and a gallon of gasoline costs about \$2.50, so a barrel is roughly 40 gallons. The tree with numbers is:

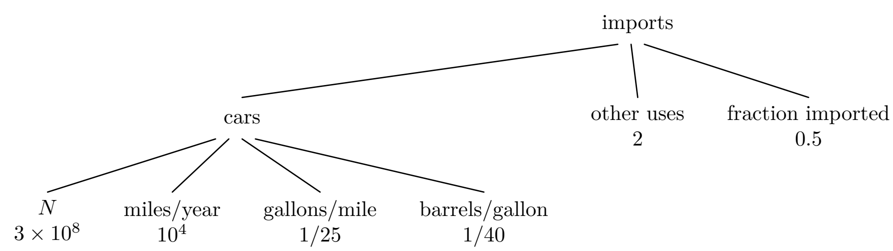

All the leaves have values, so I can propagate upward to the root. The main operation is multiplication. For the 'cars' node:

$$3 \times 10^8 \text{ cars} \times \frac{10^4 \text{ miles}}{1 \text{ car–year}} \times \frac{1 \text{ gallon}}{25 \text{ miles}} \times \frac{1 \text{ barrel}}{40 \text{ gallons}} \sim 3 \times 10^9 \text{ barrels/year}.$$

The two adjustment leaves contribute a factor of $2 \times 0.5 = 1$, so the import estimate is

$$3 \times 10^9 \text{ barrels/year}.$$

For 2006, the true value (from the US Dept of Energy) is 3.7×10^9 barrels/year – only 25 higher than the estimate!

2.5 Theory 3: Estimating accuracy

How does divide-and-conquer reasoning produce such accurate estimates? Alas, this problem is hard to analyze directly because we do not know accuracy in advance. But we can analyze a related problem: how divide-and-conquer reasoning increases our confidence in an estimate or, more precisely, decreases our uncertainty.

The answer is that it works by subdividing a quantity about which we know little into several quantities about which we know more. Even if we need many subdivisions before we reach reliable information, the increased certainty outweighs the small penalty for combining many quantities.

To explain that telegraphic answer, I will analyze a short estimation problem using divide-and-conquer done in slow motion, then apply the lessons to the oil-imports estimate.

The slow-motion problem is to estimate area of a sheet of A4 paper. On first thought, even looking at a sheet I have no clue about its area! On second thought, I know something. For example, the area is certainly more than

1 cm^2 and less than 10^5 cm^2. That wide range makes it hard to be wrong, but it is also too wide to be useful. To narrow the range, I drew a small square with an area of roughly 1 cm^2 and guessed how many squares fit on the sheet: probably at least a few hundred and probably at most a few thousand. Turning 'few' into 3, I offer 300 cm^2 to 3000 cm^2 as a plausible range for the area.

Now compare that range to the range after doing divide and conquer. So, subdivide the area into the width and height: two quantities about which my knowledge is more precise than it is about area itself. The extra precision has a general reason and a reason specific to this problem. The general reason is that we have more experience with lengths than areas: Which is the more familiar quantity, your height or your cross-sectional area? So our length estimates are usually more accurate than our area estimates.

The reason specific to this problem is that A4 paper is the European equivalent of standard American paper, known to computers and laser printers as 'letter' paper and known commonly in the United States as 'eight-and-a-half by eleven' (inches!). In metric units, the dimensions are 21.59 cm × 27.94 cm. If A4 paper were identical to letter paper, I could now compute its exact area. However, A4 paper is, I remember from living in England, slightly thinner and longer than letter paper. I forget the exact differences between the dimensions of A4 and letter paper, hence the remaining uncertainty: I'll guess that the width lies in the range 19 . . . 21 cm and the length lies in the range 28 . . . 32 cm.

The next problem is to combine the plausible ranges for the height and width into the plausible range for the area. A first guess, because the area is the product of the width and height, is to multiply the endpoints of the width and height ranges:

$$A_{min} = 19\,\text{cm} \times 28\,\text{cm} = 532\,\text{cm}^2;$$
$$A_{max} = 21\,\text{cm} \times 32\,\text{cm} = 672\,\text{cm}^2.$$

This method turns out to overextend the range – a mistake that I correct later – but even the overextended range spans only a factor of 1.26 whereas the starting range of 300 . . . 3000 cm^2 spans a factor of 10. Divide and conquer significantly narrowed the range by replacing quantities about which we have little knowledge, such as the area, with quantities about which we have more knowledge.

The second bonus, which I now quantify correctly, is that subdividing into many quantities carries only a small penalty, smaller than suggested by naively multiplying endpoints. The naive method overestimates the range because it assumes the worst. To see how, imagine an extreme case: estimating a quantity that is the product of ten factors, each that you know to within a factor of 2 (in other words, each plausible range is a factor of 4). Is your plausible range for the final quantity a factor of $4^{10} \approx 10^6$?! That conclusion is terribly pessimistic. A more likely case is that a few of the ten estimates will be too large and a few too small, and therefore that several errors will cancel.

To quantify and fix this pessimism, I will explain plausible ranges using probabilities. Probabilities are the tool for this purpose. They reflect incomplete knowledge *not* frequencies in a random experiment; see [16] for a book-length discussion and application of this fundamental point.

To illustrate the probabilistic description, start with the proposition

$$\mathsf{H} \equiv \text{The area of A4 lies in the range } 300\ldots3000\,\text{cm}^2.$$

and information

$$\mathrm{I} \equiv \text{What I know about the area } \textit{before} \text{ using divide and conquer.}$$

Now I want to know the conditional probability $\mathsf{P}(\mathsf{H}|\mathrm{I})$: the probability of H given my knowledge before trying divide and conquer. There is no algorithm known for computing this probability in such a complicated problem situation. How, for example, do we represent my state of knowledge? The best we can do in these cases is to introspect or, in plain English, to talk to our gut.

My gut is the organ with the most access to my knowledge and its incompleteness, and it tells me that I would feel mild surprise but not shock if I learned that the true area lay outside the range $300\ldots3000\,\text{cm}^2$. The surprise suggests that $\mathsf{P}(\mathsf{H}|\mathrm{I})$ is larger than $1/2$. The mildness of the surprise suggests that $\mathsf{P}(\mathsf{H}|\mathrm{I})$ is not much larger than $1/2$. I'll quantify it as $\mathsf{P}(\mathsf{H}|\mathrm{I}) = 2/3$: I would give 2-to-1 odds that the true area is within the plausible range. Throughout this book I'll use a rough 2-to-1 odds range to quantify a plausible range. I could have used a 1-to-1 odds range instead, but the 2-to-1 odds range will help give plausible ranges an intuitive interpretation as a region on a log-normal distribution. That interpretation will help quantify how to combine plausible ranges.

For the moment, I need only the idea that the plausible range contains roughly $2/3$ of the probability. With a further assumption of symmetry, the plausible range $300\ldots3000\,\text{cm}^2$ represents the following probabilities:

$$P(A < 300\,\text{cm}^2) = 1/6;$$
$$P(300\,\text{cm} \leqslant A \leqslant 3000\,\text{cm}^2) = 2/3;$$
$$P(A > 3000\,\text{cm}^2) = 1/6.$$

Here is the corresponding picture with width proportional to probability:

	$p \approx 1/6$	$p \approx 2/3$	$p \approx 1/6$
A	$< 300\,\text{cm}^2$	$300 \ldots 3000\,\text{cm}^2$	$> 3000\,\text{cm}^2$

For the height h and width w, after doing divide and conquer and using the similarity between A4 and letter paper, the plausible ranges are $28 \ldots 32$ cm and $19 \ldots 21$ cm respectively. Here are their probability interpretations:

	$p \approx 1/6$	$p \approx 2/3$	$p \approx 1/6$
h	$< 28\,\text{cm}$	$28 \ldots 32\,\text{cm}$	$> 32\,\text{cm}$

	$p \approx 1/6$	$p \approx 2/3$	$p \approx 1/6$
w	$< 19\,\text{cm}$	$19 \ldots 21\,\text{cm}$	$> 21\,\text{cm}$

Computing the plausible range for the area requires a complete probabilistic description of a plausible range. There is a correct answer to this question – at least if, like me, you are an objective Bayesian – and it depends on the information available to the person giving the range. But no one knows the exact recipe to deduce probabilities from the complex, diffuse, seemingly contradictory information lodged in a human mind.

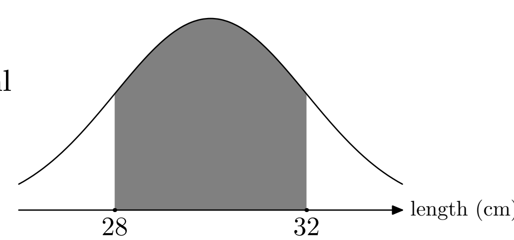

The best that we can do for now is to guess a reasonable and convenient probability distribution. I will use a log-normal distribution meaning that the uncertainty in the quantity's logarithm has a normal (or Gaussian) distribution. As an example, the figure shows the probability distribution for the length of A4 length (after taking into account the similarity to letter paper). The shaded range is the the so-called one-sigma range $\mu - \sigma$ to $\mu + \sigma$. It contains 68% of the probability – a figure conveniently close to 2/3. So to convert a plausible range to a log-normal distribution, use the lower and upper endpoints of the plausible range as $\mu - \sigma$ to $\mu + \sigma$. The peak of the distribution – the most likely value – occurs midway between the endpoints. Since 'midway' is on a logarithmic scale, the midpoint is at $\sqrt{28 \times 32}$ cm or approximately 29.93 cm.

Problem 2.5 Midpoints

The midpoint on the log scale is also known as the geometric mean. Show that it is never greater than the midpoint on the usual scale (which is also known as the arithmetic mean). Can the two midpoints ever be equal?

The log-normal distribution supplies the missing information required to combine plausible ranges. When adding independent quantities, you add their means and their variances. So when multiplying independent quantities, add the means and variances in the logarithmic space. Here is the resulting recipe. Let the plausible range for h be $l_1 \ldots u_1$ and the plausible range for w be $l_2 \ldots u_2$. First compute the geometric mean (midpoint) of each range:

$$\mu_1 = \sqrt{l_1 u_1};$$
$$\mu_2 = \sqrt{l_2 u_2}.$$

The midpoint of the range for $A = hw$ is the product of the two midpoints: $\mu = \mu_1 \mu_2$.

To compute the plausible range, first compute the ratios measuring the width of the ranges:

$$r_1 = u_1/l_1;$$
$$r_2 = u_2/l_2.$$

These ratios measure the width of the ranges. The combined ratio – that is, the ratio of endpoints for the combined plausible range – is

$$r = \exp\left(\sqrt{(\ln r_1)^2 + (\ln r_2)^2}\right).$$

For approximate range calculations, the following contour graphs often provide enough accuracy:

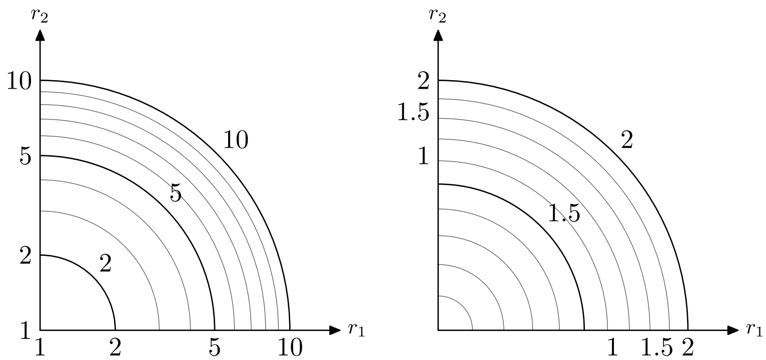

After finding the range, choose the lower and upper endpoints l and u to make $u/l = r$ and $\sqrt{lu} = \mu$. In other words, the plausible range is

$$\frac{\mu}{\sqrt{r}} \ldots \mu\sqrt{r}.$$

▶ *Use a simple example to check that this method produces reasonable results.*

▶ *Apply the method to find the plausible range for the area of an A4 sheet.*

Problem 2.6 Deriving the ratio
Use Bayes theorem to confirm this method for combining plausible ranges.

Let's check this method in a simple example: width and height ranges of $1 \ldots 2\,\mathrm{m}$. What is the plausible range for the area? The naive approach of multiplying endpoints produces a plausible range of $1 \ldots 4\,\mathrm{m}^2$ – a span of a factor of 4. The correct range should be narrower. Indeed, assuming the log-normal distribution, the range spans a factor of

$$\exp\left(\sqrt{2 \times (\ln 2)^2}\right) \approx 2.67.$$

This span and the midpoint determine the range. The area midpoint is the product of the width and height midpoints, each of which is $\sqrt{2}\,\mathrm{m}$. So the midpoint is $2\,\mathrm{m}^2$. The correct endpoints are therefore

$$\frac{2\,\mathrm{m}^2}{\sqrt{2.67}} \ldots 2\,\mathrm{m}^2 \times \sqrt{2.67}$$

or $1.23 \ldots 3.27\,\mathrm{m}^2$. In other words, I assign roughly a $1/6$ probability that the area is less than $1.23\,\mathrm{m}^2$ and roughly a $1/6$ probability that it is greater than $3.27\,\mathrm{m}^2$. Those conclusions seem reasonable when using such uncertain knowledge of length and width.

Having checked that the method is reasonable, it is time to test it in the original illustrative problem: the plausible area range for an A4 sheet. The naive plausible range was $532\ldots672\,\mathrm{cm}^2$, and the correct plausible range will be narrower. Indeed, the log-normal method gives the narrower area range of $550\ldots650\,\mathrm{cm}^2$ with a best guess (most likely value) of $598\,\mathrm{cm}^2$. How did we do? The true area is exactly $2^{-4}\,\mathrm{m}^2$ or $625\,\mathrm{cm}^2$ because – I remembered only after doing this calculation! – An paper is constructed to have one-half the area of A$(n-1)$ paper, with A0 paper having an area of $1\,\mathrm{m}^2$. The true area is only 5% larger than the best guess, suggesting that we used accurate information about the length and width; and it falls within the plausible range but not right at the center, suggesting that the method for computing the plausible range is neither too daring nor too conservative.

Problem 2.7 Volume of a room

Estimate the volume of your favorite room, giving your plausible range before and after using divide and conquer.

The analysis of combining ranges illustrates the two crucial points about divide-and-conquer reasoning. First, the main benefit comes from subdividing vague knowledge (such as the area itself) into pieces about which our knowledge is accurate. Second, this benefit swamps the small accuracy penalty from combining many quantities into one.

To confirm these lessons, examine the benefit of divide-and-conquer reasoning in the example from **Section 2.4**: estimating the annual US oil imports. To quantify the benefit, I compare my plausible ranges before and after using divide and conquer.

Before I use divide and conquer, I have almost no idea what the oil imports are, and I am scared even to guess. To nudge me along, I imagine a mugger demanding, 'Your guess or your life!' In which case I counteroffer with, 'Can I give you a range instead of a number? I'd be surprised if the annual imports are less than10^7 barrels/yr or more than 10^{12} barrels/yr.' The imaginary mugger, being my own creation, always accepts my offer.

Problem 2.8 Your range
What is your plausible range for the annual oil imports?

I need little prodding to narrow my plausible range using divide-and-conquer reasoning. It required making several estimates:

1. N_{people}: US population;
2. f_{car}: cars per person;
3. l: average distance that a car is driven
4. m: average gas mileage;
5. V: volume of a barrel;
6. f_{other}: factor to multiply auto consumption to include all other consumption;
7. $f_{imported}$: fraction of oil that is imported.

Problem 2.9 Your ranges
Give your plausible range for each quantity, i.e. the range for which you assign a two-thirds probability that the true value lies within the range.

Here are my plausible ranges with a few notes of explanation:

1. N_{people}: 290–310 million. I recently read in the newspaper that the US population just reached the milestone of 300 million. How much should I believe what I read in the paper? The media lie when it serves the powerful, but I cannot find any reason to lie about the US population, so I trust the figure, and throw in a bit of uncertainty to reflect the difficulties encountered in counting the population (e.g. what about undocumented immigrants, who are unlikely to fill out census forms?).
2. f_{car}: 0.5–1.5.
3. l: $7 \cdot 10^3$–$20 \cdot 10^3$ mi. Some books assessing used cars consider a low-mileage car to have less than 10^4 mi per year of age. So I guess that the average is somewhat larger than 10^4 mi/yr. But I am not confident of my recollection or the deduction, so my plausible range spans a factor of 3.
4. m: 15–40 miles/gallon;

5. V: 30–60 gallons;
6. f_{other}: 1.5–3;
7. f_{imported}: 0.3–0.8.

▶ *What is the resulting plausible range for the oil imports?*

Now combine the ranges using the method we used for the area of a sheet of A4 paper. That method produces the following plausible range:

$$1.0 \ldots 3.1 \ldots 9.6 \cdot 10^9 \text{ barrels/year}.$$

Compare this range to the range for the off-the-cuff guess $10^7 \ldots 10^{12}$ barrels/yr. That range spanned a factor of 10^5 whereas the improved range spans a mere factor of 10 – thanks to divide-and-conquer reasoning.

2.6 Example 3: Gold or bills?

The chapter's final estimation example is dedicated to readers who forgo careers in the financial industry for less lucrative careers in teaching and research:

▶ *Having broken into a bank vault, should we take the $100 bills or the gold?*

The answer depends partly on the ease and costs of fencing the loot – an analysis beyond the scope of this book. But within our scope is the following question: Which choice lets us carry out the most money? Our carrying capacity is limited by weight and volume. In this analysis, I assume that the lowest limit comes from weight (or mass). The mass subdivides into two subproblems – the value per mass for $100 bills and the value per mass for gold – each of which subdivides into two subproblems:

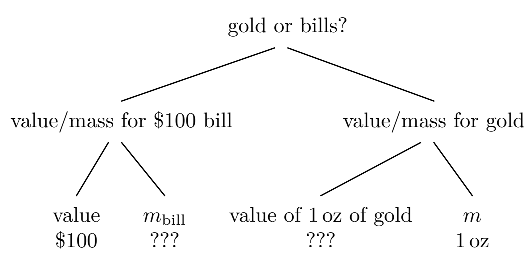

Two leaves have defined values: the value of a $100 bill and the mass of 1 oz (1 ounce) of gold. The two other leaves need divide-and-conquer estimates. In the first round of analysis I make point estimates; in the second round, I account for uncertainty by using the plausible ranges of **Section 2.5**.