Sustainable Energy — without the hot air

David JC MacKay

UIT
CAMBRIDGE, ENGLAND

Now, all people are created equal, but we don't all emit 5½ tons of CO_2 per year. We can break down the emissions of the year 2000, showing how the 34-billion-ton rectangle is shared between the regions of the world:

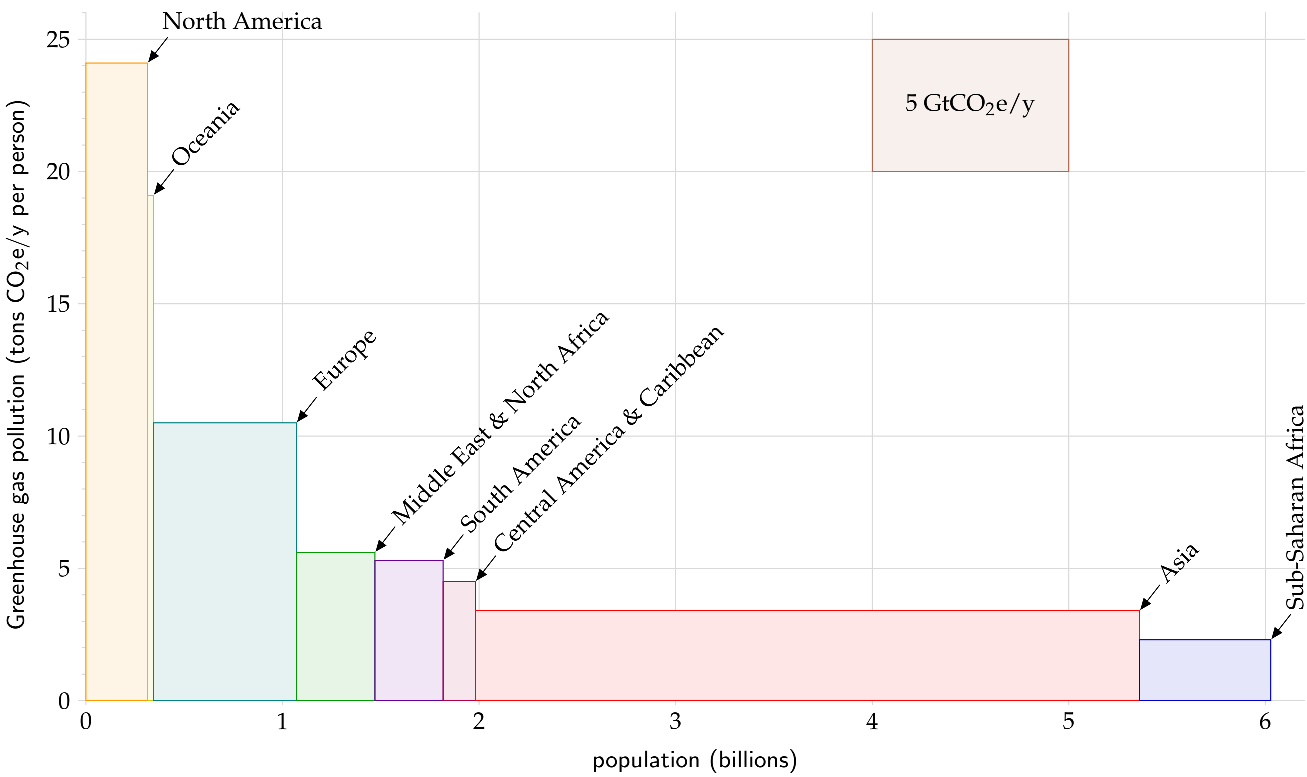

This picture, which is on the same scale as the previous one, divides the world into eight regions. Each rectangle's area represents the greenhouse gas emissions of one region. The width of the rectangle is the population of the region, and the height is the average per-capita emissions in that region.

In the year 2000, Europe's per-capita greenhouse gas emissions were twice the world average; and North America's were four times the world average.

We can continue subdividing, splitting each of the regions into countries. This is where it gets really interesting:

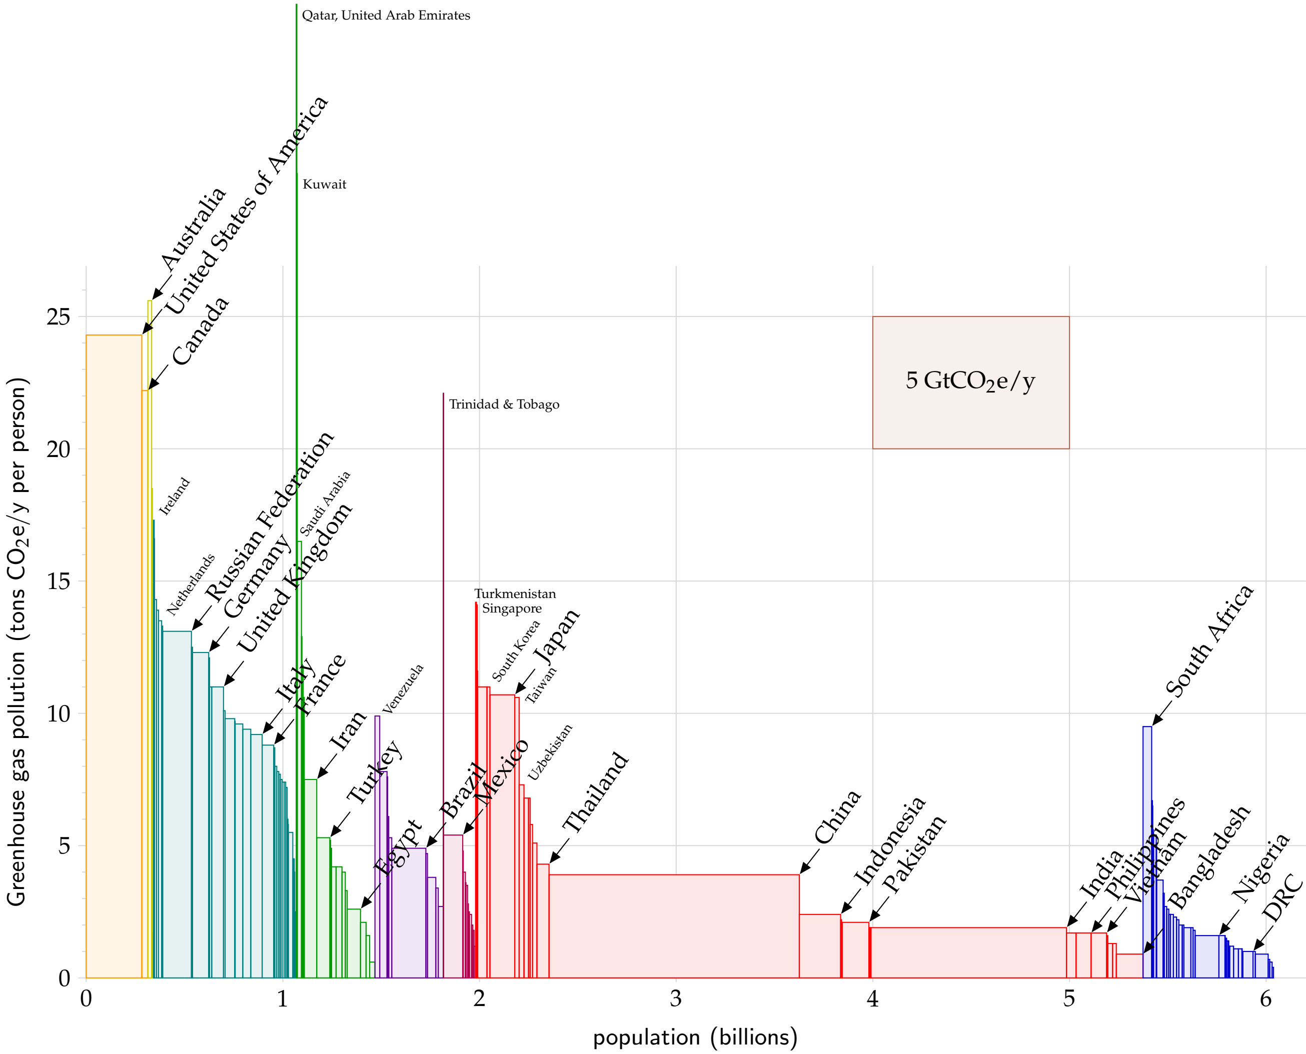

The major countries with the biggest per-capita emissions are Australia, the USA, and Canada. European countries, Japan, and South Africa are notable runners up. Among European countries, the United Kingdom is resolutely average. What about China, that naughty "out of control" country? Yes, the area of China's rectangle is about the same as the USA's, but the fact is that their per-capita emissions are *below* the world average. India's per-capita emissions are less than *half* the world average. Moreover, it's worth bearing in mind that much of the industrial emissions of China and India are associated with the manufacture of *stuff for rich countries*.

So, assuming that "something needs to be done" to reduce greenhouse gas emissions, who has a special responsibility to do something? As I said, that's an ethical question. But I find it hard to imagine any system of ethics that denies that the responsibility falls especially on the countries to the left hand side of this diagram – the countries whose emissions are two, three, or four times the world average. Countries that are most able to pay. Countries like Britain and the USA, for example.

Historical responsibility for climate impact

If we assume that the climate has been damaged by human activity, and that someone needs to fix it, who should pay? Some people say "the polluter should pay." The preceding pictures showed who's doing the polluting today. But it isn't the *rate* of CO_2 pollution that matters, it's the cumulative *total* emissions; much of the emitted carbon dioxide (about one third of it) will hang around in the atmosphere for at least 50 or 100 years. If we accept the ethical idea that "the polluter should pay" then we should ask how big is each country's historical footprint. The next picture shows each country's cumulative emissions of CO_2, expressed as an average emission rate over the period 1880–2004.

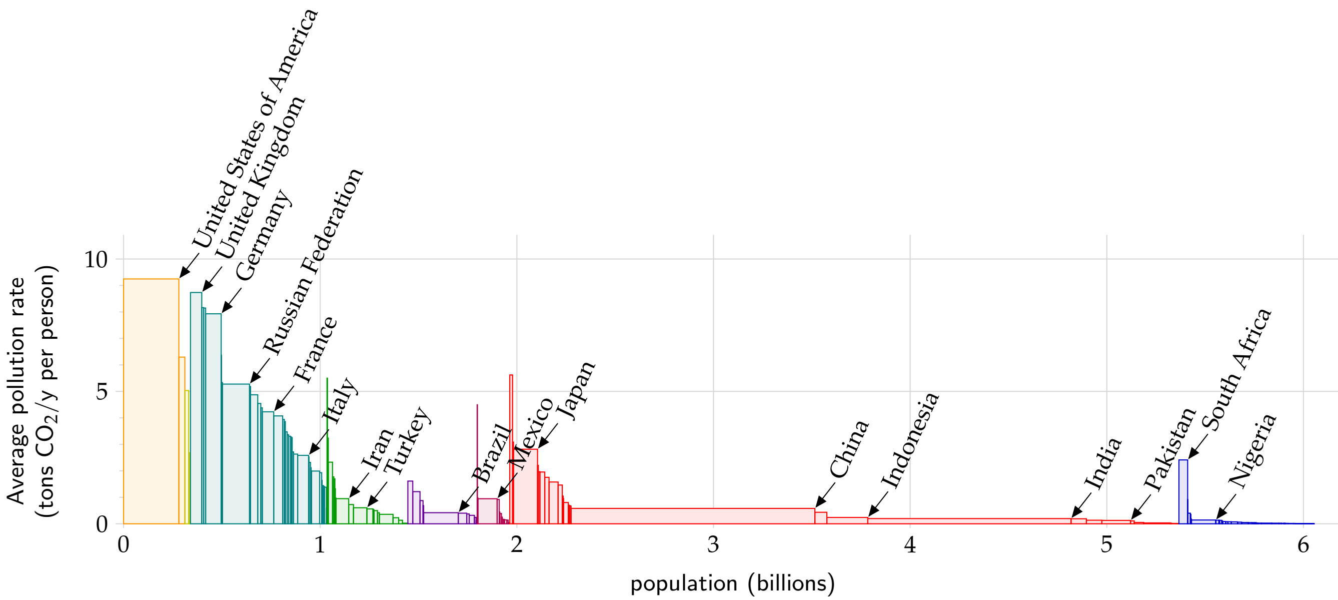

Congratulations, Britain! The UK has made it onto the winners' podium. We may be only an average European country today, but in the table of historical emitters, per capita, we are second only to the USA.

OK, that's enough ethics. What do scientists reckon needs to be done, to avoid a risk of giving the earth a 2 °C temperature rise (2 °C being the rise above which they predict lots of bad consequences)? The consensus is clear. We need to get off our fossil fuel habit, and we need to do so fast. Some countries, including Britain, have committed to at least a 60% reduction in greenhouse-gas emissions by 2050, but it must be emphasized that 60% cuts, radical though they are, are unlikely to cut the mustard. If the world's emissions were gradually reduced by 60% by 2050, climate scientists reckon it's more likely than not that global temperatures will rise by more than 2 °C. The sort of cuts we need to aim for are shown in figure 1.8. This figure shows two possibly-safe emissions scenarios presented by Baer and Mastrandrea (2006) in a report from the Institute for Public Policy Research. The lower curve assumes that a decline in emissions started in 2007, with total global emissions falling at roughly 5% per year. The upper curve assumes a brief delay in the start of the decline, and a 4% drop per year in global emissions. Both scenarios are believed to offer a modest chance of avoiding a 2 °C temperature rise above the pre-industrial level. In the lower scenario, the chance that the temperature rise will *exceed* 2 °C is estimated to be 9–26%. In the upper scenario, the chance of exceeding 2 °C is estimated to be 16–43%. These possibly-safe emissions trajectories, by the way, involve significantly sharper reductions in emissions than any of the scenarios presented by the Intergovernmental Panel on Climate Change (IPCC), or by the Stern Review (2007).

These possibly-safe trajectories require global emissions to fall by 70% or 85% by 2050. What would this mean for a country like Britain? If we subscribe to the idea of "contraction and convergence," which means that all countries aim eventually to have equal per-capita emissions, then Britain needs to aim for cuts greater than 85%: it should get down from its current 11 tons of CO_2e per year per person to roughly 1 ton per year per

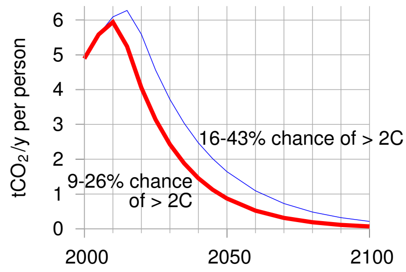

Figure 1.8. Global emissions for two scenarios considered by Baer and Mastrandrea, expressed in tons of CO_2 per year per person, using a world population of six billion. Both scenarios are believed to offer a modest chance of avoiding a 2 °C temperature rise above the pre-industrial level.

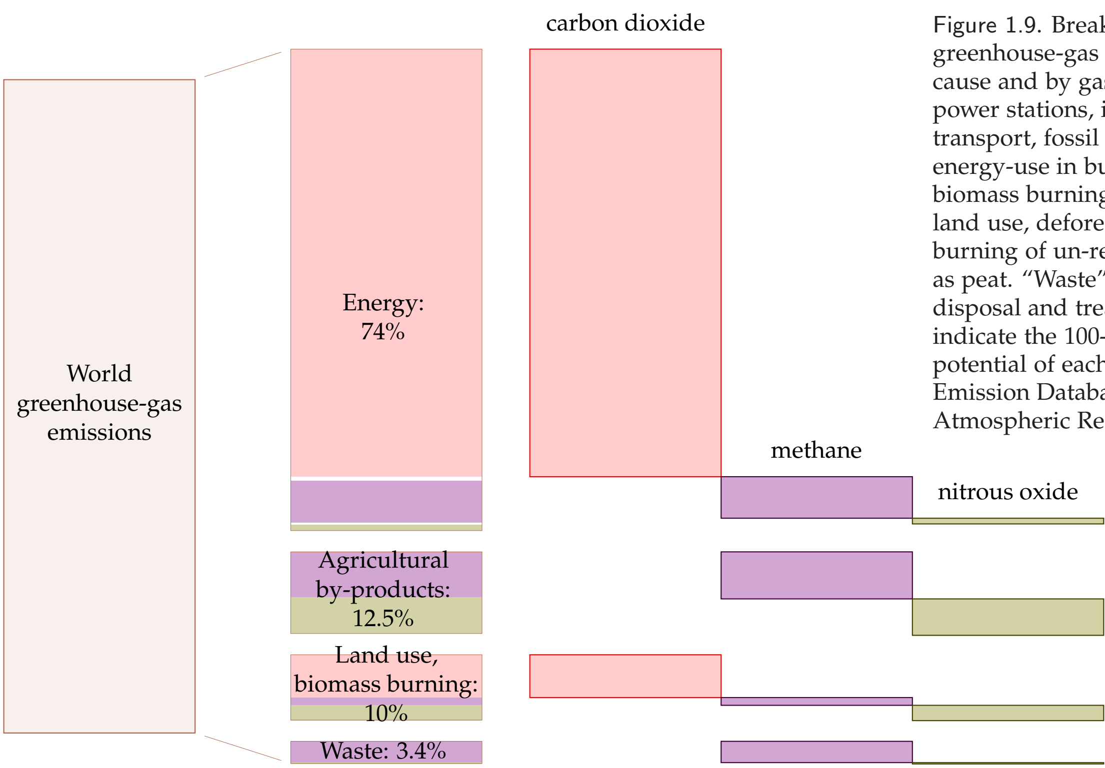

Figure 1.9. Breakdown of world greenhouse-gas emissions (2000) by cause and by gas. "Energy" includes power stations, industrial processes, transport, fossil fuel processing, and energy-use in buildings. "Land use, biomass burning" means changes in land use, deforestation, and the burning of un-renewed biomass such as peat. "Waste" includes waste disposal and treatment. The sizes indicate the 100-year global warming potential of each source. Source: Emission Database for Global Atmospheric Research.

advantage that it doesn't monopolize the land. Burning methane gas from landfill sites is a similar way of getting energy, but it's sustainable only as long as we have a sustainable source of junk to keep putting into the landfill sites. (Most of the landfill methane comes from wasted food; people in Britain throw away about 300 g of food per day per person.) Incinerating household waste is another slightly less roundabout way of getting power from solar biomass.

4. We can grow plants and feed them directly to energy-requiring humans or other animals.

For all of these processes, the first staging post for the energy is in a chemical molecule such as a carbohydrate in a green plant. We can therefore estimate the power obtainable from any and all of these processes by estimating how much power could pass through that first staging post. All the subsequent steps involving tractors, animals, chemical facilities, landfill sites, or power stations can only lose energy. So the power at the first staging post is an upper bound on the power available from all plant-based power solutions.

So, let's simply estimate the power at the first staging post. (In Chapter D, we'll go into more detail, estimating the maximum contribution of each process.) The average harvestable power of sunlight in Britain is 100 W/m^2. The most efficient plants in Europe are about 2%-efficient at turning solar energy into carbohydrates, which would suggest that plants might deliver 2 W/m^2; however, their efficiency drops at higher light levels, and the best performance of any energy crops in Europe is closer to 0.5 W/m^2. Let's cover 75% of the country with quality green stuff. That's 3000 m^2 per person devoted to bio-energy. This is the same as the British land area

Figure 6.10. Some *Miscanthus* grass enjoying the company of Dr Emily Heaton, who is 5'4" (163 cm) tall. In Britain, *Miscanthus* achieves a power per unit area of 0.75 W/m^2. Photo provided by the University of Illinois.

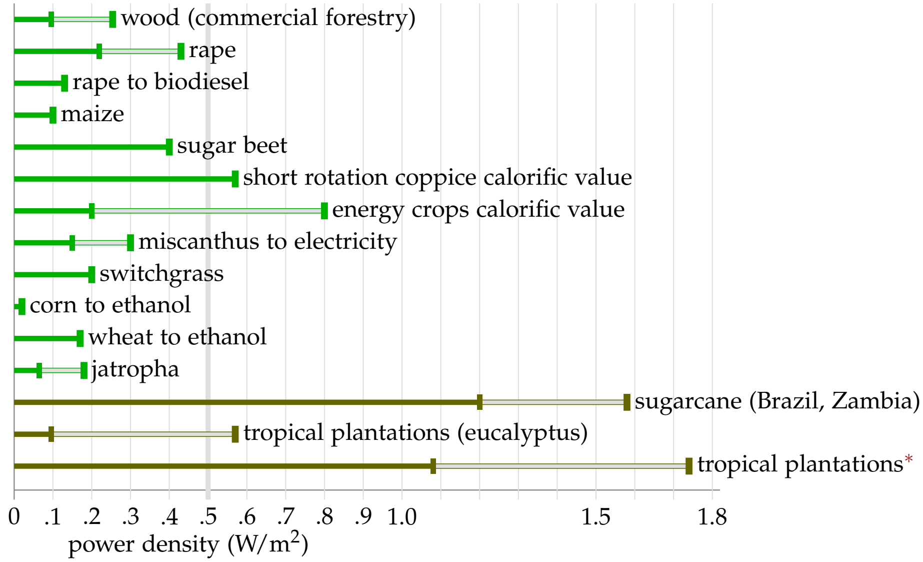

Figure 6.11. Power production, per unit area, achieved by various plants. For sources, see the end-notes. These power densities vary depending on irrigation and fertilization; ranges are indicated for some crops, for example wood has a range from 0.095–0.254 W/m^2. The bottom three power densities are for crops grown in tropical locations. The last power density (tropical plantations*) assumes genetic modification, fertilizer application, and irrigation. In the text, I use 0.5 W/m^2 as a summary figure for the best energy crops in NW Europe.

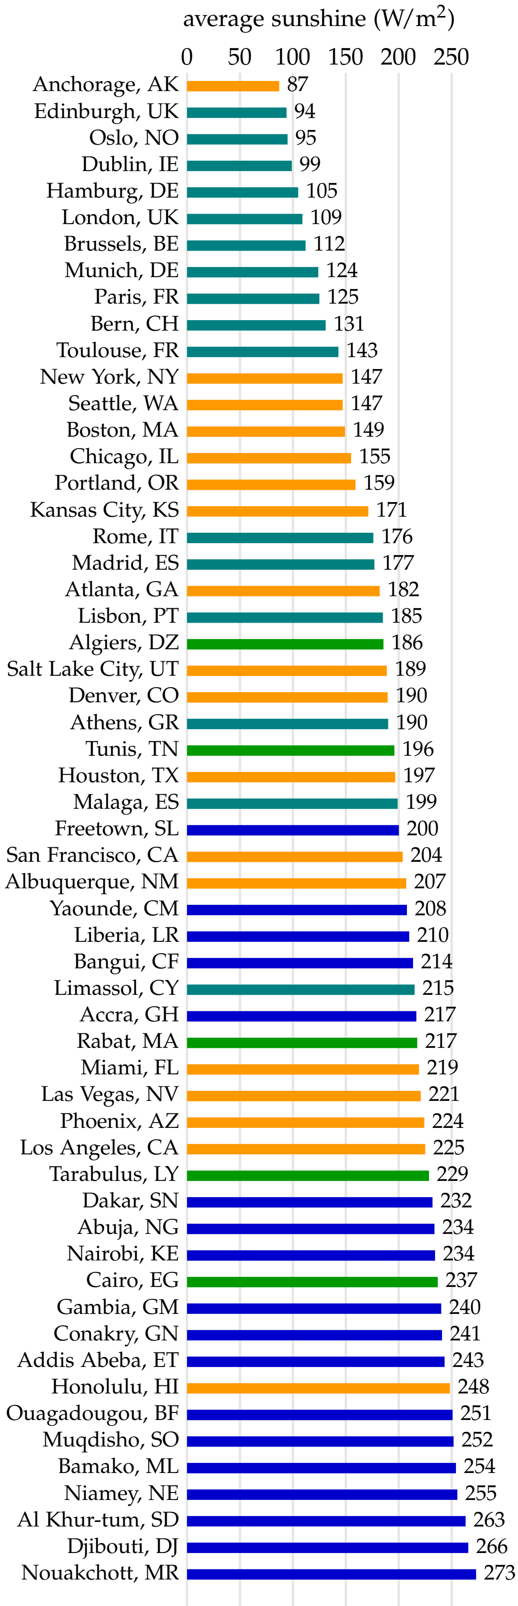

Figure 6.16. Average power of sunshine falling on a horizontal surface in selected locations in Europe, North America, and Africa.

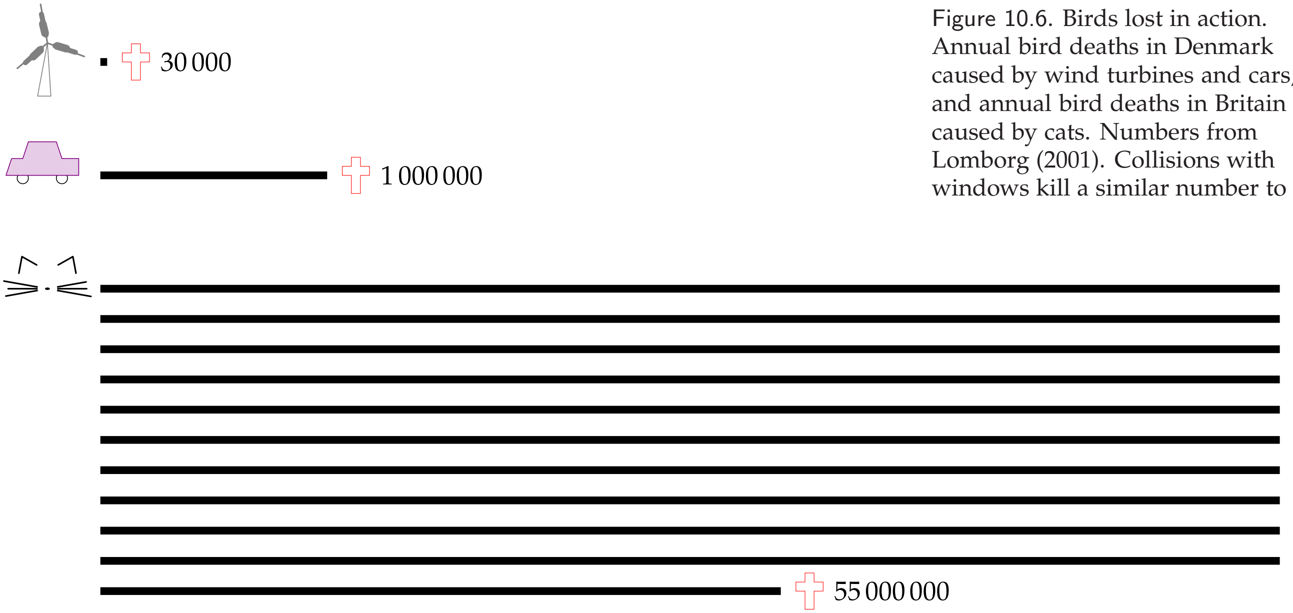

Figure 10.6. Birds lost in action. Annual bird deaths in Denmark caused by wind turbines and cars, and annual bird deaths in Britain caused by cats. Numbers from Lomborg (2001). Collisions with windows kill a similar number to cats.

Notes and further reading

page no.

60 *The Kentish Flats wind farm in the Thames Estuary...*
See `www.kentishflats.co.uk`. Its 30 Vestas V90 wind turbines have a total peak output of 90 MW, and the predicted average output was 32 MW (assuming a load factor of 36%). The mean wind speed at the hub height is 8.7 m/s. The turbines stand in 5 m-deep water, are spaced 700 m apart, and occupy an area of 10 km^2. The power density of this offshore wind farm was thus predicted to be 3.2 W/m^2. In fact, the average output was 26 MW, so the average load factor in 2006 was 29% [`wbd8o`]. This works out to a power density of 2.6 W/m^2. The North Hoyle wind farm off Prestatyn, North Wales, had a higher load factor of 36% in 2006. Its thirty 2 MW turbines occupy 8.4 km^2. They thus had an average power density of 2.6 W/m^2.

– *...shallow offshore wind, while roughly twice as costly as onshore wind, is economically feasible, given modest subsidy.* Source: Danish wind association `windpower.org`.

– *...deep offshore wind is at present not economically feasible.*
Source: British Wind Energy Association briefing document, September 2005, `www.bwea.com`. Nevertheless, a deep offshore demonstration project in 2007 put two turbines adjacent to the Beatrice oil field, 22 km off the east coast of Scotland (figure 10.8). Each turbine has a "capacity" of 5 MW and sits in a water depth of 45 m. Hub height: 107 m; diameter 126 m. All the electricity generated will be used by the oil platforms. Isn't that special! The 10 MW project cost £30 million – this price-tag of £3 per watt (peak) can be

per unit length of exposed coastline, and multiplying by the length of coastline. We ignore the question of what mechanism could collect all this power, and start by working out how much power it is.

The power of Atlantic waves has been measured: it's about 40 kW per metre of exposed coastline. That sounds like a lot of power! If everyone owned a metre of coastline and could harness their whole 40 kW, that would be plenty of power to cover modern consumption. However, *our population is too big*. There is not enough Atlantic-facing coastline for everyone to have their own metre.

As the map on p73 shows, Britannia rules about 1000 km of Atlantic coastline (one million metres), which is $^1/_{60}$ m per person. So the total raw incoming power is 16 kWh per day per person. If we extracted all this power, the Atlantic, at the seaside, would be as flat as a millpond. Practical systems won't manage to extract all the power, and some of the power will inevitably be lost during conversion from mechanical energy to electricity. Let's assume that brilliant wave-machines are 50%-efficient at turning the incoming wave power into electricity, and that we are able to pack wave-machines along 500 km of Atlantic-facing coastline. That would mean we could deliver 25% of this theoretical bound. That's 4 kWh per day per person. As usual, I'm intentionally making pretty extreme assumptions to boost the green stack – I expect the assumption that we could line *half of the Atlantic coastline* with wave absorbers will sound bananas to many readers.

How do the numbers assumed in this calculation compare with today's technology? As I write, there are just three wave machines working in deep water: three Pelamis wave energy collectors (figure 12.1) built in Scotland and deployed off Portugal. No actual performance results have been published, but the makers of the Pelamis ("designed with survival as the key objective before power capture efficiency") predict that a two-kilometre-long wave-farm consisting of 40 of their sea-snakes would deliver 6 kW per metre of wave-farm. Using this number in the previous calculation, the power delivered by 500 kilometres of wave-farm is reduced to 1.2 kWh per day per person. While wave power may be useful for small communities on remote islands, I suspect it can't play a significant role in the solution to Britain's sustainable energy problem.

What's the weight of a Pelamis, and how much steel does it contain? One snake with a maximum power of 750 kW weighs 700 tons, including 350 tons of ballast. So it has about 350 tons of steel. That's a weight-to-power ratio of roughly 500 kg per kW (peak). We can compare this with the steel requirements for offshore wind: an offshore wind-turbine with a maximum power of 3 MW weighs 500 tons, including its foundation. That's a weight-to-power ratio of about 170 kg per kW, one third of the wave machine's. The Pelamis is a first prototype; presumably with further investment and development in wave technology, the weight-to-power ratio would fall.

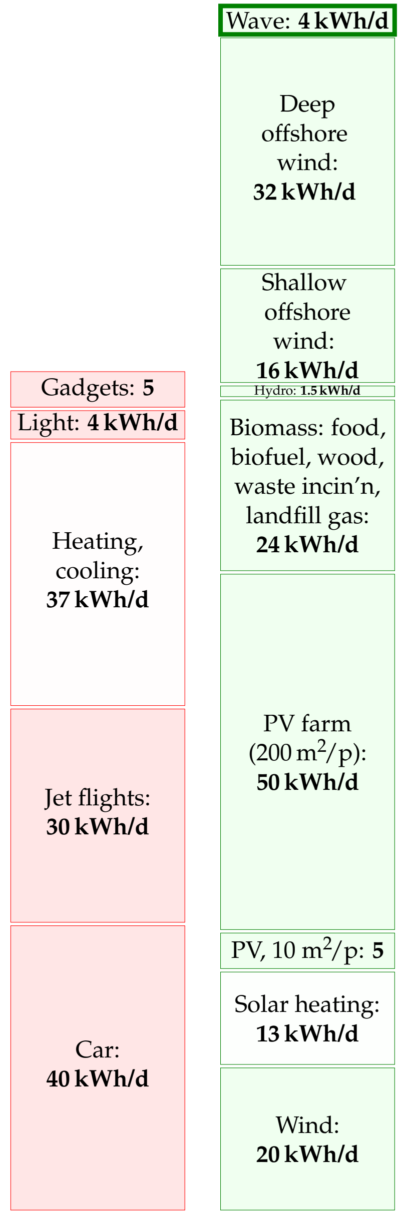

Figure 12.2. Wave.

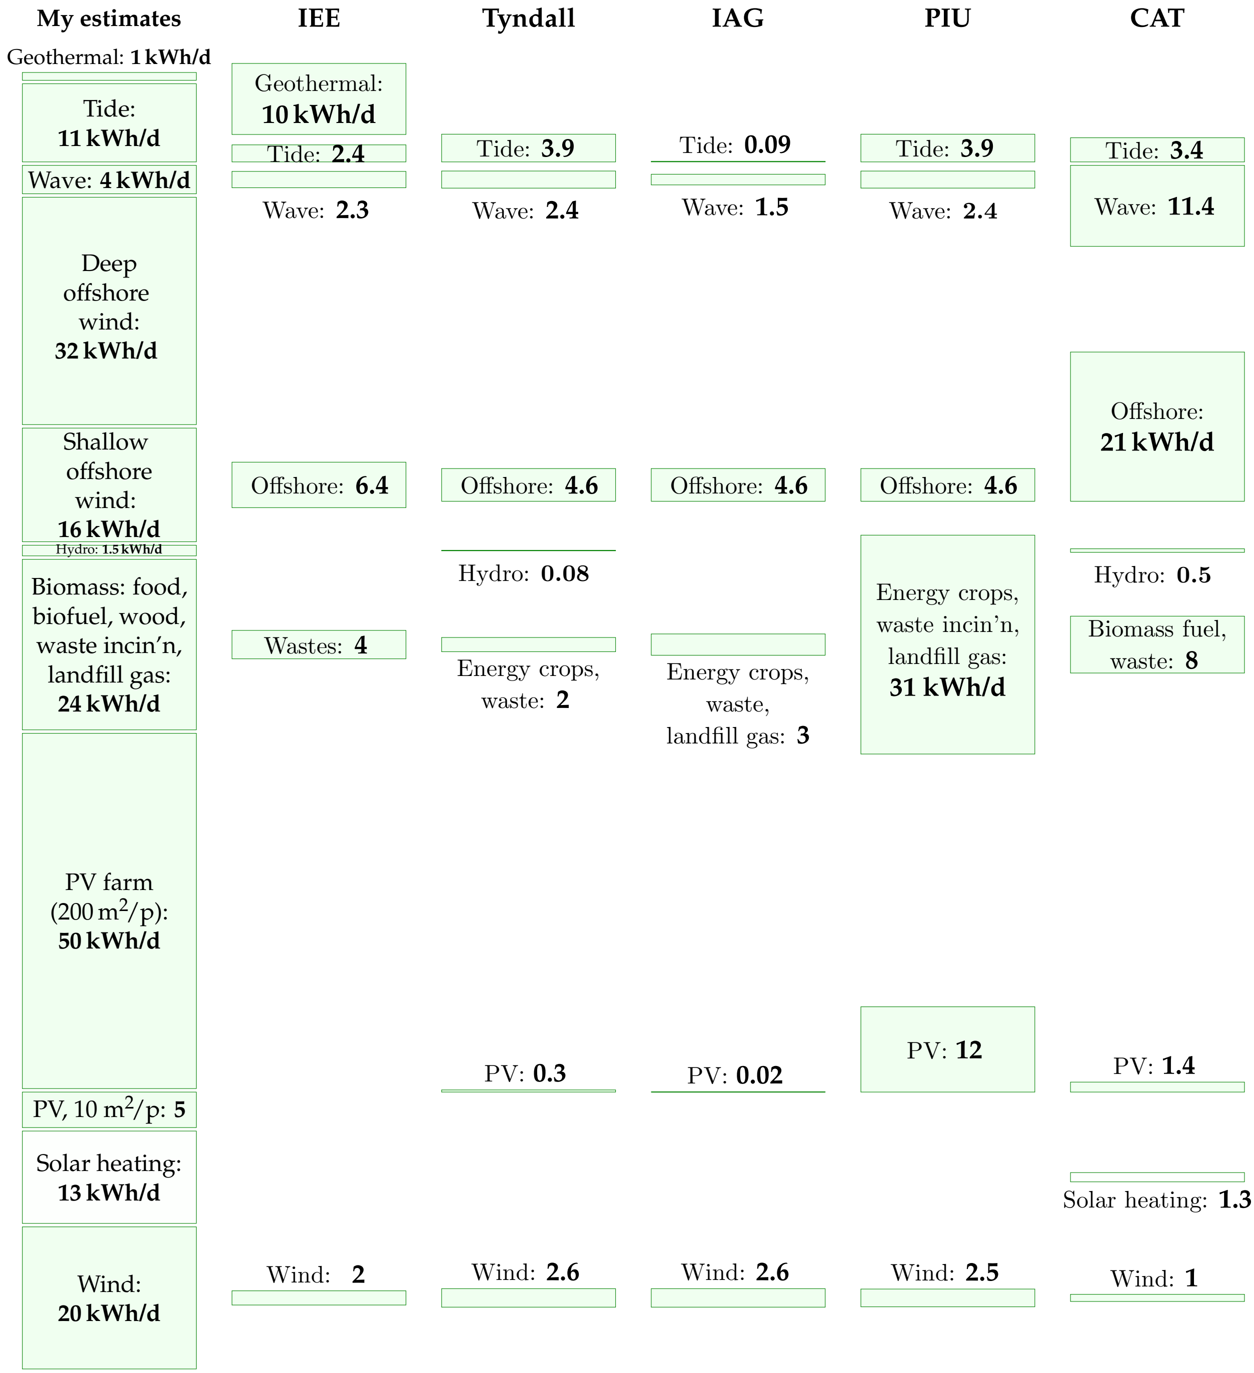

Figure 18.6. Estimates of theoretical or practical renewable resources in the UK, by the Institute of Electrical Engineers, the Tyndall Centre, the Interdepartmental Analysts Group, and the Performance and Innovation Unit; and the proposals from the Centre for Alternative Technology's "Island Britain" plan for 2027.

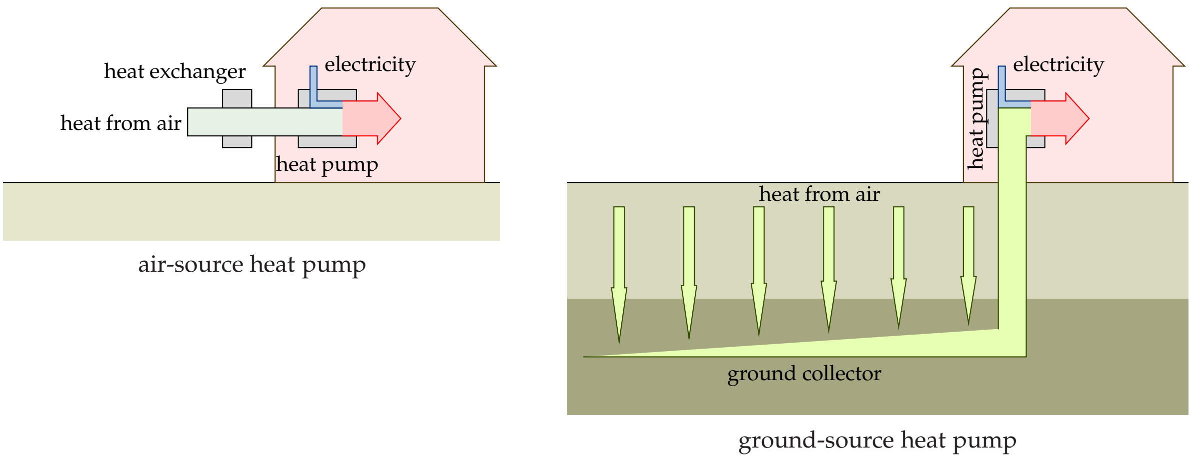

Figure 21.10. Heat pumps.

station could be captured for a useful purpose *without impairing the power station's electricity production*. This sadly is not true, as the numbers will show. Delivering useful heat to a customer always reduces the electricity produced to some degree. The true net gains from combined heat and power are often much smaller than the hype would lead you to believe.

A final impediment to rational discussion of combined heat and power is a myth that has grown up recently, that decentralizing a technology somehow makes it greener. So whereas big centralized fossil fuel power stations are "bad," flocks of local micro-power stations are imbued with goodness. But if decentralization is actually a good idea then "small is beautiful" should be evident in the numbers. Decentralization should be able to stand on its own two feet. And what the numbers actually show is that *centralized* electricity generation has many benefits in both economic and energy terms. Only in large buildings is there any benefit to local generation, and usually that benefit is only about 10% or 20%.

The government has a target for growth of combined heat and power to 10 GW of electrical capacity by 2010, but I think that growth of gas-powered combined heat and power would be a mistake. Such combined heat and power is not green: it uses fossil fuel, and it locks us into continued use of fossil fuel. Given that heat pumps are a better technology, I believe we should leapfrog over gas-powered combined heat and power and go directly for heat pumps.

Heat pumps

Like district heating and combined heat and power, heat pumps are already widely used in continental Europe, but strangely rare in Britain. Heat pumps are back-to-front refrigerators. Feel the back of your refrigerator: it's *warm*. A refrigerator moves heat from one place (its inside) to ton of coal (which is what you might use to heat a house over a year). Now imagine everyone on the planet burning one ton of coal per year: that's 6 Gt C per year, because the planet has 6 billion people.

Where is the carbon?

Where is all the carbon? We need to know how much is in the oceans, in the ground, and in vegetation, compared to the atmosphere, if we want to understand the consequences of CO_2 emissions.

Figure 31.2 shows where the carbon is. Most of it – 40 000 Gt – is in the ocean (in the form of dissolved CO_2 gas, carbonates, living plant and

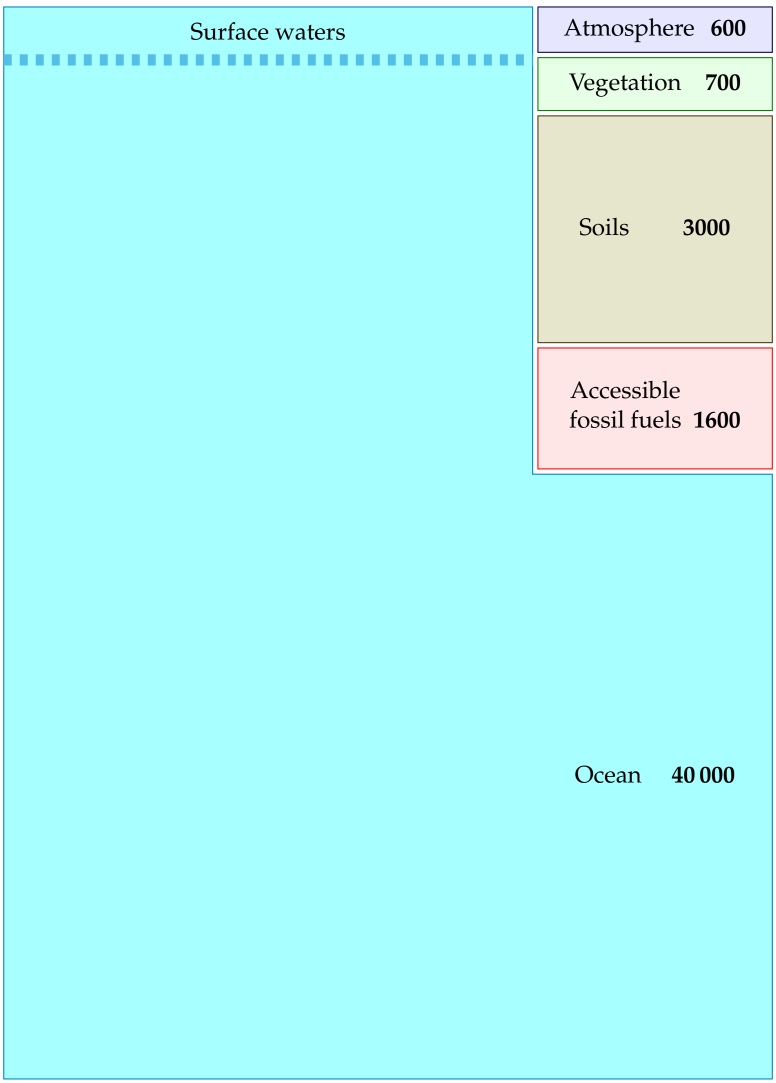

Figure 31.2. Estimated amounts of carbon, in gigatons, in accessible places on the earth. (There's a load more carbon in rocks too; this carbon moves round on a timescale of millions of years, with a long-term balance between carbon in sediment being subducted at tectonic plate boundaries, and carbon popping out of volcanoes from time to time. For simplicity I ignore this geological carbon.)

animal life, and decaying materials). Soils and vegetation together contain about 3700 Gt. Accessible fossil fuels – mainly coal – contain about 1600 Gt. Finally, the atmosphere contains about 600 Gt of carbon.

Until recently, all these pools of carbon were roughly in balance: all flows of carbon out of a pool (say, soils, vegetation, or atmosphere) were balanced by equal flows into that pool. The flows into and out of the fossil fuel pool were both negligible. Then humans started burning fossil fuels. This added two extra *unbalanced* flows, as shown in figure 31.3.

The rate of fossil fuel burning was roughly 1 Gt C/y in 1920, 2 Gt C/y in 1955, and 8.4 Gt C in 2006. (These figures include a small contribution from cement production, which releases CO_2 from limestone.)

How has this significant extra flow of carbon modified the picture shown in figure 31.2? Well, it's not exactly known. Figure 31.3 shows the key things that *are* known. Much of the extra 8.4 Gt C per year that we're putting into the atmosphere stays in the atmosphere, raising the atmospheric concentration of carbon-dioxide. The atmosphere equilibrates fairly rapidly with the surface waters of the oceans (this equilibration takes only five or ten years), and there is a net flow of CO_2 from the atmosphere into the surface waters of the oceans, amounting to 2 Gt C per year. (Recent research indicates this rate of carbon-uptake by the oceans may be reducing, however.) This unbalanced flow into the surface waters causes ocean acidification, which is bad news for coral. Some extra carbon is moving into vegetation and soil too, perhaps about 1.5 Gt C per year, but these flows are less well measured. Because roughly half of the carbon emissions are staying in the atmosphere, continued carbon pollution at a rate of 8.4 Gt C per year will continue to increase CO_2 levels in the atmosphere, and in the surface waters.

What is the long-term destination of the extra CO_2? Well, since the amount in fossil fuels is so much smaller than the total in the oceans, "in the long term" the extra carbon will make its way into the ocean, and the amounts of carbon in the atmosphere, vegetation, and soil will return to normal. However, "the long term" means thousands of years. Equilibration between atmosphere and the *surface* waters is rapid, as I said, but figures 31.2 and 31.3 show a dashed line separating the surface waters of the ocean from the rest of the ocean. On a time-scale of 50 years, this boundary is virtually a solid wall. Radioactive carbon dispersed across the globe by the atomic bomb tests of the 1960s and 70s has penetrated the oceans to a depth of only about 400 m. In contrast the average depth of the oceans is about 4000 m.

The oceans circulate slowly: a chunk of deep-ocean water takes about 1000 years to roll up to the surface and down again. The circulation of the deep waters is driven by a combination of temperature gradients and salinity gradients, so it's called the thermohaline circulation (in contrast to the circulations of the surface waters, which are wind-driven).

This slow turn-over of the oceans has a crucial consequence: we have

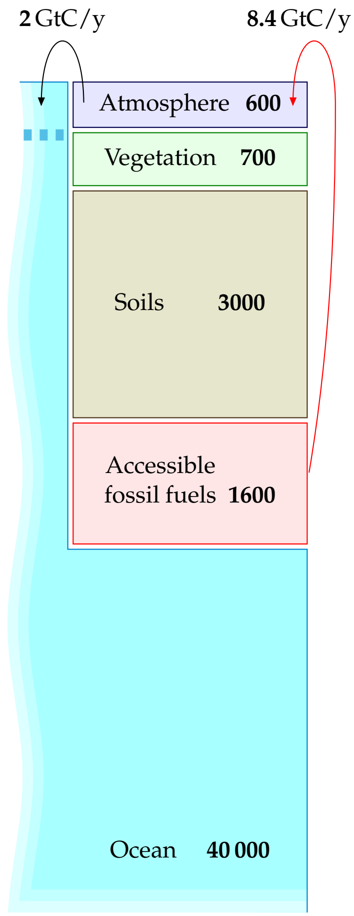

Figure 31.3. The arrows show two extra carbon flows produced by burning fossil fuels. There is an imbalance between the 8.4 Gt C/y emissions into the atmosphere from burning fossil fuels and the 2 Gt C/y take-up of CO_2 by the oceans. This cartoon omits the less-well quantified flows between atmosphere, soil, vegetation, and so forth.

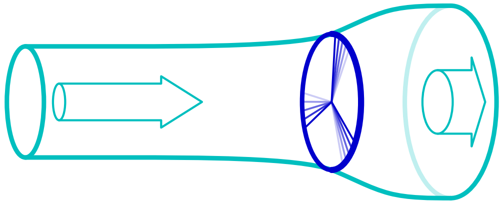

Figure B.2. Flow of air past a windmill. The air is slowed down and splayed out by the windmill.

normal; the speed of such a wind is therefore comparable to the typical speed of the cyclist, which is, let's say, 21 km per hour (13 miles per hour, or 6 metres per second). In Cambridge, the wind is only occasionally this big. Nevertheless, let's use this as a typical British figure (and bear in mind that we may need to revise our estimates).

The density of air is about 1.3 kg per m^3. (I usually round this to 1 kg per m^3, which is easier to remember, although I haven't done so here.) Then the typical power of the wind per square metre of hoop is

$$\frac{1}{2}\rho v^3 = \frac{1}{2} 1.3\,\mathrm{kg/m^3} \times (6\,\mathrm{m/s})^3 = 140\,\mathrm{W/m^2}. \quad \text{(B.3)}$$

Not all of this energy can be extracted by a windmill. The windmill slows the air down quite a lot, but it has to leave the air with *some* kinetic energy, otherwise that slowed-down air would get in the way. Figure B.2 is a cartoon of the actual flow past a windmill. The maximum fraction of the incoming energy that can be extracted by a disc-like windmill was worked out by a German physicist called Albert Betz in 1919. If the departing wind speed is one third of the arriving wind speed, the power extracted is 16/27 of the total power in the wind. 16/27 is 0.59. In practice let's guess that a windmill might be 50% efficient. In fact, real windmills are designed with particular wind speeds in mind; if the wind speed is significantly greater than the turbine's ideal speed, it has to be switched off.

As an example, let's assume a diameter of $d = 25$ m, and a hub height of 32 m, which is roughly the size of the lone windmill above the city of Wellington, New Zealand (figure B.3). The power of a single windmill is

$$\begin{aligned} & \text{efficiency factor} \times \text{power per unit area} \times \text{area} \\ = \; & 50\% \times \frac{1}{2}\rho v^3 \times \frac{\pi}{4} d^2 && \text{(B.4)} \\ = \; & 50\% \times 140\,\mathrm{W/m^2} \times \frac{\pi}{4}(25\,\mathrm{m})^2 && \text{(B.5)} \\ = \; & 34\,\mathrm{kW}. && \text{(B.6)} \end{aligned}$$

Indeed, when I visited this windmill on a very breezy day, its meter showed it was generating 60 kW.

To estimate how much power we can get from wind, we need to decide how big our windmills are going to be, and how close together we can pack them.

Figure B.3. The Brooklyn windmill above Wellington, New Zealand, with people providing a scale at the base. On a breezy day, this windmill was producing 60 kW, (1400 kWh per day). Photo by Philip Banks.

How densely could such windmills be packed? Too close and the up-wind ones will cast wind-shadows on the downwind ones. Experts say that windmills can't be spaced closer than 5 times their diameter without losing significant power. At this spacing, the power that windmills can generate per unit land area is

$$\frac{\text{power per windmill (B.4)}}{\text{land area per windmill}} = \frac{\frac{1}{2}\rho v^3 \frac{\pi}{8} d^2}{(5d)^2} \quad \text{(B.7)}$$

$$= \frac{\pi}{200}\frac{1}{2}\rho v^3 \quad \text{(B.8)}$$

$$= 0.016 \times 140\,\text{W/m}^2 \quad \text{(B.9)}$$

$$= 2.2\,\text{W/m}^2. \quad \text{(B.10)}$$

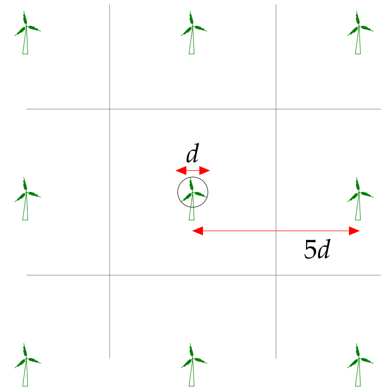

Figure B.4. Wind farm layout.

This number is worth remembering: a wind farm with a wind speed of 6 m/s produces a power of 2 W per m^2 of land area. Notice that our answer does not depend on the diameter of the windmill. The *d*s cancelled because bigger windmills have to be spaced further apart. Bigger windmills might be a good idea in order to catch bigger windspeeds that exist higher up (the taller a windmill is, the bigger the wind speed it encounters), or because of economies of scale, but those are the only reasons for preferring big windmills.

Power per unit area	
wind farm (speed 6 m/s)	$2\,\text{W/m}^2$

Table B.5. Facts worth remembering: wind farms.

This calculation depended sensitively on our estimate of the windspeed. Is 6 m/s plausible as a long-term typical windspeed in windy parts of Britain? Figures 4.1 and 4.2 showed windspeeds in Cambridge and Cairngorm. Figure B.6 shows the mean winter and summer windspeeds in eight more locations around Britain. I fear 6 m/s was an overestimate of the typical speed in most of Britain! If we replace 6 m/s by Bedford's

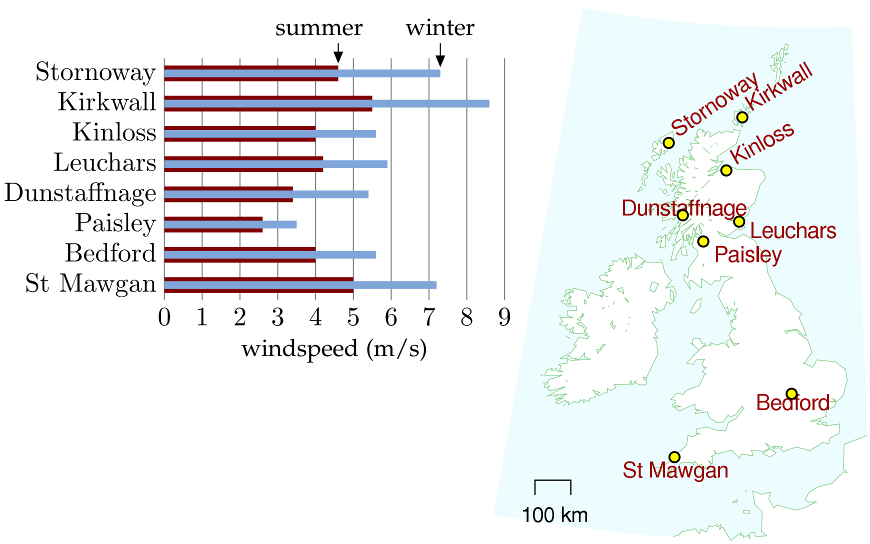

Figure B.6. Average summer windspeed (dark bar) and average winter windspeed (light bar) in eight locations around Britain. Speeds were measured at the standard weatherman's height of 10 metres. Averages are over the period 1971–2000.

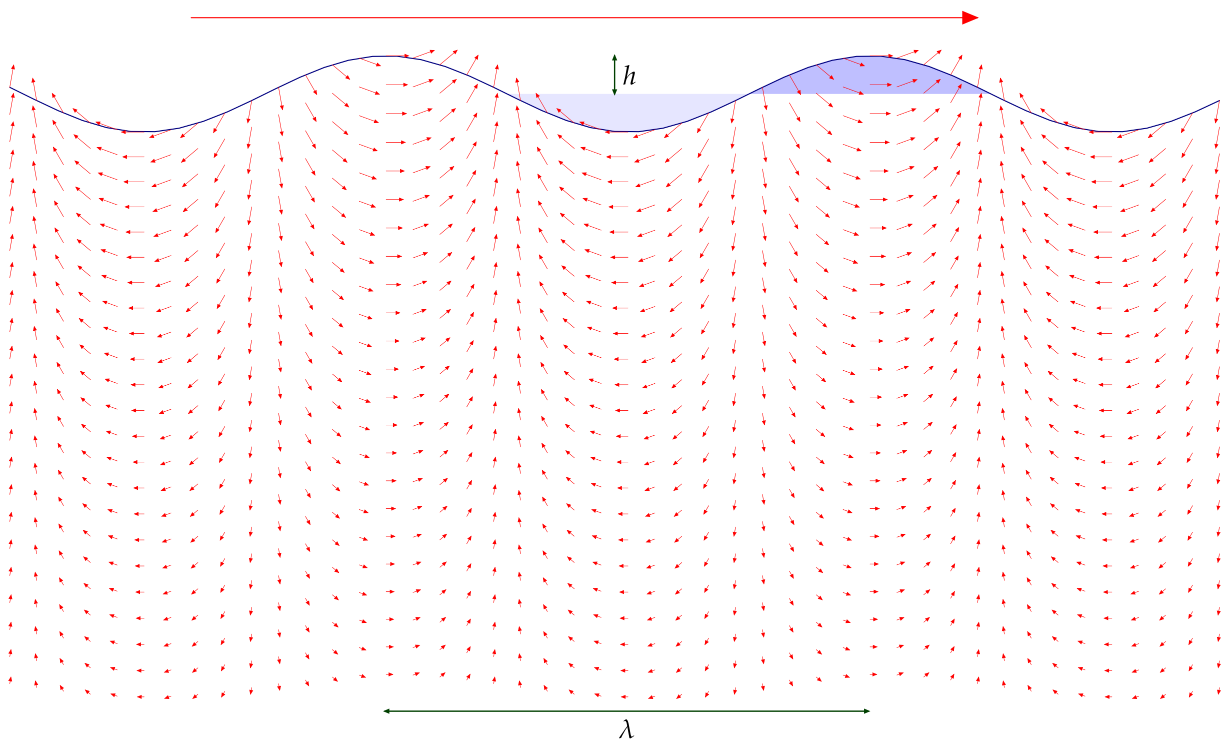

Figure F.2. A wave has energy in two forms: potential energy associated with raising water out of the light-shaded troughs into the heavy-shaded crests; and kinetic energy of all the water within a few wavelengths of the surface – the speed of the water is indicated by the small arrows. The speed of the wave, travelling from left to right, is indicated by the much bigger arrow at the top.

For a wave of wavelength λ and period T, if the height of each crest and depth of each trough is $h = 1\,\text{m}$, the potential energy passing per unit time, per unit length, is

$$P_{\text{potential}} \simeq m^* g\bar{h}/T, \tag{F.1}$$

where m^* is the mass per unit length, which is roughly $\frac{1}{2}\rho h(\lambda/2)$ (approximating the area of the shaded crest in figure F.2 by the area of a triangle), and $\bar{h}$ is the change in height of the centre-of-mass of the chunk of elevated water, which is roughly h. So

$$P_{\text{potential}} \simeq \frac{1}{2}\rho h \frac{\lambda}{2} g h/T. \tag{F.2}$$

(To find the potential energy properly, we should have done an integral here; it would have given the same answer.) Now λ/T is simply the speed at which the wave travels, v, so:

$$P_{\text{potential}} \simeq \frac{1}{4}\rho g h^2 v. \tag{F.3}$$

Waves have kinetic energy as well as potential energy, and, remarkably, these are exactly equal, although I don't show that calculation here; so the total power of the waves is double the power calculated from potential

G Tide II

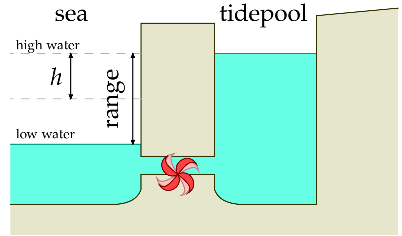

Figure G.1. A tide-pool in cross section. The pool was filled at high tide, and now it's low tide. We let the water out through the electricity generator to turn the water's potential energy into electricity.

Power density of tidal pools

To estimate the power of an artificial tide-pool, imagine that it's filled rapidly at high tide, and emptied rapidly at low tide. Power is generated in both directions, on the ebb and on the flood. (This is called two-way generation or double-effect generation.) The change in potential energy of the water, each six hours, is mgh, where h is the change in height of the centre of mass of the water, which is half the range. (The range is the difference in height between low and high tide; figure G.1.) The mass per unit area covered by tide-pool is $\rho \times (2h)$, where ρ is the density of water ($1000\,\text{kg/m}^3$). So the power per unit area generated by a tide-pool is

$$\frac{2\rho h g h}{6\,\text{hours}},$$

assuming perfectly efficient generators. Plugging in $h = 2\,\text{m}$ (i.e., range 4 m), we find the power per unit area of tide-pool is $3.6\,\text{W/m}^2$. Allowing for an efficiency of 90% for conversion of this power to electricity, we get

$$\text{power per unit area of tide-pool} \simeq 3\,\text{W/m}^2.$$

So to generate 1 GW of power (on average), we need a tide-pool with an area of about $300\,\text{km}^2$. A circular pool with diameter 20 km would do the trick. (For comparison, the area of the Severn estuary behind the proposed barrage is about $550\,\text{km}^2$, and the area of the Wash is more than $400\,\text{km}^2$.

If a tide-pool produces electricity in one direction only, the power per unit area is halved. The average power density of the tidal barrage at La Rance, where the mean tidal range is 10.9 m, has been $2.7\,\text{W/m}^2$ for decades (p87).

The raw tidal resource

The tides around Britain are genuine tidal waves. (Tsunamis, which are called "tidal waves," have nothing to do with tides: they are caused by underwater landslides and earthquakes.) The location of the high tide (the crest of the tidal wave) moves much faster than the tidal flow – 100 miles per hour, say, while the water itself moves at just 1 mile per hour.

The energy we can extract from tides, using tidal pools or tide farms, can never be more than the energy of these tidal waves from the Atlantic. We can estimate the total power of these great Atlantic tidal waves in the same way that we estimate the power of ordinary wind-generated waves. The next section describes a standard model for the power arriving in