5 Taking out the big part

Taking out the big part, the technique of this chapter, is a species of successive approximation. First do the most important part of the analysis: the big part. Then estimate changes relative to this big part. This hygienic approach keeps calculations clean enough to do mentally. Here are a few examples beginning with products, powers, and roots, then moving to exponentials and fierce integrals.

5.1 Multiplication

Suppose you have to estimate 31.5×721. A first estimate comes from rounding 31.5 to 30 and 721 to 700:

$$31.5 \times 721 \approx 30 \times 700 = 21000.$$

This product is the big part whose estimation is the first step. In the second step, estimate the correction. You could estimate the correction directly by expanding the product:

$$31.5 \times 721 = (30 + 1.5) \times (700 + 21).$$

Expanding produces four terms:

$$30 \times 700 + 1.5 \times 700 + 30 \times 21 + 1.5 \times 21.$$

What a mess! Using fractional or relative changes cleans up the calculation. The first step is to estimate the fractional change in each factor: 31.5 is 5% more than 30, and 721 is 3% more than 700. So

$$31.5 \times 721 = \underbrace{30 \times (1 + 0.05)}_{31.5} \times \underbrace{700 \times (1 + 0.03)}_{721}.$$

Reorder the pieces to combine the fractional changes:

$$\underbrace{30 \times 700}_{\text{big part}} \times \underbrace{(1 + 0.05) \times (1 + 0.03)}_{\text{correction factor}}.$$

The big part is already evaluated, so the problem reduces to estimating the correction factor. An algebraic method gives

$$(1 + 0.05) \times (1 + 0.03) = 1 \times 1 + 0.05 \times 1 + 1 \times 0.03 + \underbrace{0.05 \times 0.03}_{\text{tiny}}.$$

Because the last term is the product of two corrections, each small, it is smaller than the other terms. Ignoring it gives

$$(1 + 0.05) \times (1 + 0.03) \approx 1 + 0.05 + 0.03 = 1.08.$$

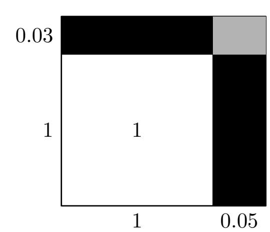

This algebra has an elegant picture. The unit square represents the 1×1 product. Enlarge its width by 0.05 to $1+0.05$, and enlarge its height by 0.03 to $1 + 0.03$. The new rectangle has area $(1 + 0.05) \times (1 + 0.03)$, which is the sought-after product. The four pieces of the figure correspond to the four terms in the expansion of $(1 + 0.05) \times (1 + 0.03)$. Relative to the unit square, the new rectangle has a thin rectangle on the right that has area 0.05 and a thin rectangle on top that has area 0.03. There's also an adjustment of 0.05×0.03 for the gray rectangle. It is tiny compared to the long rectangles, so neglect it. Then the area is roughly $1 + 0.05 + 0.03$, which is a geometric proof that the correction factor is roughly

$$1 + 0.05 + 0.03 = 1.08.$$

It represents an 8% increase. The uncorrected product is 21000, and 8% of it is 1680, so

$$31.5 \times 721 = 21000 \times \text{correction factor} \approx 21000 + 1680 = 22680.$$

The true value is 22711.5, so the estimate is low by 0.15%, which is the area of the tiny, gray rectangle.

This numerical example illustrates a general pattern. Suppose that you can easily find the product xy, as in the preceding example with $x = 30$ and $y = 700$, and you want a nearby product $(x + \Delta x)(y + \Delta y)$, where $\Delta x \ll x$ and $\Delta y \ll y$. Call $\Delta(xy)$ the change in the product xy due to the changes in x and y:

$$(x + \Delta x)(y + \Delta y) = xy + \Delta(xy).$$

To find the new product, you could find $\Delta(xy)$ (since xy is easy). But do not expand the product directly:

$$(x + \Delta x)(y + \Delta y) = xy + x\Delta y + y\Delta x + xy.$$

Instead, extract the big part of the product and study the correction factor. The big part is xy, so extract xy by extracting x from the first factor and y from the second factor. The correction factor that remains is

$$\left(1 + \frac{\Delta x}{x}\right)\left(1 + \frac{\Delta y}{y}\right) = 1 + \underbrace{\frac{\Delta x}{x} + \frac{\Delta y}{y} + \frac{\Delta x}{x}\frac{\Delta y}{y}}_{\text{frac. change in } xy}.$$

The $\Delta x/x$ is the fractional change in x. The $\Delta y/y$ is the fractional change in y. And the $(\Delta x/x)(\Delta y/y)$, the product of two tiny factors, is tiny compared to fractional changes containing one tiny factor. So, for small changes:

$$\begin{aligned}\begin{pmatrix}\text{fractional}\\ \text{change}\\ \text{in } xy\end{pmatrix} &\simeq \frac{\Delta x}{x} + \frac{\Delta y}{y}\\ &= \begin{pmatrix}\text{fractional}\\ \text{change}\\ \text{in } x\end{pmatrix} + \begin{pmatrix}\text{fractional}\\ \text{change}\\ \text{in } y\end{pmatrix}.\end{aligned}$$

In other words, for small changes:

> The fractional change in a product is the sum of fractional changes in its factors.

The simplicity of this rule means that fractional changes simplify computations.

5.2 Squares

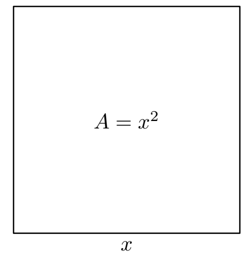

Squares are a particular kind of product, so we could approximate squares using the preceding algebra or pictures. Instead I derive the picture from scratch, to practice with pictures and to introduce the notion of low-entropy expressions. Let A be the area of a square and x be the length of its side, so $A = x^2$. Now imagine increasing x to $x + \Delta x$, producing an area $(x + \Delta x)^2$. This analysis is useful if you can choose x to be a number whose square you know; then Δx is the change to get to the number whose square you want to compute. For example, if you want to compute 9.2^2, set $x = 9$ and $\Delta x = 0.2$ and find how much the area increases. The algebraic approach is to expand

$$(x + \Delta x)^2 = x^2 + 2x\Delta x + (\Delta x)^2.$$

An alternative approach is to elaborate the picture.

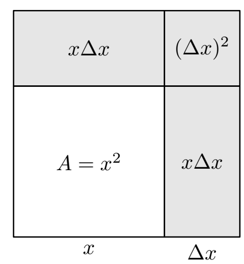

The new area is shaded and has three parts. As long as $\Delta x \ll x$, the tiny corner square is small compared to the two rectangles. So the change in area is

$$\Delta A \approx \underbrace{x\Delta x}_{\text{top rect.}} + \underbrace{x\Delta x}_{\text{right rect.}} = 2x\Delta x.$$

But this result is difficult to remember because it is has *high-entropy* [6]. The combination of x and Δx seem arbitrary. If ΔA had turned out to be x^2 or $(\Delta x)^2$, it would also have seemed reasonable. A high-entropy form has variables scattered all over, in a seemingly unconstrained arrangement. A low-entropy form groups together relevant variables to make a form that is easy to understand and therefore to remember.

To turn $\Delta A = 2x\Delta x$ into low-entropy form, divide by $A = x^2$. This choice has two reasons. The first reason is the theme of this chapter: take out the big part. You know how to square x, so A or x^2 is the big part. To take it out, divide the left side ΔA by A and the right side $2x\Delta x$ by x^2. The second reason comes the method of **Chapter 1**: dimensions. There are many dimensions in the world, so requiring an expression to be dimensionless eliminates this freedom and reduces the entropy:

> Expressions with dimensions have higher entropy than expressions without dimensions.

The high-entropy result has dimensions of area; to make it dimensionless, divide both sides by an area. For the left side ΔA, the natural, realted quantity is the area A. For the right side $2x\Delta x$, the natural, related quantity is the area x^2. So two reasons – taking out the big part and dimensions – suggest dividing by $A = x^2$. A method with two justifications is probably sound, and here is the result:

$$\frac{\Delta A}{A} \approx \frac{2x\Delta x}{x^2} = 2\frac{\Delta x}{x}.$$

Each side has a simple interpretation. The left side, $\Delta A/A$, is the fractional change in area. The right side contains $\Delta x/x$, which is the fractional change in side length. So

$$\begin{pmatrix}\text{fractional}\\ \text{change}\\ \text{in } x^2\end{pmatrix} \approx 2 \times \begin{pmatrix}\text{fractional}\\ \text{change}\\ \text{in } x\end{pmatrix}.$$

This statement of the result is easier to understand than the high-entropy form. It says that fractional changes produce fractional changes. The only seemingly arbitrary datum to remember is the factor of 2, but it too will make sense after studying cubes and square roots.

Meanwhile you might be tempted into guessing that, because $A = x^2$, the fractional changes follow the same pattern:

$$\begin{pmatrix}\text{fractional}\\ \text{change}\\ \text{in } A\end{pmatrix} \approx \begin{pmatrix}\text{fractional}\\ \text{change}\\ \text{in } x\end{pmatrix}^2.$$

That reasonable conjecture is wrong! Try a numerical example. Imagine a 10% increase in x, from 1 to 1.1. Then x^2 increases to roughly 1.2, a fractional increase of 0.2. If the candidate formula above were correct, the fractional increase would be only 0.01.

Let's finish the study of squares with 9.2^2, the numerical example mentioned before. Its big part is $9^2 = 81$. Going from 9 to 9.2 is a fractional increase of 2/90, so 9.2^2 should increase by $2 \times 2/90 = 4/90$:

$$9.2^2 \approx 81 \times \left(1 + \frac{4}{90}\right) \approx 81 + 3.6 = 84.6.$$

The exact answer is 84.64, a mere 0.05% higher.

5.3 Fuel efficiency

Section 2.7 used dimensional analysis and an experiment of dropping paper cones to show that drag force is proportional to v^2, where v is the speed that an object moves through a fluid. This result applied in the limit of high Reynolds number, which is the case for almost all flows in our everyday experience. Highway driving is at a roughly steady speed, so gasoline is burned in fighting drag rather than in lossy, stop-and-go changes of speed. The energy required for a car to travel a distance d at speed v is then

$$E = Fd \propto v^2 d,$$

where F is the drag force. In the 1970's, oil became expensive in Western countries for reasons that were widely misunderstood and often misexplained (maybe intentionally). For a thorough analysis, see [7]. Whatever the causes, the results were hard to avoid. The United States reduced oil consumption by mandating a speed limit of 55 mph on highways. For the sake of this problem, imagine that cars drove at 65 mph before the speed limit was imposed. *By what fraction does the gasoline consumption fall due to the change in speed from 65 to 55 mph?* Pretend that the speed limit does not affect how far people drive. It may be a dubious assumption, since people regulate their commuting by total time rather than distance, but that twist can be the subject of a subsequent analysis (do the big part first).

Fractional changes keep the analysis hygienic. The drag force and the energy consumption are proportional to $v^2 d$, and the distance d is, by assumption, constant. So $E \propto v^2$ and

$$\begin{pmatrix}\text{fractional}\\ \text{change}\\ \text{in } E\end{pmatrix} = 2 \times \begin{pmatrix}\text{fractional}\\ \text{change}\\ \text{in } v\end{pmatrix}.$$

A drop in v from 65 to 55 mph is a drop of roughly 15% so the energy consumption drops by $2 \times 15\% = 30\%$. It is a large reduction in automotive oil consumption. Considering the large fraction of oil consumed by car travel, this 30% drop in highway oil consumption produces a substantial reduction in total oil consumption.

5.4 Third powers

The next example extends the analysis to the volume of a cube with side length x. The usual question recurs: If x increases by Δx, what happens to the volume V? If you do not use fractional changes, you can try to guess what happens by analogy with the change in area. Perhaps

$$\Delta V \sim x^2 \Delta x$$

or maybe

$$\Delta V \sim x(\Delta x)^2?$$

Both choices have a volume on each sides, so their dimensions are correct, and dimensions do not favor either choice. In short, it's a pain to remember how to distribute the three powers of length on the right side. Should the x get all of them, two of them, one of them, or none?

Instead of trying to remember the high-entropy form, work it out from scratch, rewrite it as a fractional change, and see how simple and low-entropy it becomes. The full ΔV is

$$\Delta V = (x + \Delta x)^3 - x^3 = 3x^2 \Delta x + 3x(\Delta x)^2 + (\Delta x)^3.$$

The terms with the higher powers of Δx are the smallest, so ignore them. This approximation leaves

$$\Delta V \approx 3x^2 \Delta x.$$

The fractional change is

$$\frac{\Delta V}{V} \approx \frac{3x^2 \Delta x}{x^3} = 3\,\frac{\Delta x}{x}.$$

This result has the same form as the fractional change in area but with a factor of 3. In words:

$$\begin{pmatrix}\text{fractional}\\ \text{change}\\ \text{in } x^3\end{pmatrix} \approx 3 \times \begin{pmatrix}\text{fractional}\\ \text{change}\\ \text{in } x\end{pmatrix}.$$

The factor of 3 comes from the exponent of x in $V = x^3$, just as the 2 came from the exponent of x in $A = x^2$. Let's look at two examples.

For the first example, estimate 6.3^3. The big part is $6^3 = 216$. Since 0.3 is 5% larger than 6, its cube is $3 \times 5\% = 15\%$ larger than 6^3:

$$6.3^3 \approx 216 \times (1 + 0.15).$$

To calculate 216×0.15, first calculate the big part 200×0.15, which is 30. Then increase the result by 8% of 30, because 216 is 8% larger than 200. Since 8% of 30 is 2.4:

$$216 \times 0.15 = 30 + 2.4 = 32.4$$

Then

$$6.3^3 \approx 216 + 32.4 = 248.4.$$

The true value is 250.047, which is only 0.7% larger.

The second example comes from the physics of wind energy. The power produced by a wind turbine is related to the force exerted by the wind, which is (like the drag force) proportional to v^2. Since power is force times velocity, it should be proportional to v^3. Therefore a 10% increase in wind speed increases generated power by 30%! The hunt for fast winds is one reson that wind turbines are placed high in the atmosphere (for example, on cliffs) or at sea, where winds are faster than near land surfaces.

5.5 Reciprocals

The preceding examples used positive exponents. To explore fractional changes in new territory, try a negative exponent. This example is about the simplest one: reciprocals, where $n = -1$. Suppose that you want to estimate $1/13$ mentally. The big part is $1/10$ because 10 is a nearby factor of 10, which means its reciprocal is easy. So $1/13 \approx 0.1$. To get a more accurate approximation, take out the big part $1/10$ and approximate the correction factor:

$$\frac{1}{13} = \frac{1}{10} \times \frac{1}{1+0.3}.$$

The correction factor is close to 1, reflecting that most of the result is in the big part $1/10$. The correction factor has the form $(1+x)^{-1}$, where $x = 0.3$. It is therefore approximately $1 - x$ as I hope the following example and picture will convince you. If a book is discounted 10% and shipping costs add 10% of the discounted price, the final total is almost exactly the original price. Try an example with a \$20 book. It gets reduced to \$18 but shipping adds \$1.80, for a total of \$19.80. Except for the tiny error of \$0.20, a 10% increase and a 10% decrease cancel each other. In general

$$\underbrace{(1-x)}_{\text{decrease}} \times \underbrace{(1+x)}_{\text{increase}} \approx 1.$$

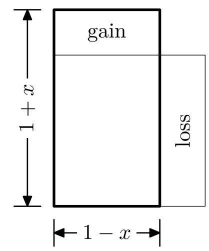

The picture confirms the algebra. Relative to the original unit square, the new $(1-x)\times(1+x)$ rectangle loses a rectangle on the right with area x and gains a rectangle on the top, also with area x. So the gain cancels the loss, keeping the area at 1. The error in this tally is the tiny square with area x^2; however, as long as x^2 is small, do not worry. That pictorial approximation leads to

$$\frac{1}{1+x} \approx 1-x.$$

In words,

$$\begin{pmatrix}\text{fractional}\\ \text{change}\\ \text{in } z^{-1}\end{pmatrix} = -1 \times \begin{pmatrix}\text{fractional}\\ \text{change}\\ \text{in } z\end{pmatrix}.$$

If z increases by 30%, from 1 to 1.3, then z^{-1} decreseas by 30%, from 1 to 0.7. So $1/1.3 = 0.7$ and

$$\frac{1}{13} = \frac{1}{10} \times \frac{1}{1.3} \approx 0.1 \times 0.7 = 0.07.$$

The error in the approximation comes from the neglected x^2 term in the reciprocal $(1+x)^{-1}$. To reduce the error, reduce x by making the big part a close approximation. Massage the original fraction to make the denominator close to 1/100:

$$\frac{1}{13} \times \frac{8}{8} = \frac{8}{104} = \frac{8}{100} \times \frac{1}{1.04}.$$

The big part $8/100 = 0.08$ is still easy, and the correction factor 1.04 has a smaller x: only 0.04. A 4% increase in a denominator produces a 4% decrease in the quantity itself, so

$$\frac{1}{13} \approx 0.08 - 4\%,$$

where the −4% means 'subtract 4% of the previous quantity'. To find the 4%, mentally rewrite 0.08 as 0.0800. Since 4% of 800 is 32, reduce the 0.08 by 0.0032:

$$\frac{1}{13} \approx 0.0800 - 0.0032 = 0.0768.$$

To make an even more accurate value, multiply 1/13 by 77/77 to get 77/1001. The big part is 0.077 and the correction factor is a reduction by 0.1%, which is 0.00077. The result is 0.076923. For comparison, the true value is .0769230769....

The second application follows up the reduction in gasoline consumption due to a 55-mph speed limit, analyzed in **Section 5.3**. How much does the reduction in energy consumption increase fuel efficiency? Fuel efficiency is inversely proportional to energy consumption, so the −30% change in energy consumption produces a +30% change in fuel efficiency. It is often measured in miles per gallon, and a typical value for highway driving may be 35 mph. The 55 mph speed limit would increase it to roughly 45 mph, a larger increase than the legally mandated engineering increases over the last few decades.

5.6 Square roots

After positive and negative integer exponents, the next frontier is fractional exponents. The most common example is square roots, so let's apply these methods to $\sqrt{10}$. First take out the big part from $\sqrt{10}$. The big part is from the number whose square root is easy, which is 9. So factor out $\sqrt{9}$:

$$\sqrt{10} = \sqrt{9} \times \sqrt{1 + \frac{1}{9}}.$$

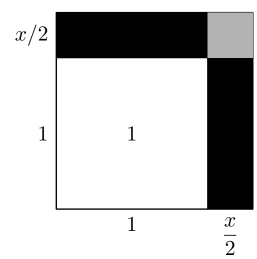

The problem reduces to estimating $\sqrt{1+x}$ with $x = 1/9$ in this case. Reversing the analysis for squaring in **Section 5.2** produces a recipe for square roots. For squaring, the problem was to find the area given the side length. Here the problem is to find the side length $\sqrt{1+x}$ given that the area is $1 + x$. Relative to the unit square, the three shaded areas that make an L contribute the extra area x. The width of the vertical rectangle, or the height of the horizontal rectangle, is the change in side length. To find those dimensions, study the areas. Most of the contribution comes from the two dark rectangles, so ignore the tiny gray square. In that approximation, each rectangle contributes an area $x/2$. The rectangles measure $1 \times \Delta x$ or $\Delta x \times 1$, so their small dimension is roughly $\Delta x = x/2$. Thus the side length of the enclosing square is $1 + x/2$. This result produces the first square-root approximation:

$$\sqrt{1+x} \approx 1 + \frac{x}{2}.$$

The right side represents a fractional increase of $x/2$, so

$$\begin{pmatrix}\text{fractional}\\ \text{change}\\ \text{in } \sqrt{z}\end{pmatrix} \simeq \frac{1}{2} \times \begin{pmatrix}\text{fractional}\\ \text{change}\\ \text{in } z\end{pmatrix},$$

or in words

> A fractional change in z produces one-half the fractional change in $\sqrt{z}$.

This result is the missing piece in estimating $\sqrt{10}$. The missing step was $\sqrt{1+x}$ with $x = 1/9$. Using the approximation,

$$\sqrt{1+\frac{1}{9}} \approx 1 + \frac{1}{18}.$$

So increase the big part (which is 3) by 1/18:

$$\sqrt{10} \approx 3 \times \left(1 + \frac{1}{18}\right) = 3\frac{1}{6} = 3.166\ldots.$$

The true value is $3.1622\ldots$; the estimate is accurate to 0.14%, a reasonable trade for three lines of work.

A few more lines and a refined picture increase the accuracy. The previous analysis ignored the tiny gray square. But now we know enough about the diagram to account for it, or at least to account for most of it. Neglecting the tiny square produced a square of side $1 + x/2$, which has area $1 + x$ plus the area of the tiny square. The tiny square is $x/2$ on each side so its area is $x^2/4$. The error in the first approximation $\sqrt{1+x} = 1 + x/2$ arises from this extra area.

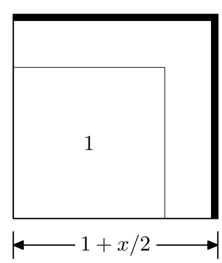

To fix the approximation, shrink the big square slightly, just enough to remove an L-shaped shaded piece with area $x^2/4$. The dimensions of the L cannot be determined exactly – or else we could take square roots exactly – but it is solvable almost exactly using the knowledge from the earlier approximations. The analysis is by successive approximations. The L has two arms, each almost a thin rectangle that is as long or tall as the whole square, which means a length of $1 + x/2$. The 'almost' comes from ignoring the miniscule corner square where the two arms overlap. In this approximation, each arm has area $x^2/8$ in order that the L have area $x^2/4$. Since each sliver has length $1 + x/2$, the widths are

$$\text{width} = \frac{\text{area}}{\text{length}} = \frac{x^2/8}{1 + x/2}.$$

The $1 + x/2$ in the denominator is a fractional increase in the denominator of $x/2$, so it is a fractional decrease of $x/2$ in the numerator:

$$\frac{x^2/8}{1 + x/2} \approx \frac{x^2}{8}\left(1 - \frac{x}{2}\right) = \frac{x^2}{8} - \frac{x^3}{16}.$$

This result is the thin width of the either rectangle arm. So shrink each side of the old square by $x^2/8 - x^3/16$, giving the next approximation to $\sqrt{1+x}$:

$$\sqrt{1+x} = 1 + \frac{x}{2} - \frac{x^2}{8} + \frac{x^3}{16}.$$

The cubic term $x^3/16$ is a bonus. We tried to compute the approximation after $1 + x/2$, which presumably would give the coefficient of the $x^2/8$ term, yet we get the x^3 coefficient for free!

For mental calculation, I often neglect the cubic term. And, consistent with taking out the big part, I represent the $x^2/8$ as an adjustment on the next biggest part, which is the $x/2$ term:

$$\sqrt{1+x} = 1 + \frac{x}{2}\left(1 - \frac{x}{4}\right).$$

This formula gives the next approximation for $\sqrt{10}$. The zeroth approximation is $\sqrt{10} = 3$, which is the big part. The next approximation includes the $x/2$ term to give

$$\sqrt{10} = 3 + \frac{1}{6}.$$

The correction is $1/6$. With $x = 1/9$, the correction needs reducing by $x/4 = 1/36$. Because $1/36$ of $1/6$ is $1/216$, the next approximation is

$$\sqrt{10} = 3 + \frac{1}{6} - \frac{1}{216}.$$

For $1/216$ use fractional changes to approximate it: 216 is 8% larger than 200, so

$$\frac{1}{216} \approx \underbrace{\frac{1}{200}}_{0.0050} - 8\%.$$

The percentage is not hard: $8\% \times 50 = 4$, so

$$\frac{1}{216} \approx \underbrace{0.0050} - 0.0004 = 0.0046.$$

Thus

$$\sqrt{10} \approx 3 + 0.1666 - 0.0046 \approx 3.1626.$$

The true value is $3.162277\ldots$, so the estimate is accurate to 0.01%.

Estimating square roots often benefits from a trick to speed convergence of the series. To see the need for the trick, try to estimate $\sqrt{2}$ using the preceding approximations. The big part is $\sqrt{1}$, which is no help. What remains is the whole problem: $\sqrt{1+x}$ with $x = 1$. Its first approximation is

$$\sqrt{2} \approx 1 + \frac{x}{2} = \frac{3}{2}.$$

Compared to the true value $1.414\ldots$ this approximation is large by 6%. The next approximation includes the $x^2/8$ term:

$$\sqrt{2} \approx 1 + \frac{x}{2} - \frac{x^2}{8} = \frac{11}{8} = 1.375,$$

which is small by roughly 3%. The convergence is slow because $x = 1$, so successive terms do not shrink much despite the growing powers of x. If only I could shrink x! The following trick serves this purpose:

$$\sqrt{2} = \frac{\sqrt{4/3}}{\sqrt{2/3}}.$$

Each square root has the form $\sqrt{1+x}$ where $x = \pm 1/3$. Retain up to the $x/2$ term:

$$\sqrt{2} = \frac{\sqrt{4/3}}{\sqrt{2/3}} \approx \frac{1 + 1/6}{1 - 1/6} = \frac{7}{5} = 1.4.$$

This quick approximation is low by only 1%! With the $x^2/8$ correction for each square root, the approximation becomes $\sqrt{2} \approx 83/59 = 1.406\ldots$, which is low by 0.5%. The extra effort to include the quadratic term is hardly worth only a factor of 2 in accuracy.

5.7 In general

Look at the patterns for fractional changes. Here they are, in the order that we studied them:

$$\begin{pmatrix}\text{fractional}\\ \text{change}\\ \text{in } z^2\end{pmatrix} \simeq 2 \times \begin{pmatrix}\text{fractional}\\ \text{change}\\ \text{in } z\end{pmatrix},$$

$$\begin{pmatrix}\text{fractional}\\ \text{change}\\ \text{in } z^3\end{pmatrix} \simeq 3 \times \begin{pmatrix}\text{fractional}\\ \text{change}\\ \text{in } z\end{pmatrix},$$

$$\begin{pmatrix}\text{fractional}\\ \text{change}\\ \text{in } z^{-1}\end{pmatrix} \simeq -1 \times \begin{pmatrix}\text{fractional}\\ \text{change}\\ \text{in } z\end{pmatrix},$$

$$\begin{pmatrix}\text{fractional}\\ \text{change}\\ \text{in } z^{1/2}\end{pmatrix} \simeq 1/2 \times \begin{pmatrix}\text{fractional}\\ \text{change}\\ \text{in } z\end{pmatrix}.$$

The general pattern is

$$\begin{pmatrix}\text{fractional}\\ \text{change}\\ \text{in } z^n\end{pmatrix} \simeq n \times \begin{pmatrix}\text{fractional}\\ \text{change}\\ \text{in } z\end{pmatrix}.$$

Before trying to prove it, check an easy case that was not part of the data used to make the generalization: $n = 1$. The fractional changes in z and z^1 are identical, so the pattern works. You can also check it when n is a nonnegative integer. In that case, z^n is a product of n factors of z. The product principle from **Section 5.1** is that the fractional change in a product is the sum of fractional changes in its factors. With n identical factors, the sum is indeed n times the fractional change in each factor.

The shortest proof for general n is by logarithmic differentiation. As the name says: First take the logarithm and then differentiate. The logarithm of $f = z^n$ is $n \log z$. Differentiating, or rather taking the differential, gives

$$\frac{df}{f} = n\frac{dz}{z}.$$

That result is exact for infinitesimal changes ($dz = 0$). For finite changes, use Δz instead of dz and turn the equals sign into an $\approx$:

$$\frac{\Delta f}{f} \approx n\frac{\Delta z}{z},$$

which is the symbolic expression of the general pattern:

> The fractional change in z^n is n times the fractional change in z.

5.8 Seasons

An application of these results is to evaluate a common explanation for seasons. It is often said that, because the earth is closer to the sun in the summer than in the winter, summers are warmer than winters. The earth–sun distance does vary throughout the year because the earth orbits in an ellipse rather than a circle. As the distance varies, so does the solar flux, which is the amount of solar energy per unit area hitting the surface. The flux radiates back to space as blackbody radiation, the subject of numerous physics textbooks. The blackbody flux is related to the surface temperature. So the changing the earth–sun distance changes the earth's surface temperature. How large is the effect and is it enough to account for the seasons?

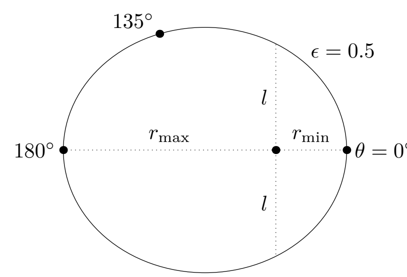

The cleanest analysis is, not surprisingly, via fractional changes starting with the fractional change in earth–sun distance. In polar coordinates, the equation of an ellipse is

$$r = \frac{l}{1 + \epsilon \cos\theta},$$

where ϵ is the eccentricity, θ is the polar angle, and l is the semi-latus rectum (proportional to the angular momentum of the orbit). The diagram shows an orbit with eccentricity of 0.5, much exaggerated compared to the earth's orbit in order to show the elliptical nature of the orbit. The distance varies from $r_{\text{min}} = l/(1+\epsilon)$ to $r_{\text{max}} = l/(1-\epsilon)$. Going from $r_{\text{min}} = l/(1+\epsilon)$ to l is a fractional increase of roughly ϵ. Going from l to $r_{\text{max}} = l/(1-\epsilon)$ is another fractional increase of ϵ, so the earth–sun distance varies by roughly 2ϵ. The earth's orbit has $\epsilon = 0.016$ or 1.6%, meaning that the distance varies by 3.2%. As a check on that number, here is the relevant orbital data:

$$r_{\text{min}} = 1.471 \cdot 10^8\,\text{km},$$
$$r_{\text{max}} = 1.521 \cdot 10^8\,\text{km}.$$

These distances differ by roughly 3.2%.

The second step is to estimate the fractional change in flux produced by this fractional change in distance. The total solar power P spreads over a giant sphere with surface area $A = 4\pi d^2$. The power per area, which is flux, is $P/A \propto d^{-2}$. Because of the -2 exponent, a distance increase of 3.2% produces a flux decrease of 6.4%.

The third step is to estimate the fractional change in temperature produced by this fractional change in incoming flux. The outgoing flux is blackbody radiation, and it equals the incoming flux. So the outgoing flux also changes by 6.4%. Statistical mechanics – the Stefan–Boltzmann law – says that blackbody flux F is proportional to T^4, where T is the surface temperature:

$$F = \sigma T^4.$$

The σ is the Stefan–Boltzmann constant, a ghastly combination of the quantum of action $\hbar$, the speed of light c, Boltzmann's constant k_B, and $\pi^2/60$. But its composition is not relevant, because we are interested only in the fractional change in T. The freedom comes from using fractional changes, and is one of the most important reasons to use them. Since $T \propto F^{1/4}$, if flux changes by 6.4%, then T changes by 6.4%/4 or 1.6%. To find the actual change in temperature, multiply this percentage by the surface temperature T. Do not fall into the trap of thinking that, in winter anyway, the temperature is often 0 °C, so the change ΔT is also 0 °C! The blackbody flux $F \propto T^4$ depends on T being an absolute temperature: measured relative to absolute zero. On one such scale, the Kelvin scale, $T = 300\,\mathrm{K}$ so a 1.6% variation is about 5 K. The reference points of the Celsius and Kelvin scales are different, but their degrees are the same size, so a 5 K difference is also a 5 °C difference. This change is too small to account for the difference between summer and winter, making the proposed explanation for seasons implausible. The explanation has other flaws, such as not explaining how Australia and Europe have opposite seasons despite being almost exactly equidistant from the sun. If orbital distance changes do not produce seasons, what does?

5.9 Exponentials

The preceding examples investigated the approximation

$$(1 + x)^n \simeq 1 + nx$$

where the exponent n was a positive integer, negative integer, and even a fraction. The examples used moderate exponents: 1/2 for the square roots, −1 for reciprocals, and −2 and 1/4 for the seasons. Now push n to an extreme, but skillfully. If you simply make n huge, then you end up evaluating quantities like 1.1^{800}, which is not instructive. Instead, let n grow but shrink

x in parallel to keep nx fixed. An intuitive value for nx is 1, and these examples keep $nx = 1$ while increasing n:

$$1.1^{10} = 2.59374\ldots,$$
$$1.01^{100} = 2.70481\ldots,$$
$$1.001^{1000} = 2.71692\ldots.$$

In each case, $nx = 1$ so the usual approximation is

$$(1 + x)^n = 2 \approx 1 + nx = 2,$$

which is significantly wrong. The problem lies in nx growing too large. In the examples with moderate n, the product nx was much smaller than 1. So new mathematics happens when nx grows beyond that limited range.

To explain what happens, guess features of the solution and then find an explanation related to those features. The sequence starting with 1.1^{10} seems to approach $e = 2.718\ldots$, the base of the natural logarithms. That limit suggests that we study not $(1 + x)^n$ but rather its logarithm:

$$\ln(1 + x)^n = n \ln(1 + x).$$

As long as x itself is not large (nx can still be large), then $\ln(1 + x) \approx 1 + x$. So $n \ln(1 + x) \approx nx$ and

$$(1 + x)^n \approx e^{nx}.$$

When $nx \ll 1$, then e^{nx} approximates to $1+nx$, which reproduces the familiar approximation $1+nx$. When nx grows large, the approximation $e^{nx} = 1+nx$ fails, and you have to use e^{nx} itself.

5.10 Extreme cases

The general n^{th} power $(1 + x)^n$ has several extreme cases depending on n, x, and nx. One limit is taking $n \to 0$. Then $(1 + x)^n$ turns into $1 + n \ln x$, whose proof is left as an exercise for you. The other two limits have been the subject of the preceding analyses. When $x \to 0$, the limit is e^{nx}. If $nx \to 0$ in additon $x \to 0$, then e^{nx} limits to $1 + nx$, which is the result from the first examples in this chapter. Here is a pictoral summary:

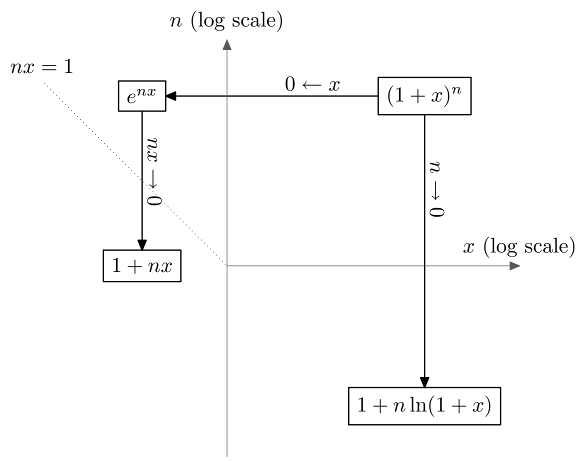

Here are a few numerical examples of these limits:

limit	x	n	$(1+x)^n \approx$
$n \to 0$	1	0.1	$1 + 0.1 \ln 2$
$x \to 0$	0.1	30	e^3
$x, nx \to 0$	0.1	3	1.3

These limits come in handy in the next problem.

5.11 Daunting integral

As a physics undergraduate, I spent many late nights in the department library eating pizza while doing problem sets. The graduate students, in the same boat for their courses, would share their favorite mathematics and physics problems, which included the following from the former USSR. The Landau institute for theoretical physics required an entrance exam of 'mathematical preliminaries'. One preliminary was to evaluate

$$\int_{-\pi/2}^{\pi/2} \cos^{100} t \, dt$$

to within 5% in less than 5 minutes, without a calculator or computer! That $\cos^{100} t$ looks frightening. Normal techniques for trigonometric functions do not help. For example, this identity is useful when integrating $\cos^2 t$:

$$\cos^2 t = \frac{1}{2}(\cos 2t - 1).$$

Here it would produce

$$\cos^{100} t = \left(\frac{\cos 2t - 1}{2}\right)^{50},$$

which becomes a trigonometric monster after expanding the 50^{th} power. The answer is to approximate; after all, we need an answer accurate only to 5%. An approximation for $\cos t$ is $\cos t = 1 - t^2/2$. So

$$\cos^{100} t \simeq \left(1 - \frac{t^2}{2}\right)^{100},$$

which looks like $(1+x)^n$ with $x = -t^2/2$ and $n = 100$. In the range $t \approx 0$ where the approximation for cosine is valid, it is the extreme case $x \to 0$ of $(1+x)^n$, which is e^{nx}. So

$$\cos^{100} t = \left(1 - \frac{t^2}{2}\right)^{100} = e^{-50t^2}.$$

The integrand has the general form $e^{-\alpha t^2}$, which is the Gaussian analyzed in **Section 2.2** and **Section 3.4**. This simple conclusion, that a high power of a cosine becomes a Gaussian, seems hard to believe, but the computer-generated plots of $\cos^n t$ for $n = 1 \ldots 5$ show the cosine curve turning into the Gaussian bell shape as n increases. A plot is not a proof, but it increases confidence in a surprising result.

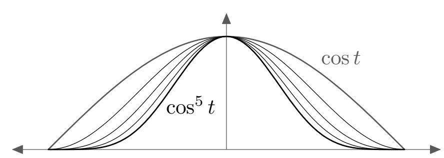

The argument has a few flaws but do not concern yourself with them now. Follow Bob Marley: Don't worry, be happy. In other words, approximate first and (maybe) ask questions later *after* getting an answer. To promote this *sang froid* or courage, I practice what I preach and defer the analysis of the flaws. If the limits were infinite, the integral would be

$$\int_{-\infty}^{\infty} e^{-\alpha t^2}\, dt,$$

which is doable. Alas, our limits are $-\pi/2$ to $\pi/2$ rather than from $-\infty$ to ∞. Do not worry; just extend the limits and justify it at the end. The infinite-range integral of the Gaussian is

$$\int_{-\infty}^{\infty} e^{-\alpha t^2}\,dt = \sqrt{\frac{\pi}{\alpha}}.$$

For $\cos^{100} t$, the parameter is $\alpha = 50$ so the original integral becomes

$$\int_{-\pi/2}^{\pi/2} \cos^{100} t\,dt \approx \int_{-\infty}^{\infty} e^{-50t^2}\,dt = \sqrt{\frac{\pi}{50}}.$$

Since $50 \approx 16\pi$, the integral is $\sqrt{1/16} = 0.25$. The exact answer is

$$\int_{-\pi/2}^{\pi/2} \cos^{n} t\,dt = 2^{-n}\binom{n}{n/2}\pi,$$

whose proof I leave as a fun exercise for you. For $n = 100$, the result is

$$\frac{12611418068195524166851562157\pi}{158456325028528675187087900672} = 0.25003696348037\ldots$$

The `maxima` program, which computed this exact rational-fraction multiple of π, is free software originally written at MIT as the `Macsyma` project. Using a recent laptop (circa 2006) with an Intel 1.83 GHz Core Duo CPU, `maxima` required roughly 20 milliseconds to compute the exact result. Our estimate of 1/4 used a method that requires less than, say, thirty seconds of human time (with practice), and it is accurate to almost 0.01%. Not a bad showing for wetware.

In order to estimate accurately the computation times for such integrals, I tried a higher exponent:

$$\int_{-\pi/2}^{\pi/2} \cos^{10000} t\,dt$$

In 0.26 seconds, `maxima` returned a gigantic rational-fractional multiple of π. Converting it to a floating-point number gave $0.025065\ldots$, which is almost exactly one-tenth of the previous answer. That rescaling makes sense: Increasing the exponent by a factor of 100 increases the denominator in the integral by $\sqrt{100} = 10$.

Now look at the promised flaws in the argument. Here are the steps in slow motion, along with their defects:

1. Approximate $\cos t$ by $1 - t^2/2$. This approximation is valid as long as $t \approx 0$. However, the integral ranges from $t = -\pi/2$ to $t = \pi/2$, taking t beyond the requirement $t \approx 0$.

2008-03-06 13:24:47 / rev ebd336097912+

2. Approximate $(1-t^2/2)^n$ as $e^{-nt^2/2}$. This approximation is valid when $t^2/2 \approx 0$. Again, however, t ranges beyond that limited domain.
3. Replace the difficult limits $-\pi/2 \ldots \pi/2$ with the easier ones $-\infty \ldots \infty$. The infinite limits permit the polar-coordinates trick of **Section 2.2** – which I call a trick because I've never seen a different problem that uses it. However, what justifies extending the limits?

All three flaws have an justification in the high exponent (100 in this case). Raising $\cos t$ to a high power means that the result is close to zero when $\cos t$ drops even slightly below 1. For example, when $t = 0.5$, its cosine is $0.877\ldots$ and $\cos^{100} t \approx 2 \cdot 10^{-6}$. The exponential approximation e^{-50t^2} is roughly $3.7 \cdot 10^{-6}$, which seems inaccurate: The error is almost 100%! But that error is a relative error or fractional error. The absolute error is roughly $2 \cdot 10^{-6}$. it is fine to make large relative errors where the integrand is tiny. In the region where the integrand contributes most of the area, which is $t \approx 0$, steps 1 and 2 of the approximation are valid. In the other regions, who cares?!

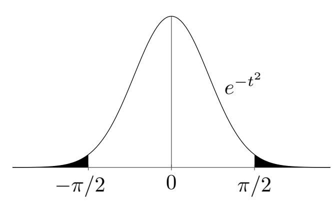

The same argument justifies the third step: extending the limits to infinity. It would be foolhardy to extend the limits in the original integral to give

$$\int_{-\infty}^{\infty} \cos^{100} t \, dt.$$

because each hump of $\cos^{100} t$ contributes equal area and the extended limits enclose an infinity of humps. But this objection disappears if you extend the limits after making the first two approximations. Those approximations give

$$\int_{-\pi/2}^{\pi/2} e^{-50t^2} \, dt.$$

Because the Gaussian e^{-50t^2} is miniscule at and beyond $t = \pm\pi/2$, it is safe to extend the limits to $-\infty \ldots \infty$. The figure shows the tails of $e^{-t^2/2}$, and they are already small. In the faster-decaying function e^{-50t^2}, the tails are so miniscule that they would be invisible at any feasible printing resolution.

I do not want to finish the example with a verification. So try a small additional investigation. It arose because of the high accuracy of the approximation when 100 or 10000 is the exponent of the cosine. I wondered how well the approximation does in the other extreme case, when the exponent is small. To study the accuracy, define

$$f(n) \equiv \int_{-\pi/2}^{\pi/2} \cos^n t \, dt.$$

The preceding approximations produce the approximation

$$f_0(n) = \sqrt{\frac{2\pi}{n}},$$

as you can check by trying the exponents $n = 100$ and $n = 10000$. The fractional error is

$$\frac{f_0(n)}{f(n)} - 1.$$

Here are a few values computed by `maxima`:

n	$f_0(n)/f(n) - 1$
1	0.2533141373155
2	0.1283791670955
3	0.0854018818374
4	0.0638460810704
5	0.0509358530746
6	0.0423520253928
7	0.0362367256182
8	0.0316609527730
9	0.0281092532666
10	0.0252728978367
100	0.0025030858398
1000	0.0002500312109
10000	0.0000250003124

Particularly interesting is the small fractional error when $n = 1$, a case where you can confirm `maxima`'s calculation by hand. The exact integral is

$$f(1) = \int_{-\pi/2}^{\pi/2} \cos^1 t \, dt.$$

So $f(1) = 2$, which compares to the approximation $f_0(1) = \sqrt{2\pi} \approx 2.5$. Even with an exponent as small as $n = 1$, which invalidates each step in the approximation, the error is only 25%. With $n = 2$, the error is only 13% and from there it is, so to speak, all downhill.

5.12 What you have learned

Take out the big part, and use fractional changes to adjust the answer. Using tha procedure keeps calculations hygienic. The fundamental formula is

$$(1+x)^n \simeq 1+nx,$$

or

$$\begin{pmatrix}\text{fractional}\\ \text{change}\\ \text{in } z^n\end{pmatrix} \simeq n \times \begin{pmatrix}\text{fractional}\\ \text{change}\\ \text{in } z\end{pmatrix}.$$

When the exponent n times the fractional change x grows too large (becomes comparable with 1), you need a more accurate approximation:

$$(1+x)^n \simeq e^{nx}.$$