occasional articles that appear in newspapers when rent skyrockets and interest in the subject increases, roughly 20 percent of rent goes toward landlords' property taxes. We therefore estimate that \$1000 is the annual property tax per person.

Pasadena has roughly $2 \cdot 10^5$ people, as stated on the road signs that grace the entries to Pasadena. So the annual tax collected is \$200 million. If we add federal subsidies to the budget, the total budget is probably double that, or \$400 million. A rule of thumb in these financial calculations is to double any estimate that you make, to correct for costs or revenues that you forgot to include. This rule of thumb is not infallible. We can check its validity in this case by estimating the federal contribution. The federal budget is roughly \$2 trillion, or \$6000 for every person in the United States (any recent almanac tells us the federal budget and the US population). One-half of the \$6000 funds defense spending and interest on the national debt; it would be surprising if fully one-half of the remaining \$3000 went to the cities. Perhaps \$1000 per person goes to cities, which is roughly the amount that the city collects from property taxes. Our doubling rule is accurate in this case.

For practice, we cross-check the local-tax estimate of \$200 million, by estimating the total land value in Pasadena, and guessing the tax rate. The area of Pasadena is $100\,\mathrm{km}^2 \sim 36\,\mathrm{mi}^2$, and $1\,\mathrm{mi}^2 = 640\,\mathrm{acres}$. You can look up this acre–square-mile conversion, or remember that, under the Homestead Act, the US government handed out land in 160-acre parcels—known as *quarter lots* because they were $0.25\,\text{mi}^2$. Land prices are exorbitant in southern California (sun, sand, surf, and mountains, all within a few hours drive); the cost is roughly \$1 million per acre (as you can determine by looking at the homes-for-sale section of the newspaper). We guess that property tax is 1 percent of property value. You can determine a more accurate value by asking anyone who owns a home, or by asking City Hall. The total tax is

$$\begin{aligned} W &\sim \underbrace{36\,\text{mi}^2}_{\textit{area}} \times \frac{640\,\text{acres}}{1\,\text{mi}^2} \times \underbrace{\frac{\$1\text{ million}}{1\,\text{acre}}}_{\textit{land price}} \times \underbrace{0.01}_{\textit{tax}} \\ &\sim \$200\text{ million.} \end{aligned} \tag{1.6}$$

This revenue is identical to our previous estimate of local revenue; the equality increases our confidence in the estimates. As a check on our estimate, we looked up the budget of Pasadena. In 1990, it was \$350 million; this value is much closer to our estimate of \$400 million than we have a right to expect!

The cost of lighting, calculated in Section 1.2, consumes only 0.2 percent of the city's budget. Astronomers should not wait for Pasadena to turn out the lights.

1.4 Diaper production

How many disposable diapers are manufactured in the United States every year?

We begin with a guess. The number must be in the millions—say, 10 million—because of the huge outcry when environmentalists suggested banning disposable diapers to conserve landfill space and to reduce disposed plastic. To estimate such a large number, we divide and conquer. We estimate the number of diaper users—babies, assuming that all babies use diapers, and that no one else does—and the number of diapers that each baby uses in 1 year. These assumptions are not particularly accurate, but they provide a start for our estimation. How many babies are there? We hereby define a baby as a child under 2 years of age. What fraction of the population are babies? To estimate this fraction, we begin by assuming that everyone lives exactly 70 years—roughly the life expectancy in the United States—and then abruptly dies. (The life expectancy is more like 77 years, but an error of 10 percent is not significant given the inaccuracies in the remaining estimates.)

How could we have figured out the average age, if we did not already know it? In the United States, government retirement (Social Security) benefits begin at age 65 years, the canonical retirement age. If the life expectancy were less than 65 years—say, 55 years—then so many people would complain about being short-changed by Social

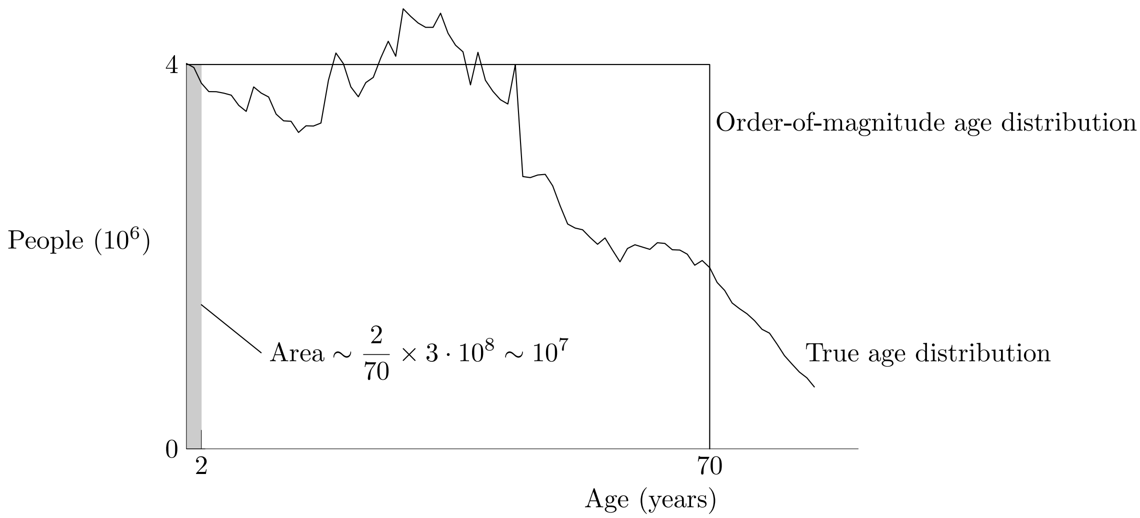

Figure 1.4. *Number of people versus age (in the United States). The true age distribution is irregular and messy; without looking it up, we cannot know the area between ages 0.0 years and 2.0 years (to estimate the number of babies). The rectangular graph—which has the same area and similar width—immediately makes clear what the fraction under 2 years is: It is roughly* $2/70 \sim 0.03$. *The population of the United States is roughly* $3 \cdot 10^8$, *so the number of babies is* $\sim 0.03 \times 3 \cdot 10^8 \sim 10^7$.

Security that the system would probably be changed. If the life expectancy were much longer than 65 years—say, if it were 90 years—then Social Security would cost much more: It would have to pay retirement benefits for $90 - 65 = 25$ years instead of for $75 - 65 = 10$ years, a factor of 2.5 increase. It would have gone bankrupt long ago. So, the life expectancy must be around 70 or 80 years; if it becomes significantly longer, expect to see the retirement age increased accordingly. For definiteness, we choose one value: 70 years. Even if 80 years is a more accurate estimate, we would be making an error of only 15 percent, which is probably smaller than the error that we made in guessing the cutoff age for diaper use. It would hardly improve the accuracy of the final estimate to agonize over this 15 percent.

To compute how people are between the ages of 0 and 2.0 years, consider an analogous problem. In a 4-year university (which graduates everyone in 4 years and accepts no transfer students) with 1000 students, how many students graduate in each year's class? The answer is 250, because $1000/4 = 250$. We can translate this argument into the following mathematics. Let τ be lifetime of a person. We assume that the population is steady: The birth and death rates are equal. Let the rates be $\dot{N}$. Then the total population is $N = \dot{N}\tau$, and the population between ages τ_1 and τ_2 is

$$N\frac{\tau_2 - \tau_1}{\tau} = \dot{N}(\tau_2 - \tau_1). \tag{1.7}$$

So, if everyone lives for 70 years exactly, then the fraction of the population whose age is between 0 and 2 years is $2/70$ or ~ 0.03 (Figure 1.4). There are roughly $3 \cdot 10^8$ people in the United States, so

$$N_{\text{babies}} \sim 3 \cdot 10^8 \times 0.03 \sim 10^7 \text{ babies}. \tag{1.8}$$

We have just seen another example of skillful lying. The jagged curve in Figure 1.4 shows a cartoon version of the actual mortality curve for the United States. We simplified this curve into the boxcar shape (the rectangle), because we know how to deal with rectangles. Instead of integrating the complex, jagged curve, we integrate a simple, civilized curve: a rectangle of the same area and similar width. This procedure is **order-of-magnitude integration**. Similarly, when we studied the Brinks armored-car example (Section 1.1), we pretended that the cargo space was a cube; that procedure was **order-of-magnitude geometry**.

How many diapers does each baby use per year? This number is large—maybe 100, maybe 10,000—so a wild guess is not likely to be accurate. We divide and conquer, dividing 1 year into 365 days. Suppose that each baby uses 8 diapers per day; newborns use many more, and older toddlers use less; our estimate is a reasonable compromise. Then, the annual use per baby is ~ 3000, and all 10^7 babies use $3{\cdot}10^{10}$ diapers. The actual number manufactured is $1.6{\cdot}10^{10}$ per year, so our initial guess is low, and our systematic estimate is high.

This example also illustrates how to deal with **flows**: People move from one age to the next, leaving the flow (dying) at different ages, on average at age 70 years. From that knowledge alone, it is difficult to estimate the number of children under age 2 years; only an actuarial table would give us precise information. Instead, we invent a table that makes the calculation simple: Everyone lives to the life expectancy, and then dies abruptly. The calculation is simple, and the approximation is at least as accurate as the approximation that every child uses diapers for exactly 2 years. In a product, the error is dominated by the most uncertain factor; you waste your time if you make the other factors more accurate than the most uncertain factor.

1.5 Meteorite impacts

How many large meteorites hit the earth each year?

This question is not yet clearly defined: What does *large* mean? When you explore a new field, you often have to estimate such ill-defined quantities. The real world is messy. You have to constrain the question before you can answer it. After you answer it, even with crude approximations, you will understand the domain more clearly, will know which constraints were useful, and will know how to improve them. If your candidate set of assumptions produce a wildly inaccurate estimate—say, one that is off by a factor of 100,000—then you can be sure that your assumptions contain a fundamental flaw. Solving such an inaccurate model exactly is a waste of your time. An order-of-magnitude analysis can prevent this waste, saving you time to create more realistic models. After you are satisfied with your

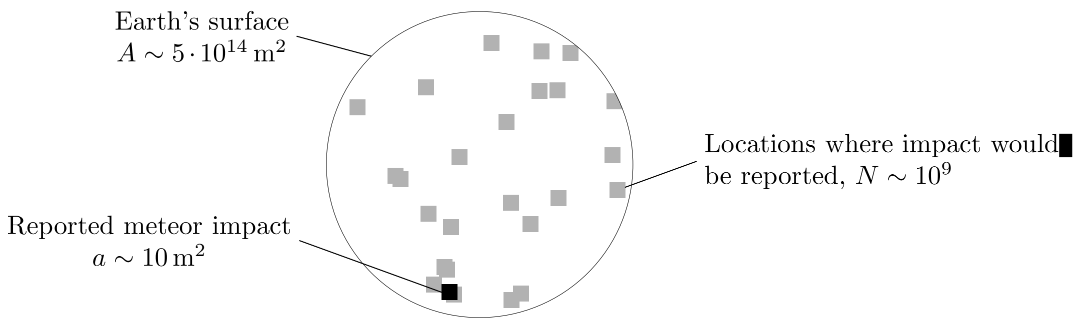

Figure 1.5. *Large-meteorite impacts on the surface of the earth. Over the surface of the earth, represented as a circle, every year one meteorite impact (black square) causes sufficient damage to be reported by* Sky&Telescope*. The gray squares are areas where such a meteorite impact would have been reported—for example, a house or car in an industrial country; they have total area* $Na \sim 10^{10}\,\mathrm{m}^2$*. The gray squares cover only a small fraction of the earth's surface. The expected number of large impacts over the whole earth is* $1 \times A/Na \sim 5{\cdot}10^4$*, where* $A \sim 5 \cdot 10^{14}\,\mathrm{m}^2$ *is the surface area of the earth.*

assumptions, you can invest the effort to refine your model.

Sky&Telescope magazine reports approximately one meteorite impact per year. However, we cannot simply conclude that only one large meteorite falls each year, because *Sky&Telescope* presumably does not report meteorites that land in the ocean or in the middle of corn fields. We must adjust this figure upward, by a factor that accounts for the cross-section (effective area) that *Sky&Telescope* reports cover (Figure 1.5). Most of the reports cite impacts on large, expensive property such as cars or houses, and are from industrial countries, which have $N \sim 10^9$ people. How much target area does each person's car and living space occupy? Her car may occupy $4\,\mathrm{m}^2$, and her living space (portion of a house or apartment) may occupy $10\,\mathrm{m}^2$. [A country dweller living in a ranch house presents a larger target than $10\,\mathrm{m}^2$, perhaps $30\,\mathrm{m}^2$. A city dweller living in an apartment presents a smaller target than $10\,\mathrm{m}^2$, as you can understand from the following argument. Assume that a meteorite that lands in a city crashes through 10 stories. The target area is the area of the building roof, which is one-tenth the total apartment area in the building. In a city, perhaps $50\,\mathrm{m}^2$ is a typical area for a two-person apartment, and $3\,\mathrm{m}^2$ is a typical target area per person. Our estimate of $10\,\mathrm{m}^2$ is a compromise between the rural value of $30\,\mathrm{m}^2$ and the city value of $3\,\mathrm{m}^2$.]

Because each person presents a target area of $a \sim 10\,\mathrm{m}^2$, the total area covered by the reports is $Na \sim 10^{10}\,\mathrm{m}^2$. The surface area of the earth is $A \sim 4\pi \times (6{\cdot}10^6\,\mathrm{m})^2 \sim 5{\cdot}10^{14}\,\mathrm{m}^2$, so the reports of one impact per year cover a fraction $Na/A \sim 2 \cdot 10^{-5}$ of the earth's surface. We multiply our initial estimate of impacts by the reciprocal, A/Na, and estimate $5 \cdot 10^4$ large-meteorite impacts per year. In the solution, we defined large implicitly, by the criteria that *Sky & Telescope* use.

1.6 What you have learned

You now know a basic repertoire of order-of-magnitude techniques:

- *Divide and conquer:* Split a complicated problem into manageable chunks, especially when you must deal with tiny or huge numbers, or when a formula naturally factors into parts (such as $V \sim l \times w \times h$).

- *Guess:* Make a guess before solving a problem. The guess may suggest a method of attack. For example, if the guess results in a tiny or huge number, consider using divide and conquer. The guess may provide a rough estimate; then you can remember the final estimate as a correction to the guess. Furthermore, guessing—and checking and modifying your guess—improves your intuition and guesses for future problems.
- *Talk to your gut:* When you make a guess, ask your gut how it feels. Is it too high? Too low? If the guess is both, then it's probably reliable.
- *Lie skillfully:* Simplify a complicated situation by assuming what you need to know to solve it. For example, when you do not know what shape an object has, assume that it is a sphere or a cube.
- *Cross-check:* Solve a problem in more than one way, to check whether your answers correspond.
- *Use guerrilla warfare:* Dredge up related facts to help you make an estimate.
- *Punt:* If you're worried about a physical effect, do not worry about it in your first attempt at a solution. The productive strategy is to start estimating, to explore the problem, and then to handle the exceptions once you understand the domain.
- *Be an optimist:* This method is related to *punt*. If an assumption allows a solution, make it, and worry about the damage afterward.
- *Lower your standards:* If you cannot solve the entire problem as asked, solve those parts of it that you can, because the subproblem might still be interesting. Solving the subproblem also clarifies what you need to know to solve the original problem.
- *Use symbols:* Even if you do not know a certain value—for example, the energy density stored in muscle—define a symbol for it. It may cancel later. If it does not, and the problem is still too complex, then lower your standards.

We apply these techniques, and introduce a few more, in the chapters to come. With a little knowledge and a repertoire of techniques, you can estimate many quantities.

1.7 Exercises

► **1.1** *Rewriting*
Estimate the radius of the earth. Prove that the earth is (a) huge and (b) tiny, by choosing appropriate units for the radius.

► **1.2** *Batteries*
What is the cost of energy from a 9V battery? From a wall socket (the mains)?

► **1.3** *Human warmth*
How much heat do you generate just sitting around?

► **1.4** *Fuel economy*
What is the fuel consumption, in passenger–miles per gallon, of a 747 jumbo jet?

► **1.5** *Bandwidth*
What is the data rate (bits/s) of a 747 filled with DVD's crossing the Atlantic?

► **1.6** *Pit spacing*
What is the spacing of the pits on a CD-ROM disc? *Extra:* Test your estimate with a simple experiment.