AROUND THE WORLD IN 80 WAYS

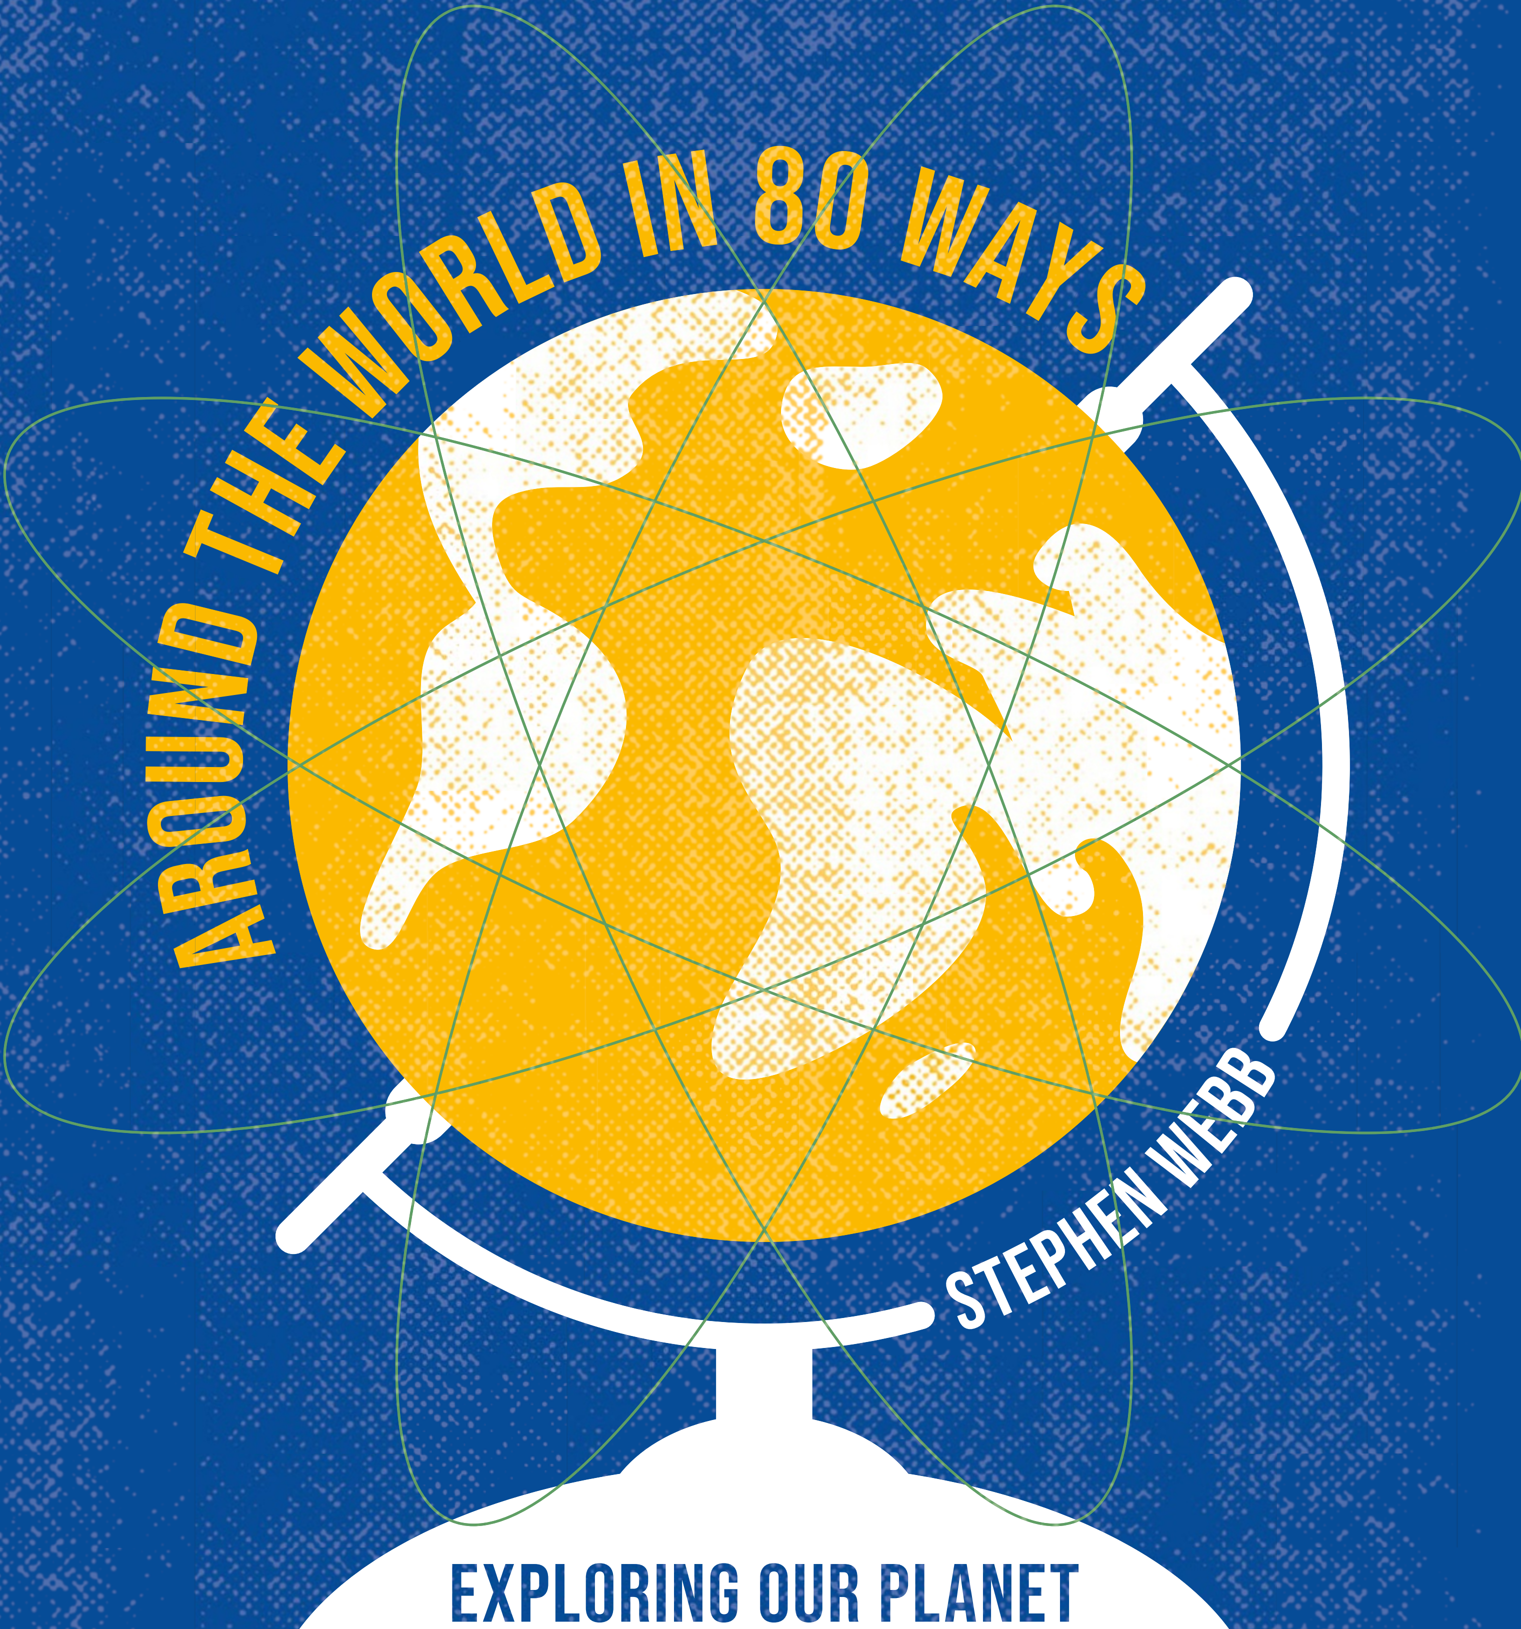

STEPHEN WEBB

EXPLORING OUR PLANET THROUGH MAPS AND DATA

Springer

Around the World in 80 Ways

Stephen Webb

Around the World in 80 Ways

Exploring Our Planet Through Maps and Data

Springer

Stephen Webb
DCQE
University of Portsmouth
Portsmouth, UK

ISBN 978-3-031-02439-9   ISBN 978-3-031-02440-5 (eBook)
https://doi.org/10.1007/978-3-031-02440-5

This Springer imprint is published by the registered company Springer Nature Switzerland AG
The registered company address is: Gewerbestrasse 11, 6330 Cham, Switzerland

To Chris Caron, Angela Lahee, and Lisa Scalone
who have always been a delight to work with

Acknowledgments

I would like to thank the Open Source and the Creative Commons communities. Without their generosity, and the giving of their time and knowledge, the writing of this book would have been a difficult and arduous task. Thanks to them, this was a fun project for me. I would also like to thank Alina Maniar and Carl Berndtsson for kindly allowing the reproduction of photographs.

I am particularly grateful to Brian Clegg for his extensive, expert, and insightful comments on an early draft of the manuscript. Any merits this book might possess will be in large part due to Brian's input. The mistakes, of course, are all my own. Lisa Scalone has once again provided brilliant editorial advice.

Finally, as always, I would like to thank Heike and Jessica for everything.

Contents

Mapping the World

Over the years I've had discussions with people who hold, shall we say, opinions that contradict mainstream scientific thought. Often, I felt, my interlocutors—anti-vaxxers, climate-change deniers, flat-Earthers, and others who discern conspiracy where the rest of us see science—refused to accept that human activity can impact the globe ("we're just one species; how can we cause climate change?") or even that there *is* a globe ("if Earth is curved, how come we can sometimes see very distant skyscrapers?"). How can they be so *ignorant* of the world, I'd wonder.

And then I read Hans Rosling's book *Factfulness*, the introduction to which contains a dozen simple questions about the world. Alongside each question Rosling offered three possible answers. I got four answers correct: the same as a chimpanzee would score were it to point randomly at the page. I consoled myself with the thought that Rosling posed those same questions to almost 12,000 people in 14 countries and not a single person got all of them right. The average score was just two correct answers (worse, then, than a chimpanzee) and 1800 people got every question wrong. It seems that *all* of us, Nobel prize winners included, are ignorant about many aspects of our world. But if my own knowledge was shaky, how could I in good conscience argue with the anti-vaxxers, the climate-change deniers, and the flat-Earthers?

I got into the habit of testing my intuition against the most robust data I could find. So I asked myself questions such as: What is the commonest cause of death in different countries? Do different nationalities search Google for different terms or are people across the globe interested in the same things? Which nations generate the most trash? At the same time I discovered that, using the power of modern GIS software, it is easy to create maps of these data—and a map is far easier to grasp than a table of numbers.

After making a few maps I hoped they would help me prove my case when I got into discussions with people who have an allergy to science. I didn't realise that some people set the bar low when it comes to evidence that might support their beliefs but impossibly high for evidence that might challenge them. My charts changed no minds.

Nevertheless, I enjoyed the process of making the maps—so I made more of them. I made 80, in fact, some quirky, others heartbreaking. They fall naturally into eight themes, and I've collected them here. However, as anti-vaxxers, climate-change deniers, and flat-Earthers would quickly point out, my maps deceive. How do I deceive thee? Let me count the ways. . .

1 Map Projections

All maps project the surface of a three-dimensional sphere (in other words, Earth's surface) onto a two-dimensional space (such as the page of a book). It's mathematically impossible to do that without distorting something. This fact is presumably of no interest to flat-Earthers, but it should bother the rest of us. One can project Earth's curved surface onto a plane in an infinite number of different ways, and all of them distort the 'truth' in some way. The question is: what distortion are we willing to accept and what features do we want to preserve?

A map has four basic characteristics—area, direction, distance, and shape. Map projections differ in how they try to preserve these characteristics. In my primary school the map adorning the wall was a Mercator projection. A Mercator map is useful for sailors: it preserves directions, so any course of constant bearing appears as a straight-line segment on the map. The downside is that, while areas are accurate close to the equator, areas inflate as one heads towards the poles. For years, thanks to the Mercator map hanging on my school wall, I believed Antarctica is Earth's largest continent and Greenland rivals America in size.

S. Webb, *Around the World in 80 Ways*, https://doi.org/10.1007/978-3-031-02440-5_1

The American inventor Buckminster Fuller wanted a map of the world that better preserved the relative size of areas and the shapes of areas. He projected the world map onto the surface of an icosohedron, a three-dimensional object consisting of 20 equilateral triangles, and then cut the icosohedron so it could be laid out flat. His Dymaxion map (see Fig. 1, left), better known as the Fuller projection, butchers the world; good luck using it for navigation. But since one can cut the projection in many different ways, one can use it to illustrate themes that are difficult to show with other map projections.

Or consider the Werner projection (see Fig. 1, right), developed by the German mathematician Johannes Werner in the sixteenth century, which turns the world into a heart. I can think of few practical uses for a map based on the Werner projection. But the projection is not *wrong*. And isn't there something uplifting about seeing the world in the shape of a heart?

There is no 'right' map projection. From the infinite options on offer, you simply use the one that suits your task.

For the purposes of this book a map projection that seeks to retain the relative size of areas makes sense. I have chosen the Equal Earth projection, developed in 2018 by three geodetic engineers, Bojan Šavrič, Tom Patterson, and Bernhard Jenny. On an Equal Earth map Greenland no longer competes against America in size. Be aware, however, that this projection distorts and stretches shapes, directions, and distances north–south. I believe these maps can be *useful* and I try throughout to be *honest*, but the maps are not *truthful*—nor can they be. Lie number one.

2 Map Conventions

We bring our prejudices to the act of making and reading maps. For example, the prime meridian, the line of longitude 0°, runs through Greenwich. My maps are thus centred on a line that takes in London. But that particular definition of the prime meridian is an accident of history. In the latter half of the nineteenth century, much of the world's commerce relied on sea-charts that had a prime meridian centred on Greenwich; furthermore the USA used it as the basis of its national system of time zones. So formalising Greenwich as the centre of world time, something agreed at the International Meridian Conference in 1884, was convenient for the largest number of people *at that time*. Our maps do not have to look that way. The map that used to hang on my old school wall could just as reasonably have been centred on a line taking in New York, or Moscow, or Shanghai. In Fig. 2, the world map is centred on the 150th meridian east, a line that takes in Queensland and New South Wales in Australia.

Although this map might seem unfamiliar, it is recognisable. But another convention is that north should appear at the top of a map. We could just as easily adopt the opposite convention. This 'upside-down' view of the world (see Fig. 3) is *entirely* unfamiliar. It looks *wrong*. And yet it is as valid as our usual representation. If we wished we could represent our maps with east, say, at the top. In medieval Europe, before explorers adopted the magnetic compass, most mapmakers did just that: the rising of the sun provided an important bearing.

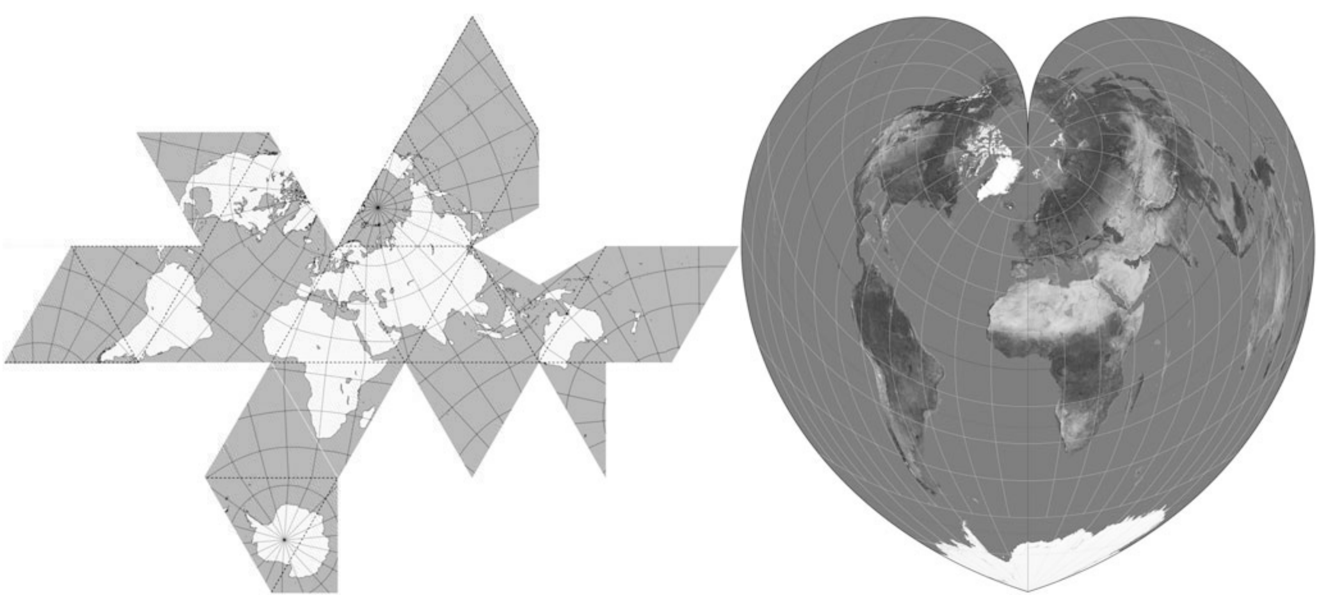

Fig. 1 Left: the Dymaxion, or Fuller, projection of the world centred on Europe and Africa. (Thomasee73, CC BY-SA 4.0). Right: the Werner projection of the world, with imagery derived from NASA's Blue Marble summer month composite. (Strebe, CC BY-SA 3.0)

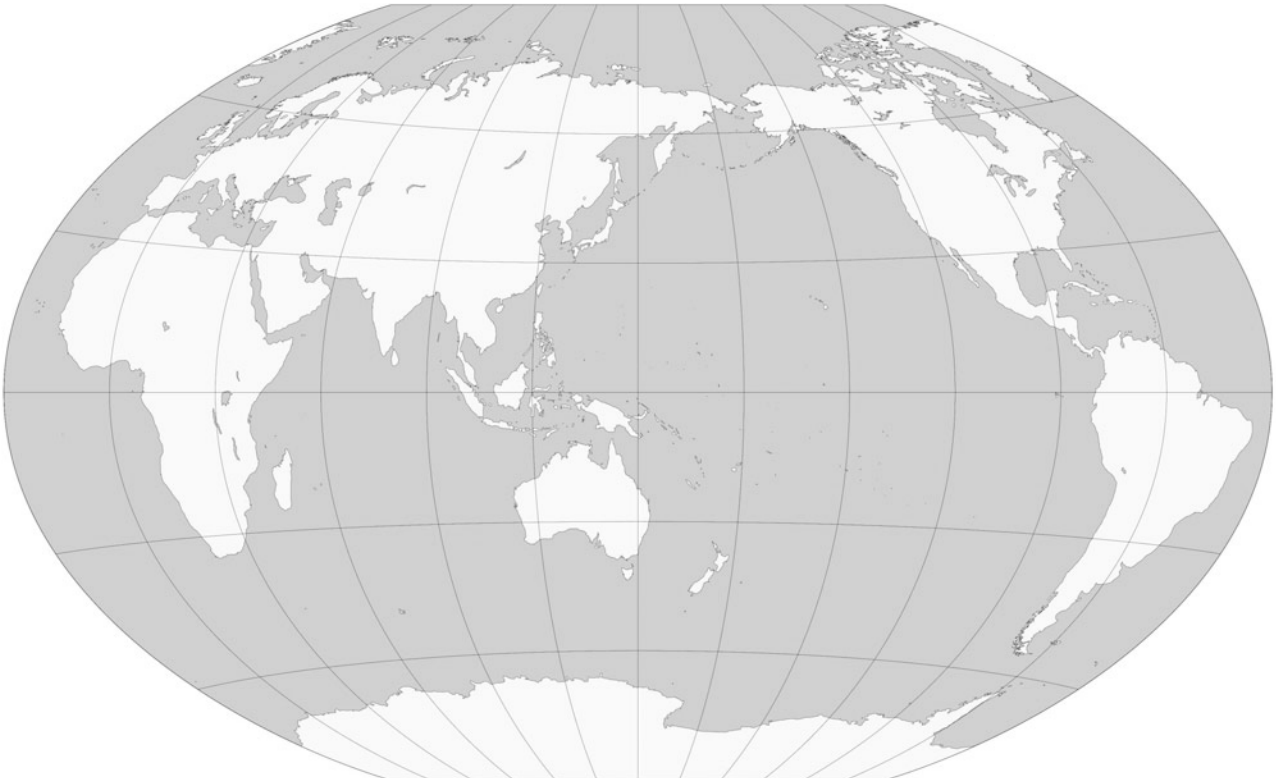

Fig. 2 The world map centred on the 150 meridian east, which takes in the Pacific Ocean. This map uses the Winkel III projection. (Milenioscuro, CC BY-SA 3.0)

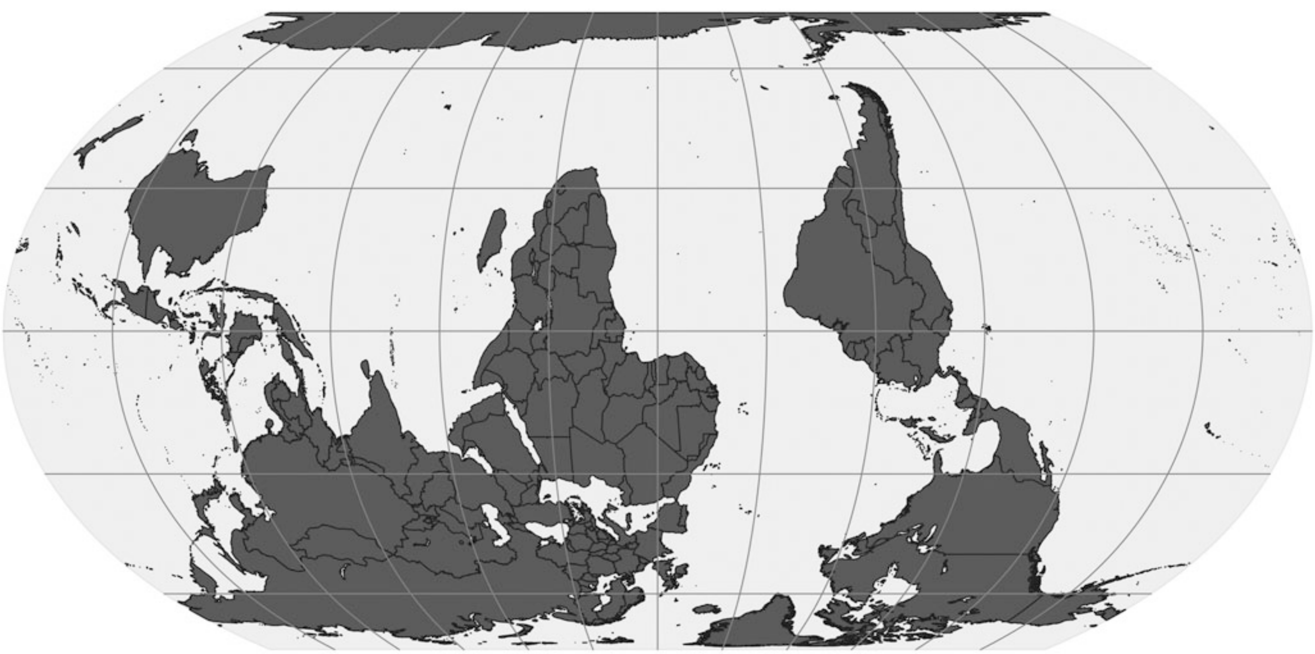

Fig. 3 The world map with south at the top. This map uses the Equal Earth projection. (Own work)

If we never stop to note that our maps enshrine arbitrary conventions then we will find it difficult to ever see the world anew. Lie number two.

3 What Is a Country Anyway?

My maps often compare countries according to some statistic. But in this context even the straightforward notion of 'country' is open to debate.

Take, for example, French Guiana. It has Brazil bordering to the east and south, Suriname bordering to the west. If you ask about the pattern of its flag or the colour of its passport then the answer is clear: the flag and passport of French Guiana are identical to those of France. And that is because French Guiana is a department and region of France; it just happens to be overseas. (This gives rise to an interesting trivia question: with which country does France have its longest border? Answer: Brazil.) On the other hand, if you happen to be interested in, say, the proportion of a country's land covered by forest then it makes no sense to argue that the figure for mainland France (which has forestry cover of about 30%) should be carried over to French Guiana (where forests cover about 99% of the land). So whether France and French Guiana should be considered identical depends upon context. It is another choice.

Or consider the UK. In many cases of interest we gain insight by comparing the figures for England, Scotland, Wales, and Northern Ireland. But differences between Northern Ireland and Wales, say, are hard to perceive on a world map. Besides, combined figures for the UK are often easier to obtain. So UK it is. Except there are occasions where the data *require* a separation of the Home Nations: England has won the World Cup, Scotland has not.

Or consider Hong Kong and Macau. Both these places are special administrative regions of China. There are occasions when it makes sense not to distinguish between China and Hong Kong/Macau (particularly when these places are too small to show up on a world map). But in some contexts it makes sense to draw attention to these special administrative regions. Many other special regions, dependent territories, and autonomous areas exist. We need to look at these on a case-by-case basis.

Even the term 'China' is ambiguous. There are two 'Chinas': the Republic of China, islands that lie about 800 km east of Hong Kong, which we usually call Taiwan; and the People's Republic of China, the most populous nation on Earth, which we usually refer to simply as 'China'. Behind that ambiguity sits decades of political controversy.

The control and ownership of a surprising number of places remain a matter of dispute. Kosovo, for example, is currently recognised as an independent state by 97 of the 193 UN member states (that's 50.3%)—so is Kosovo an independent state or not? Some states, recognised as independent by much of the international community, do not control their territory. At time of writing, for example, 138 UN member states have recognised the State of Palestine—but in practice most of the territory it claims is under the control of Israel. And some states receive little recognition by the international community despite being in effective control of all or part of a disputed territory. The Sahrawi Arab Democratic Republic, for example, controls about a fifth of Western Sahara (Morocco administers the rest) but at the time of writing only 40 UN member states recognise the claim. And then there are places such as Transnistria, a breakaway state from Moldova that is currently recognised only by other non-recognised states—in this case by Abkhazia (which most countries recognise as part of Georgia), Artsakh (which most countries recognise as part of Azerbaijan), and South Ossetia (another region that most countries recognise as part of Georgia).

My favourite example of the complexities involved in questions of territorial ownership is a piece of land in Kazakhstan, an ellipse of size 90 km in the east–west direction and 85 km north–south, at the centre of which is the Baikonur Cosmodrome. Although it lies well inside Kazakhstan it is, until a lease runs out in 2050, formally a part of Russia. I have chosen to leave this plot of land on my maps, not as a mountweazel but simply as a reminder that national boundaries are the product of often labyrinthine historical flows.

The maps in this book typically refer to about 250 nations, territories and dependencies—a number significantly greater than the number of member states of the United Nations. The comparisons I make are driven by the availability of data, never by political choice. But data availability (or the lack of it; the maps often include regions with 'No data') is always driven by *someone's* choice. Ignore that and you fall for lie number three.

4 Visualising the Data

Most of the maps in this book are choropleths. The choropleth, or 'colour map', has a long history. The mathematician Charles Dupin created the first choropleth in 1826. Dupin was interested in levels of literacy across France, and he shaded different regions of his country according to a colour scale running from white (high levels of literacy, the light colour symbolising 'enlightened' France) to black (low levels of literacy, the inky colour representing 'dark' France). Dupin's map had a practical use: at a glance it showed a clear divide between the north and south of the country, which hinted at a disparity of education between the two parts of

France. Once a problem becomes visible like this, politicians can start to develop policies to address the issue.

So choropleths can be valuable: they are an easy way of representing a large amount of data in a succinct, easy-to-digest, visually appealing way. And because a wide variety of data is collected at national level, it is a trivial task for a modern GIS system to spit out a choropleth illustrating that data. But one must be careful when interpreting a choropleth.

For example, suppose two countries possess the same value of the topic under discussion. Following on from Dupin, we might find they have the same rate of literacy, say. In this case we assign the same colour to both countries. But if one country covers a large area and the other is small, the larger country will appear more prominent: the eye can't help but attach greater significance to a larger block of colour.

And then there's the question of which colour to use. Quite apart from the cultural significance of colour, there is a practical point that has its roots in the technology of printing. Suppose you represent some quantity with a colour progression running from white (for small values of the quantity) to red (for large values). Well, printers will struggle to distinguish more than five or six shades on the choropleth. Furthermore, the human eye will struggle to discriminate one country from another if too many shades of red exist. (Personally, I struggle to distinguish between more than four shades of red.)

Then there are the choices one makes when classifying the data. Suppose we have some measurement, which runs from 1 to 100, on 250 countries. And suppose we want to put each country into one of five data groups, or 'bins', depending on its measurement. How should we proceed? We could ensure that we have 50 countries in each of the five bins. That seems reasonable. But, depending on the underlying distribution of the measurement, we might need bins of different sizes. We might end up with bins of size 1–2.7, 2.8–3.6, 3.7–22.1, 22.2–22.5, 22.5–100, say. Besides leading to bin sizes that are awkward to work with, it appears unfair to group a country with a measurement of 22.5 with one that measures 100. Or we could decide to have bins of equal size: 1–20; 21–40; 41–60; 61–80; 81–100. That, too, seems reasonable. But, depending on the underlying distribution of the measurement, we might end up with very different numbers of countries in the different bins: 145, 85, 17, 2, 1, say. That, too, appears unfair. Or we could try to minimise variation within each bin while also rounding the boundaries up to whole numbers for ease of reading. Those three approaches (and there are other options I haven't mentioned) would lead to maps with a different appearance in each case *but the underlying data would be the same*. Map appearance depends on our choice of how to classify the data. Remember, when you look at these maps, that one could create other valid visualisations. Otherwise you fall for lie number four.

5 Collecting and Interpreting the Data

The Covid-19 pandemic has given us all a lesson in the importance of data, but also in how difficult it can be to collect data in a consistent way and in how fiendishly hard it can be to interpret data.

The SARS-CoV-2 virus exists in order to find a host in which it can replicate and, once that host is no longer useful, move on to another. The illness caused by the virus is a by-product of that drive to replicate. So we have a situation in which the virus spreads through the population, and governments in turn respond by restricting the movement of people and by implementing public health measures (with varying degrees of rigour). Now consider how difficult it is to answer the most basic question about this virus: how deadly is the disease, Covid-19, that it causes?

Journalists often write about the case fatality rate. It *should* be easy to determine, right? You just need two numbers: the number of deaths from disease divided by the number of diagnosed cases over some period of time. But the world is not that simple.

Take the number of cases. Some people infected with the SARS-CoV-2 virus can be asymptomatic: they don't know they have had the infection. If a person is feeling fine then it's unlikely he or she would be tested for Covid-19; but that creates an uncertainty in our estimate for the number of cases. When people *are* tested some (admittedly small) proportion of tests return a wrong result because no diagnostic test is ever perfect: we get false positives and false negatives. Test results can even go missing, which again generates an uncertainty. Other factors can be at play, too: in some states of the USA, for example, politicians have asked scientists to 'massage' data in order to align with a particular view. And in many countries, the public health infrastructure is too fragile to collect reliable data. All this and more means case numbers, even confirmed case numbers, come with some uncertainty attached.

But we can at least count the number of deaths caused by Covid-19, right? Well, no. Public Health England provides a daily count of those who died within 28 days of testing positive for coronavirus, regardless of cause of death. On the other hand, particularly early on in the pandemic, the data did not include people who almost certainly died as a result of the virus but who were never tested. Even if everyone were tested with a perfectly sensitive test there would still be ambiguity: if someone dies of a stroke five weeks after testing positive, should that be counted as a Covid-19 death? Perhaps; perhaps not. It would be a judgement call for a doctor. (And I am ignoring here a death caused by a heart attack, say, that would have been prevented had the victim sought assistance—but did not do so because of worries of contracting Covid-19 in the health care system. In that case the SARS-CoV-2 virus would not be in the patient's body, but the patient died because the virus was in the body of the community.) Maybe the only way to measure the number of deaths is to count excess deaths over the most recent five-year average and attribute the excess to Covid-19. Even that approach is problematic because the response to the pandemic might have reduced deaths from other causes (fewer cars on the road, for example, means fewer road traffic fatalities).

Determining the impact of Covid-19 in a single jurisdiction is hard. The difficulty increases when one attempts to compare nations. Different countries have reported Covid-19 deaths in different ways: some countries adopted a broad definition of what constitutes a coronavirus fatality, others tried to cover up the number of deaths. Healthcare systems and social support differ between countries. Furthermore, one thing we know for certain about Covid-19 is that it kills older people more readily than it kills younger people, so comparisons should really take into account the age profile of a country.

This is not a primer on epidemiology, so I'll leave the discussion there. The point I want to make is that collecting data, even for something as straightforward as the number of people who die from a particular disease, is difficult. It can be even more difficult to divine the meaning of that data. Experts can—and often do—disagree over the meaning of data. (And bear in mind, as you read, that I often express an opinion on data in areas in which I have no particular expertise.)

This is not to say we should avoid attempts at comparing countries. In the case of Covid-19, for example, surely we *should* try to compare countries, even if we know such comparisons are imperfect, because we might be able to learn more about the virus. Nevertheless, we all need to question the collection and interpretation of data—or we fall prey to lie number five.

6 A Moment in Time

The maps in this book refer to the latest data to which I have access. But things change. Remember, as you read, that these maps represent snapshots of activity rather than immutable truths. Otherwise, lie number six will deceive you.

7 Questions, Questions

Lie number seven is subtle: even the questions we choose to ask can be a source of bias. In *Factfulness*, for example, one of Rosling's questions was: "In 1996, tigers, giant pandas, and black rhinos were all listed as endangered. How many of these three species are more critically endangered today?" We feel better when we learn the surprising answer: "None of them". Had Rosling asked instead about South China tigers, polar bears, and northern white rhinos—well, we wouldn't feel so good. It matters, which questions we choose to ask.

* * *

Seven deadly deceptions. You might conclude that anti-vaxxers, climate-change deniers, and flat-Earthers have a point. If maps are so misleading, why bother engaging with them? Indeed, there is a deeper question here: why should any of us attempt to understand the world in this way if the truth is so elusive?

I can think of a number of reasons why the attempt is worth it.

The first is it reminds us science is not a body of facts nor a collection of truths. Rather, it is a *process*. Scientists understand that science does not achieve complete certainty, but it *is* our best route to finding robust and reliable knowledge. We need to learn to live with uncertainty. If nothing else, accepting uncertainty is healthier than having complete faith in answers that are wrong. Particularly in these times, when many of our political leaders start with a gut feeling and then look for 'alternative facts' to justify their feeling, this is an important lesson.

Even where we have uncertain data we can still learn things, still draw conclusions, still compare one part of the world with another. We don't know, for example, the precise number of judicial executions taking place in China each year. But we *do* know the number is greater than that of any other country, and we *do* know many countries have abolished the death penalty. A map of capital punishment, illustrating the number of executions by country in a given year, cannot be completely accurate. But it can still provide insights.

There are other reasons for working with these maps. It's *interesting* to ponder the location of world heritage sites, say, or the world's tallest buildings. It's *fun* to reflect on national success at the World Cup Finals, say, or the distribution of medals at the Olympics. It's *important* to contemplate different peoples' standard of living, say, or their access to electricity.

Finally, if you disagree with the maps here—perhaps you think my colour scheme is misleading, or I put the data into too many or too few bins, or the size of the bins is confusing . . . well, you can create your own! I provide details of the sources I used so you can generate different versions of any of these maps or, as new data become available, create updated versions. The creation of world maps was once a task for professional cartographers. Nowadays, as I explain in an appendix, the widespread availability of open source geographical information systems mean anyone who can use a computer can generate a choropleth. Give it a try!

The World Itself

Astronomers now know of thousands of exoplanets, planets that orbit distant stars. None of those planets are much like ours. This might be because our techniques favour the detection of giant planets and planets that orbit close to their parent star. But it is also possible that rocky, tectonically active, water-rich planets in possession of a large satellite—in other words, life-bearing places like Earth—are rare. In this first chapter we take a look at our planetary home.

1 Earth's Poles (Map 1)

Just over 4.54 billion years ago, part of a giant molecular cloud began to collapse under the force of gravity. The collapse led to the formation of a protostar, the progenitor of our Sun, around which a disc of gas and dust began to orbit. Over time, dust grains would occasionally collide and clump together, eventually forming rocks as large as 200 m. In turn, these rocks collided and formed planetesimals as large as 10 km. The collisions continued for several million years, and led to the formation of four terrestrial planets (Mercury, Venus, Earth, Mars) in the inner solar system and four giant planets (Jupiter, Saturn, Uranus, Neptune) in the outer solar system, along with numerous smaller bodies. Collisions in the solar system never stopped (they continue still) and some astronomers hypothesise that in one such collision Theia, an object the size of Mars, struck the infant Earth. That impact formed the Moon. Ever since, Earth and Moon have danced together in their yearly lap of the Sun.

If the Theia hypothesis is correct, that Moon-forming collision caused Earth's rotational axis to tilt. We feel the results to this day. Our planet's axis of rotation, which meets the surface at the North and South Poles, is tilted at just over 23° with respect to the plane of Earth's orbit around the Sun. And because this tilt is fixed, regardless of where Earth is in its annual orbit, we experience seasons. Summer in the northern hemisphere sees the south polar regions freeze in darkness; when it is summer in the southern hemisphere the north polar regions experience the dark. In total, Earth's polar regions receive less solar radiation than its temperate and torrid zones—and so surface temperatures are lower at the poles. Over time, therefore, the polar regions develop ice caps.

As a child, I believed the ice cap at Antarctica must be Earth's biggest continent: the bottom of any map showed a sprawling land mass across Earth's full width. But Antarctica is only the fifth largest of the seven continents. Its apparent size is the result of the typical methods used to represent Earth's curvature on a flat page. As we saw in the Introduction, when you depict a curved surface on a flat map you must distort some element of the map. For most map makers, and for most map readers, it makes sense to distort the polar regions (since few people ever visit them) and maintain an accurate representation of the temperate and torrid zones (since people live there). But such distortion is a problem when a map illustrates some form of human activity because the question of what happens at the poles is irrelevant: it is not worth considering the incidence of road traffic accidents at the North Pole, say, or deaths from malaria at the South Pole. The Equal Earth projection used in this book compresses features in the north–south direction near the poles, so the polar regions do not dominate and the relative size of Antarctica is shown correctly. Since this small continent often appears in this book in grey (in order to denote 'No data') the reader's eye might slip over the area at the bottom of a map. There is then a danger that we *underestimate* the importance of the poles. The polar regions are essential to humanity's future. So, before we start focus on the rest of Earth, let's look at the polar regions. First, the North Pole.

Maps that 'look down' from above often show the North Pole as a point in the blue of the Arctic Ocean. The North Pole does indeed lay in the middle of the Arctic Ocean. The nearest land, Kaffeklubben (Coffee Club) Island, is about 690 km away just off the northern tip of Greenland. The nearest permanently inhabited settlement is 810 km away,

S. Webb, *Around the World in 80 Ways*, https://doi.org/10.1007/978-3-031-02440-5_2

1 Earth's poles

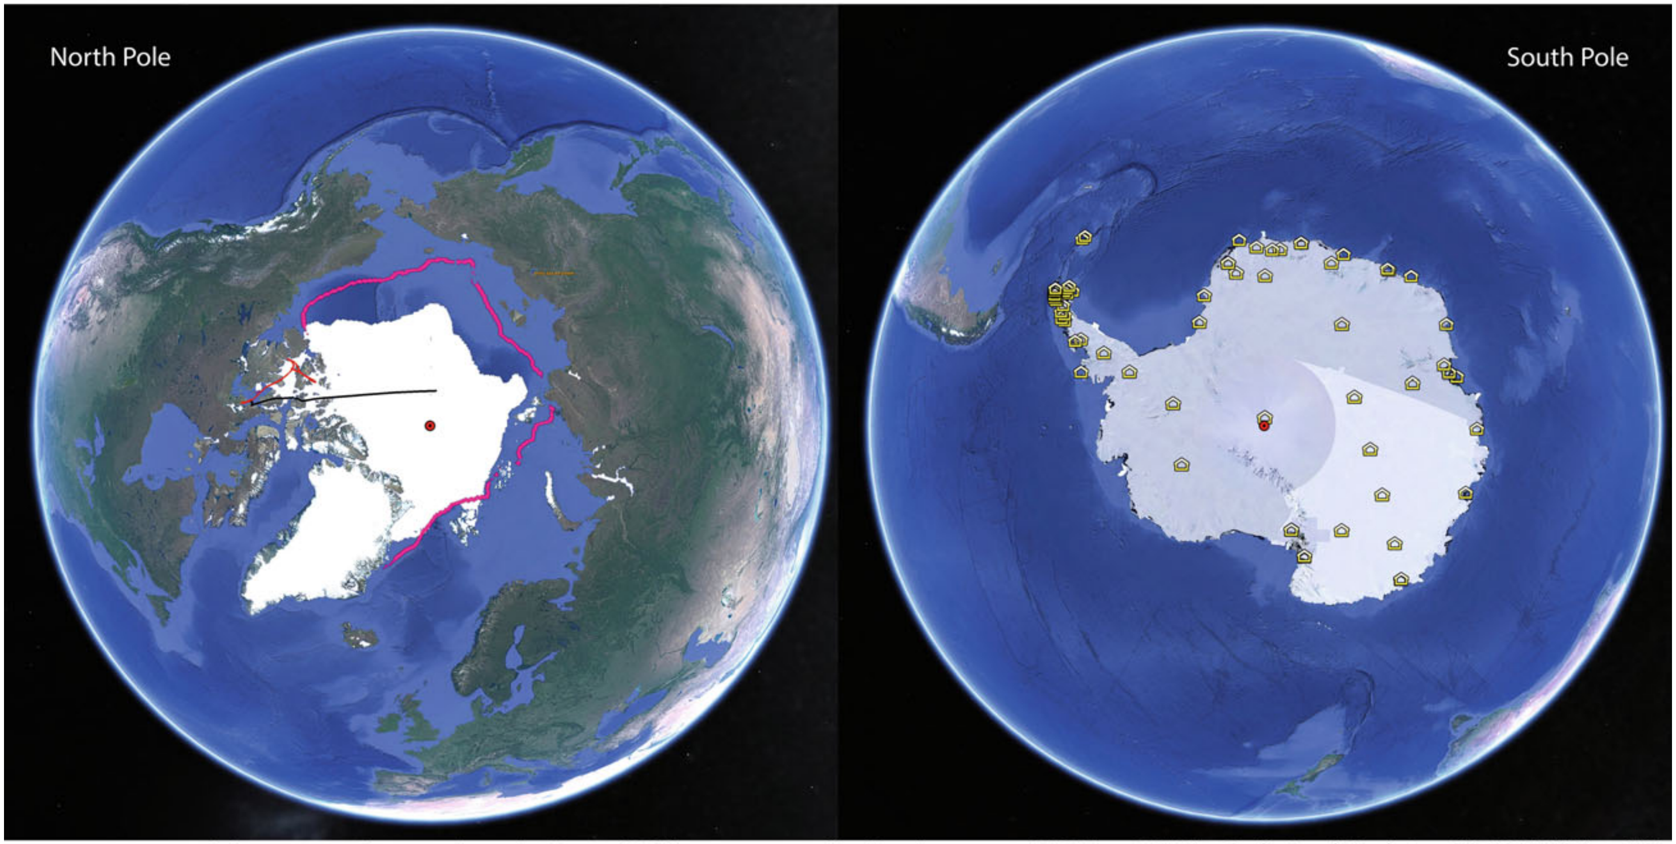

• Pole — Movement of magnetic pole (since 1831) — Movement of magnetic pole (1590 – 1831) — Median ice extent (1981 – 2010) Scientific base

Credit: Google Earth; Data SIO, NOAA, US Navy, NGA, GEBCO, NSIDC; Image Landsat/Copernicus; Image US Geological Survey; Image PGC/NASA

Map 1

Fig. 4 To reach the North Pole with a surface ship requires something like this Russian vessel, a nuclear-powered icebreaker called the '50 Years of Victory'. The ship can break through ice up to 2.8 m thick, whether sailing backwards or forwards. Nuclear-powered submarines are another option: they can sail beneath the ice, then punch a hole through it when they surface. (Christopher Michel, CC BY 3.0)

in the Canadian territory of Nanavut. But such maps mislead because they suggest that you or I could sail across the top of the world. With a normal boat, that journey is impossible (see Fig. 4). The high Arctic waters have an almost permanent cover of shifting sea ice. Ice presents such a challenge to exploration that humans first reached the North Pole as recently as 1926, and that was in an airship. (The claims of earlier expeditions, which used wooden sleds and dog teams to attempt to reach the North Pole, have not withstood scrutiny.)

If the polar icepack were easily navigable, a European ship bound for east Asia could shave about 4000 km from its journey. Huge economic benefits would flow. Small wonder, then, that explorers began searching for a so-called 'Northwest Passage' as early as the sixteenth century. As late as 1845, British explorers were still dying in the search: Sir John Franklin led an expedition of two ships, HMS *Erebus* and HMS *Terror*, neither of which returned. Today, global warming is causing the icepack to shrink and thin: the magenta line on the image shows the median ice extent for the month of October in the 30-year period 1981–2010 and the latest icepack, as can be clearly seen, is much smaller than the median. A reliable passage for commercial shipping might soon be available for a few months each year. Some economies will benefit but scientists view the creation of a Northwest Passage as being of small reward compared to the dangers associated with our climate emergency.

The North Pole is a fixed geographic feature. The North *Magnetic* Pole—the location at which a compass needle, if allowed to rotate freely, points vertically down—moves in response to happenings deep in Earth's core. In 1831 James Clark Ross became the first to reach the North Magnetic Pole. Since then the Pole has moved, and is currently drifting towards Siberia at a rate of 50 km/yr. (The black line on the image shows how it has drifted since 1831; the red line is an estimate of its position dating back to 1590.) Eventually, Earth's magnetic field will flip: north and south will swap, an event that last happened 780,000 years ago.

The South Pole, unlike the North, is not located on sea: Antarctica is a continental land mass. The ice here is about 2.7 km thick. That this ice sits on land is a cause for concern in a warming world. Arctic sea ice, when it melts, has no effect on global sea level (for the same reason an ice cube floating in a glass of water does not change the water level when it melts). A melting Antarctic would put *additional* water into the oceans and global sea level would rise. If *all* Antarctic ice were to melt, sea levels would rise by about 60 m. That extreme case will almost certainly not happen, but even a modest rise in sea levels would be enough to make some coastal cities uninhabitable (see Map 31).

Although Antarctica is a colder, harsher environment than anything found in the northern hemisphere, explorers reached the South Pole before the North. The great Norwegian polar explorer Roald Amundsen claimed the honour in 1911. Scott and four others got there 34 days after Amundsen, then perished on the return journey. (In 1914 the indefatigable explorer Ernest Shackleton attempted the first crossing of Antarctica via the South Pole. He did not make it, but his story of *Endurance* is astonishing.)

The ever-shifting nature of the Arctic icecap makes it difficult to construct permanent scientific bases there. Antarctica, in contrast, is home to about 75 permanent research bases and its population—scientists, technicians, and support staff—exceeds 1100 during winter and swells to more than 4200 during summer. One such research base, the Amundsen–Scott South Pole Station, is the southernmost structure on Earth: it lies directly at the Pole. Scientists at the Station experience just one 'day' and one 'night' during the year: for six months the Sun is above the horizon and for six months it is below. This period of prolonged darkness, combined with the dry atmosphere (Antarctica is one of the driest deserts on Earth), makes the Station a good place for astronomy. The Station is also home to some of the most esoteric experiments on the planet. IceCube, for example, hunts for neutrinos—mysterious particles which, although abundant in the universe, are reluctant to interact with the rest of the world. IceCube snares high-energy neutrinos from the depths of the cosmos. See Fig. 5.

2 Meteorites (Map 2)

Our knowledge of how Earth came into being and of how our solar system evolved comes in large part from a study of meteorites. Certain types of meteorite that we find here on Earth developed in the same dust that gave rise to the Sun and planets, in the same protoplanetary disk from which Earth formed. The thought inspires awe: a pristine meteoroid can lap the Sun without incident for four billion years until, one day, its orbit intersects that of our planet. The meteoroid blazes through the atmosphere (at which point we call it a meteor) and then a snapshot of our solar system's pre-history falls to Earth as a meteorite. You can imagine how fortunate I felt, then, when I was given a behind-the-scenes tour of the

Fig. 5 The IceCube Laboratory, at the Amundsen–Scott South Pole Station, underneath a night sky in winter. This building hosts computers that are used to collect and process data from the experiment's detectors. The detectors themselves are buried deep underneath Antarctic ice, and look 'up' through Earth's bulk in search of neutrinos. (John Hardin, CC BY 4.0)

2 Meteorites

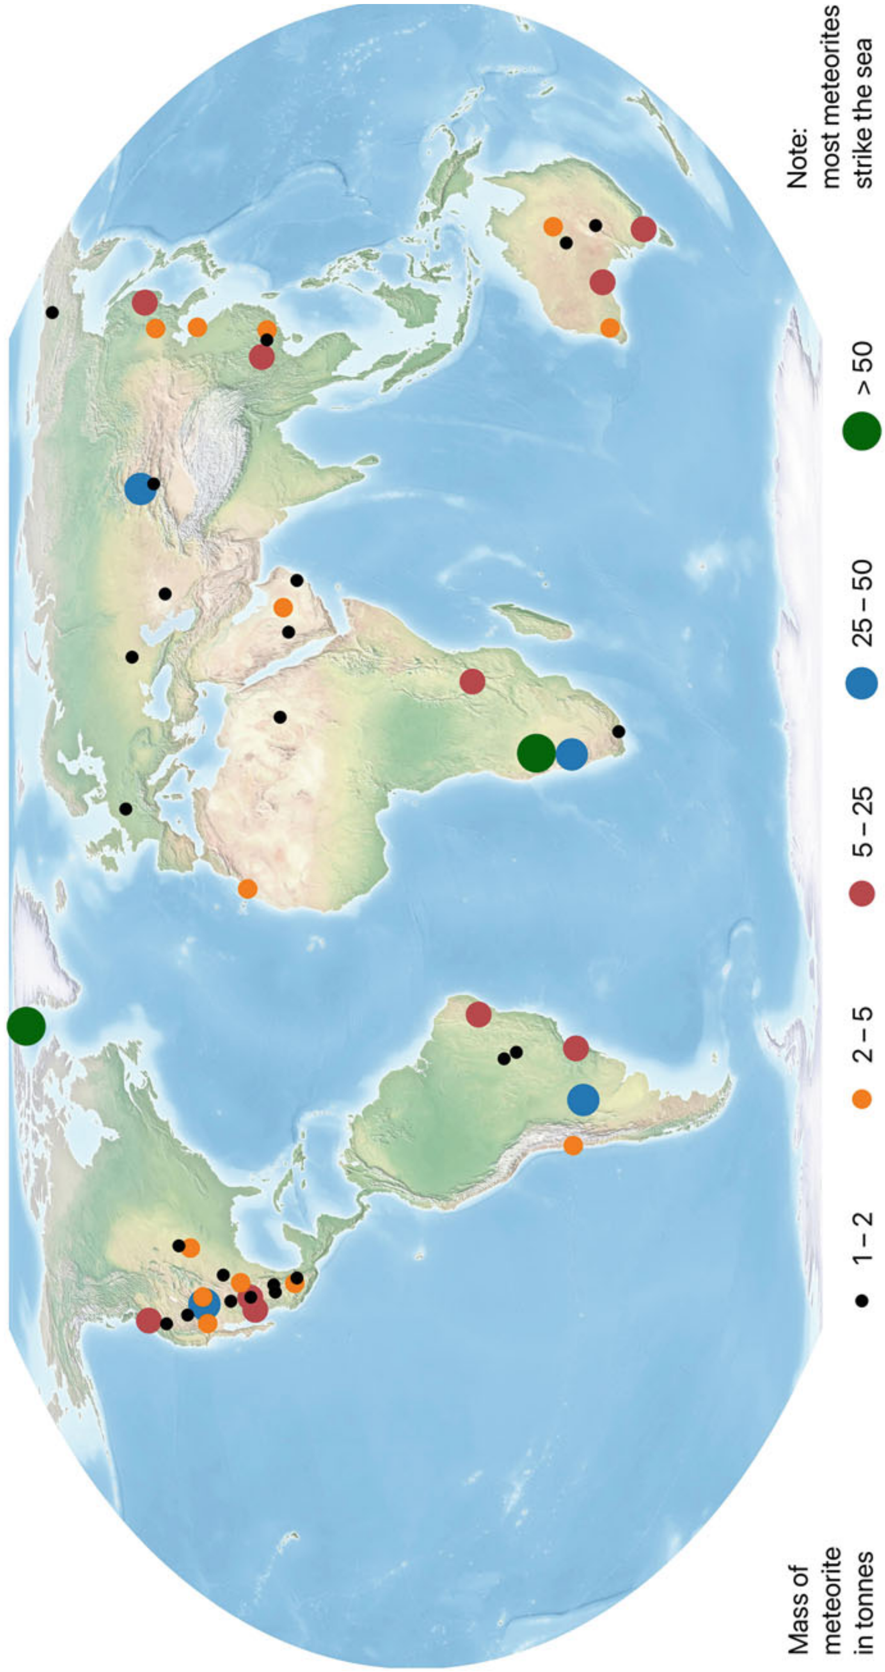

Map 2

Fig. 6 Left: the South Corston fragment of the Strathmore meteorite. The Corston fragment, which weighs just over 1 kg, landed in a field about 50 m from a farm house. The impact made a hole 15 cm deep. Right: The Easter Essendy fragment. Weighing slightly over 10 kg, this was the largest of the four meteorite fragments to be recovered. The impact made a hole 45 cm deep. The other two fragments are called Carsie and Keithick. The Carsie fragment is interesting in that an eyewitness, Mrs. Grace Walsh, saw it hit the ground. The Keithick fragment became the most famous of the four when it drilled a hole through the roof of a house occupied by the Hill family. (Geni, CC BY 4.0)

National Museums Scotland—the highlight of which was the chance to hold the largest of the four pieces of the Strathmore meteorite.

On 3 December 1917, just after 1 pm, a fireball flashed across the clear winter skies of southern Scotland. Within seconds, people in the central region of Strathmore heard an explosion as four objects crashed to the ground. Contemporary accounts describe a range of responses to the fireball. One witness said he "heard it fizzin'"; a more pompous report declared the "community [was] thrown into a state of consternation". The four objects were soon recovered, and represent parts of the largest meteorite fall ever recorded in Scotland. The piece I held weighs 10.2 kg. Mineralogists have studied the meteorite, and the prevailing hypothesis is it was once part of an ancient stony asteroid that was hit by a larger body some 468 million years ago; the impact smashed the asteroid into tiny pieces, with orbits that cross Earth's. The Strathmore meteorite is nothing special to look at, just a dull lump of rock; see Fig. 6. Nevertheless, it's quite a feeling to hold an object you know spent 468 million years in the cold of space before smashing into the cold of Scotland.

It is not every day you get to hold a space rock, but meteorites are not uncommon. As Earth follows its orbit it encounters lots of material. Most of this stuff is small: grains of dust and pea-sized bits of rock. This type of material burns up in the atmosphere—we see a "shooting star"—and never makes it to the ground. The larger bits of rock, though, can survive the shock of a high-speed journey through the atmosphere. We have two possible ways to estimate how many of them hit Earth each year. First, we can use an all-sky camera to monitor the meteorites that fall in a particular area and then (assuming all areas receive the same number of meteorites) extrapolate the total for Earth. Second, we can go to a place where there is little vegetation or erosion (a desert does nicely) and then count the number of meteorites lying around. One can estimate how long a meteorite has been there by measuring the amount of weathering. From these data one can determine how many meteorites fall in that place each year and then extrapolate the total for Earth. Although both methods generate estimates that come with large uncertainties attached to them, it seems reasonable to conclude that Earth receives between 40,000–80,000 tonnes of meteoritic mass every year. One estimate suggests between 18,000–84,000 meteorites with mass greater than 10 g hit Earth each year. The majority of meteorites fall into the ocean rather than onto land, for the simple reason that there is more than twice as much ocean cover than land cover. Ocean-falling meteorites are typically lost from view. Some

Fig. 7 A visitor centre has been created around the Hoba meteorite in Groonfontein, Namibia, which is classed as a national monument. Hoba is the largest known single-piece meteorite. (Sergio Conti, CC BY 2.0)

meteorites, though, hit land and can be recovered. The Strathmore meteorite was one such example.

The Meteoritical Society maintains a database of meteorite finds. At the time of writing the database contains 61,865 records. The location of a large fraction of these meteorites is Antarctica. This is not because Antarctica is particularly large (as we saw in Map 1, it isn't) nor is it because it attracts meteorites (it doesn't). Rather, scientists have determined that Antarctica is the ideal place to go meteorite hunting—so some of the scientific bases shown in Map 1 are home to geologists who go looking for space rocks. The conditions are ideal: the dry, cold climate preserves the meteorites that land there; the ice sheets create a "conveyor belt" that concentrates them in areas where they can be easily collected; and high-speed winds scour the ice surface and expose buried rocks.

In addition to the location field in the Meteoritical Society database, each record contains a wealth of other information. Each record, for example, contains an estimate of the mass of the meteorite. The overwhelming majority of meteorites in the database are low-mass objects, of a few grammes or less. This is exactly what one would expect: the solar system contains many more low-mass objects than high-mass objects. But the database also contains details of substantial meteorites, objects much bigger than the Strathmore meteorite. The database has records of 53 objects with a mass greater than 1 tonne.

The most massive object in the database, weighing in at 60 tonnes, is the Hoba meteorite. Jacobus Hermanus Brits, the owner of a farm in the Otjozondjupa region of Namibia, discovered it in 1920, quite by chance, when he heard a loud metallic screech as he ploughed one of his fields. Brits had the obstruction excavated and uncovered a tablet of metal 2.7 m by 2.7 m by 0.9 m in size. Workers estimated its mass to be 66 tonnes; the combined effects of weathering, sampling, and vandalism have reduced its weight to its present 60 tonnes. Scientists soon identified the tablet as an iron meteorite that fell within the past 80,000 years. Given its weight and size this object—the most massive naturally occurring piece of iron anywhere on Earth's surface—has never moved in all its time here.

The Hoba meteorite is now a tourist attraction (see Fig. 7), receiving thousands of visitors each year, and the commonest question those tourists ask is: why is there no crater? Surely a 66 tonne lump of metal barrelling into Earth's surface would leave a sizeable hole? The answer is simple: our planet's atmosphere, acting on the object's unusually flat shape, slowed the meteorite's descent. When the iron slab crashed