

Order-of-Magnitude Physics

Understanding the World with Dimensional Analysis,
Educated Guesswork, and White Lies

Sanjoy Mahajan

University of Cambridge

Sterl Phinney

California Institute of Technology

Peter Goldreich

Institute for Advanced Study

Copyright © 1995–2006

Send comments to sanjoy@mrao.cam.ac.uk

Draft of 2006-03-20 23:51:19 [rev 22987c8b4860]

Contents

1	Wetting your feet	1
1.1	Armored cars	1
1.2	Cost of lighting Pasadena, California	9
1.3	Pasadena's budget	11
1.4	Diaper production	12
1.5	Meteorite impacts	14
1.6	What you have learned	16
1.7	Exercises	17
2	Some financial math	18
2.1	Rule of 72	18
2.2	Mortgages: A first approximation	19
2.3	Realistic mortgages	21
2.4	Short-term limit	22
2.5	Long-term limit	23
2.6	What you have learned	24
2.7	Exercises	25
	Bibliography	26

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 of the book, **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.

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

$$\text{kinetic energy} \sim Mv^2. \quad (1.1)$$

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

$$V \sim \text{length} \times \text{width} \times \text{height} \quad (1.2)$$

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

$$V \sim 2 \text{ m} \times 2 \text{ m} \times 2 \text{ m} \sim 10 \text{ m}^3 = 10^7 \text{ cm}^3. \quad (1.3)$$

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 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.¹ 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 an inkjet printer; next to the printer is a ream of 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:

$$v \sim 6 \text{ cm} \times 15 \text{ cm} \times 10^{-2} \text{ cm} \sim 1 \text{ cm}^3. \quad (1.4)$$

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.

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




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 $V \sim 2 \text{ m} \times 2 \text{ m} \times 2 \text{ m} \sim 10 \text{ m}^3$.

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

\$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. Here’s why. Suppose your estimate for a quantity Q turns into this multiplication:

$$Q = 2 \times 5 \times y, \quad (1.5)$$

and you still need to estimate y . Let’s say that you think that y is roughly 12 – and in fact y is 12 – but you are unsure by a factor of 3: a value of 4 or of 36 also seem plausible. Then your upper and lower estimates for Q are:

$$\begin{aligned} Q_1 &= 2 \times 5 \times 4 \\ Q_2 &= 2 \times 5 \times 36. \end{aligned} \quad (1.6)$$

If you estimate Q by using the arithmetic mean of the lower and upper estimates Q_1 and Q_2 , then your estimate is

$$\frac{Q_1 + Q_2}{2} = 2 \times 5 \times \frac{4 + 36}{2} = 200, \quad (1.7)$$

which is much greater than the true value $2 \times 5 \times 12 = 120$. The estimate using the geometric mean is

$$\sqrt{Q_1 Q_2} = 2 \times 5 \times \sqrt{4 \times 36} = 2 \times 5 \times 12, \quad (1.8)$$

which is the true value.

The uncertainty in y is given as a factor (here, a factor of 3) rather than as a difference (for example, 12 ± 8). So this example intrinsically favors the geometric mean, and its conclusion – that one should use the geometric mean – is not surprising given this starting point. So how reasonable is the starting point, that one’s uncertainty in y is more accurately characterized by a factor than by a difference? To test it, imagine that you are unsure whether iron is more or less dense than rock and that an estimate, say of the earth’s mass, depends on the ratio of densities $r = \rho_{\text{iron}}/\rho_{\text{rock}}$. Suppose that your knowledge of r is as follows: ‘I don’t know which substance is denser. Perhaps they are comparable ($r = 1$). However, if they are not comparable, then their densities are probably not more than a factor of 2 different.’ In other words, $r = 1/2$ and $r = 2$ also seem plausible. This characterization using factors matches our intuition for densities. In contrast, using differences does not express our intuitions. If r could range as high as 2, then as a difference $r = 1 \pm 1$ and $r = 0$ would be a reasonable value, as reasonable as $r = 2$. However, to say the conclusion is to doubt it. Your intuition rebels because it knows that $r = 0$ is nonsense: Iron is dense, much denser than air or feathers!

To make the example more extreme, imagine a situation with more uncertainty, say where your uncertainty is a factor of 10. In terms of the density ratio, both $r = 10$ and $r = 1/10$ seem plausible. This level of uncertainty might be appropriate when comparing substances with which you have little experience; for example, the crust of a neutron star with the core of a white dwarf. The expression $r = 10 = 1 + 9$ expresses a reasonable intuition for the upper bound. However, the lower bound – if you use differences rather than factors – does not express a reasonable intuition about densities for it would say that $r = 1 - 9 = -8$. Since densities must be positive, this result is nonsense.

This point about positive and negative is essential to this discussion. The quantities that we multiply to make an estimate almost always have a known sign. For the sake of argument, suppose that the sign of a particular quantity is positive. If you are completely uncertain about its value, it could range between 0 and ∞ . This range is not symmetric (where is the midpoint?), so differences do not accurately characterize your uncertainty on this range. However, the *logarithm* of the quantity ranges from $-\infty$ to ∞ , which is a symmetric range. Here, differences can accurately characterize your uncertainty. Differences on a logarithmic scale become factors on a linear scale; logarithms were invented because of this property: They convert multiplication into addition. So the moral of these examples is: On a logarithmic scale, use *differences* to characterize uncertainty; and on a linear scale, use *factors* to characterize uncertainty.

Returning to the example that prompted this discussion of arith-

metic and geometric means, let's estimate the typical denomination in the armored car, by asking our gut about nearby estimates. It is noncommittal when asked about \$10 or \$100 bills; both denominations sound reasonable. It has strong feelings when we ask it about \$1 bills ('who would bother to transport them?') or \$1000 ('never seen one'): both seem unreasonable. The geometric mean of 10 and 100 is roughly 30, so we can imagine an armored car filled with \$30 bills. Because US money does not come in \$30 bills, we instead use a nearby actual denomination of \$20.

If the car is filled with \$20 bills, it would contain \$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, $\rho \sim 1 \text{ g cm}^{-3}$. A bill, which has volume $v \sim 1 \text{ cm}^3$, has mass $m \sim 1 \text{ g}$. And 10^7 cm^3 of bills would have a mass of $10^7 \text{ g} = 10 \text{ tons}$.²

This cargo is large. [Metric tons are 10^6 g ; English tons are roughly $0.9 \cdot 10^6 \text{ g}$, 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

2. 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.

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 many books (and a lot of the rest of this book).

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). If the 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, often reduces the error in the estimate. If you guess a number of the order of 10^{10} in one step, you might be in error by a factor of 10. For example, you might estimate the number of stars in our galaxy as $N \sim 10^{10}$, but 10^9 or 10^{11} might feel equally plausible. Now break the estimate into two pieces:

$$N \sim A \times B, \quad (1.9)$$

where A and B are each of order 10^5 . Now you estimate A and B . What is the typical error in your estimate of N ? There probably is a general rule about guessing, that the *logarithm* of a number is in error by a fixed fraction. If estimates of order 10^{10} are often in error by a factor of 10, which is one unit on a log-base-10 scale, then estimates of order 10^5 would be in error by one-half of a unit or by a factor of 3. In the one-shot estimate for N , its logarithm $\log_{10} N$ could be 9, 10, or 11 – the three values feeling equally plausible:

$$\begin{array}{rcccl} \log_{10} N & 9 & 10 & 11 & \\ p & \frac{1}{3} & \frac{1}{3} & \frac{1}{3} & \end{array} \quad (1.10)$$

For A or B , each of order 10^5 , their logarithms could be 4.5, 5, or 5.5, each value feeling equally plausible. Now see what happens when you multiply A and B or, equivalently, when you add their logarithms:

$$\begin{aligned} \log_{10} N &= \log_{10} A + \log_{10} B \\ &= \left\{ \begin{array}{c} 4.5 \\ 5 \\ 5.5 \end{array} \right\} + \left\{ \begin{array}{c} 4.5 \\ 5 \\ 5.5 \end{array} \right\}, \end{aligned} \quad (1.11)$$

where the curly braces list equally plausible values. Since each of A and B have three possibilities, their sum has nine possibilities including duplicates. Here are the possible sums and their probabilities, assuming that the errors in A and B are uncorrelated:

$$\begin{array}{rcccl} \log_{10} N & 9 & 9.5 & 10 & 10.5 & 11 \\ p & \frac{1}{9} & \frac{2}{9} & \frac{3}{9} & \frac{2}{9} & \frac{1}{9} \end{array} \quad (1.12)$$

Compare this distribution to the one-shot distribution (1.10). This distribution, which results from breaking the estimate into two parts, is more likely to produce a $\log_{10} N$ close to 10. In other words, breaking the estimate into two parts reduces the expected error, because the errors in the parts have a chance to cancel.

Generalizing this example, we break N into k roughly equal parts; the example used $k = 2$. The estimate of each part is then in error by a factor of $\gamma = 10^{1/k}$. If these errors are uncorrelated, their logarithms combine as steps in a random walk. So the error in $\log_{10} N$ arises from a random walk of k steps, each with step size $\log_{10} \gamma = 1/k$. In such a random walk, the expected root-mean-squared distance from the origin is

$$x_{\text{rms}} = (\text{number of steps})^{1/2} \times \text{step size}, \quad (1.13)$$

which is

$$x_{\text{rms}} = k^{1/2} \times \frac{1}{k} = k^{-1/2}. \quad (1.14)$$

This distance from the origin is the typical error in the logarithm of N . By increasing k , you decrease this distance and therefore decrease the error in N . The moral is: Divide and conquer!




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 $a \sim (50 \text{ m})^2$.

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 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 (all land is lighted). 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 mi², or 60 km².) How much area does each lamp govern? In a car—say, at 1 km/min or $\sim 20 \text{ m s}^{-1}$ —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 \times 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 $p \sim 300$ W. All N lamps consume $P \sim Np \sim 4 \cdot 10^4 \times 300$ W $\sim 1.2 \cdot 10^4$ 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_1). As all pedestrians in New York city know, 10 north-south blocks or 20 east-west blocks make 1 mile (or 1600 m); so $d \sim 100$ m. In street layout, Pasadena is closer to New York city than to the desert. So we use $d \sim 100$ 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 $f_L \sim 1$; in Pasadena $f_L \sim 0.3$ 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.

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

$$\begin{aligned} \text{cost} &\sim \underbrace{100 \text{ km}^2}_A \times \frac{10^6 \text{ m}^2}{1 \text{ km}^2} \times \underbrace{\frac{1 \text{ lamp}}{2.5 \cdot 10^3 \text{ m}^2}}_a \times \underbrace{\frac{8 \text{ hrs}}{1 \text{ day}}}_{\text{night}} \\ &\times \frac{365 \text{ days}}{1 \text{ year}} \times \underbrace{\frac{\$0.08}{1 \text{ kW-hour}}}_{\text{price}} \times 0.3 \text{ kW} \times 0.25 \quad (1.15) \\ &\sim \$1 \text{ million.} \end{aligned}$$

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




Figure 1.3. Fraction of Pasadena that is lighted. The streets (thick lines) are spaced $d \sim 100$ m apart. Each lamp, spaced 50 m apart, lights a $50 \text{ m} \times 50 \text{ m}$ 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, $f = 0.75$.

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 allotted 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 – such as county taxes, rental of city property, local sales taxes (value-added taxes), and so on. 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 $100 \text{ km}^2 \sim 36 \text{ mi}^2$, and $1 \text{ mi}^2 = 640 \text{ acres}$. 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 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

$$\begin{aligned}
 W &\sim \underbrace{36 \text{ mi}^2}_{\text{area}} \times \frac{640 \text{ acres}}{1 \text{ mi}^2} \times \underbrace{\frac{\$1 \text{ million}}{1 \text{ acre}}}_{\text{land price}} \times \underbrace{0.01}_{\text{tax}} \\
 &\sim \$200 \text{ million.}
 \end{aligned} \tag{1.16}$$


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 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

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 $\sim 0.03 \times 3 \cdot 10^8 \sim 10^7$.

the population between ages τ_1 and τ_2 is

$$N \frac{\tau_2 - \tau_1}{\tau} = \dot{N}(\tau_2 - \tau_1). \quad (1.17)$$

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

$$N_{\text{babies}} \sim 3 \cdot 10^8 \times 0.03 \sim 10^7 \text{ babies}. \quad (1.18)$$

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 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 m^2 , and her living space (portion of a house or apartment) may occupy 10 m^2 . [A country dweller living in a ranch house presents a larger target than 10 m^2 , perhaps 30 m^2 . A city dweller living in an apartment presents a smaller target than 10 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 m^2 is a typical area for a two-person apartment, and 3 m^2 is a typical target area per person. Our estimate of 10 m^2 is a compromise between the rural value of 30 m^2 and the city value of 3 m^2 .]

Because each person presents a target area of $a \sim 10 \text{ m}^2$, the total area covered by the reports is $Na \sim 10^{10} \text{ m}^2$. The surface area of the earth is $A \sim 4\pi \times (6 \cdot 10^6 \text{ m})^2 \sim 5 \cdot 10^{14} \text{ m}^2$, 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

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 $Na \sim 10^{10} \text{ m}^2$. 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 $A \sim 5 \cdot 10^{14} \text{ m}^2$ is the surface area of the earth.

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. The symbol may cancel later in the solution to the problem. If it does not cancel, and the problem is still too complex, lower your standards. By solving a simpler problem, you begin to understand the area. From that understanding, you may learn enough to solve the original problem.

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 quantities that occur in fluid dynamics, biophysics, and many other areas.

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.

2 Some financial math

This chapter, a portion of a longer draft, discusses growth rates and mortgages.

2.1 Rule of 72

If the world population grows at 1% annually, how long until the population doubles? Answer: 72 years. If inflation is 3% annually, how long before the \$1000 under your mattress is worth only \$250 in today's dollars? Answer: 48 years. Such problems, where a quick back-of-the-envelope answer is doable mentally, are amenable to the **rule of 72**:

If a quantity grows at $q\%$ annually, then it doubles in $72/q$ years.

Before we justify this rule, let's apply it to the two preceding examples. In the world-population example, $q = 1$ so the doubling time is 72 years. In the mattress-money example, the money has fallen in worth from \$1000 to \$250, so prices have quadrupled. With annual inflation being 3%, prices double in $72/3 = 24$ years. Quadrupling requires two doublings, so the fall to \$250 takes $2 \times 24 = 48$ years.

The rule arises as follows. Let r be the fractional growth in one period. In the inflation example, the period is one year and $r = 3\%$ or 0.03. After n periods, the quantity has grown by a factor

$$f = (1 + r)^n. \quad (2.1)$$

Take the natural logarithm of both sides:

$$\ln f = n \ln(1 + r). \quad (2.2)$$

As long as $r \ll 1$, you can approximate the right side: $\ln(1 + r) \approx r$. So

$$\ln f \approx nr. \quad (2.3)$$

Doubling means $f = 2$, so the number of periods is given by

$$n \approx \frac{\ln 2}{r} = \frac{0.69}{r}. \quad (2.4)$$

You can express r as a percentage instead of as a fraction: as $q\%$ where $q = 100r$. Then the number of periods required for the quantity to double is:

$$n \approx \frac{69}{q}. \quad (2.5)$$

The doubling time is n periods:

$$t \approx \frac{69 \dots}{q} \times \text{one period.} \quad (2.6)$$

For the inflation example, the period is one year and $r = 0.03$, so the denominator becomes 3% and $t = 23$ years. However, the first analysis of this example used 24 years as doubling time. The discrepancy is in the 69 versus 72. The more exact rule is the ‘rule of 69’ of (2.6). The problem is that 69 has very few factors: only 3 and 23. On the other hand, 72 has a zillion factors: 2, 3, 4, 6, 8, 12, 18, 24, and 36. For most percentage growth rates, the mental divisions are easier if you use 72 rather than 69.

The requirement for $r \ll 1$ means that the quantity does not grow appreciably in one period. In Weimar Germany, inflation was so severe that shoppers brought baskets filled with (paper) money to pay for their groceries. Others, it is said, stole the baskets rather than the money – because baskets were worth more than the money. Or, and this story is probably apocryphal: A taxi journey cost less than a bus journey, because you paid the bus fare at the beginning of the trip, whereas you paid the taxi fare at the end of the trip, by which time the money had sufficiently inflated to make it cheaper than the bus fare. Except for such special circumstances, the approximation $r \ll 1$ is usually accurate.

2.2 Mortgages: A first approximation

Mortgage calculations provide an example of interpreting and simplifying mathematical expressions. A fixed-rate mortgage is the standard mortgage in the United States. In many countries (the United Kingdom and South Africa, for example), the standard is a variable-rate mortgage. For this exposition, fixed-rate mortgages have a large advantage: They are deterministic and therefore easier to analyze.

Here’s how they work. The bank lends you, say, \$180,000 to buy a house (or, in Manhattan, a closet). You pay it back in **equal** monthly payments for 30 years. As a first guess, you’d pay \$500 every month for a total of \$180,000. The bank, however, will not like this arrangement. It gets all its money back but only after a long time. Instead of lending it to you, it could have invested the money – even in a savings account at another bank – and made more money. In other words, it would rather have money now than later. And so would you, which is why you are asking the bank for money to buy a house. Hence the bank charges you **interest** to account for the increased value of money now compared to later. A typical interest rate might be quoted at 6%. That value has no notion of time, whereas we are looking for a quantity that specifies the rate at which money declines in value. But all is well: The 6% shorthand means 6% per year.

So, you could hold on to the money for 30 years and include interest when you pay it back. For the bank to be happy, they would

want the money that you give them, after discounting its value, to equal the loan amount of \$180,000. You would then make a payment, in 30 years, of

$$P_{\text{final}} = (1 + r\tau)^n P, \quad (2.7)$$

where τ is the compounding period, r is the interest rate, and n is the number of periods in the term of the loan. When $r\tau \ll 1$, you can approximate the factor

$$f = (1 + r\tau)^n \quad (2.8)$$

by taking the logarithm of both sides, the same method that led to (2.3). Then

$$\ln f = n \underbrace{\ln(1 + r\tau)}_{\approx r\tau} \approx r\tau n, \quad (2.9)$$

so

$$(1 + r\tau)^n \approx e^{r\tau n}. \quad (2.10)$$

The repayment (2.7) becomes

$$P_{\text{final}} \approx P e^{r\tau n}. \quad (2.11)$$

For the hypothetical loan above, $n = 360$, $\tau = 1$ month, and $r = 0.005/\text{month}$ (from $0.06/\text{year}$). Equivalently, and this version makes the mental calculations easier, you can consider $\tau = 1$ year and use $r = 0.06/\text{year}$. Either way, $r\tau n = 1.8$. So you make a payment of roughly $e^{1.8}P$ or $6P$. If you spread this payment over 30 years, the payments are:

$$\frac{\text{total payment}}{\text{loan term}} = \frac{(1 + r\tau)^n P}{n\tau} = \$3000 \text{ month}^{-1}. \quad (2.12)$$

The first approach (pay back P after 30 years) cheats the bank, whereas this refiguring cheats you. Your calculation of present value already accounts for the declining value of money, so why should you return the money before 30 years? In other words, if you make more than one payment, you should pay less than \$3000/month. How much less depends on how many payments you make.

So the true monthly payment will be between \$500 and \$3000. To make a rough estimate, compute the geometric mean of these upper and lower bounds; we choose the geometric mean rather than the normal (arithmetic) mean for the reasons discussed on page 4. The geometric mean of 500 and 3000 is

$$\sqrt{500 \times 3000} = 1000 \times \sqrt{1.5} = 1225. \quad (2.13)$$

This estimate for the monthly payment, \$1225, is close to the true value of \$1079.19/month. So the geometric-mean estimate is reasonably accurate, more accurate than we have a right to expect. In Section 2.5 we show you a more principled and even more accurate method.

2.3 Realistic mortgages

The previous estimates, except for the geometric-mean estimate, assumed that you return the money in one lump. Let's consider the opposite, more realistic limit: You make many smaller payments. In other words, the number of periods n is large. Then how much is the payment amount? You get the principal at $t = 0$ and make the first payment at $t = \tau$, the second payment at $t = 2\tau$ and so on. To find the payment amount p , add the present values of each payment and set the total present value equal to the principal. With those values equal, you and the bank are making an equal trade: You get money now, and they get a steady stream of payments with the same present value as the lump sum that they gave you.

The payment at $t = \tau$ has present value

$$\frac{\text{payment}}{\text{discount factor}} = \frac{p}{1 + r\tau}, \quad (2.14)$$

and the k th payment, at $t = k\tau$, has present value

$$\frac{p}{(1 + r\tau)^k}. \quad (2.15)$$

As long as $r\tau \ll 1$, we can approximate $(1 + r\tau)^k$ by $e^{kr\tau}$. So payment k has present value $pe^{-kr\tau}$ and the total present value of the n payments, which is also the principal P , is:

$$P = \sum_1^n pe^{-kr\tau} = p \sum_1^n e^{-kr\tau}. \quad (2.16)$$

Let's approximate this sum. The first approximation, which makes only a small error when n is large, is to pretend that the payments start at $t = 0$ rather than at $t = \tau$, thereby reducing the lower limit from 1 to 0:

$$P = p \sum_0^n e^{-kr\tau}. \quad (2.17)$$

The second approximation, also accurate when n is large, is to change the sum into an integral:

$$\begin{aligned} P &= p \int_0^n e^{-kr\tau} dk \\ &= \frac{p}{r\tau} (1 - e^{-nr\tau}). \end{aligned} \quad (2.18)$$

The factor $nr\tau$ is the loan term, which we will call T . With this notation, the payment is

$$p = \frac{Pr\tau}{1 - e^{-rT}}. \quad (2.19)$$

This formula is at least plausible. The numerator $Pr\tau$ is the interest in the first period before any principal has been returned. The

denominator then corrects this payment because you also return the principal. As you return principal, the portion of p devoted to interest falls, and the portion devoted to principal rises – keeping their sum constant (and equal to p).

The dimensionless parameter here is rT , the factor in the exponent. The existence of a dimensionless parameter suggests the next step, which is to examine the formula in two limits: $rT \rightarrow 0$ and $rT \rightarrow \infty$.

2.4 Short-term limit

We first look at the limit $rT \rightarrow 0$. Roughly speaking, this limit describes a short-term loan or a low-interest loan. Speaking more precisely, the important value is not r or T itself – for they have dimensions and therefore no intrinsic scale – but their dimensionless product rT . In this limit, $e^{-rT} \approx 1$ and

$$p = \frac{Pr\tau}{1 - 1} \rightarrow \infty. \quad (2.20)$$

We've approximated too boldly! The next-best approximation for e^{-rT} is $e^{-rT} \approx 1 - rT$. Then the payment (2.19) is

$$p = \frac{Pr\tau}{rT} = P \frac{\tau}{T} = \frac{P}{n}. \quad (2.21)$$

So you pay back the principal in n equal installments, each of size P/n : You are not paying interest! Those glad tidings are not surprising, since this calculation assumed that $rT \approx 0$, which is the dimensionally correct way of saying ‘negligible interest’. This conclusion, while more informative than the infinite payment (2.20), is not news but it is a useful first approximation.

For the next approximation, take the approximation for e^{-rT} one more step:

$$e^{-rT} \approx 1 - rT + \frac{(rT)^2}{2} + \dots \quad (2.22)$$

Then the payment (2.19) becomes

$$p = \frac{Pr\tau}{rT - (rT)^2/2} = \frac{Pr\tau}{rT} \frac{1}{1 - rT/2}. \quad (2.23)$$

The first factor is the previous approximation P/n from (2.21). The second factor, since $rT \ll 1$, is approximately $1 + rT/2$. So the payment is

$$p = \frac{P}{n} \left(1 + \frac{rT}{2} \right) = \frac{P}{n} + \frac{1}{2} Pr\tau. \quad (2.24)$$

This result has been called the ‘Persian folk method of figuring interest’ [1].

Its main term is P/n , the principal repaid per period on a zero-interest loan. The correction term $Pr\tau/2$ also has an intuitive explanation. The main factor $Pr\tau$ is the interest accumulated during the

first period, before any principal has been returned. So the correction term somehow accounts for a non-zero interest rate. Its factor of one-half arises because the outstanding principal is dwindling from P to 0, so on average you should pay less than $Pr\tau$ in interest each period. The following argument shows how much less. In the first approximation, the interest rate is zero, so the payment is P/n and goes all toward principal. Thus the outstanding principal declines linearly from P to 0. The average of the outstanding principal, over the term of the loan, is the (arithmetic!) mean of 0 and P , or $P/2$.

Having this information enables the next approximation, where interest rate is slightly greater than zero. Now you need to pay interest on the outstanding principal. The average outstanding principal is hardly affected by the tiny interest rate, since most of the payment is still toward principal. So the interest will be charged on $P/2$, the average outstanding principal in the first approximation. In one period, this interest is $Pr\tau/2$, which is the correction term in (2.24). One of the exercises (which one?) contains an example using this formula.

2.5 Long-term limit

Now for the other limit: the term is long or $rT \gg 1$. Then the usual exponential factor e^{-rT} is nearly zero. If we assume that it is exactly zero, perhaps an overeager approximation, then the payment expression (2.19) simplifies to:

$$p = \frac{Pr\tau}{1 - e^{-rT}} \approx Pr\tau. \quad (2.25)$$

The quantity $Pr\tau$ is the interest in one period before the principal P falls. Since the payment exactly covers the interest, with nothing left over to reduce the outstanding principal, the payment remains constant and equal to the interest: The loan is an interest-only loan. Such a loan is also called an annuity. In this case you sold the bank an annuity: It paid you a lump sum, the principal, and receives forever a stream of payments.

Many house mortgages fall toward this limit. Here are typical terms: United States, 30 years; South Africa, 25 years; England, 25 years. Typical rates might be 6% fixed (US in 2005) to 13% variable (South Africa in 2004). For the typical American mortgage,

$$rT = \frac{6\%}{\text{year}} \times 30 \text{ years} = 1.8, \quad (2.26)$$

which, although larger than 1, is not much larger. So we have to improve the approximation $e^{-rT} = 0$ that led to the interest-only repayment schedule (2.25). The next improvement is:

$$p = \frac{Pr\tau}{1 - e^{-rT}} \approx Pr\tau(1 + e^{-rT}), \quad (2.27)$$

where the last step uses the binomial approximation to $(1 - x)^{-1}$ for $x \ll 1$.

This approximation increases the interest-only estimate of $Pr\tau$ by the small fraction e^{-rT} . For the mortgage in (2.26), the parameter rT is 1.8 and $e^{rT} \approx 6$. So, relative to an annuity, the payment increases by e^{-rT} or by one-sixth. For the loan used before, with $P = \$180,000$, the monthly interest is, with $\tau = 1$ month:

$$Pr\tau = \$180,000 \times \frac{6\%}{\text{year}} \times \frac{1 \text{ year}}{12}. \quad (2.28)$$

The years upstairs and downstairs cancel, and the 6 and 12 combine to make a gentle fraction:

$$\begin{aligned} Pr\tau &= \$180,000 \times \frac{6}{12} \times 1\% \\ &= \$180,000 \times \frac{1}{2} \times \frac{1}{100} \\ &= \$900. \end{aligned} \quad (2.29)$$

This example shows that 6% (like 12%) is an easy rate for which to figure monthly interest mentally. So if you have another rate, such as 9%, an easy way to calculate mentally is to pretend that the monthly interest at 6%, then adjust your answer (for 9% you would multiply it by 1.5). The monthly interest (2.29) needs to be adjusted upwards slightly to include repaying the principal. The correction factor in (2.27) is roughly $1 + 1/6$, so the payment becomes

$$p = \$900 \times \left(1 + \frac{1}{6}\right) = \$1050. \quad (2.30)$$

The correct value is \$1079.19, so this estimate is quite accurate, significantly more accurate than the geometric-mean estimate of \$1225.

2.6 What you have learned

You now know a basic repertoire of approximations for growth rates and loans.

- *Rule of 72:* If a quantity grows at $q\%$ annually, then it doubles in $72/q$ years.
- *Short-term loans:* In a short-term loan, most of the payment goes toward principal. If all of it went toward the principal, originally P , then you would pay P/n each period, where n is the number of periods. In the next approximation, you include interest of $Pr\tau/2$ in the payment. The $Pr\tau$ is the interest if you never repay principal, and the factor of one-half accounts for the (nearly) linearly declining principal.
- *Long-term loans:* In a long-term loan, the opposite limit, most of the initial payment goes toward interest. If all of it went toward

interest, then the principal would never decline, and you would pay forever. In that situation, you would have sold the bank an annuity. In the next approximation – for a finite but long loan term – you increase the interest-only payment by the fraction e^{-rT} , where r is the rate and T is the loan term.

2.7 Exercises

► 2.1 *Credit cards*

American credit cards often promise a rate such as $r = 4\%$ for the ‘life of the loan’, but set a repayment term of $T = 5$ yr. Estimate the monthly payment on a loan of \$12000, if possible mentally.

► 2.2 *House mortgage*

Estimate the monthly payments for a 25-year fixed-rate mortgage at 9% interest, where $P = \$120000$.

Bibliography

- [1] Peyman Milanfar. A Persian folk method of figuring interest. *Mathematics Magazine*, 69:376, December 1996.
- [2] William L. Riordan. *Plunkitt of Tammany Hall*. E. P. Dutton, New York, 1963.